

An ℓ_∞ Eigenvector Perturbation Bound and Its Application to Robust Covariance Estimation

Jianqing Fan
Weichen Wang
Yiqiao Zhong

*Department of Operations Research and Financial Engineering
Princeton University
Princeton, NJ 08544, USA*

JQFAN@PRINCETON.EDU
NICKWEICHWANG@GMAIL.COM
YIQIAOZ@PRINCETON.EDU

Editor: Nicolai Meinshausen

Abstract

In statistics and machine learning, we are interested in the eigenvectors (or singular vectors) of certain matrices (e.g. covariance matrices, data matrices, etc). However, those matrices are usually perturbed by noises or statistical errors, either from random sampling or structural patterns. The Davis-Kahan $\sin \theta$ theorem is often used to bound the difference between the eigenvectors of a matrix A and those of a perturbed matrix $\tilde{A} = A + E$, in terms of ℓ_2 norm. In this paper, we prove that when A is a low-rank and incoherent matrix, the ℓ_∞ norm perturbation bound of singular vectors (or eigenvectors in the symmetric case) is smaller by a factor of $\sqrt{d_1}$ or $\sqrt{d_2}$ for left and right vectors, where d_1 and d_2 are the matrix dimensions. The power of this new perturbation result is shown in robust covariance estimation, particularly when random variables have heavy tails. There, we propose new robust covariance estimators and establish their asymptotic properties using the newly developed perturbation bound. Our theoretical results are verified through extensive numerical experiments.

Keywords: Matrix perturbation theory, Incoherence, Low-rank matrices, Sparsity, Approximate factor model

1. Introduction

The perturbation of matrix eigenvectors (or singular vectors) has been well studied in matrix perturbation theory (Wedin, 1972; Stewart, 1990). The best known result of eigenvector perturbation is the classic Davis-Kahan theorem (Davis and Kahan, 1970). It originally emerged as a powerful tool in numerical analysis, but soon found its widespread use in other fields, such as statistics and machine learning. Its popularity continues to surge in recent years, which is largely attributed to the omnipresent data analysis, where it is a common practice, for example, to employ PCA (Jolliffe, 2002) for dimension reduction, feature extraction, and data visualization.

The eigenvectors of matrices are closely related to the underlying structure in a variety of problems. For instance, principal components often capture most information of data and extract the latent factors that drive the correlation structure of the data (Bartholomew et al., 2011); in classical multidimensional scaling (MDS), the centered squared distance matrix encodes the coordinates of data points embedded in a low dimensional subspace

(Borg and Groenen, 2005); and in clustering and network analysis, spectral algorithms are used to reveal clusters and community structure (Ng et al., 2002; Rohe et al., 2011). In those problems, the low dimensional structure that we want to recover, is often ‘perturbed’ by observation uncertainty or statistical errors. Besides, there might be a sparse pattern corrupting the low dimensional structure, as in approximate factor models (Chamberlain et al., 1983; Stock and Watson, 2002) and robust PCA (De La Torre and Black, 2003; Candès et al., 2011).

A general way to study these problems is to consider

$$\tilde{A} = A + S + N, \tag{1}$$

where A is a low rank matrix, S is a sparse matrix, and N is a random matrix regarded as random noise or estimation error, all of which have the same size $d_1 \times d_2$. Usually A is regarded as the ‘signal’ matrix we are primarily interested in, S is some sparse contamination whose effect we want to separate from A , and N is the noise (or estimation error in covariance matrix estimation).

The decomposition (1) forms the core of a flourishing literature on robust PCA (Chandrasekaran et al., 2011; Candès et al., 2011), structured covariance estimation (Fan et al., 2008, 2013), multivariate regression (Yuan et al., 2007) and so on. Among these works, a standard condition on A is matrix incoherence (Candès et al., 2011). Let the singular value decomposition be

$$A = U\Sigma V^T = \sum_{i=1}^r \sigma_i u_i v_i^T, \tag{2}$$

where r is the rank of A , the singular values are $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$, and the matrices $U = [u_1, \dots, u_r] \in \mathbb{R}^{d_1 \times r}$, $V = [v_1, \dots, v_r] \in \mathbb{R}^{d_2 \times r}$ consist of the singular vectors. The coherences $\mu(U), \mu(V)$ are defined as

$$\mu(U) = \frac{d_1}{r} \max_i \sum_{j=1}^r U_{ij}^2, \quad \mu(V) = \frac{d_2}{r} \max_i \sum_{j=1}^r V_{ij}^2, \tag{3}$$

where U_{ij} and V_{ij} are the (i, j) entry of U and V , respectively. It is usually expected that $\mu_0 := \max\{\mu(U), \mu(V)\}$ is not too large, which means the singular vectors u_i and v_i are incoherent with the standard basis. This incoherence condition (3) is necessary for us to separate the sparse component S from the low rank component A ; otherwise A and S are not identifiable. Note that we do not need any incoherence condition on UV^T , which is different from Candès et al. (2011) and is arguably unnecessary (Chen, 2015).

Now we denote the eigengap $\gamma_0 = \min\{\sigma_i - \sigma_{i+1} : i = 1, \dots, r\}$ where $\sigma_{r+1} := 0$ for notational convenience. Also we let $E = S + N$, and view it as a perturbation matrix to the matrix A in (1). To quantify the perturbation, we define a rescaled measure as $\tau_0 := \max\{\sqrt{d_2/d_1}\|E\|_1, \sqrt{d_1/d_2}\|E\|_\infty\}$, where

$$\|E\|_1 = \max_j \sum_{i=1}^{d_1} |E_{ij}|, \quad \|E\|_\infty = \max_i \sum_{j=1}^{d_2} |E_{ij}|, \tag{4}$$

which are commonly used norms gauging sparsity (Bickel and Levina, 2008). They are also operator norms in suitable spaces (see Section 2). The rescaled norms $\sqrt{d_2/d_1}\|E\|_1$ and $\sqrt{d_1/d_2}\|E\|_\infty$ are comparable to the spectral norm $\|E\|_2 := \max_{\|u\|_2=1} \|Eu\|_2$ in many cases; for example, when E is an all-one matrix, $\sqrt{d_2/d_1}\|E\|_1 = \sqrt{d_1/d_2}\|E\|_\infty = \|E\|_2$.

Suppose the perturbed matrix \tilde{A} also has the singular value decomposition:

$$\tilde{A} = \sum_{i=1}^{d_1 \wedge d_2} \tilde{\sigma}_i \tilde{u}_i \tilde{v}_i^T, \quad (5)$$

where $\tilde{\sigma}_i$ are nonnegative and in the decreasing order, and the notation \wedge means $a \wedge b = \min\{a, b\}$. Denote $\tilde{U} = [\tilde{u}_1, \dots, \tilde{u}_r]$, $\tilde{V} = [\tilde{v}_1, \dots, \tilde{v}_r]$, which are counterparts of top r singular vectors of A .

We will present an ℓ_∞ matrix perturbation result that bounds $\|\tilde{u}_i - u_i\|_\infty$ and $\|\tilde{v}_i - v_i\|_\infty$ up to sign.¹ This result is different from ℓ_2 bounds, Frobenius-norm bounds, or the $\sin \Theta$ bounds, as the ℓ_∞ norm is not orthogonal invariant. The following theorem is a simplified version of our main results in Section 2.

Theorem 1 *Let $\tilde{A} = A + E$ and suppose the singular decomposition in (2) and (5). Denote $\gamma_0 = \min\{\sigma_i - \sigma_{i+1} : i = 1, \dots, r\}$ where $\sigma_{r+1} := 0$. Then there exists $C(r, \mu_0) = O(r^4 \mu_0^2)$ such that, if $\gamma_0 > C(r, \mu_0)\tau_0$, up to sign,*

$$\max_{1 \leq i \leq r} \|\tilde{u}_i - u_i\|_\infty \leq C(r, \mu_0) \frac{\tau_0}{\gamma_0 \sqrt{d_1}} \quad \text{and} \quad \max_{1 \leq i \leq r} \|\tilde{v}_i - v_i\|_\infty \leq C(r, \mu_0) \frac{\tau_0}{\gamma_0 \sqrt{d_2}}, \quad (6)$$

where $\mu_0 = \max\{\mu(U), \mu(V)\}$ is the coherence given after (3) and $\tau_0 := \max\{\sqrt{d_2/d_1}\|E\|_1, \sqrt{d_1/d_2}\|E\|_\infty\}$.

When A is symmetric, $\tau_0 = \|E\|_\infty$ and the condition on the eigengap is simply $\gamma_0 > C(r, \mu_0)\|E\|_\infty$. The incoherence condition naturally holds for a variety of applications, where the low rank structure emerges as a consequence of a few factors driving the data matrix. For example, in Fama-French factor models, the excess returns in a stock market are driven by a few common factors (Fama and French, 1993); in collaborative filtering, the ratings of users are mostly determined by a few common preferences (Rennie and Srebro, 2005); in video surveillance, A is associated with the stationary background across image frames (Oliver et al., 2000). We will have a detailed discussion in Section 2.3.

The eigenvector perturbation was studied by Davis and Kahan (1970), where Hermitian matrices were considered, and the results were extended by Wedin (1972) to general rectangular matrices. To compare our result with these classical results, assuming $\gamma_0 \geq 2\|E\|_2$, a combination of Wedin's theorem and Mirsky's inequality (Mirsky, 1960) (the counterpart of Weyl's inequality for singular values) implies

$$\max_{1 \leq k \leq r} \{\|v_k - \tilde{v}_k\|_2 \vee \|u_k - \tilde{u}_k\|_2\} \leq \frac{2\sqrt{2}\|E\|_2}{\gamma_0}. \quad (7)$$

where $a \vee b := \max\{a, b\}$.

1. 'Up to sign' means we can appropriately choose an eigenvector or singular vector u to be either u or $-u$ in the bounds. This is because eigenvectors and singular vectors are not unique.

Yu et al. (2015) also proved a similar bound as in (7), and that result is more convenient to use. If we are interested in the ℓ_∞ bound but naively use the trivial inequality $\|x\|_\infty \leq \|x\|_2$, we would have a suboptimal bound $O(\|E\|_2/\gamma_0)$ in many situations, especially in cases where $\|E\|_2$ is comparable to $\|E\|_\infty$. Compared with (6), the bound is worse by a factor of $\sqrt{d_1}$ for u_k and $\sqrt{d_2}$ for v_k . In other words, converting the ℓ_2 bound from Davis-Kahan theorem directly to the ℓ_∞ bound does not give a sharp result in general, in the presence of incoherent and low rank structure of A . Actually, assuming $\|E\|_2$ is comparable with $\|E\|_\infty$, for square matrices, our ℓ_∞ bound (6) matches the ℓ_2 bound (7) in terms of dimensions d_1 and d_2 . This is because $\|x\|_2 \leq \sqrt{n} \|x\|_\infty$ for any $x \in \mathbb{R}^n$, so we expect to gain a factor $\sqrt{d_1}$ or $\sqrt{d_2}$ in those ℓ_∞ bounds. The intuition is that, when A has an incoherent and low-rank structure, the perturbation of singular vectors is not concentrated on a few coordinates.

To understand how matrix incoherence helps, let us consider a simple example with no matrix incoherence, in which (7) is tight up to a constant. Let $A = d(1, 0, \dots, 0)^T(1, 0, \dots, 0)$ be a d -dimensional square matrix, and $E = d(0, 1/2, 0, \dots, 0)^T(1, 0, \dots, 0)$ of the same size. It is apparent that $\gamma_0 = d, \tau_0 = d/2$, and that $v_1 = (1, 0, \dots, 0)^T, \tilde{v}_1 = (2/\sqrt{5}, 1/\sqrt{5}, 0, \dots, 0)^T$ up to sign. Clearly, the perturbation $\|\tilde{v}_1 - v_1\|_\infty$ is not vanishing as d tends to infinity in this example, and thus, there is no hope of a strong upper bound as in (6) without the incoherence condition.

The reason that the factor $\sqrt{d_1}$ or $\sqrt{d_2}$ comes into play in (7) is that, the error $u_k - \tilde{u}_k$ (and similarly for v_k) spreads out evenly in d_1 (or d_2) coordinates, so that the ℓ_∞ error is far smaller than the ℓ_2 error. This, of course, hinges on the incoherence condition, which in essence precludes eigenvectors from aligning with any coordinate.

Our result is very different from the sparse PCA literature, in which it is usually assumed that the leading eigenvectors are sparse. In Johnstone and Lu (2009), it is proved that there is a threshold for p/n (the ratio between the dimension and the sample size), above which PCA performs poorly, in the sense that $\langle \tilde{v}_1, v_1 \rangle$ is approximately 0. This means that the principal component computed from the sample covariance matrix reveals nothing about the true eigenvector. In order to mitigate this issue, in Johnstone and Lu (2009) and subsequent papers (Vu and Lei, 2013; Ma, 2013; Berthet and Rigollet, 2013), sparse leading eigenvectors are assumed. However, our result is different, in the sense that we require a stronger eigengap condition $\gamma_0 > C(r, \mu_0)\|E\|_\infty$ (i.e. stronger signal), whereas in Johnstone and Lu (2009), the eigengap of the leading eigenvectors is a constant times $\|E\|_2$. This explains why it is plausible to have a strong uniform eigenvector perturbation bound in this paper.

We will illustrate the power of this perturbation result using robust covariance estimation as one application. In the approximate factor model, the true covariance matrix admits a decomposition into a low rank part A and a sparse part S . Such models have been widely applied in finance, economics, genomics, and health to explore correlation structure.

However, in many studies, especially financial and genomics applications, it is well known that the observations exhibit heavy tails (Gupta et al., 2013). This problem can be resolved with the aid of recent results of concentration bounds in robust estimation (Catoni, 2012; Hsu and Sabato, 2014; Fan et al., 2017a), which produces the estimation error N in (1) with an optimal entry-wise bound. It nicely fits our perturbation results, and we can tackle it easily by following the ideas in Fan et al. (2013).

Here are a few notations in this paper. For a generic d_1 by d_2 matrix, the matrix max-norm is denoted as $\|M\|_{\max} = \max_{i,j} |M_{ij}|$. The matrix operator norm induced by vector ℓ_p norm is $\|M\|_p = \sup_{\|x\|_p=1} \|Mx\|_p$ for $1 \leq p \leq \infty$. In particular, $\|M\|_1 = \max_j \sum_{i=1}^{d_1} |M_{ij}|$; $\|M\|_{\infty} = \max_i \sum_{j=1}^{d_2} |M_{ij}|$; and $\|\cdot\|$ denotes the spectral norm, or the matrix 2-norm $\|\cdot\|_2$ for simplicity. We use $\sigma_j(M)$ to denote the j^{th} largest singular value. For a symmetric matrix M , denote $\lambda_j(M)$ as its j^{th} largest eigenvalue. If M is a positive definite matrix, then $M^{1/2}$ is the square root of M , and $M^{-1/2}$ is the square root of M^{-1} .

2. The ℓ_{∞} perturbation results

2.1 Symmetric matrices

First, we study ℓ_{∞} perturbation for symmetric matrices (so $d_1 = d_2$). The approach we study symmetric matrices will be useful to analyze asymmetric matrices, because we can always augment a $d_1 \times d_2$ rectangular matrix into a $(d_1 + d_2) \times (d_1 + d_2)$ symmetric matrix, and transfer the study of singular vectors to the eigenvectors of the augmented matrix. This augmentation is called *Hermitian dilation*. (Tropp, 2012; Paulsen, 2002)

Suppose that $A \in \mathbb{R}^{d \times d}$ is an d -dimensional symmetric matrix. The perturbation matrix $E \in \mathbb{R}^{d \times d}$ is also d -dimensional and symmetric. Let the perturbed matrix be $\tilde{A} := A + E$. Suppose the spectral decomposition of A is given by

$$A = [V, V_{\perp}] \begin{pmatrix} \Lambda_1 & 0 \\ 0 & \Lambda_2 \end{pmatrix} [V, V_{\perp}]^T = \sum_{i=1}^r \lambda_i v_i v_i^T + \sum_{i>r} \lambda_i v_i v_i^T, \quad (8)$$

where $\Lambda_1 = \text{diag}\{\lambda_1, \dots, \lambda_r\}$, $\Lambda_2 = \text{diag}\{\lambda_{r+1}, \dots, \lambda_n\}$, and where $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_n|$. Note the best rank- r approximation of A under the Frobenius norm is $A_r := \sum_{i \leq r} \lambda_i v_i v_i^T$.² Analogously, the spectral decomposition of \tilde{A} is

$$\tilde{A} = \sum_{i=1}^r \tilde{\lambda}_i \tilde{v}_i \tilde{v}_i^T + \sum_{i>r} \tilde{\lambda}_i \tilde{v}_i \tilde{v}_i^T,$$

and write $\tilde{V} = [\tilde{v}_1, \dots, \tilde{v}_r] \in \mathbb{R}^{d \times r}$, where $|\tilde{\lambda}_1| \geq |\tilde{\lambda}_2| \geq \dots \geq |\tilde{\lambda}_n|$. Recall that $\|E\|_{\infty}$ given by (4) is an operator norm in the ℓ_{∞} space, in the sense that $\|E\|_{\infty} = \sup_{\|u\|_{\infty} \leq 1} \|Eu\|_{\infty}$. This norm is the natural counterpart of the spectral norm $\|E\|_2 := \sup_{\|u\|_2 \leq 1} \|\tilde{E}u\|_2$.

