Distributed estimation of principal eigenspaces

Jianqing Fan, **Dong Wang**, **Kaizheng Wang**, **Ziwei Zhu**

Department of Operations Research and Financial Engineering Princeton University

Abstract

Principal component analysis (PCA) is fundamental to statistical machine learning. It extracts latent principal factors that contribute to the most variation of the data. When data are stored across multiple machines, however, communication cost can prohibit the computation of PCA in a central location and distributed algorithms for PCA are thus needed. This paper proposes and studies a distributed PCA algorithm: each node machine computes the top K eigenvectors and transmits them to the central server; the central server then aggregates the information from all the node machines and conducts a PCA based on the aggregated information. We investigate the bias and variance for the resulting distributed estimator of the top K eigenvectors. In particular, we show that for distributions with symmetric innovation, the empirical top eigenspaces are unbiased and hence the distributed PCA is "unbiased". We derive the rate of convergence for distributed PCA estimators, which depends explicitly on the effective rank of covariance, eigen-gap, and the number of machines. We show that when the number of machines is not unreasonably large, the distributed PCA performs as well as the whole sample PCA, even without full access of whole data. The theoretical results are verified by an extensive simulation study. We also extend our analysis to the heterogeneous case where the population covariance matrices are different across local machines but share similar top eigen-structures.

Keywords

Distributed Learning; PCA; One-shot Approach; Communication Efficiency; Unbiasedness of Empirical Eigenspaces; Heterogeneity

1   Introduction

Principal component analysis (PCA) (Pearson, 1901; Hotelling, 1933) is one of the most fundamental tools in statistical machine learning. The past century has witnessed great efforts on establishing consistency and asymptotic distribution of empirical eigenvalues and eigenvectors. The early classical work of Anderson (1963) studied the asymptotic normality of eigenvalues and eigenvectors of sample covariances from multivariate Gaussian distribution with dimension d fixed and sample size n going to infinity. Recent focus moves on to the high-dimensional regimes, i.e., both n and d go to infinity. A partial list of such literatures are Johnstone (2001); Baik et al. (2005); Paul (2007); Johnstone and Lu (2009); Jung and Marron (2009); Onatski (2012); Shen et al. (2016); Wang and Fan (2017). As demonstrated by these papers, asymptotic behaviors of empirical eigenvalues and

eigenvectors depend on the scaling of n, d and also the spikiness of the covariance. When $n \ll d$, the empirical leading eigenvector $\hat{\mathbf{v}}_1$ is inconsistent in estimating the true top eigenvector $\mathbf{v}_1$ unless the top eigenvalue λ_1 diverges fast. This phenomenon inspires another line of research on sparse PCA where certain sparsity on top eigenvectors is imposed to overcome the noise accumulation due to high dimensionality; see e.g., Johnstone and Lu (2009); Vu et al. (2013); Shen et al. (2013); Cai et al. (2013). Besides the asymptotic study, there are also non-asymptotic results on PCA, for example, Nadler (2008) and Reiss and Wahl (2016).

With rapid developments of information and technology, massive datasets are now ubiquitous. Statistical analysis such as regression or PCA on such enormous data is unprecedentedly desirable. However, large datasets are usually scattered across distant places such that to fuse or aggregate them is extremely difficult due to communication cost, privacy, data security and ownerships, among others. Consider giant IT companies that collect data simultaneously from places all around the world. Constraints on communication budget and network bandwidth make it nearly impossible to aggregate and maintain global data in a single data center. Another example is that health records are scattered across many hospitals or countries. It is hard to process the data in a central location due to privacy and ownership concerns. To resolve these issues, efforts have been made to exploiting distributed computing architectures and developing distributed estimators or testing statistics based on data scattered around different locations. A typical distributed statistical method first calculates local statistics based on each sub-dataset and then combines all the subsample-based statistics to produce an aggregated statistic. Such distributed methods fully adapt to the parallel data collection procedures and thus significantly reduce the communication cost. Many distributed regression methods follow this fashion (Zhang et al., 2013; Chen and Xie, 2014; Battey et al., 2015; Lee et al., 2017; Blanchard and Mucke, 2017; Guo et al., 2017). The last two papers study distributed kernel regression with spectral regularization using eigen-decomposition of Gram matrices, which is relevant to but different from our distributed PCA.