We will use notations $O(\cdot)$ and $\Omega(\cdot)$ to hide absolute constants.³ The next theorem bounds the perturbation of eigenspaces up to a rotation.

Theorem 2 *Suppose $|\lambda_r| - \varepsilon = \Omega(r^3 \mu^2 \|E\|_{\infty})$, where $\varepsilon = \|A - A_r\|_{\infty}$, which is the approximation error measured under the matrix ∞ -norm and $\mu = \mu(V)$ is the coherence of V*

2. This is a consequence of Wielandt-Hoffman theorem.

3. We write $a = O(b)$ if there is a constant $C > 0$ such that $a < Cb$; and $a = \Omega(b)$ if there is a constant $C' > 0$ such that $a > C'b$.

defined in (3). Then, there exists an orthogonal matrix $R \in \mathbb{R}^{r \times r}$ such that

$$\|\tilde{V}R - V\|_{\max} = O\left(\frac{r^{5/2}\mu^2\|E\|_{\infty}}{(|\lambda_r| - \varepsilon)\sqrt{d}}\right).$$

This result involves an unspecified rotation R , due to the possible presence of multiplicity of eigenvalues. In the case where $\lambda_1 = \dots = \lambda_r > 0$, the individual eigenvectors of V are only identifiable up to rotation. However, assuming an eigengap (similar to Davis-Kahan theorem), we are able to bound the perturbation of individual eigenvectors (up to sign).

Theorem 3 *Assume the conditions in Theorem 2. In addition, suppose δ satisfies $\delta > \|E\|_2$, and for any $i \in [r]$, the interval $[\lambda_i - \delta, \lambda_i + \delta]$ does not contain any eigenvalues of A other than λ_i . Then, up to sign,*

$$\max_{i \in [r]} \|\tilde{v}_i - v_i\|_{\infty} = \|\tilde{V} - V\|_{\max} = O\left(\frac{r^4\mu^2\|E\|_{\infty}}{(|\lambda_r| - \varepsilon)\sqrt{d}} + \frac{r^{3/2}\mu^{1/2}\|E\|_2}{\delta\sqrt{d}}\right).$$

To understand the above two theorems, let us consider the case where A has exactly rank r (i.e., $\varepsilon = 0$), and r and μ are not large (say, bounded by a constant). Theorem 2 gives a uniform entrywise bound $O(\|E\|_{\infty}/|\lambda_r|\sqrt{d})$ on the eigenvector perturbation. As a comparison, the Davis-Kahan $\sin \Theta$ theorem (Davis and Kahan, 1970) gives a bound $O(\|E\|_2/|\lambda_r|)$ on $\|\tilde{V}R - V\|_2$ with suitably chosen rotation R .⁴ This is an order of \sqrt{d} larger than the bound given in Theorem 2 when $\|E\|_{\infty}$ is of the same order as $\|E\|_2$. Thus, in scenarios where $\|E\|_2$ is comparable to $\|E\|_{\infty}$, this is a refinement of Davis-Kahan theorem, because the max-norm bound in Theorem 2 provides an entry-wise control of perturbation. Although $\|E\|_{\infty} \geq \|E\|_2$,⁵ there are many settings where the two quantities are comparable; for example, if E has a submatrix whose entries are identical and has zero entries otherwise, then $\|E\|_{\infty} = \|E\|_2$.

Theorem 3 provides the perturbation of individual eigenvectors, under a usual eigengap assumption. When r and μ are not large, we incur an additional term $O(\|E\|_2/\delta\sqrt{d})$ in the bound. This is understandable, since $\|\tilde{v}_i - v_i\|_2$ is typically $O(\|E\|_2/\delta)$.

When the rank of A is not exactly r , we require that $|\lambda_r|$ is larger than the approximation error $\|A - A_r\|_{\infty}$. It is important to state that this assumption is more restricted than the eigengap assumption in the Davis-Kahan theorem, since $\|A - A_r\|_{\infty} \geq \|A - A_r\|_2 = |\lambda_{r+1}|$. However, different from the matrix max-norm, the spectral norm $\|\cdot\|_2$ only depends on the eigenvalues of a matrix, so it is natural to expect ℓ_2 perturbation bounds that only involve λ_r and λ_{r+1} . It is not clear whether we should expect an ℓ_{∞} bound that involves λ_{r+1} instead of ε . More discussions can be found in Section 5.

We do not pursue the optimal bound in terms of r and $\mu(V)$ in this paper, as the two quantities are not large in many applications, and the current proof is already complicated.

4. To see how the Davis-Kahan $\sin \Theta$ theorem relates to this form, we can use the identity $\|\sin \Theta(\tilde{V}, V)\|_2 = \|\tilde{V}\tilde{V}^T - VV^T\|_2$ (Stewart, 1990), and the (easily verifiable) inequality $2\min_R \|\tilde{V}R - V\|_2 \geq \|\tilde{V}\tilde{V}^T - VV^T\|_2 \geq \min_R \|\tilde{V}R - V\|_2$ where R is an orthogonal matrix.

5. Since $\|E\|_1\|E\|_{\infty} \geq \|E\|_2^2$ (Stewart, 1990), the inequality follows from $\|E\|_1 = \|E\|_{\infty}$ by symmetry.

2.2 Rectangular matrices

Now we establish ℓ_∞ perturbation bounds for general rectangular matrices. The results here are more general than those in Section 1, and in particular, we allow the matrix A to be of approximate low rank. Suppose that both A and E are $d_1 \times d_2$ matrices, and $\tilde{A} := A + E$. The rank of A is at most $d_1 \wedge d_2$ (where $a \wedge b = \min\{a, b\}$). Suppose an integer r satisfies $r \leq \text{rank}(A)$. Let the singular value decomposition of A be

$$A = \sum_{i=1}^r \sigma_i u_i v_i^T + \sum_{i=r+1}^{d_1 \wedge d_2} \sigma_i u_i v_i^T,$$

where the singular values are ordered as $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{d_1 \wedge d_2} \geq 0$, and the unit vectors $u_1, \dots, u_{d_1 \wedge d_2}$ (or unit vectors $v_1, \dots, v_{d_1 \wedge d_2}$) are orthogonal to each other. We denote $U = [u_1, \dots, u_r] \in \mathbb{R}^{d_1 \times r}$ and $V = [v_1, \dots, v_r] \in \mathbb{R}^{d_2 \times r}$. Analogously, the singular value decomposition of \tilde{A} is

$$\tilde{A} = \sum_{i=1}^r \tilde{\sigma}_i \tilde{u}_i \tilde{v}_i^T + \sum_{i=r+1}^{d_1 \wedge d_2} \tilde{\sigma}_i \tilde{u}_i \tilde{v}_i^T,$$

where $\tilde{\sigma}_1 \geq \dots \geq \tilde{\sigma}_{d_1 \wedge d_2}$. Similarly, columns of $\tilde{U} = [\tilde{u}_1, \dots, \tilde{u}_r] \in \mathbb{R}^{d_1 \times r}$ and $\tilde{V} = [\tilde{v}_1, \dots, \tilde{v}_r] \in \mathbb{R}^{d_2 \times r}$ are orthonormal.

Define $\mu_0 = \max\{\mu(V), \mu(U)\}$, where $\mu(U)$ (resp. $\mu(V)$) is the coherence of U (resp. V). This μ_0 will appear in the statement of our results, as it controls both the structure of left and right singular spaces. When, specially, A is a symmetric matrix, the spectral decomposition of A is also the singular value decomposition (up to sign), and thus μ_0 coincides with μ defined in Section 2.1.

Recall the definition of matrix ∞ -norm and 1-norm of a rectangular matrix (4). Similar to the matrix ∞ -norm, $\|\cdot\|_1$ is an operator norm in the ℓ_1 space. An obvious relationship between matrix ∞ -norm and 1-norm is $\|E\|_\infty = \|E^T\|_1$. Note that the matrix ∞ -norm and 1-norm have different number of summands in their definitions, so we are motivated to consider $\tau_0 := \max\{\sqrt{d_1/d_2}\|E\|_\infty, \sqrt{d_2/d_1}\|E\|_1\}$ to balance the dimensions d_1 and d_2 .

Let $A_r = \sum_{i \leq r} \sigma_i u_i v_i^T$ be the best rank- r approximation of A under the Frobenius norm, and let $\varepsilon_0 = \sqrt{d_1/d_2}\|A - A_r\|_\infty \vee \sqrt{d_2/d_1}\|A - A_r\|_1$, which also balances the two dimensions. Note that in the special case where A is symmetric, this approximation error ε_0 is identical to ε defined in Section 2.1. The next theorem bounds the perturbation of singular spaces.

Theorem 4 *Suppose that $\delta_0 - \varepsilon_0 = \Omega(r^3 \mu_0^2 \tau_0)$. Then, there exists orthogonal matrices $R_U, R_V \in \mathbb{R}^{r \times r}$ such that,*

$$\|\tilde{U} R_U - U\|_{\max} = O\left(\frac{r^{5/2} \mu_0^2 \tau_0}{(\sigma_r - \varepsilon_0) \sqrt{d_1}}\right), \quad \|\tilde{V} R_V - V\|_{\max} = O\left(\frac{r^{5/2} \mu_0^2 \tau_0}{(\sigma_r - \varepsilon_0) \sqrt{d_2}}\right).$$

Similar to Theorem 3, under an assumption of gaps between singular values, the next theorem bounds the perturbation of individual singular vectors.

Theorem 5 *Suppose the same assumptions in Theorem 4 hold. In addition, suppose δ_0 satisfies $\delta_0 > \|E\|_2$, and for any $i \in [r]$, the interval $[\sigma_i - \delta_0, \sigma_i + \delta_0]$ does not contain any eigenvalues of A other than σ_i . Then, up to sign,*

$$\max_{i \in [r]} \|\tilde{u}_i - u_i\|_\infty = O\left(\frac{r^4 \mu_0^2 \tau_0}{(\sigma_r - \varepsilon_0) \sqrt{d_1}} + \frac{r^{3/2} \mu_0^{1/2} \|E\|_2}{\delta_0 \sqrt{d_1}}\right), \quad (9)$$

$$\max_{i \in [r]} \|\tilde{v}_i - v_i\|_\infty = O\left(\frac{r^4 \mu_0^2 \tau_0}{(\sigma_r - \varepsilon_0) \sqrt{d_2}} + \frac{r^{3/2} \mu_0^{1/2} \|E\|_2}{\delta_0 \sqrt{d_2}}\right). \quad (10)$$

As mentioned in the beginning of this section, we will use dilation to augment all $d_1 \times d_2$ matrices into symmetric ones with size $d_1 + d_2$. In order to balance the possibly different scales of d_1 and d_2 , we consider a weighted max-norm. This idea will be further illustrated in Section 5.

2.3 Examples: which matrices have such structure?

In many problems, low-rank structure naturally arises due to the impact of pervasive latent factors that influence most observed data. Since observations are imperfect, the low-rank structure is often ‘perturbed’ by an additional sparse structure, gross errors, measurement noises, or the idiosyncratic components that can not be captured by the latent factors. We give some motivating examples with such structure.

Panel data in stock markets. Consider the excess returns from a stock market over a period of time. The driving factors in the market are reflected in the covariance matrix as a low rank component A . The residual covariance of the idiosyncratic components is often modeled by a sparse component S . Statistical analysis including PCA is usually conducted based on the estimated covariance matrix $\tilde{A} = \hat{\Sigma}$, which is perturbed from the true covariance $\Sigma = A + S$ by the estimation error N (Stock and Watson, 2002; Fan et al., 2013). In Section 3.1, we will develop a robust estimation method in the presence of heavy-tailed return data.

Video surveillance. In image processing and computer vision, it is often desired to separate moving objects from static background before further modeling and analysis (Oliver et al., 2000; Hu et al., 2004). The static background corresponds to the low rank component A in the data matrix, which is a collection of video frames, each consisting of many pixels represented as a long vector in the data matrix. Moving objects and noise correspond to the sparse matrix S and noise matrix N . Since the background is global information and reflected by many pixels of a frame, it is natural for the incoherence condition to hold.

Wireless sensor network localization. In wireless sensor networks, we are usually interested in determining the location of sensor nodes with unknown position based on a few (noisy) measurements between neighboring nodes (Doherty et al., 2001; Biswas and Ye, 2004). Let \mathbb{X} be an r by n matrix such that each column x_i gives the coordinates of each node in a plane ($r = 2$) or a space ($r = 3$). Assume the center of the sensors has been relocated at origin. Then the low rank matrix $A = \mathbb{X}^T \mathbb{X}$, encoding the true distance information, has to satisfy distance constraints given by the measurements. The noisy distance matrix \tilde{A} after centering, equals to the sum of A and a matrix N consisting of measurement

errors. Suppose that each node is a random point uniformly distributed in a rectangular region. It is not difficult to see that with high probability, the top r eigenvalues of $\mathbb{X}^T\mathbb{X}$ and their eigengap scales with the number of sensors n and the leading eigenvectors have a bounded coherence.

In our theorems, we require that the coherence μ is not too large. This is a natural structural condition associated with the low rank matrices. Consider the following very simple example: if the eigenvectors v_1, \dots, v_r of the low rank matrix A are uniform unit vectors in a sphere, then with high probability, $\max_i \|v_i\|_\infty = O(\sqrt{\log n})$, which implies $\mu = O(\log n)$. An intuitive way to understand the incoherence structure is that no coordinates of v_1 (or v_2, \dots, v_r) are dominant. In other words, the eigenvectors are not concentrated on a few coordinates.

In all our examples, the incoherence structure is natural. The factor model satisfies such structure, which will be discussed in Section 3. In the video surveillance example, ideally, when the images are static, A is a rank one matrix $x\mathbf{1}^T$. Since usually a majority of pixels (coordinates of x) help to display an image, the vector x often has dense coordinates with comparable magnitude, so A also has an incoherence structure in this example. Similarly, in the sensor localization example, the coordinates of all sensor nodes are comparable in magnitude, so the low rank matrix A formed by $\mathbb{X}^T\mathbb{X}$ also has the desired incoherence structure.

2.4 Other perturbation results

Although the eigenvector perturbation theory is well studied in numerical analysis, there is a renewed interest among statistics and machine learning communities recently, due to the wide applicability of PCA and other eigenvector-based methods. In Cai and Zhang (2016); Yu et al. (2015), they obtained variants or improvements of Davis-Kahan theorem (or Wedin's theorem), which are user-friendly in the statistical contexts. These results assume the perturbation is deterministic, which is the same as Davis-Kahan theorem and Wedin's theorem. In general, these results are sharp, even when the perturbation is random, as evidenced by the BBP transition (Baik et al., 2005).

However, these classical results can be suboptimal, when the perturbation is random and the smallest eigenvalue gap $\lambda_1 - \lambda_2$ does not capture particular spectrum structure. For example, Vu (2011); O'Rourke et al. (2013) showed that with high probability, there are bounds sharper than the Wedin's theorem, when the signal matrix is low-rank and satisfies certain eigenvalue conditions.

In this paper, our perturbation results are deterministic, thus the bound can be suboptimal when the perturbation is random with certain structure (e.g. the difference between sample covariance and population one for i.i.d. samples). However, the advantage of a deterministic result is that it is applicable to any random perturbation. This is especially useful when we cannot make strong random assumptions on the perturbation (e.g., the perturbation is an unknown sparse matrix). In Section 3, we will see examples of this type.

3. Application to robust covariance estimation

We will study the problem of robust estimation of covariance matrices and show the strength of our perturbation result. Throughout this section, we assume both rank r and the coher-

ence $\mu(V)$ are bounded by a constant, though this assumption can be relaxed. We will use C to represent a generic constant, and its value may change from line to line.

3.1 PCA in spiked covariance model

To initiate our discussions, we first consider sub-Gaussian random variables. Let $X = (X_1, \dots, X_d)$ be a random d -dimensional vector with mean zero and covariance matrix

$$\Sigma = \sum_{i=1}^r \lambda_i v_i v_i^T + \sigma^2 I_d := \Sigma_1 + \Sigma_2, \quad (\lambda_1 \geq \dots \geq \lambda_r > 0), \quad (11)$$

and \mathbb{X} be an n by d matrix, whose rows are independently sampled from the same distribution. This is the spiked covariance model that has received intensive study in recent years. Let the empirical covariance matrix be $\widehat{\Sigma} = \mathbb{X}^T \mathbb{X} / n$. Viewing the empirical covariance matrix as its population version plus an estimation error, we have the decomposition

$$\widehat{\Sigma} = \Sigma_1 + \Sigma_2 + \left(\frac{1}{n} \mathbb{X}^T \mathbb{X} - \Sigma \right),$$

which is a special case of the general decomposition in (1). Here, Σ_2 is the sparse component, and the estimation error $\mathbb{X}^T \mathbb{X} / n - \Sigma$ is the noise component. Note that v_1, \dots, v_r are just the top r leading eigenvectors of Σ and we write $V = [v_1, \dots, v_r]$. Assume the top r eigenvectors of $\widehat{\Sigma}$ are denoted by $\widehat{v}_1, \dots, \widehat{v}_r$. We want to find an ℓ_∞ bound on the estimation error $\widehat{v}_i - v_i$ for all $i \in [r]$.

When the dimension d is comparable to or larger than n , it has been shown by Johnstone and Lu (2009) that the leading empirical eigenvector \widehat{v}_1 is not a consistent estimate of the true eigenvector v_1 , unless we assume larger eigenvalues. Indeed, we will impose more stringent conditions on λ_i 's in order to obtain good ℓ_∞ bounds.

Assuming the coherence $\mu(V)$ is bounded, we can easily see $\text{Var}(X_j) \leq \sigma^2 + C\lambda_1/d$ for some constant C . It follows from the standard concentration result (e.g., Vershynin (2010)) that if rows of \mathbb{X} contains i.i.d sub-Gaussian vectors and $\log d = O(n)$, then with probability greater than $1 - d^{-1}$,

$$\left\| \frac{1}{n} \mathbb{X}^T \mathbb{X} - \Sigma \right\|_{\max} \leq C \left(\sigma^2 + \frac{\lambda_1}{d} \right) \sqrt{\frac{\log d}{n}}. \quad (12)$$

To apply Theorem 3, we treat Σ_1 as A and $\widehat{\Sigma} - \Sigma_1$ as E . If the conditions in Theorem 3 are satisfied, we will obtain

$$\max_{1 \leq k \leq r} \|\widehat{v}_k - v_k\|_\infty = O(\|E\|_\infty / (\lambda_r \sqrt{d}) + \|E\|_2 / (\delta \sqrt{d})). \quad (13)$$

Note there are simple bounds on $\|E\|_\infty$ and $\|E\|_2$:

$$\|E\|_2 \leq \|E\|_\infty \leq \sigma^2 + d \left\| \frac{1}{n} \mathbb{X}^T \mathbb{X} - \Sigma \right\|_{\max} \leq C \left\{ 1 + (d\sigma^2 + \lambda_1) \sqrt{\frac{\log d}{n}} \right\}.$$

By assuming a strong uniform eigengap, the conditions in Theorem 3 are satisfied, and the bound in (13) can be simplified. Define the uniform eigengap as

$$\gamma = \min\{\lambda_i - \lambda_{i+1} : 1 \leq i \leq r\}, \quad \lambda_{r+1} := 0.$$

Note that $\gamma \leq \min\{\lambda_r, \delta\}$, so if $\gamma > C(1 + (d\sigma^2 + \lambda_1)\sqrt{\log d/n})$, we have

$$\max_{1 \leq k \leq r} \|\widehat{v}_k - v_k\|_\infty = O_P\left(\frac{\|E\|_\infty}{\gamma\sqrt{d}}\right) = O_P\left(\frac{1 + (d\sigma^2 + \lambda_1)\sqrt{\log d/n}}{\gamma\sqrt{d}}\right),$$

In particular, when $\lambda_1 \asymp \gamma$ and $\gamma \gg \max\{1, \sigma^2 d\sqrt{\log d/n}\}$, we have

$$\max_{1 \leq k \leq r} \|\widehat{v}_k - v_k\|_\infty = o_P\left(\frac{1}{\sqrt{d}}\right).$$

The above analysis pertains to the structure of sample covariance matrix. In the following subsections, we will estimate the covariance matrix using more complicated robust procedure. Our perturbation theorems in Section 2 provide a fast and clean approach to obtain new results.

3.2 PCA for robust covariance estimation

The usefulness of Theorem 3 is more pronounced when the random variables are heavy-tailed. Consider again the covariance matrix Σ with structure (11). Instead of assuming sub-Gaussian distribution, we assume there exists a constant $C > 0$ such that $\max_{j \leq d} EX_j^4 < C$, i.e. the fourth moments of the random variables are uniformly bounded.

Unlike sub-Gaussian variables, there is no concentration bound similar to (12) for the empirical covariance matrix. Fortunately, thanks to recent advances in robust statistics (e.g., Catoni (2012)), robust estimate of Σ with guaranteed concentration property becomes possible. We shall use the method proposed in Fan et al. (2017a). Motivated by the classical M -estimator of Huber (1964), Fan et al. (2017a) proposed a robust estimator for each element of $\widehat{\Sigma}$, by solving a Huber loss based minimization problem

$$\widehat{\Sigma}_{ij} = \operatorname{argmin}_\mu \sum_{t=1}^n l_\alpha(X_{ti}X_{tj} - \mu), \quad (14)$$

where l_α is the Huber loss defined as

$$l_\alpha(x) = \begin{cases} 2\alpha|x| - \alpha^2, & |x| \geq \alpha, \\ x^2, & |x| \leq \alpha. \end{cases}$$

The parameter α is suggested to be $\alpha = \sqrt{nv^2/\log(\epsilon^{-1})}$ for $\epsilon \in (0, 1)$, where v is assumed to satisfy $v \geq \max_{ij} \sqrt{\operatorname{Var}(X_i X_j)}$. If $\log(\epsilon^{-1}) \leq n/8$, Fan et al. (2017a) showed

$$P\left(|\widehat{\Sigma}_{ij} - \Sigma_{ij}| \leq 4v\sqrt{\frac{\log(\epsilon^{-1})}{n}}\right) \geq 1 - 2\epsilon.$$

From this result, the next proposition is immediate by taking $\epsilon = d^{-3}$.

Proposition 1 *Suppose that there is a constant C with $\max_{j \leq d} EX_j^4 < C$. Then with probability greater than $1 - d^{-1}(1 + d^{-1})$, the robust estimate of covariance matrix with $\alpha = \sqrt{3nv^2 \log(d)}$ satisfies*

$$\|\widehat{\Sigma} - \Sigma\|_{\max} \leq 4v \sqrt{\frac{3 \log d}{n}},$$

where v is a pre-determined parameter assumed to be no less than $\max_{ij} \sqrt{\text{Var}(X_i X_j)}$.

This result relaxes the sub-Gaussianity assumption by robustifying the covariance estimate. It is apparent that the ℓ_∞ bound in the previous section is still valid in this case. To be more specific, suppose $\mu(V)$ is bounded by a constant. Then, (13) holds for the PCA based on the robust covariance estimation. When $\lambda_1 \asymp \gamma$ and $\gamma \gg \max\{1, \sigma^2 d \sqrt{\log d/n}\}$, we again have

$$\max_{1 \leq k \leq r} \|\widehat{v}_k - v_k\|_\infty = O_P\left(\frac{1 + (d\sigma^2 + \lambda_1) \sqrt{\log d/n}}{\gamma \sqrt{d}}\right) = o_P\left(\frac{1}{\sqrt{d}}\right).$$

Note that an entrywise estimation error $o_p(1/\sqrt{d})$ necessarily implies consistency of the estimated eigenvectors, since we can easily convert an ℓ_∞ result into an ℓ_2 result. The minimum signal strength (or magnitude of leading eigenvalues) for such consistency is shown to be $\sigma^2 d/n$ under the sub-Gaussian assumption (Wang and Fan, 2017).

If the goal is simply to prove consistency of \widehat{v}_k , the strategy of using our ℓ_∞ perturbation bounds is not optimal. However, there are also merits: our result is nonasymptotic; it holds for more general distributions (beyond sub-Gaussian distributions); and its entrywise bound gives stronger guarantee. Moreover, the ℓ_∞ perturbation bounds provide greater flexibility for analysis, since it is straightforward to adapt analysis to problems with more complicated structure. For example, the above discussion can be easily extended to a general Σ_2 with bounded $\|\Sigma_2\|_\infty$ rather than a diagonal matrix.

3.3 Robust covariance estimation via factor models

In this subsection, we will apply Theorem 3 to robust large covariance matrix estimation for approximate factor models in econometrics. With this theorem, we are able to extend the data distribution in factor analysis beyond exponentially decayed distributions considered by Fan et al. (2013), to include heavy-tailed distributions.

Suppose the observation y_{it} , say, the excess return at day t for stock i , admits a decomposition

$$y_{it} = b_i^T f_t + u_{it}, \quad i \leq d, t \leq n, \quad (15)$$

where $b_i \in \mathbb{R}^r$ is the unknown but fixed loading vector, $f_t \in \mathbb{R}^r$ denotes the unobserved factor vector at time t , and u_{it} 's represent the idiosyncratic noises. Let $y_t = (y_{1t}, \dots, y_{dt})^T$ and $u_t = (u_{1t}, \dots, u_{dt})^T$ so that $y_t = B f_t + u_t$, where $B = (b_1, \dots, b_d)^T \in \mathbb{R}^{d \times r}$. Suppose that f_t and u_t are uncorrelated and centered random vectors, with bounded fourth moments, i.e., the fourth moments of all entries of f_t and u_t are bounded by some constant. We assume $\{f_t, u_t\}$ are independent for t , although it is possible to allow for weak temporal

dependence as in Fan et al. (2013). From (15), we can decompose $\Sigma = \text{Cov}(y_t)$ into a low rank component and a residual component:

$$\Sigma = BB^T + \Sigma_u, \quad (16)$$

where $\Sigma_u := \text{Cov}(u_t)$. To circumvent the identifiability issue common in latent variable models, here we also assume, without loss of generality, $\text{Cov}(f_t) = I_r$ and that $B^T B$ is a diagonal matrix, since rotating B will not affect the above decomposition (16).

We will need two major assumptions for our analysis: (1) the factors are *pervasive* in the sense of Definition 2, and (2) there is a constant $C > 0$ such that $\|\Sigma_u^{-1}\|_2, \|\Sigma_u\|_2 \leq C$, which are standard assumptions in the factor model literature. The pervasive assumption is reasonable in financial applications, since the factors have impacts on a large fraction of the outcomes (Chamberlain et al., 1983; Bai, 2003). If the factor loadings $\{b_i\}_{i=1}^d$ are regarded as random realizations from a bounded random vector, the assumption holds (Fan et al., 2013).

Definition 2 *In the factor model (15), the factors are called pervasive if there is a constant $C > 0$ such that $\|B\|_{\max} \leq C$ and the eigenvalues of the r by r matrix $B^T B/d$ are distinct and bounded away from zero and infinity.*

Let $\{\lambda_i, v_i\}_{i=1}^r$ be the top r eigenvalues and eigenvectors of Σ , and similarly, $\{\bar{\lambda}_i, \bar{v}_i\}_{i=1}^r$ for BB^T . In the following proposition, we show that pervasiveness is naturally connected to the incoherence structure. This connects the econometrics and machine learning literature and provide a good interpretation on the concept of the incoherence. Its proof can be found in the appendix.

Proposition 3 *Suppose there exists a constant $C > 0$ such that $\|\Sigma_u\| \leq C$. The factors f_t are pervasive if and only if the coherence $\mu(V)$ for $V = (v_1, \dots, v_r) \in \mathbb{R}^{d \times r}$ is bounded by some constant, and $\lambda_i = \lambda_i(\Sigma) \asymp d$ for $i \leq r$ so that $\min_{1 \leq i \neq j \leq r} |\lambda_i - \lambda_j|/\lambda_j > 0$.*

Our goal is to obtain a good covariance matrix estimator by exploiting the structure (16). Our strategy is to use a generalization of the principal orthogonal complement thresholding (POET) method proposed in Fan et al. (2013). The generic POET procedure encompasses three steps:

- (1) Given three pilot estimators $\widehat{\Sigma}, \widehat{\Lambda} = \text{diag}(\widehat{\lambda}_1, \dots, \widehat{\lambda}_r), \widehat{V} = (\widehat{v}_1, \dots, \widehat{v}_r)$ respectively for true covariance Σ , leading eigenvalues $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_r)$ and leading eigenvectors $V = (v_1, \dots, v_r)$, compute the principal orthogonal complement $\widehat{\Sigma}_u$:

$$\widehat{\Sigma}_u = \widehat{\Sigma} - \widehat{V}\widehat{\Lambda}\widehat{V}^T. \quad (17)$$

- (2) Apply the correlation thresholding to $\widehat{\Sigma}_u$ to obtain thresholded estimate $\widehat{\Sigma}_u^\top$ defined as follows:

$$\widehat{\Sigma}_{u,ij}^\top = \begin{cases} \widehat{\Sigma}_{u,ij}, & i = j \\ s_{ij}(\widehat{\Sigma}_{u,ij})I(|\widehat{\Sigma}_{u,ij}| \geq \tau_{ij}), & i \neq j \end{cases}, \quad (18)$$

where $s_{ij}(\cdot)$ is the generalized shrinkage function (Antoniadis and Fan, 2001; Rothman et al., 2009) and $\tau_{ij} = \tau(\widehat{\sigma}_{u,ii}\widehat{\sigma}_{u,jj})^{1/2}$ is an entry-dependent threshold. τ will be determined later in Theorem 6. This step exploits the sparsity of Σ_u .

(3) Construct the final estimator $\widehat{\Sigma}^\top = \widehat{V}\widehat{\Lambda}\widehat{V}^\top + \widehat{\Sigma}_u^\top$.

The key feature in the above procedure lies in the flexibility of choosing the pilot estimators in the first step. We will choose $\widehat{\Sigma}$ according to data generating distribution. Typically we can use $\widehat{\lambda}_i, \widehat{v}_i$ for $i \leq r$ as the eigenvalues/vectors of $\widehat{\Sigma}$. However, $\widehat{\Lambda}$ and \widehat{V} in general do not have to come from the spectral information of $\widehat{\Sigma}$ and can be obtained separately via different methods.

To guide the selection of proper pilot estimators, Fan et al. (2017+) provided a high level sufficient condition for this simple procedure to be effective, and its performance is gauged, in part, through the sparsity level of Σ_u , defined as $m_d := \max_{i \leq d} \sum_{j \leq d} |\Sigma_{u,ij}|^q$. When $q = 0$, m_d corresponds to the maximum number of nonzero elements in each row of Σ_u . For completeness, we present the theorem given by Fan et al. (2017+) in the following.

Theorem 6 *Let $w_n = \sqrt{\log d/n} + 1/\sqrt{d}$. Suppose there exists $C > 0$ such that $\|\Sigma_u^{-1}\|, \|\Sigma_u\| \leq C$ and we have pilot estimators $\widehat{\Sigma}, \widehat{\Lambda}, \widehat{V}$ satisfying*

$$\|\widehat{\Sigma} - \Sigma\|_{\max} = O(\sqrt{\log d/n}), \quad (19)$$

$$|\widehat{\lambda}_i/\lambda_i - 1| = O(\sqrt{\log d/n}), \quad (20)$$

$$\|\widehat{v}_i - v_i\|_\infty = O(w_n/\sqrt{d}). \quad (21)$$

Under the pervasiveness condition of the factor model (15), with $\tau \asymp w_n$, if $m_d w_n^{1-q} = o(1)$, the following rates of convergence hold with the generic POET procedure:

$$\|\widehat{\Sigma}_u^\top - \Sigma_u\|_2 = O(m_d w_n^{1-q}) = \|(\widehat{\Sigma}_u^\top)^{-1} - \Sigma_u^{-1}\|_2, \quad (22)$$

and

$$\begin{aligned} \|\widehat{\Sigma}^\top - \Sigma\|_{\max} &= O(w_n), \\ \|\widehat{\Sigma}^\top - \Sigma\|_\Sigma &= O\left(\frac{\sqrt{d} \log d}{n} + m_d w_n^{1-q}\right), \\ \|(\widehat{\Sigma}^\top)^{-1} - \Sigma^{-1}\|_2 &= O(m_d w_n^{1-q}), \end{aligned} \quad (23)$$

where $\|A\|_\Sigma = d^{-1/2} \|\Sigma^{-1/2} A \Sigma^{-1/2}\|_F$ is the relative Frobenius norm.

We remark that the additional term $1/\sqrt{d}$ in w_n , is due to the estimation of unobservable factors and is negligible when the dimensional d is high. The optimality of the above rates of convergence is discussed in details in Fan et al. (2017+). Theorem 6 reveals a profound deterministic connection between the estimation error bound of the pilot estimators with the rate of convergences of the POET output estimators. Notice that the eigenvector estimation error is under the ℓ_∞ norm, for which our ℓ_∞ perturbation bounds will prove to be useful.

In this subsection, since we assume only bounded fourth moments, we choose $\widehat{\Sigma}$ to be the robust estimate of covariance matrix Σ defined in (14). We now invoke our ℓ_∞ bounds to show that the spectrum properties (eigenvalues and eigenvectors) are stable to perturbation.

Let us decompose $\widehat{\Sigma}$ into a form such that Theorem 3 can be invoked:

$$\widehat{\Sigma} = \sum_{i=1}^r \bar{\lambda}_i \bar{v}_i \bar{v}_i^T + \Sigma_u + (\widehat{\Sigma} - \Sigma),$$

where $\widehat{\Sigma}$ is viewed as \widetilde{A} , the low-rank part $\sum_{i=1}^r \bar{\lambda}_i \bar{v}_i \bar{v}_i^T$, which is also BB^T , is viewed as A , and the remaining terms are treated as E . The following results follow immediately.

Proposition 4 *Assume that there is a constant $C > 0$ such that $\|\Sigma_u\| \leq C$. If the factors are pervasive, then with probability greater than $1 - d^{-1}$, we have (19) – (21) hold with $\widehat{\lambda}_i, \widehat{v}_i$ as the leading eigenvalues/vectors of $\widehat{\Sigma}$ for $i \leq r$. In addition, (22) and (23) hold.*

The inequality (19) follows directly from Proposition 1 under the assumption of bounded fourth moments. It is also easily verifiable that (20), (21) follow from (19) by Weyl's inequality and Theorem 3 (noting that $\|\Sigma_u\|_\infty \leq \sqrt{d}\|\Sigma_u\|$). See Section 3.2 for more details.