Among all the efforts towards creating accurate and efficient distributed statistical methods, there has been rapid advancement on distributed PCA over the past two decades. Unlike the traditional PCA where we have the complete data matrix $\mathbf{X} \in \mathbb{R}^{N \times d}$ with d features of N samples at one place, the distributed PCA needs to handle data that are partitioned and stored across multiple servers. There are two data partition regimes: "horizontal" and "vertical". In the horizontal partition regime, each server contains all the features of a subset of subjects, while in the vertical partition regime, each server has a subset of features of all the subjects. To conduct distributed PCA in the horizontal regime, Qu et al. (2002) proposes that each server computes several top eigenvalues and eigenvectors on its local data and then sends them to the central server that aggregates the information together. Yet there is no theoretical guarantee on the approximation error of the proposed algorithm. Liang et al. (2014), Kannan et al. (2014) and Boutsidis et al. (2016) aim to find a good rank-K approximation of $\widehat{\mathbf{X}}$ of $\mathbf{X}$. To assess the approximation quality, they compare $\left\| \widehat{\mathbf{X}} - \mathbf{X} \right\|_F$ against $\min_{rank(\mathbf{B}) \leq K} \left\| \mathbf{B} - \mathbf{X} \right\|_F$ and study the excess risk. For the distributed PCA in the

vertical data partition regime, there is also a great amount of literature, for example, Kargupta et al. (2001), Li et al. (2011), Bertrand and Moonen (2014), Schizas and Aduroja (2015), etc. This line of research is often motivated from sensor networks and signal processing where the vertically partitioned data are common. Our work focuses on the horizontal partition regime, i.e., we have partitions over the samples rather than the features.

Despite these achievements, very few papers establish rigorous statistical error analysis of the proposed distributed PCA methods. To our best knowledge, the only works that provide statistical analysis so far are El Karoui and d'Aspremont (2010) and Chen et al. (2016). To estimate the leading singular vectors of a large target matrix, both papers propose to aggregate singular vectors of multiple random approximations of the original matrix. El Karoui and d'Aspremont (2010) adopts sparse approximation of the matrix by sampling the entries, while Chen et al. (2016) uses Gaussian random sketches. The works are related to ours, since we can perceive sub-datasets in the distributed PCA problem as random approximations. However, our analysis is more general, since it does not rely on any matrix incoherence assumption as required by El Karoui and d'Aspremont (2010) and it explicitly characterizes how the probability distribution affects the final statistical error in finite sample error bounds. Besides, our aggregation algorithm is much simpler than the one in Chen et al. (2016). The manuscript Garber et al. (2017) came out after we submitted the first draft of our work. The authors focused on estimation of the first principal component rather than the multi-dimensional eigenspaces, based on very different approaches.