Note that in the case of sub-Gaussian variables, sample covariance matrix and its leading eigenvalues/vectors will also serve the same purpose due to (12) and Theorem 3 as discussed in Section 3.1.

We have seen that the ℓ_∞ perturbation bounds are useful in robust covariance estimation, and particularly, they resolve a theoretical difficulty in the generic POET procedure for factor model based covariance matrix estimation.

4. Simulations

4.1 Simulation: the perturbation result

In this subsection, we implement numerical simulations to verify the perturbation bound in Theorem 3. We will show that the error behaves in the same way as indicated by our theoretical bound.

In the experiments, we let the matrix size d run from 200 to 2000 by an increment of 200. We fix the rank of A to be 3 ($r = 3$). To generate an incoherence low rank matrix, we sample a $d \times d$ random matrix with iid standard normal variables, perform singular value decomposition, and extract the first r right singular vectors v_1, v_2, \dots, v_r . Let $V = (v_1, \dots, v_r)$ and $D = \text{diag}(r\gamma, (r-1)\gamma, \dots, \gamma)$ where γ as before represents the eigengap. Then, we set $A = VDVT^T$. By orthogonal invariance, v_i is uniformly distributed on the unit sphere \mathbb{S}^{d-1} . It is not hard to see that with probability $1 - O(d^{-1})$, the coherence of V $\mu(V) = O(\sqrt{\log d})$.

We consider two types of sparse perturbation matrices E : (a) construct a $d \times d$ matrix E_0 by randomly selecting s entries for each row, and sampling a uniform number in $[0, L]$ for each entry, and then symmetrize the perturbation matrix by setting $E = (E_0 + E_0^T)/2$; (b) pick $\rho \in (0, 1)$, $L' > 0$, and let $E_{ij} = L'\rho^{|i-j|}$. Note that in (b) we have $\|E\|_\infty \leq 2L'/(1-\rho)$, and thus we can choose suitable L' and ρ to control the ℓ_∞ norm of E . This covariance structure is common in cases where correlations between random variables depend on their “distance” $|i-j|$, which usually arises from autoregressive models.




Figure 1: The left plot shows the perturbation error of eigenvectors against matrix size d ranging from 200 to 2000, with different eigengap γ . The right plot shows $\log(\text{err})$ against $\log(d)$. The slope is around -0.5 . Blue lines represent $\gamma = 10$; red lines $\gamma = 50$; green lines $\gamma = 100$; and black lines $\gamma = 500$. We report the largest error over 100 runs.

The perturbation of eigenvectors is measured by the element-wise error:

$$\text{err} := \max_{1 \leq i \leq r} \min_{\eta_i \in \{\pm 1\}} \|\eta_i \tilde{v}_i - v_i\|_\infty,$$

where $\{\tilde{v}_i\}_{i=1}^r$ are the eigenvectors of $\tilde{A} = A + E$ in the descending order.

To investigate how the error depends on γ and d , we generate E according to mechanism (a) with $s = 10, L = 3$, and run simulations in different parameter configurations: (1) let the matrix size d range from 200 to 2000, and choose the eigengap γ in $\{10, 50, 100, 500\}$ (Figure 1); (2) fix the product $\gamma\sqrt{d}$ to be one of $\{2000, 3000, 4000, 5000\}$, and let the matrix size d run from 200 to 2000 (Figure 2).

To find how the errors behave for E generated from different methods, we run simulations as in (1) but generate E differently. We construct E through mechanism (a) with $L = 10, s = 3$ and $L = 0.6, s = 50$, and also through mechanism (b) with $L' = 1.5, \rho = 0.9$ and $L' = 7.5, \rho = 0.5$ (Figure 3). The parameters are chosen such that $\|E\|_\infty$ is about 30.

In Figure 1 – 3, we report the largest error based on 100 runs. Figure 1 shows that the error decreases as d increases (the left plot); and moreover, the logarithm of the error is linear in $\log(d)$, with a slope -0.5 , that is, $\text{err} \propto 1/\sqrt{d}$ (the right plot). We can take the eigengap γ into consideration and characterize the relationship in a more refined way. In Figure 2, it is clear that err almost falls on the same horizontal line for different configurations of d and γ , with $\gamma\sqrt{d}$ fixed. The right panel clearly indicates that $\text{err} \times \gamma\sqrt{d}$ is a constant, and therefore $\text{err} \propto 1/(\gamma\sqrt{d})$. In Figure 3, we find that the errors behave almost the




Figure 2: The left plot shows the perturbation error of eigenvectors against matrix size d ranging from 200 to 2000, when $\gamma\sqrt{d}$ is kept fixed, with different values. The right plot shows the error multiplied by $\gamma\sqrt{d}$ against d . Blue lines represent $\gamma\sqrt{d} = 2000$; red lines $\gamma\sqrt{d} = 3000$; green lines $\gamma\sqrt{d} = 4000$; and black lines $\gamma\sqrt{d} = 5000$. We report the largest error over 100 runs.

same regardless of how E is generated. These simulation results provide stark evidence supporting the ℓ_∞ perturbation bound in Theorem 3.

4.2 Simulation: robust covariance estimation

We consider the performance of the generic POET procedure in robust covariance estimation in this subsection. Note that the procedure is flexible in employing any pilot estimators $\widehat{\Sigma}, \widehat{\Lambda}, \widehat{V}$ satisfying the conditions (19) – (21) respectively.

We implemented the robust procedure with four different initial trios: (1) the sample covariance $\widehat{\Sigma}^S$ with its leading r eigenvalues and eigenvectors as $\widehat{\Lambda}^S$ and \widehat{V}^S ; (2) the Huber's robust estimator $\widehat{\Sigma}^R$ given in (14) and its top r eigen-structure estimators $\widehat{\Lambda}^R$ and \widehat{V}^R ; (3) the marginal Kendall's tau estimator $\widehat{\Sigma}^K$ with its corresponding $\widehat{\Lambda}^K$ and \widehat{V}^K ; (4) lastly, we use the spatial Kendall's tau estimator to estimate the leading eigenvectors instead of the marginal Kendall's tau, so \widehat{V}^K in (3) is replaced with \widetilde{V}^K . We need to briefly review the two types of Kendall's tau estimators here, and specifically give the formula for $\widehat{\Sigma}^K$ and \widetilde{V}^K .

Kendall's tau correlation coefficient, for estimating pairwise comovement correlation, is defined as

$$\widehat{\tau}_{jk} := \frac{2}{n(n-1)} \sum_{t < t'} \text{sgn}((y_{tj} - y_{t'j})(y_{tk} - y_{t'k})). \quad (24)$$




Figure 3: These plots show $\log(\text{err})$ against $\log(d)$, with matrix size d ranging from 200 to 2000 and different eigengap γ . The perturbation E is generated from different ways. Top left: $L = 10, s = 3$; top right: $L = 0.6, s = 50$; bottom left: $L' = 1.5, \rho = 0.9$; bottom right: $L' = 7.5, \rho = 0.5$. The slopes are around -0.5 . Blue lines represent $\gamma = 10$; red lines $\gamma = 50$; green lines $\gamma = 100$; and black lines $\gamma = 500$. We report the largest error over 100 runs.

Its population expectation is related to the Pearson correlation via the transform $r_{jk} = \sin\left(\frac{\pi}{2} E[\hat{\tau}_{jk}]\right)$ for elliptical distributions (which are far too restrictive for high-dimensional applications). Then $\hat{r}_{jk} = \sin\left(\frac{\pi}{2} \hat{\tau}_{jk}\right)$ is a valid estimation for the Pearson correlation r_{jk} . Letting $\hat{R} = (\hat{r}_{jk})$ and $\hat{D} = \text{diag}(\sqrt{\hat{\Sigma}_{11}^R}, \dots, \sqrt{\hat{\Sigma}_{dd}^R})$ containing the robustly estimated standard deviations, we define the marginal Kendall's tau estimator as

$$\hat{\Sigma}^K = \hat{D} \hat{R} \hat{D}. \quad (25)$$

In the above construction of \hat{D} , we still use the robust variance estimates from $\hat{\Sigma}^R$.

The spatial Kendall's tau estimator is a second-order U-statistic, defined as

$$\tilde{\Sigma}^K := \frac{2}{n(n-1)} \sum_{t < t'} \frac{(y_t - y_{t'})(y_t - y_{t'})^T}{\|y_t - y_{t'}\|_2^2}. \quad (26)$$

Then \tilde{V}^S is constructed by the top r eigenvectors of $\tilde{\Sigma}^K$. It has been shown by Fan et al. (2017+) that under elliptical distribution, $\hat{\Sigma}^K$ and its top r eigenvalues $\hat{\Lambda}^K$ satisfy (19) and (20) while \tilde{V}^S suffices to conclude (21). Hence Method (4) indeed provides good initial estimators if data are from elliptical distribution. However, since $\hat{\Sigma}^K$ attains (19) for elliptical distribution, by similar argument for deriving Proposition 4 based on our ℓ_∞ perturbation bound, \hat{V}^K consisting of the leading eigenvectors of $\hat{\Sigma}^K$ is also valid for the generic POET procedure. For more details about the two types of Kendall's tau, we refer the readers to Fang et al. (1990); Choi and Marden (1998); Han and Liu (2014); Fan et al. (2017+) and references therein.

In summary, Method (1) is designed for the case of sub-Gaussian data; Method (3) and (4) work under the situation of elliptical distribution; while Method (2) is proposed in this paper for the general heavy-tailed case with bounded fourth moments without further distributional shape constraints.

We simulated n samples of $(f_t^T, u_t^T)^T$ from two settings: (a) a multivariate t-distribution with covariance matrix $\text{diag}\{I_r, 5I_d\}$ and various degrees of freedom ($\nu = 3$ for very heavy tail, $\nu = 5$ for medium heavy tail and $\nu = \infty$ for Gaussian tail), which is one example of the elliptical distribution (Fang et al., 1990); (b) an element-wise iid one-dimensional t distribution with the same covariance matrix and degrees of freedom $\nu = 3, 5$ and ∞ , which is a non-elliptical heavy-tailed distribution.

Each row of coefficient matrix B is independently sampled from a standard normal distribution, so that with high probability, the pervasiveness condition holds with $\|B\|_{\max} = O(\sqrt{\log d})$. The data is then generated by $y_t = Bf_t + u_t$ and the true population covariance matrix is $\Sigma = BB^T + 5I_d$.

For d running from 200 to 900 and $n = d/2$, we calculated errors of the four robust estimators in different norms. The tuning for α in minimization (14) is discussed more thoroughly in Fan et al. (2017b). For the thresholding parameter, we used $\tau = 2\sqrt{\log d/n}$. The estimation errors are gauged in the following norms: $\|\hat{\Sigma}_u^T - \Sigma_u\|$, $\|(\hat{\Sigma}^T)^{-1} - \Sigma^{-1}\|$ and $\|\hat{\Sigma}^T - \Sigma\|_\Sigma$ as shown in Theorem 6. The two different settings are separately plotted in Figures 4 and 5. The estimation errors of applying sample covariance matrix $\hat{\Sigma}^S$ in Method (1) are used as the baseline for comparison. For example, if relative Frobenius norm is used to measure performance, $\|(\hat{\Sigma}^T)^{(k)} - \Sigma\|_\Sigma / \|(\hat{\Sigma}^T)^{(1)} - \Sigma\|_\Sigma$ will be depicted for $k = 2, 3, 4$, where $(\hat{\Sigma}^T)^{(k)}$ are generic POET estimators based on Method (k). Therefore if the ratio curve moves below 1, the method is better than naive sample estimator (Fan et al., 2013) and vice versa. The more it gets below 1, the more robust the procedure is against heavy-tailed randomness.

The first setting (Figure 4) represents a heavy-tailed elliptical distribution, where we expect Methods (2), (3), (4) all outperform the POET estimator based on the sample covariance, i.e. Method (1), especially in the presence of extremely heavy tails (solid lines for $\nu = 3$). As expected, all three curves under various measures show error ratios visibly smaller than 1. On the other hand, if data are indeed Gaussian (dotted line for $\nu = \infty$), Method (1) has better behavior under most measures (error ratios are greater than 1). Nevertheless, our robust Method (2) still performs comparably well with Method (1), whereas the median error ratios for the two Kendall's tau methods are much worse. In addition, the IQR (interquartile range) plots reveal that Method (2) is indeed more stable than two Kendall's tau Methods (3) and (4). It is also noteworthy that Method (4), which




Figure 4: *Error ratios of robust estimates against varying dimension. Blue lines represent errors of Method (2) over Method (1) under different norms; black lines errors of Method (3) over Method (1); red lines errors of Method (4) over Method (1). (f_t^T, u_t^T) is generated by multivariate t -distribution with $df = 3$ (solid), 5 (dashed) and ∞ (dotted). The median errors and their IQR's (interquartile range) over 100 simulations are reported.*

leverages the advantage of spatial Kendall's tau, performs more robustly than Method (3), which solely base its estimation of the eigen-structure on marginal Kendall's tau.

The second setting (Figure 5) provides an example of non-elliptical distributed data. We can see that the performance of the general robust Method (2) dominates the other three methods, which verifies the benefit of robust estimation for a general heavy-tailed distribu-




Figure 5: *Error ratios of robust estimates against varying dimension. Blue lines represent errors of Method (2) over Method (1) under different norms; black lines errors of Method (3) over Method (1); red lines errors of Method (4) over Method (1). (f_t^T, u_t^T) is generated by element-wise iid t -distribution with $df = 3$ (solid), 5 (dashed) and ∞ (dotted). The median errors and their IQR's (interquartile range) over 100 simulations are reported.*

tion. Note that Kendall's tau methods do not apply to distributions outside the elliptical family, excluding even the element-wise iid t distribution in this setting. Nonetheless, even in the first setting where the data are indeed elliptical, with proper tuning, the proposed robust method can still outperform Kendall's tau by a clear margin.

5. Proof organization of main theorems

5.1 Symmetric Case

For shorthand, we write $\tau = \|E\|_\infty$, and $\kappa = \sqrt{d}\|EV\|_{\max}$. An obvious bound for κ is $\kappa \leq \sqrt{r\mu}\tau$ (by Cauchy-Schwarz inequality). We will use these notations throughout this subsection.

Recall the spectral decomposition of A in (8). Expressing E in terms of the column vectors of V and V_\perp , which form an orthogonal basis in \mathbb{R}^n , we write

$$[V, V_\perp]^T E [V, V_\perp] =: \begin{pmatrix} E_{11} & E_{12} \\ E_{21} & E_{22} \end{pmatrix}. \quad (27)$$

Note that $E_{12} = E_{21}^T$ since E is symmetric. Conceptually, the perturbation results in a rotation of $[V, V_\perp]$, and we write a candidate orthogonal basis as follows:

$$\bar{V} := (V + V_\perp Q)(I_r + Q^T Q)^{-1/2}, \quad \bar{V}_\perp := (V_\perp - V Q^T)(I_{d-r} + Q Q^T)^{-1/2}, \quad (28)$$

where $Q \in \mathbb{R}^{(d-r) \times r}$ is to be determined. It is straightforward to check that $[\bar{V}, \bar{V}_\perp]$ is an orthogonal matrix. We will choose Q in a way such that $(\bar{V}, \bar{V}_\perp)^T \tilde{A} (\bar{V}, \bar{V}_\perp)$ is a block diagonal matrix, i.e., $\bar{V}_\perp^T \tilde{A} \bar{V} = 0$. Substituting (28) and simplifying the equation, we obtain

$$Q(\Lambda_1 + E_{11}) - (\Lambda_2 + E_{22})Q = E_{21} - Q E_{12} Q. \quad (29)$$

The approach of studying perturbation through a quadratic equation is known. See Stewart (1990) for example. Yet, to the best of our knowledge, existing results study perturbation under orthogonal-invariant norms (or unitary-invariant norms in the complex case), which includes a family of matrix operator norms and Frobenius norm, but excludes the matrix max-norm. The advantages of orthogonal-invariant norms are pronounced: such norms of a symmetric matrix only depend on its eigenvalues regardless of eigenvectors; moreover, with suitable normalization they are *consistent* in the sense $\|AB\| \leq \|A\| \cdot \|B\|$. See Stewart (1990) for a clear exposition.

The max-norm, however, does not possess these important properties. An imminent issue is that it is not clear how to relate Q to $V_\perp Q$, which will appear in (29) after expanding E according to (27), and which we want to control. Our approach here is to study $\bar{Q} := V_\perp Q$ directly through a transformed quadratic equation, obtained by left multiplying V_\perp to (29). Denote $H = V_\perp E_{21}$, $\bar{Q} = V_\perp Q$, $\bar{L}_1 = \Lambda_1 + E_{11}$, $\bar{L}_2 = V_\perp(\Lambda_2 + E_{22})V_\perp^T$. If we can find an appropriate matrix \bar{Q} with $\bar{Q} = V_\perp Q$, and it satisfies the quadratic equation

$$\bar{Q} \bar{L}_1 - \bar{L}_2 \bar{Q} = H - \bar{Q} H^T \bar{Q}, \quad (30)$$

then Q also satisfies the quadratic equation (29). This is because multiplying both sides of (30) by V_\perp^T yields (29), and thus any solution \bar{Q} to (30) with the form $\bar{Q} = V_\perp Q$ must result in a solution Q to (29).

Once we have such \bar{Q} (or equivalently Q), then $(\bar{V}, \bar{V}_\perp)^T \tilde{A} (\bar{V}, \bar{V}_\perp)$ is a block diagonal matrix, and the span of column vectors of \bar{V} is a candidate space of the span of first r

eigenvectors, namely $\text{span}\{\tilde{v}_1, \dots, \tilde{v}_r\}$. We will verify the two spaces are identical in Lemma 7. Before stating that lemma, we first provide bounds on $\|\bar{Q}\|_{\max}$ and $\|\bar{V} - V\|_{\max}$.

Lemma 5 *Suppose $|\lambda_r| - \varepsilon > 4r\mu(\tau + 2r\kappa)$. Then, there exists a matrix $Q \in \mathbb{R}^{(d-r) \times r}$ such that $\bar{Q} = V_{\perp}Q \in \mathbb{R}^{d \times r}$ is a solution to the quadratic equation (30), and \bar{Q} satisfies $\|\bar{Q}\|_{\max} \leq \omega/\sqrt{d}$. Moreover, if $r\omega < 1/2$, the matrix \bar{V} defined in (28) satisfies*

$$\|\bar{V} - V\|_{\max} \leq 2\sqrt{\mu}\omega r/\sqrt{d}. \quad (31)$$

Here, ω is defined as $\omega = 8(1 + r\mu)\kappa/(|\lambda_r| - \varepsilon)$.

The second claim of the lemma (i.e., the bound (31)) is relatively easy to prove once the first claim (i.e., the bound on $\|\bar{Q}\|_{\max}$) is proved. To understand this, note that we can rewrite \bar{V} as $\bar{V} = (V + \bar{Q})(I_r + \bar{Q}^T\bar{Q})^{-1/2}$, and $\|\bar{Q}^T\bar{Q}\|_{\max}$ can be controlled by a trivial inequality $\|\bar{Q}^T\bar{Q}\|_{\max} \leq d\|\bar{Q}\|_{\max}^2 \leq w^2$. To prove the first claim, we construct a sequence of matrices through recursion that converges to the fixed point \bar{Q} , which is a solution to the quadratic equation (30). For all iterates of matrices, we prove a uniform max-norm bound, which leads to a max-bound on $\|\bar{Q}\|_{\max}$ by continuity. To be specific, we initialize $\bar{Q}^0 = 0$, and given \bar{Q}^t , we solve a linear equation:

$$\bar{Q}\bar{L}_1 - \bar{L}_2\bar{Q} = H - \bar{Q}^t H^T \bar{Q}^t, \quad (32)$$

and the solution is defined as \bar{Q}^{t+1} . Under some conditions, the iterate \bar{Q}^t converges to a limit \bar{Q} , which is a solution to (30). The next general lemma captures this idea. It follows from Stewart (1990) with minor adaptations.

Lemma 6 *Let T be a bounded linear operator on a Banach space \mathcal{B} equipped with a norm $\|\cdot\|$. Assume that T has a bounded inverse, and define $\beta = \|T^{-1}\|^{-1}$. Let $\varphi : \mathcal{B} \rightarrow \mathcal{B}$ be a map that satisfies*

$$\|\varphi(x)\| \leq \eta\|x\|^2, \quad \text{and} \quad \|\varphi(x) - \varphi(y)\| \leq 2\eta \max\{\|x\|, \|y\|\}\|x - y\| \quad (33)$$

for some $\eta \geq 0$. Suppose that \mathcal{B}_0 is a closed subspace of \mathcal{B} such that $T^{-1}(\mathcal{B}_0) \subseteq \mathcal{B}_0$ and $\varphi(\mathcal{B}_0) \subseteq \mathcal{B}_0$. Suppose $y \in \mathcal{B}_0$ that satisfies $4\eta\|y\| < \beta^2$. Then, the sequence initialized with $x_0 = 0$ and iterated through

$$x_{k+1} = T^{-1}(y + \varphi(x_k)), \quad k \geq 0 \quad (34)$$

converges to a solution x^* to $Tx = y + \varphi(x)$. Moreover, we have $x^* \subseteq \mathcal{B}_0$, and $\|x^*\| \leq 2\|y\|/\beta$.

To apply this lemma to the equation (30), we view \mathcal{B} as the space of matrices $\mathbb{R}^{d \times r}$ with the max-norm $\|\cdot\|_{\max}$, and \mathcal{B}_0 as the subspace of matrices of the form $V_{\perp}Q$ where $Q \in \mathbb{R}^{(d-r) \times r}$. The linear operator T is set to be the $T(\bar{Q}) = \bar{Q}\bar{L}_1 - \bar{L}_2\bar{Q}$, and the map φ is set to be the quadratic function $\varphi(\bar{Q}) = -\bar{Q}H^T\bar{Q}$. Roughly speaking, under the assumption of Lemma 6, the nonlinear effect caused by φ is weak compared with the linear operator T . Therefore, it is crucial to show T is invertible, i.e. to give a good lower bound on $\|T^{-1}\|_{\max}^{-1} = \inf_{\|\bar{Q}\|_{\max}=1} \|T(\bar{Q})\|_{\max}$. Since the norm is not orthogonal-invariant, a subtle

issue arises when A is not of exact low rank, which will be discussed at the end of the subsection.

If there is no perturbation (i.e., $E = 0$), all the iterates \bar{Q}^t are simply 0, so \bar{V} is identical to V . If the perturbation is not too large, the next lemma shows that the column vectors of \bar{V} span the same space as $\text{span}\{\tilde{v}_1, \dots, \tilde{v}_r\}$.

In other words, with a suitable orthogonal matrix R , the columns of $\bar{V}R$ are $\tilde{v}_1, \dots, \tilde{v}_r$.

Lemma 7 *Suppose $|\lambda_r| - \varepsilon > \max\{3\tau, 64(1 + r\mu)r^{3/2}\mu^{1/2}\kappa\}$. Then, there exists an orthogonal matrix $R \in \mathbb{R}^{r \times r}$ such that the column vectors of $\bar{V}R$ are $\tilde{v}_1, \dots, \tilde{v}_r$.*

Proof of Theorem 2 It is easy to check that under the assumption of Theorem 2, the conditions required in Lemma 5 and Lemma 7 are satisfied. Hence, the two lemmas imply Theorem 2. \blacksquare

To study the perturbation of individual eigenvectors, we assume, in addition to the condition on $|\lambda_r|$, that $\lambda_1, \dots, \lambda_r$ satisfy a uniform gap, (namely $\delta > \|E\|_2$). This additional assumption is necessary, because otherwise, the perturbation may lead to a change of relative order of eigenvalues, and we may be unable to match eigenvectors from the order of eigenvalues. Suppose $R \in \mathbb{R}^{r \times r}$ is an orthogonal matrix such that $\bar{V}R$ are eigenvectors of \bar{A} . Now, under the assumption of Theorem 2, the column vectors of \tilde{V} and $\bar{V}R$ are identical up to sign, so we can rewrite the difference $\tilde{V} - V$ as

$$\tilde{V} - V = \bar{V}(R - I_r) + (\bar{V} - V). \quad (35)$$

We already provided a bound on $\|\bar{V} - V\|_{\max}$ in Lemma 5. By the triangular inequality, we can derive a bound on $\|\tilde{V}\|_{\max}$. If we can prove a bound on $\|R - I_r\|_{\max}$, it will finally leads to a bound on $\|\tilde{V} - V\|_{\max}$. In order to do so, we use the Davis-Kahan theorem to obtain an bound on $\langle \tilde{v}_i, v_i \rangle$ for all $i \in [r]$. This will lead to a max-norm bound on $R - I_r$ (with the price of potentially increasing the bound by a factor of r). The details about the proof of Theorem 3 are in the appendix.

We remark that the conditions on $|\lambda_r| - \varepsilon$ in Theorem 2 and Theorem 3 are only useful in cases where $|\lambda_r| > \|A - A_r\|_{\infty}$. Ideally, we would like to have results with assumptions only involving λ_r and λ_{r+1} , like in the Davis-Kahan theorem. Unfortunately, unlike orthogonal-invariant norms that only depend on the eigenvalues of a matrix, the max-norm $\|\cdot\|_{\max}$ is not orthogonal-invariant, and thus it also depends on the eigenvectors of a matrix. For this reason, it is not clear whether we could obtain a lower bound on $\|T^{-1}\|_{\max}^{-1}$ using only the eigenvalues λ_r and λ_{r+1} so that Lemma 6 could be applied. The analysis appears to be difficult if we do not have a bound on $\|T^{-1}\|_{\max}^{-1}$, considering that even in the analysis of linear equations, we need invertibility and a well-controlled condition number.

5.2 Asymmetric case

Let A^d, E^d be $d_1 + d_2$ square matrices defined as

$$A^d := \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}, \quad E^d := \begin{pmatrix} 0 & E \\ E^T & 0 \end{pmatrix}.$$

Also denote $\tilde{A}^d := A^d + E^d$. This augmentation of an asymmetric matrix into a symmetric one is called Hermitian dilation. Here the superscript d means the Hermitian dilation. We also use this notation to denote quantities corresponding to A^d and \tilde{A}^d .

An important observation is that

$$\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} \begin{pmatrix} u_i \\ \pm v_i \end{pmatrix} = \pm \sigma_i \begin{pmatrix} u_i \\ \pm v_i \end{pmatrix}.$$

From this identity, we know that A^d have nonzero eigenvalues $\pm \sigma_i$ where $1 \leq i \leq \text{rank}(A)$, and its corresponding eigenvectors are $(u_i^T, \pm v_i^T)^T$. For a given r , we stack these (normalized) eigenvectors with indices $i \in [r]$ into a matrix $V^d \in \mathbb{R}^{(d_1+d_2) \times 2r}$:

$$V^d := \frac{1}{\sqrt{2}} \begin{pmatrix} U & U \\ V & -V \end{pmatrix}.$$

Through the augmented matrices, we can transfer eigenvector results for symmetric matrices to singular vectors of asymmetric matrices. However, we cannot directly invoke the results proved for symmetric matrices, due to an issue about the coherence of V^d : when d_1 and d_2 are not comparable, the coherence $\mu(V^d)$ can be very large even when $\mu(V)$ and $\mu(U)$ are bounded. To understand this, consider the case where $r = 1$, $d_1 \gg d_2$, and all entries of U are $O(1/\sqrt{d_1})$, and all entries of V are $O(1/\sqrt{d_2})$. Then, the coherences $\mu(U)$ and $\mu(V)$ are $O(1)$, but $\mu(V^d) = O((d_1 + d_2)/d_2) \gg 1$.

This unpleasant issue about the coherence, nevertheless, can be tackled if we consider a different matrix norm. In order to deal with the different scales of d_1 and d_2 , we define the weighted max-norm for any matrix M with $d_1 + d_2$ rows as follows:

$$\|M\|_w := \left\| \begin{pmatrix} \sqrt{d_1} I_{d_1} & 0 \\ 0 & \sqrt{d_2} I_{d_2} \end{pmatrix} M \right\|_{\max}. \quad (36)$$

In other words, we rescale the top d_1 rows of M by a factor of $\sqrt{d_1}$, and rescale the bottom d_2 rows by $\sqrt{d_2}$. This weighted norm serves to balance the potential different scales of d_1 and d_2 .

The proofs of theorems in Section 2.2 will be almost the same with those in the symmetric case, with the major difference being the new matrix norm. Because the derivation is slightly repetitive, we will provide concise proofs in the appendix. Similar to the decomposition in (2.1),

$$A^d = \begin{pmatrix} 0 & A_r \\ A_r^T & 0 \end{pmatrix} + \begin{pmatrix} 0 & A - A_r \\ A^T - A_r^T & 0 \end{pmatrix} =: A_r^d + (A^d - A_r^d),$$

where A_r^d has rank $2r$. Equivalently,

$$A_r^d = \sum_{i=1}^r \sigma_i (u_i^T, v_i^T)^T (u_i^T, v_i^T) - \sum_{i=1}^r \sigma_i (u_i^T, -v_i^T)^T (u_i^T, -v_i^T).$$

Analogously, we define notations in (28)–(30) and use d in the superscript to signify that they are augmented through Hermitian dilation. It is worthwhile to note that $\Lambda_1^d = \text{diag}\{\sigma_1, \dots, \sigma_r, -\sigma_r, \dots, -\sigma_1\}$, and that $\min\{|\pm \sigma_i| : i \in [r]\} = \sigma_r$ (a simi-

lar quantity as $|\lambda_r|$). Recall $\mu_0 = \mu(U) \vee \mu(V)$, $\tau_0 = \sqrt{d_1/d_2}\|E\|_\infty \vee \sqrt{d_2/d_1}\|E\|_1$ and $\varepsilon_0 = \sqrt{d_1/d_2}\|A - A_r\|_\infty \vee \sqrt{d_2/d_1}\|A - A_r\|_1$. In the proof, we will also use $\kappa_0 = \max\{\sqrt{d_1}\|EV\|_{\max}, \sqrt{d_2}\|E^T U\|_{\max}\}$, which is a quantity similar to κ .

The next key lemma, which is parallel to Lemma 5, provides a bound on the solution \bar{Q}^d to the quadratic equation

$$\bar{Q}^d \bar{L}_1^d - \bar{L}_2^d \bar{Q}^d = H^d - \bar{Q}^d (H^d)^T \bar{Q}^d. \quad (37)$$

Lemma 8 *Suppose $\sigma_r - \varepsilon_0 > 16r\mu_0(\tau_0 + r\kappa_0)$. Then, there exists a matrix $Q^d \in \mathbb{R}^{(d_1+d_2-2r) \times 2r}$ such that $\bar{Q}^d = V_\perp^d Q^d \in \mathbb{R}^{(d_1+d_2) \times 2r}$ is a solution to the quadratic equation (37), and \bar{Q}^d satisfies $\|\bar{Q}^d\|_w \leq \omega_0$. Moreover, if $r\omega_0 < 1/2$, the matrix \bar{V}^d defined in (28) satisfies*

$$\|\bar{V}^d - V^d\|_w \leq 6\sqrt{\mu_0} r\omega_0. \quad (38)$$

Here, ω_0 is defined as $\omega_0 = 8(1 + r\mu_0)\kappa_0/3(\sigma_r - \varepsilon_0)$.

In this lemma, the bound (38) bears a similar form to (31): if we consider the max-norm, the first d_1 rows of $\bar{V}^d - V^d$ correspond to the left singular vectors u_i 's, and they scale with $1/\sqrt{d_1}$; and the last d_2 rows correspond to the right singular vectors v_i 's, which scale with $1/\sqrt{d_2}$. Clearly, the weighted max-norm $\|\cdot\|_w$ indeed helps to balance the two dimensions.

The rest of the proofs can be found in the appendix.

Acknowledgments

We would like to acknowledge support for this project from the following grants: NSF grants DMS-1406266, DMS-1662139, DMS-1712591 and NIH grant R01-GM072611-13.

A. Proofs for Section 2.1

Denote the column span of a matrix M by $\text{span}(M)$. Suppose two matrices $M_1, M_2 \in \mathbb{R}^{n \times m}$ ($m \leq n$) have orthonormal column vectors. It is known that (Stewart, 1990)

$$d(M_1, M_2) := \|M_1 M_1^T - M_2 M_2^T\|_2 = \|\sin \Theta(M_1, M_2)\|_2. \quad (39)$$

where $\Theta(M_1, M_2)$ are the canonical angles between $\text{span}(M_1)$ and $\text{span}(M_2)$. Recall the notations defined in (27), and also recall $\kappa = \sqrt{d}\|EV\|_{\max}$, $\Lambda_1 = \text{diag}\{\lambda_1, \dots, \lambda_r\}$, $\Lambda_2 = \text{diag}\{\lambda_{r+1}, \dots, \lambda_n\}$, $\bar{L}_1 = \Lambda_1 + E_{11}$, $\bar{L}_2 = V_{\perp}(\Lambda_2 + E_{22})V_{\perp}^T$ and $H = V_{\perp}E_{21}$. The first lemma bounds $\|H\|_{\max}$.

Lemma 9 *We have the following bound on $\|H\|_{\max}$:*

$$\|H\|_{\max} \leq (1 + r\mu)\kappa/\sqrt{d}.$$

Proof Using the definition $E_{21} = V_{\perp}^T EV$ in (27), we can write $H = V_{\perp}V_{\perp}^T EV$. Since the columns of V and V_{\perp} form an orthogonal basis in \mathbb{R}^d , clearly

$$VV^T + V_{\perp}V_{\perp}^T = I_d. \quad (40)$$

By Cauchy-Schwarz inequality and the definition of μ , for any $i, j \in [d]$,

$$|(VV^T)_{ij}| = \sum_{k=1}^r |V_{ik}V_{jk}| \leq \left(\sum_{k=1}^r V_{ik}^2\right)^{1/2} \cdot \left(\sum_{k=1}^r V_{jk}^2\right)^{1/2} \leq \frac{r\mu}{d}.$$

Using the identity (40) and the above inequality, we derive

$$\begin{aligned} \|H\|_{\max} &\leq \|EV\|_{\max} + \|VV^T EV\|_{\max} \\ &\leq (1 + d\|VV^T\|_{\max})\|EV\|_{\max} \leq (1 + r\mu)\|EV\|_{\max}, \end{aligned}$$

which completes the proof. ■

Lemma 10 *If $|\lambda_r| > \kappa r\sqrt{\mu}$, then \bar{L}_1 is an invertible matrix. Furthermore,*

$$\inf_{\|Q_0\|_{\max}=1} \|Q_0\bar{L}_1 - \bar{L}_2 Q_0\|_{\max} \geq |\lambda_r| - 3r\mu(\tau + r\kappa) - \varepsilon, \quad (41)$$

where Q_0 is an $d \times r$ matrix.

Proof Let Q_0 be any $d \times r$ matrix with $\|Q_0\|_{\max} = 1$. Note

$$Q_0\bar{L}_1 - \bar{L}_2 Q_0 = Q_0\Lambda_1 + Q_0E_{11} - \bar{L}_2 Q_0.$$

We will derive upper bounds on Q_0E_{11} and $\bar{L}_2 Q_0$, and a lower bound on $Q_0\Lambda_1$. Since $E_{11} = V^T EV$ by definition, we expand Q_0E_{11} and use a trivial inequality to derive

$$\|Q_0E_{11}\|_{\max} \leq d\|Q_0V^T\|_{\max}\|EV\|_{\max}. \quad (42)$$

By Cauchy-Schwarz inequality and the definition of μ in (3), for $i, j \in [d]$,

$$|(Q_0 V^T)_{ij}| \leq \sum_{k=1}^r |(Q_0)_{ik} V_{jk}| \leq \left(\sum_{k=1}^r (Q_0)_{ik}^2 \right)^{1/2} \left(\sum_{k=1}^r V_{jk}^2 \right)^{1/2} \leq \sqrt{r} \cdot \sqrt{\frac{r\mu}{d}},$$

Substituting $\|EV\|_{\max} = \kappa/\sqrt{d}$ into (42), we obtain an upper bound:

$$\|Q_0 E_{11}\|_{\max} \leq \kappa r \sqrt{\mu}. \quad (43)$$

To bound $\bar{L}_2 Q_0 = (V_{\perp} E_{22} V_{\perp}^T + (A - A_r)) Q_0$, we use the identity (40) and write

$$V_{\perp} E_{22} V_{\perp}^T Q_0 = V_{\perp} V_{\perp}^T E V_{\perp} V_{\perp}^T Q_0 = (I_d - V V^T) E (I_d - V V^T) Q_0.$$

Using two trivial inequalities $\|EQ_0\|_{\max} \leq \|E\|_{\infty} \|Q_0\|_{\max} = \|E\|_{\infty}$ and $\|V^T Q_0\|_{\max} \leq \|V^T\|_{\infty} \|Q_0\|_{\max} \leq \sqrt{d}$, we have

$$\begin{aligned} \|E(I_d - V V^T) Q_0\|_{\max} &\leq \|EQ_0\|_{\max} + r \|EV\|_{\max} \|V^T Q_0\|_{\max} \\ &\leq \|E\|_{\infty} + r \sqrt{d} \|EV\|_{\max} = \tau + r\kappa. \end{aligned}$$

In the proof of Lemma 9, we showed $\|V V^T\|_{\max} \leq r\mu/d$. Thus,

$$\|V_{\perp} E_{22} V_{\perp}^T Q_0\|_{\max} \leq (1 + d \|V V^T\|_{\max}) \cdot \|E(I_d - V V^T) Q_0\|_{\max} \leq (1 + r\mu)(\tau + r\kappa).$$

Moreover, $\|(A - A_r) Q_0\|_{\max} \leq \|A - A_r\|_{\infty} \|Q_0\|_{\max} = \varepsilon$. Combining the two bounds,

$$\|\bar{L}_2 Q_0\|_{\max} \leq (1 + r\mu)(\tau + r\kappa) + \varepsilon. \quad (44)$$

It is straightforward to obtain a lower bound on $\|Q_0 \Lambda_1\|_{\max}$: since there is an entry of Q_0 , say $(Q_0)_{ij}$, that has an absolute value of 1, we have

$$\|Q_0 \Lambda_1\|_{\max} \geq |(Q_0)_{ij} \lambda_j| \geq |\lambda_r|. \quad (45)$$

To show \bar{L}_1 is invertible, we use (42) and (45) to obtain

$$\|Q_0 \bar{L}_1\|_{\max} \geq \|Q_0 \Lambda_1\|_{\max} - \|Q_0 E_{11}\|_{\max} \geq |\lambda_r| - \kappa r \sqrt{\mu}.$$

When $|\lambda_r| - \kappa r \sqrt{\mu} > 0$, \bar{L}_1 must have full rank, because otherwise we can choose an appropriate Q_0 in the null space of \bar{L}_1^T so that $Q_0 \bar{L}_1 = 0$, which is a contradiction. To prove the second claim of the lemma, we combine the lower bound (45) with upper bounds (43) and (44) to derive

$$\begin{aligned} \|Q_0 \bar{L}_1 - \bar{L}_2 Q_0\|_{\max} &\geq \|Q_0 \Lambda_1\|_{\max} - \|Q_0 E_{11}\|_{\max} - \|\bar{L}_2 Q_0\|_{\max} \\ &\geq |\lambda_r| - \kappa r \sqrt{\mu} - (1 + r\mu)(\tau + r\kappa) - \varepsilon \\ &\geq |\lambda_r| - 3r\mu(\tau + r\kappa) - \varepsilon, \end{aligned}$$

which is exactly the desired inequality. ■

Next we prove Lemma 6. This lemma follows from Stewart (1990), with minor changes that involves \mathcal{B}_0 . We provide a proof for the sake of completeness.

Proof of Lemma 6 Let us write $\alpha = \|y\|$ for shorthand and recall $\beta = \|T^{-1}\|^{-1}$. As the first step, we show that the sequence $\{x_k\}_{k=0}^\infty$ is bounded. By construction in (34), we bound $\|x_{k+1}\|$ using $\|x_k\|$:

$$\|x_{k+1}\| \leq \|T^{-1}\|(\|y\| + \|\varphi(x_k)\|) \leq \frac{\alpha}{\beta} + \frac{\eta}{\beta}\|x_k\|^2.$$

We use this inequality to derive an upper bound on $\{x_k\}$ for all k . We define $\xi_0 = 0$ and

$$\xi_{k+1} = \frac{\alpha}{\beta} + \frac{\eta}{\beta}\xi_k^2, \quad k \geq 0,$$

then clearly $\|x_k\| \leq \xi_k$ (which can be shown by induction). It is easy to check (by induction) that the sequence $\{\xi_k\}_{k=1}^\infty$ is increasing. Moreover, since $4\alpha\eta < \beta^2$, the quadratic function

$$\phi(\xi) = \frac{\alpha}{\beta} + \frac{\eta}{\beta}\xi^2,$$

has two fixed points (namely solutions to $\phi(\xi) = \xi$), and the smaller one satisfies

$$\xi_\star = \frac{2\alpha}{\beta + \sqrt{\beta^2 - 4\eta\alpha}} < \frac{2\alpha}{\beta}.$$

If $\xi_k < \xi_\star$, then $\xi_{k+1} = \phi(\xi_k) \leq \phi(\xi_\star) = \xi_\star$. Thus, by induction, all ξ_k are bounded by ξ_\star . This implies $\|x_k\| \leq \xi_\star < 2\alpha/\beta$. The next step is to show that the sequence $\{x_k\}$ converges. Using the recursive definition (34) again, we derive

$$\begin{aligned} \|x_{k+1} - x_k\| &\leq \|T^{-1}\| \|\varphi(x_k) - \varphi(x_{k-1})\| \\ &\leq 2\beta^{-1}\eta \max\{\|x_k\|, \|x_{k-1}\|\} \|x_k - x_{k-1}\| \\ &\leq \frac{4\alpha\eta}{\beta^2} \|x_k - x_{k-1}\|. \end{aligned}$$

Since $4\alpha\eta/\beta^2 < 1$, the sequence $\{x_k\}_{k=0}^\infty$ is a Cauchy sequence, and convergence is secured. Let $x^\star \in \mathcal{B}$ be the limit. It is clear by assumption that $x_k \in \mathcal{B}_0$ implies $x_{k+1} \in \mathcal{B}_0$, so $x^\star \in \mathcal{B}_0$ and $\|x^\star\| \leq 2\alpha/\beta$ by continuity.

The final step is to show x^\star is a solution to $Tx = y + \phi(x)$. Because $\{x_k\}_{k=0}^\infty$ is bounded and ϕ satisfies (33), the sequence $\{\phi(x_k)\}_{k=0}^\infty$ converges to $\phi(x^\star)$ by continuity and compactness. The linear operator T is also continuous, so we can take limits on both sides of $Tx_{k+1} = y + \phi(x_k)$, we conclude that x^\star is a solution to $Tx = y + \phi(x)$. \blacksquare

With all the preparations, we are now ready to present the key lemma. As discussed in Section 5, we set

$$\mathcal{B}_0 := \{\bar{Q} \in \mathbb{R}^{d \times r} : \bar{Q} = V_\perp Q \text{ for some } Q \in \mathbb{R}^{(d-r) \times r}\}.$$

which is a subspace of $\mathcal{B} = \mathbb{R}^{d \times r}$. Consider the matrix max-norm $\|\cdot\|_{\max}$ in \mathcal{B} .

Lemma 11 *Suppose $|\lambda_r| - \varepsilon > 4r\mu(\tau + 2r\kappa)$. Then there exists a solution $\bar{Q} \in \mathcal{B}_0$ to the equation (30) with*

$$\|\bar{Q}\|_{\max} \leq \frac{8(1+r\mu)\kappa}{(|\lambda_r| - \varepsilon)\sqrt{d}}.$$

Proof We will invoke Lemma 6 and apply it to the quadratic equation (30). To do so, we first check the conditions required in Lemma 6.

Let the linear operator \mathcal{T} be $\mathcal{T}\bar{Q} = \bar{Q}\bar{L}_1 - \bar{L}_2\bar{Q}$. By Lemma 10, \mathcal{T} has a bounded inverse, and $\beta := \|\mathcal{T}^{-1}\|_{\max}^{-1}$ is bounded from below:

$$\beta \geq |\lambda_r| - 3r\mu(\tau + r\kappa) - \varepsilon. \quad (46)$$

Let us define φ by $\varphi(\bar{Q}) = \bar{Q}H^T\bar{Q}$. To check the inequalities in (33), observe that

$$\|\varphi(\bar{Q})\|_{\max} \leq rd\|\bar{Q}\|_{\max}\|H\|_{\max}\|\bar{Q}\|_{\max} \leq (1+r\mu)\kappa r\sqrt{d}\|\bar{Q}\|_{\max}^2$$

where we used Lemma 9. We also observe

$$\begin{aligned} \|\varphi(\bar{Q}_1) - \varphi(\bar{Q}_2)\|_{\max} &= \|\bar{Q}_1H^T(\bar{Q}_1 - \bar{Q}_2) + (\bar{Q}_1 - \bar{Q}_2)H^T\bar{Q}_2\|_{\max} \\ &\leq rd\|\bar{Q}_1\|_{\max}\|H\|_{\max}\|\bar{Q}_1 - \bar{Q}_2\|_{\max} + rd\|\bar{Q}_1 - \bar{Q}_2\|_{\max}\|H\|_{\max}\|\bar{Q}_2\|_{\max} \\ &\leq 2(1+r\mu)\kappa r\sqrt{d} \max\{\|\bar{Q}_1\|_{\max}, \|\bar{Q}_2\|_{\max}\}\|\bar{Q}_1 - \bar{Q}_2\|_{\max}. \end{aligned}$$

Thus, if we set $\eta = (1+r\mu)\kappa r\sqrt{d}$, then inequalities required in (33) are satisfied. For any \bar{Q} with $\bar{Q} = V_{\perp}Q \in \mathcal{B}_0$, obviously $\varphi(\bar{Q}) = V_{\perp}QH^T\bar{Q} \in \mathcal{B}_0$. To show $\mathcal{T}^{-1}(\bar{Q}) \in \mathcal{B}_0$, let $Q_0 = \mathcal{T}^{-1}(\bar{Q})$ and observe that

$$Q_0\bar{L}_1 - \bar{L}_2Q_0 = \bar{Q} \in \mathcal{B}_0.$$

By definition, we know $\bar{L}_2Q_0 = V_{\perp}(E_{22} + \Lambda_2)V_{\perp}^TQ_0 \in \mathcal{B}_0$, so we deduce $Q_0\bar{L}_1 \in \mathcal{B}_0$. Our assumption implies $|\lambda_r| > \kappa r\sqrt{\mu}$, so by Lemma 10, the matrix \bar{L}_1 is invertible, and thus $Q_0 \in \mathcal{B}_0$. The last condition we check is $4\eta\|H\|_{\max} < \beta^2$. By Lemma 9 and (46), this is true if

$$4(1+r\mu)^2\kappa^2r < [|\lambda_r| - 3r\mu(\tau + r\kappa) - \varepsilon]^2.$$

The above inequality holds when $|\lambda_r| > 4r\mu(V)(\tau + 2r\kappa) + \varepsilon$. Under this condition, we have, by Lemma 6,

$$\|\bar{Q}\|_{\max} \leq \frac{2(1+r\mu)\kappa}{(|\lambda_r| - 3r\mu(\tau + r\kappa) - \varepsilon)\sqrt{d}} \leq \frac{8(1+r\mu)\kappa}{(|\lambda_r| - \varepsilon)\sqrt{d}},$$

where, the second inequality is due to $3r\mu(\tau + r\kappa) \leq 3(|\lambda_r| - \varepsilon)/4$. ■

The next lemma is a consequence of Lemma 11. We define, as in Lemma 5, that $\omega = 8(1+r\mu)\kappa/(|\lambda_r| - \varepsilon)$.

Lemma 12 *If $r\omega^2 < 1/2$, then*

$$\|(I_r + \overline{Q}^T \overline{Q})^{-1/2} - I_r\|_{\max} \leq r\omega^2, \quad \|(I_r + \overline{Q}^T \overline{Q})^{-1/2}\|_{\max} \leq \frac{3}{2}.$$

Proof By the triangular inequality, the second inequality is immediate from the first one. To prove the first inequality, suppose the spectral decomposition of $\overline{Q}^T \overline{Q}$ is $\overline{Q}^T \overline{Q} = \overline{U} \overline{\Sigma} \overline{U}^T$, where $\overline{\Sigma} = \text{diag}\{\overline{\lambda}_1, \dots, \overline{\lambda}_r\}$ where $\overline{\lambda}_1 \geq \dots \geq \overline{\lambda}_r$, and $\overline{U} = [\overline{u}_1, \overline{u}_2, \dots, \overline{u}_r]$ where $\overline{u}_1, \dots, \overline{u}_r$ are orthonormal vectors in \mathbb{R}^r . Since $\overline{Q}^T \overline{Q}$ has nonnegative eigenvalues, we have $\overline{\lambda}_r \geq 0$. Using these notations, we can rewrite the matrix as

$$(I_r + \overline{Q}^T \overline{Q})^{-1/2} - I_r = \sum_{i=1}^r ((1 + \overline{\lambda}_i)^{-1/2} - 1) \overline{u}_i \overline{u}_i^T.$$

Note that $\overline{\lambda}_1 \leq \|\overline{Q}^T \overline{Q}\| \leq rd \|\overline{Q}\|_{\max}^2 \leq r\omega^2$, which implies $\overline{\lambda}_1 < 1/2$. It is easy to check that $1 + |x| \geq (1+x)^{-1/2} \geq 1 - |x|$ whenever $|x| < 1/2$. From this fact, we know $|(1 + \overline{\lambda}_i)^{-1/2} - 1| \leq \overline{\lambda}_i \leq r\omega^2$. Using Cauchy-Schwarz inequality, we deduce that for any $j, k \in [d]$,

$$\begin{aligned} |[(I_r + \overline{Q}^T \overline{Q})^{-1/2} - I_r]_{jk}| &\leq \sum_{i=1}^r |(1 + \overline{\lambda}_i)^{-1/2} - 1| \cdot |\overline{U}_{ji} \overline{U}_{ki}| \\ &\leq r\omega^2 \cdot \left(\sum_{i=1}^r \overline{U}_{ji}^2 \right)^{1/2} \left(\sum_{i=1}^r \overline{U}_{ki}^2 \right)^{1/2} \\ &\leq r\omega^2. \end{aligned}$$

This leads to the desired max-norm bound. ■

Proof of Lemma 5 The first claim of the lemma (the existence of \overline{Q} and its max-norm bound) follows directly from Lemma 11. To prove the second claim, we split $\overline{V} - V$ into two parts:

$$\begin{aligned} \overline{V} - V &= V \left((I_r + Q^T Q)^{-1/2} - I_r \right) + V_{\perp} Q (I_r + Q^T Q)^{-1/2} \\ &= V \left((I_r + \overline{Q}^T \overline{Q})^{-1/2} - I_r \right) + \overline{Q} (I_r + \overline{Q}^T \overline{Q})^{-1/2}, \end{aligned} \quad (47)$$

where we used identity $V_{\perp}^T V_{\perp} = I_{d-r}$. Note that $r\omega < 1/2$ implies $r\omega^2 = (r\omega)^2/r < 1/(4r) < 1/2$. Thus, we can use Lemma 12 and derive

$$\left\| V \left((I_r + \overline{Q}^T \overline{Q})^{-1/2} - I_r \right) \right\|_{\max} \leq \sqrt{\frac{r^2 \mu}{d}} \|(I_r + \overline{Q}^T \overline{Q})^{-1/2} - I_r\|_{\max} \leq \sqrt{\frac{\mu}{d}} r^2 \omega^2.$$

where we used Cauchy-Schwarz inequality. Using the above inequality and the bound on $\|Q\|_{\max}$ (namely, the first claim in the lemma),

$$\begin{aligned}\|\bar{V} - V\|_{\max} &\leq \sqrt{\frac{\mu}{d}} r^2 \omega^2 + r \|\bar{Q}\|_{\max} \|(I_r + \bar{Q}^T \bar{Q})^{-1/2}\|_{\max} \\ &\leq (\sqrt{\mu} \omega^2 r^2 + 3\omega r/2)/\sqrt{d}.\end{aligned}$$

Simplifying the bound using $r\omega \leq 1/2$ and a trivial bound $\mu \geq 1$, we obtain (31). \blacksquare

Proof of Lemma 7 Using the identity in (39), it follows from Davis-Kahan sin Θ theorem (Davis and Kahan, 1970) and Weyl's inequality that

$$d(\tilde{V}, V) \leq \frac{\|E\|_2}{\delta_r - \|E\|_2},$$

when $\delta_r > \|E\|_2$, where $\delta_r = |\lambda_r| - |\lambda_{r+1}|$. Since $\lambda_{r+1} \leq \|A - A_r\|_2 \leq \varepsilon$ and $\|E\|_2 \leq \tau$, the condition $|\lambda_r| - \varepsilon > 3\tau$ implies $\delta_r > 3\|E\|_2$. Hence, we have $d(\tilde{V}, V) < 1/2$. Moreover,

$$\begin{aligned}d(\bar{V}, V) &= \|\bar{V}\bar{V}^T - VV^T\|_2 \leq \|\bar{V}(\bar{V} - V)^T\|_2 + \|(\bar{V} - V)V^T\|_2 \\ &\leq 2\|\bar{V} - V\|_2 \leq 2\sqrt{rd}\|\bar{V} - V\|_{\max} \\ &\leq 4r^{3/2}\sqrt{\mu}\omega,\end{aligned}$$

where we used a trivial inequality $\|M\|_2 \leq \|M\|_F \leq \sqrt{rd}\|M\|_{\max}$ for any $M \in \mathbb{R}^{d \times r}$. Under the condition $|\lambda_r| - \varepsilon > 64(1+r\mu)r^{3/2}\mu^{1/2}\kappa$, it is easy to check that $4r^{3/2}\sqrt{\mu}\omega \leq 1/2$. Thus, we obtain $d(\bar{V}, V) < 1/2$. By the triangular inequality,

$$d(\tilde{V}, \bar{V}) \leq d(\tilde{V}, V) + d(\bar{V}, V) < 1.$$

Since $(\bar{V}, \bar{V}_\perp)^T \tilde{A}(\bar{V}, \bar{V}_\perp)$ is a block diagonal matrix, $\text{span}(\bar{V})$ is the same as the subspace spanned by r eigenvectors of \tilde{A} . We claim that $\text{span}(\bar{V}) = \text{span}(\tilde{v}_1, \dots, \tilde{v}_r)$. Otherwise, there exists an eigenvector $u \in \text{span}(\bar{V})$ whose associated eigenvalue is distinct from $\tilde{\lambda}_1, \dots, \tilde{\lambda}_r$ (since $\delta_r > 3\|E\|_2$), and thus u is orthogonal to $\tilde{v}_1, \dots, \tilde{v}_r$. Therefore,

$$\|(\tilde{V}\tilde{V}^T - \bar{V}\bar{V}^T)u\|_2 = \|\bar{V}\bar{V}^T u\|_2 = \|u\|_2.$$

This implies $d(\tilde{V}, \bar{V}) \geq 1$, which is a contradiction. \blacksquare

Proof of Theorem 3 We split $\tilde{V} - V$ into two parts—see (35). In the following, we first obtain a bound on $\|R - I_r\|_{\max}$, which then results in a bound on $\|\tilde{V} - V\|_{\max}$.

Under the assumption of the theorem, $r\omega < 1/2$, so

$$\|\bar{V}\|_{\max} \leq \|\bar{V} - V\|_{\max} + \|V\|_{\max} \leq (2\sqrt{\mu}r\omega)/\sqrt{d} + \sqrt{r\mu}/\sqrt{d} \leq 2\sqrt{r\mu/d}. \quad (48)$$

To bound $\|R - I_r\|_{\max}$, we rewrite R as $R = \bar{V}^T \bar{V} R = \bar{V}^T \tilde{V}$. Expand \bar{V} according to (28),

$$R = (I_r + \bar{Q}^T \bar{Q})^{-1/2} (V + \bar{Q})^T \tilde{V}.$$

Let us make a few observations: (a) $\|\bar{Q}^T \tilde{V}\|_{\max} \leq \sqrt{d} \|\bar{Q}\|_{\max} \leq \omega$ by Cauchy-Schwarz inequality; (b) $\|\tilde{V}^T V\|_{\max} \leq 1$ by Cauchy-Schwarz inequality again; and (c) $\|(I_r + \bar{Q}^T \bar{Q})^{-1/2} - I_r\|_{\max} \leq r\omega^2$ by Lemma 12. Using these inequalities, we have

$$\begin{aligned} \|R - (V + \bar{Q})^T \tilde{V}\|_{\max} &\leq r \|(I_r + \bar{Q}^T \bar{Q})^{-1/2} - I_r\|_{\max} \|(V + \bar{Q})^T \tilde{V}\|_{\max} \\ &\leq r^2 \omega^2 (1 + \omega). \end{aligned} \quad (49)$$

Furthermore, by Davis-Kahn $\sin \Theta$ theorem (Davis and Kahan, 1970) and Weyl's inequality, for any $i \in [r]$,

$$\sin \theta(v_i, \tilde{v}_i) = \sqrt{1 - \langle v_i, \tilde{v}_i \rangle^2} \leq \frac{\|E\|_2}{\delta - \|E\|_2}. \quad (50)$$

when $\delta > \|E\|_2$ (δ is defined in Theorem 3). This leads to the bound $\sin \theta(v_i, \tilde{v}_i) \leq 2\|E\|_2/\delta$ (which is a simplified bound). This is because when $\delta \geq 2\|E\|_2$, the bound is implied by (50); when $\delta < 2\|E\|_2$, the bound trivially follows from $\sin \theta(v_i, \tilde{v}_i) \leq 1$. We obtain, up to sign, for $i \leq r$,

$$\sqrt{1 - \langle v_i, \tilde{v}_i \rangle} \leq \sqrt{1 - \langle v_i, \tilde{v}_i \rangle^2} \leq \frac{2\|E\|_2}{\delta}. \quad (51)$$

In other words, each diagonal entry of $I_r - V^T \tilde{V}$, namely $1 - \langle v_i, \tilde{v}_i \rangle$, is bounded by $4\|E\|_2^2/\delta^2$. Since $\{\tilde{v}_i\}_{i=1}^r$ are orthonormal vectors, we have $1 - \langle v_i, \tilde{v}_i \rangle^2 \geq \sum_{i' \neq i} \langle v_i, \tilde{v}_{i'} \rangle^2 \geq \langle v_i, \tilde{v}_j \rangle^2$ for any $i \neq j$, which leads to bounds on off-diagonal entries of $V^T \tilde{V} - I_r$. We will combine the two bounds. Note that when $\delta \geq 2\|E\|_2$,

$$\|V^T \tilde{V} - I_r\|_{\max} \leq \max \left\{ \frac{4\|E\|_2^2}{\delta^2}, \frac{2\|E\|_2}{\delta} \right\} = \frac{2\|E\|_2}{\delta};$$

and when $\delta < 2\|E\|_2$, $\|V^T \tilde{V} - I_r\|_{\max}$ is trivially bounded by 1 (up to sign), which is trivially bounded by $2\|E\|_2/\delta$. In either case, we deduce

$$\|V^T \tilde{V} - I_r\|_{\max} \leq \frac{2\|E\|_2}{\delta}. \quad (52)$$

Using the bounds in (49) and (52) and $\|\bar{Q}^T \tilde{V}\|_{\max} \leq \omega$, we obtain

$$\begin{aligned} \|R - I_r\|_{\max} &\leq \|R - (V + \bar{Q})^T \tilde{V}\|_{\max} + \|V^T \tilde{V} - I_r\|_{\max} + \|\bar{Q}^T \tilde{V}\|_{\max} \\ &\leq r^2 \omega^2 (1 + \omega) + \frac{2\|E\|_2}{\delta} + \omega. \end{aligned}$$

We use the inequality $r\omega < 1/2$ to simplify the above bound:

$$\begin{aligned} \|R - I_r\|_{\max} &\leq r^2\omega^2(1 + \omega) + \omega + 2\|E\|_2/\delta \leq \left(\frac{1}{2} + \frac{1}{4} + 1\right)r\omega + 2\|E\|_2/\delta \\ &\leq 2r\omega + 2\|E\|_2/\delta. \end{aligned} \quad (53)$$

We are now ready to bound $\|\tilde{V} - V\|_{\max}$. In (35), we use the bounds (48), (53), (31) to obtain

$$\begin{aligned} \|\tilde{V} - V\|_{\max} &= \|\bar{V}(R - I_r) + (\bar{V} - V)\|_{\max} \leq r\|\bar{V}\|_{\max}\|R - I_r\|_{\max} + \|\bar{V} - V\|_{\max} \\ &\leq 2r\sqrt{r\mu/d}(2r\omega + 2\|E\|_2/\delta) + 2r\sqrt{\mu}\omega/\sqrt{d} \\ &\leq \frac{(4r^{5/2}\mu^{1/2} + 2r\mu^{1/2})\omega}{\sqrt{d}} + \frac{4r^{3/2}\mu^{1/2}\|E\|_2}{\delta\sqrt{d}} \\ &\leq \frac{48(1 + r\mu)r^{5/2}\mu^{1/2}\kappa}{(|\lambda_r| - \varepsilon)\sqrt{d}} + \frac{4r^{3/2}\mu^{1/2}\|E\|_2}{\delta\sqrt{d}}. \end{aligned}$$

Using a trivial inequality $\kappa \leq \sqrt{r\mu}\tau$, the above bound leads to

$$\|\tilde{V} - V\|_{\max} = O\left(\frac{r^4\mu^2\tau}{(|\lambda_r| - \varepsilon)\sqrt{d}} + \frac{r^{3/2}\mu^{1/2}\|E\|_2}{\delta\sqrt{d}}\right).$$

■

B. Proofs for Section 2.2

Recall the definitions of μ_0 , τ_0 , κ_0 and ε_0 in Section 5.2. Similar to the symmetric case, we will use the following easily verifiable inequalities.

$$\kappa_0 \leq \sqrt{r\mu_0}\tau_0, \quad \|E\|_2 \leq \left(\sqrt{d_1/d_2}\|E\|_{\infty} \cdot \sqrt{d_2/d_1}\|E\|_1\right)^{1/2} \leq \tau_0. \quad (54)$$

Lemma 13 *Parallel to Lemma 9, we have*

$$\|H^d\|_w \leq (1 + r\mu_0)\kappa_0,$$

where $\kappa_0 = \sqrt{d_1}\|EV\|_{\max} \vee \sqrt{d_2}\|E^T U\|_{\max}$ as defined.

Proof Recall $H^d = V_{\perp}^d(V_{\perp}^d)^T E^d V^d = E^d V^d - V^d(V^d)^T E^d V^d$. Note $V^d(V^d)^T = \text{diag}(UU^T, VV^T)$ and $\|UU^T\|_{\max} \leq r\mu(U)/d_1$, $\|VV^T\|_{\max} \leq r\mu(V)/d_2$. Thus,

$$\begin{aligned} \|H^d\|_w &\leq \|E^d V^d\|_w + \|V^d(V^d)^T E^d V^d\|_w \\ &\leq (1 + d_1\|UU^T\|_{\max} \vee d_2\|VV^T\|_{\max})\|E^d V^d\|_w \leq (1 + r\mu_0)\kappa_0. \end{aligned}$$

■

Lemma 14 *Parallel to Lemma 10, if $\sigma_r > 2\kappa_0 r \sqrt{\mu_0}$, then \bar{L}_1^d is a non-degenerate matrix. Furthermore, we have the following bound*

$$\inf_{\|Q_0^d\|_w=1} \|Q_0^d \bar{L}_1^d - \bar{L}_2^d Q_0^d\|_w \geq \sigma_r - 4r\mu_0(\tau_0 + r\kappa_0) - \varepsilon_0, \quad (55)$$

where $Q_0^d \in \mathbb{R}^{(d_1+d_2) \times 2r}$.

Proof Following similar derivations with Lemma 10, we have $\|Q_0^d E_{11}^d\|_w \leq 2\kappa_0 r \sqrt{\mu_0}$, and for any matrix $Q_0^d \in \mathbb{R}^{(d_1+d_2) \times 2r}$ with $\|Q_0^d\|_w = 1$,

$$\|\bar{L}_2^d Q_0^d\|_w = \|(A^d - A_r^d)Q_0^d + V_\perp^d (V_\perp^d)^T E^d V_\perp^d (V_\perp^d)^T Q_0^d\|_w \leq \varepsilon_0 + (1 + r\mu_0)(\tau_0 + r\kappa_0).$$

This can be checked by expressing Q_0^d as a block matrix and expand the matrix multiplication. In particular, one can verify that (i) $\|(A^d - A_r^d)Q_0^d\|_w \leq \varepsilon_0$; (ii) For any matrix M with $d_1 + d_2$ rows, $\|V^d (V^d)^T M\|_w \leq r\mu_0 \|M\|_w$; (iii) $\|E^d Q_0^d\|_w \leq \tau_0$; (iv) $\|E^d V^d (V^d)^T Q_0^d\|_w \leq r\kappa_0$. Moreover, $\|Q_0^d \Lambda_1^d\|_w \geq \sigma_r \|Q_0^d\|_w \geq \sigma_r$. Thus,

$$\inf_{\|Q_0^d\|_w=1} \|Q_0^d \bar{L}_1^d - \bar{L}_2^d Q_0^d\|_w \geq \sigma_r - 4r\mu_0(\tau_0 + r\kappa_0) - \varepsilon_0,$$

which is the desired inequality in the lemma. In addition, \bar{L}_1^d is non-degenerate if $\sigma_r > 2\kappa_0 r \sqrt{\mu_0} > 0$. \blacksquare

Lemma 15 *Parallel to Lemma 11, there is a solution $\bar{Q}^d \in \mathcal{B}_0$ to the system (37) such that if $\sigma_r - \varepsilon_0 > 16r\mu_0(\tau_0 + r\kappa_0)$, then*

$$\|\bar{Q}^d\|_w \leq \frac{8(1 + r\mu_0)\kappa_0}{3(\sigma_r - \varepsilon_0)}.$$

Proof We again invoke Lemma 6. Let \mathcal{B} be the space $\mathbb{R}^{(d_1+d_2) \times 2r}$ equipped with the weighted max-norm $\|\cdot\|_w$. We also define \mathcal{B}_0 as a subspace of \mathcal{B} consisting of matrices of the form $V_\perp^d Q^d$ where Q^d has size $(d_1 + d_2 - 2r) \times 2r$. Let the linear operator \mathcal{T}^d be $\mathcal{T}^d \bar{Q}^d := \bar{Q}^d \bar{L}_1^d - \bar{L}_2^d \bar{Q}^d$. First notice from Lemma 14, \mathcal{T}^d is a linear operator with bounded inverse, i.e., $\beta := \|(\mathcal{T}^d)^{-1}\|_w^{-1}$ is bounded from below by

$$\beta \geq \sigma_r - 4r\mu_0(\tau_0 + r\kappa_0) - \varepsilon_0.$$

Let φ be a map given by $\varphi(\bar{Q}^d) = \bar{Q}^d (H^d)^T \bar{Q}^d$. Note that $H^d \in \mathcal{B}$. Using the (easily verifiable) inequality

$$\|M_1 M_2^T M_3\|_w \leq 2r \|M_1\|_w \|M_2^T M_3\|_{\max} \leq 4r \|M_1\|_w \|M_2\|_w \|M_3\|_w \quad \forall M_1, M_2, M_3 \in \mathcal{B}, \quad (56)$$

we derive, by the bound on $\|H^d\|_w$ (Lemma 13), that

$$\|\varphi(\bar{Q}^d)\|_w \leq 4r \|H^d\|_w \|\bar{Q}^d\|_w^2 \leq 4r(1 + r\mu_0)\kappa_0 \|\bar{Q}^d\|_w^2.$$

Moreover, using the inequality (56) and the bound on $\|H^d\|_w$ (Lemma 13),

$$\begin{aligned} \|\varphi(\bar{Q}_1^d) - \varphi(\bar{Q}_2^d)\|_w &\leq 4r\|\bar{Q}_1^d\|_w\|H^d\|_w\|\bar{Q}_1^d - \bar{Q}_2^d\|_w + 4r\|\bar{Q}_1^d - \bar{Q}_2^d\|_w\|H^d\|_w\|\bar{Q}_2^d\|_w \\ &\leq 8r(1+r\mu_0)\kappa_0 \max\{\|\bar{Q}_1^d\|_w, \|\bar{Q}_2^d\|_w\}\|\bar{Q}_1^d - \bar{Q}_2^d\|_w. \end{aligned}$$

Thus, we can choose $\eta = 4r(1+r\mu_0)\kappa_0$, and the condition (33) in Lemma 6 is satisfied. To ensure $4\eta\|H^d\|_w < \beta^2$, it suffices to require (again by Lemma 13),

$$16r(1+r\mu_0)^2\kappa_0^2 < [\sigma_r - 4r\mu_0(\tau_0 + r\kappa_0) - \varepsilon_0]^2.$$

It is easily checkable that the above inequality holds when $\sigma_r - \varepsilon_0 > 16r\mu_0(\tau_0 + r\kappa_0)$. Under this condition, by Lemma 6,

$$\|\bar{Q}^d\|_w \leq \frac{2\|H^d\|_w}{\beta} \leq \frac{2(1+r\mu_0)\kappa_0}{\sigma_r - 4r\mu_0(\tau_0 + r\kappa_0) - \varepsilon_0} \leq \frac{2(1+r\mu_0)\kappa_0}{\sigma_r - \varepsilon_0 - (\sigma_r - \varepsilon_0)/4} \leq \frac{8(1+r\mu_0)\kappa_0}{3(\sigma_r - \varepsilon_0)},$$

which completes the proof. \blacksquare

Proof of Lemma 8 The first claim of the lemma (existence of \bar{Q}^d and its max-norm bound) follows from Lemma 15. To prove the second claim, we split $\bar{V}^d - V^d$ into two parts:

$$\bar{V}^d - V^d = V^d \left((I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2} - I_{2r} \right) + \bar{Q}^d (I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2}, \quad (57)$$

Note $\kappa_0 \leq \tau_0 \sqrt{r\mu_0}$ (see (54)). It can be checked that the condition $\sigma_r - \varepsilon_0 > 16r\mu_0(\tau_0 + r\kappa_0)$ implies $r\omega_0 < 1/3$. Since $\|(\bar{Q}^d)^T \bar{Q}^d\|_{\max} \leq 2\omega_0^2$ and $\|(\bar{Q}^d)^T \bar{Q}^d\|_2 \leq 2r\|(\bar{Q}^d)^T \bar{Q}^d\|_{\max} \leq 4r\omega_0^2 < 1/2$, similar to Lemma 12, we have

$$\begin{aligned} \|(I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2} - I_{2r}\|_{\max} &\leq 4r\omega_0^2, \\ \|(I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2}\|_{\max} &\leq 3/2. \end{aligned} \quad (58)$$

This yields

$$\begin{aligned} \|\bar{V}^d - V^d\|_w &= \|V^d \left((I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2} - I_{2r} \right)\|_w + \|\bar{Q}^d (I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2}\|_w \\ &\leq \sqrt{2r^2\mu_0} 4r\omega_0^2 + 2r \cdot 3/2 \cdot \|\bar{Q}^d\|_w \leq 8\sqrt{\mu_0}\omega_0^2 r^2 + 3\omega_0 r \end{aligned} \quad (59)$$

$$\leq 8\sqrt{\mu_0}\omega_0 r/3 + 3\sqrt{\mu_0}\omega_0 r \leq 6\sqrt{\mu_0}r\omega_0. \quad (60)$$

\blacksquare

Lemma 16 *Suppose $\sigma_r - \varepsilon_0 > \max\{16r\mu_0(\tau_0 + r\kappa_0), 64r^{3/2}\mu_0^{1/2}(1+r\mu_0)\kappa_0\}$. Then, there exists an orthogonal matrix $R_V \in \mathbb{R}^{r \times r}$ (or R_U) such that the column vectors of $\bar{V}R_V$ (and $\bar{U}R_U$) are the top r right (and left) singular vectors of \tilde{A} .*

Proof Similar to the proof of Lemma 7, we will prove $d(\bar{V}, V) < 1/2$ and $d(\tilde{V}, V) \leq 1/2$, which would then imply that \bar{V} and \tilde{V} are the same only up to an orthogonal transformation. The same is true for \bar{U} and \tilde{U} , and we will leave out its proof.

By Weyl's inequality for singular values (also known as Mirsky's theorem (Mirsky, 1960)), for any i , $|\tilde{\sigma}_i - \sigma_i| \leq \|E\|$. By Wedin's perturbation bounds for singular vectors (Wedin, 1972),

$$d(\tilde{V}, V) \leq \frac{\|E\|_2}{\sigma_r - \|E\|_2}.$$

Note that $\|E\|_2 \leq \tau_0$ (see (54)) Under the assumption in the lemma, clearly $\sigma_r - \varepsilon_0 > 3\tau_0$, and we have $d(\tilde{V}, V) \leq 1/2$. Moreover, by Lemma 8, we have $\|\bar{V}^d - V^d\|_w \leq 6r\omega_0\sqrt{\mu_0}$. Note that each column vector of V^d and \bar{V}^d are $(d_1 + d_2)$ -dimensional. Looking at the last d_2 dimensions, we have $\|\bar{V} - V\|_{\max} \leq 6r\omega_0\sqrt{\mu_0/d_1}$.

$$d(\bar{V}, V) \leq 2\|\bar{V} - V\| \leq 2\sqrt{rd_1}\|\bar{V} - V\|_{\max} \leq 12r^{3/2}\mu_0^{1/2}\omega_0.$$

Under the assumption of the lemma, $d(\bar{V}, V) \leq 1/2$. Therefore, we deduce $d(\tilde{V}, \bar{V}) = 0$, and conclude that there exists an orthogonal matrix $R_V \in \mathbb{R}^{r \times r}$ such that $\tilde{V} = \bar{V}R_V$. ■

Proof of Theorem 4 Lemma 8, together with Lemma 16, implies Theorem 4. ■

Proof of Theorem 5 Similar to the proof of Theorem 3, we first split the difference $\tilde{V}^d - V^d$:

$$\tilde{V}^d - V^d = \bar{V}^d(R^d - I_{2r}) + (\bar{V}^d - V^d). \quad (61)$$

To bound the first term, note that under our assumption, $r\omega_0 < 1/3$ (derived in the proof of Lemma 8), it is easy to check $\|\bar{V}^d\|_w \leq 3\sqrt{r\mu_0}$. We rewrite the matrix R^d as

$$R^d = (I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2} (V^d + \bar{Q}^d)^T \tilde{V}^d.$$

Notice that $\|(\bar{Q}^d)^T \tilde{V}^d\|_{\max} \leq \sqrt{2}\|\bar{Q}^d\|_w \leq \sqrt{2}\omega_0$, $\|(\tilde{V}^d)^T V^d\|_{\max} \leq 1$ and

$$\begin{aligned} \|R^d - (V^d + \bar{Q}^d)^T \tilde{V}^d\|_{\max} &\leq 2r\|(I_{2r} + (\bar{Q}^d)^T \bar{Q}^d)^{-1/2} - I_{2r}\|_{\max} \|(\bar{Q}^d)^T \tilde{V}^d\|_{\max} \\ &\leq 8r^2\omega_0^2(1 + \sqrt{2}\omega_0). \end{aligned}$$

where we used (58). Following the same derivations as in the proof of Theorem 3, and using the (easily verifiable) fact $\|E^d\|_2 = \|E\|_2$, we can bound $\|(V^d)^T \tilde{V}^d - I_{2r}\|_{\max}$ by $2\|E\|_2/\delta_0$. Thus, using $r\omega_0 \leq 1/3$, under $\delta_0 > 2\|E\|_2$, we have

$$\begin{aligned} \|R^d - I_{2r}\|_{\max} &\leq \|R^d - (V^d + \bar{Q}^d)^T \tilde{V}^d\|_{\max} + \|(V^d)^T \tilde{V}^d - I_{2r}\|_{\max} + \|(\bar{Q}^d)^T \tilde{V}^d\|_{\max} \\ &\leq 8r^2\omega_0^2(1 + \sqrt{2}\omega_0) + \frac{2\|E\|_2}{\delta_0} + \sqrt{2}\omega_0 < 4\sqrt{2}r\omega_0 + \frac{2\|E\|_2}{\delta_0}. \end{aligned} \quad (62)$$

Finally, in order to bound $\|\tilde{V}^d - V^d\|_w$, we use (60), (61) and (62), and derive

$$\begin{aligned} \|\tilde{V}^d - V^d\|_w &= \|\bar{V}^d(R^d - I_{2r}) + (\bar{V}^d - V^d)\|_w \leq 2r\|\bar{V}^d\|_w\|R^d - I_{2r}\|_{\max} + \|\bar{V}^d - V^d\|_w \\ &\leq 6r\sqrt{r\mu_0}\left(4\sqrt{2}r\omega_0 + \frac{2\|E\|_2}{\delta_0}\right) + 6r\sqrt{\mu_0}\omega_0 \leq \left(40r^{5/2}\omega_0 + 12r^{3/2}\frac{\|E\|_2}{\delta_0}\right) \cdot \sqrt{\mu_0} \\ &\leq \frac{107r^{5/2}\mu_0^{1/2}(1+r\mu_0)\kappa_0}{\sigma_r - \varepsilon_0} + \frac{12r^{3/2}\mu_0^{1/2}\|E\|_2}{\delta_0} \\ &= O\left(\frac{r^4\mu_0^2\tau_0}{\sigma_r - \varepsilon_0} + \frac{r^{3/2}\mu_0^{1/2}\|E\|_2}{\delta_0}\right). \end{aligned}$$

This completes the proof. ■

C. Proofs for Section 3

Proof of Proposition 3 Note first by Weyl's inequality, $|\lambda_i - \bar{\lambda}_i| \leq \|\Sigma_u\| \leq C$. So this implies that $\bar{\lambda}_i = \lambda_i(B^T B) \asymp d$ if and only if $\lambda_i = \lambda_i(\Sigma) \asymp d$ for $i \leq r$. And furthermore the eigenvalues of $B^T B/d$ are distinct if and only if $\min_{1 \leq i \neq j \leq r} |\lambda_i(\Sigma) - \lambda_j(\Sigma)|/\lambda_j(\Sigma) > 0$.

To prove the equivalency of bounded $\|B\|_{\max}$ and bounded coherence. We first prove the necessary condition. Again from Weyl's inequality, $\lambda_i(\Sigma) \leq C$ for $i \geq r+1$. If $\mu(V)$ is bounded, Σ_{ii} must also be bounded, since $\Sigma_{ii} \leq \sum_{j=1}^r v_{ij}^2 \lambda_j(\Sigma) + \lambda_{r+1}(\Sigma) \leq C(\mu(V) + 1)$. Therefore $\|b_i\|^2 \leq \|\tilde{b}_i\|^2 + (\Sigma_u)_{ii} = \Sigma_{ii}$ implies $\|B\|_{\max}$ is bounded. Namely, the factors are pervasive.

On the contrary, if pervasiveness holds, we need to prove that $\mu(V)$ is bounded. Let $B = (\tilde{b}_1, \dots, \tilde{b}_r)$. Obviously $\bar{\lambda}_i = \|\tilde{b}_i\|^2 \asymp d$ and $\bar{v}_i = \tilde{b}_i/\|\tilde{b}_i\|$. Without loss of generality, assume $\bar{\lambda}_i$'s are decreasing. So $\|\bar{v}_i\|_{\infty} \leq \|B\|_{\max}/\|\tilde{b}_i\| \leq C/\sqrt{d}$ and $\mu(\bar{V}) \leq C$ where $\bar{V} = (\bar{v}_1, \dots, \bar{v}_r)$. By Theorem 3,

$$\|\bar{v}_i - v_i\|_{\infty} \leq C \frac{\|\Sigma_u\|_{\infty}}{\bar{\gamma}\sqrt{d}},$$

where $\bar{\gamma} = \min\{\bar{\lambda}_i - \bar{\lambda}_{i+1} : 1 \leq i \leq r\} \asymp d$ with the convention $\bar{\lambda}_{r+1} = 0$. Hence, we have $\|v_i\|_{\infty} \leq C/\sqrt{d}$, which implies bounded coherence $\mu(V)$. ■

References

- Anestis Antoniadis and Jianqing Fan. Regularization of wavelet approximations. *Journal of the American Statistical Association*, 96(455), 2001.
- Jushan Bai. Inferential theory for factor models of large dimensions. *Econometrica*, 71(1): 135–171, 2003.
- Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, pages 1643–

- 1697, 2005.
- David J Bartholomew, Martin Knott, and Iriini Moustaki. *Latent variable models and factor analysis: A unified approach*, volume 904. John Wiley & Sons, 2011.
- Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*, 41(4):1780–1815, 2013.
- Peter J Bickel and Elizaveta Levina. Covariance regularization by thresholding. *The Annals of Statistics*, pages 2577–2604, 2008.
- Pratik Biswas and Yinyu Ye. Semidefinite programming for ad hoc wireless sensor network localization. In *Proceedings of the 3rd international symposium on Information processing in sensor networks*, pages 46–54. ACM, 2004.
- Ingwer Borg and Patrick JF Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2005.
- T Tony Cai and Anru Zhang. Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics. *arXiv preprint arXiv:1605.00353*, 2016.
- Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.
- Olivier Catoni. Challenging the empirical mean and empirical variance: a deviation study. In *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, volume 48, pages 1148–1185. Institut Henri Poincaré, 2012.
- Gary Chamberlain, Michael Rothschild, et al. Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*, 51(5):1281–1304, 1983.
- Venkat Chandrasekaran, Sujay Sanghavi, Pablo A Parrilo, and Alan S Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2): 572–596, 2011.
- Yudong Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- Kyungmee Choi and J Marden. A multivariate version of kendall’s τ . *Journal of Nonparametric Statistics*, 9(3):261–293, 1998.
- Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- Fernando De La Torre and Michael J Black. A framework for robust subspace learning. *International Journal of Computer Vision*, 54(1-3):117–142, 2003.
- Lance Doherty, Kristofer SJ Pister, and Laurent El Ghaoui. Convex position estimation in wireless sensor networks. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1655–1663. IEEE, 2001.

- Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993.
- Jianqing Fan, Yingying Fan, and Jinchi Lv. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, 147(1):186–197, 2008.
- Jianqing Fan, Yuan Liao, and Martina Mincheva. Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society: Series B*, 75:1–44, 2013.
- Jianqing Fan, Qiefeng Li, and Yuyan Wang. Robust estimation of high-dimensional mean regression. *Journal of Royal Statistical Society, Series B*, 79:247–265, 2017a.
- Jianqing Fan, Weichen Wang, and Yiqiao Zhong. Robust covariance estimation for approximate factor models. *Journal of Econometrics*, 2017b.
- Jianqing Fan, Han Liu, and Weichen Wang. Large covariance estimation through elliptical factor models. *The Annals of Statistics*, page to appear, 2017+.
- Kai-Tai Fang, Samuel Kotz, and Kai Wang Ng. *Symmetric multivariate and related distributions*. CRC Press, 1990.
- Arjun K Gupta, Tamas Varga, and Taras Bodnar. *Elliptically contoured models in statistics and portfolio theory*. Springer, 2013.
- Fang Han and Han Liu. Scale-invariant sparse PCA on high-dimensional meta-elliptical data. *Journal of the American Statistical Association*, 109(505):275–287, 2014.
- Daniel Hsu and Sivan Sabato. Heavy-tailed regression with a generalized median-of-means. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 37–45, 2014.
- Weiming Hu, Tieniu Tan, Liang Wang, and Steve Maybank. A survey on visual surveillance of object motion and behaviors. *Systems, Man, and Cybernetics, Part C: IEEE Transactions on Applications and Reviews*, 34(3):334–352, 2004.
- Peter J Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- Iain M. Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009. doi: 10.1198/jasa.2009.0121.
- Ian Jolliffe. *Principal Component Analysis*. Wiley Online Library, 2002.
- Zongming Ma. Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*, 41(2):772–801, 2013.
- Leon Mirsky. Symmetric gauge functions and unitarily invariant norms. *The quarterly journal of mathematics*, 11:50–59, 1960.

- Andrew Y Ng, Michael I Jordan, Yair Weiss, et al. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 2:849–856, 2002.
- Nuria M Oliver, Barbara Rosario, and Alex P Pentland. A bayesian computer vision system for modeling human interactions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):831–843, 2000.
- Sean O’Rourke, Van Vu, and Ke Wang. Random perturbation of low rank matrices: Improving classical bounds. *arXiv preprint arXiv:1311.2657*, 2013.
- Vern Paulsen. *Completely bounded maps and operator algebras*, volume 78. Cambridge University Press, 2002.
- Jasson DM Rennie and Nathan Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22nd international conference on Machine learning*, pages 713–719. ACM, 2005.
- Karl Rohe, Sourav Chatterjee, and Bin Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, pages 1878–1915, 2011.
- Adam J Rothman, Elizaveta Levina, and Ji Zhu. Generalized thresholding of large covariance matrices. *Journal of the American Statistical Association*, 104(485):177–186, 2009.
- Gilbert W Stewart. *Matrix perturbation theory*. 1990.
- James H Stock and Mark W Watson. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460):1167–1179, 2002.
- Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Van Vu. Singular vectors under random perturbation. *Random Structures & Algorithms*, 39(4):526–538, 2011.
- Vincent Q Vu and Jing Lei. Minimax sparse principal subspace estimation in high dimensions. *The Annals of Statistics*, 41(6):2905–2947, 2013.
- Weichen Wang and Jianqing Fan. Asymptotics of empirical eigenstructure for high dimensional spiked covariance. *The Annals of Statistics*, 45(3):1342–1374, 2017.
- Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.
- Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102:315–323, 2015.

Ming Yuan, Ali Ekici, Zhaosong Lu, and Renato Monteiro. Dimension reduction and coefficient estimation in multivariate linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(3):329–346, 2007.