We propose a distributed algorithm with only one-shot communication to solve for the top K eigenvectors of the population covariance matrix Σ when samples are scattered across m servers. We first calculate for each subset of data ℓ its top K eigenvectors $\left\{ \widehat{\mathbf{V}}_K^{(\ell)} = \left(\widehat{\mathbf{v}}_1^{(\ell)}, \cdots, \widehat{\mathbf{v}}_K^{(\ell)} \right) \right\}_{\ell=1}^m$ of the sample covariance matrix there, then compute the average of projection matrices of the eigenspaces $\widetilde{\Sigma} = (1/m) \sum_{i=1}^m \widehat{\mathbf{V}}_K^{(\ell)} \widehat{\mathbf{V}}_K^{(\ell)T}$, and finally take the top K eigenvectors of $\widetilde{\Sigma}$ as the final estimator $\widetilde{\mathbf{V}}_K = \left(\widetilde{\mathbf{v}}_1^{(\ell)}, \cdots, \widetilde{\mathbf{v}}_K^{(\ell)} \right)$. The communication cost of this method is of order $O(mKd)$. We establish rigorous non-asymptotic analysis of the statistical error $\left\| \widetilde{\mathbf{V}}_K \widetilde{\mathbf{V}}_K^T - \mathbf{V}_K \mathbf{V}_K^T \right\|_F$, and show that as long as we have a sufficiently large number of samples in each server, $\widetilde{\mathbf{V}}_K$ enjoys the same statistical error rate as the standard PCA over the full sample. The eigenvalues of Σ are easily estimated once we get good estimators of the eigenvectors, using another round of communication.

The rest of the paper is organized as follows. In Section 2, we introduce the problem setup of the distributed PCA. In Section 3, we elucidate our distributed algorithm for estimating the top K eigenvectors. Section 4 develops the statistical error rates of the aggregated estimator. The results are extended to heterogeneous samples in Section 5. Finally in Section 6 we present extensive simulation results to validate our theories.

2 Problem setup

We first collect all the notations that will be used. By convention we use regular letters for scalars and bold letters for both matrices and vectors. We denote the set $\{1, 2, 3,..., d\}$ by $[d]$ for convenience. For two scalar sequences $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, we say $a_n \gtrsim b_n \left(a_n \lesssim b_n\right)$ if there exists a universal constant $C > 0$ such that $a_n \geq Cb_n\left(a_n \leq Cb_n\right)$, and $a_n \asymp b_n$ if both $a_n \gtrsim b_n$ and $a_n \lesssim b_n$ hold. For a random variable $X \in \mathbb{R}$, we define $\left\|X\right\|_{\psi_2} = \sup_{p \geq 1}\left(\mathbb{E}|X|^p\right)^{\frac{1}{p}}/\sqrt{p}$ and define $\left\|X\right\|_{\psi_1} = \sup_{p \geq 1}\left(\mathbb{E}|X|^p\right)^{\frac{1}{p}}/p$. Please refer to Vershynin (2012) for equivalent definitions of ψ_2-norm and ψ_1-norm. For two random variables X and Y, we use $X \overset{d}{=} Y$ to denote that X and Y have identical distributions. Define $\mathbf{e}_i$ to be the unit vector whose components are all zero except that the i-th component equals 1. For $q \geq r$, $\mathcal{O}_{q \times r}$ denotes the space of $q \times r$ matrices with orthonormal columns. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, we use $\left\|\mathbf{A}\right\|_F$, $\left\|\mathbf{A}\right\|_*$ and $\left\|\mathbf{A}\right\|_2$ to denote the Frobenius norm, nuclear norm and spectral norm of $\mathbf{A}$, respectively. $\mathrm{Col}(\mathbf{A})$ represents the linear space spanned by column vectors of $\mathbf{A}$. We denote the Moore-Penrose pseudo inverse of a matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ by $A^\dagger$. For a symmetric matrix $\mathbf{A}$, we use $\lambda_j(\mathbf{A})$ to refer to its j-th largest eigenvalue.

Suppose we have N i.i.d random samples $\left\{\mathbf{X}_i\right\}_{i = 1}^N \subseteq \mathbb{R}^d$ with $\mathbb{E}\mathbf{X}_1 = \mathbf{0}$ and covariance matrix $\mathbb{E}\left(\mathbf{X}_1\mathbf{X}_1^T\right) = \Sigma$. By spectral decomposition, $\Sigma = \mathbf{V}\Lambda\mathbf{V}^T$, where $\Lambda = \mathrm{diag}(\lambda_1, \lambda_2..., \lambda_d)$ with $\lambda_1 \geq \lambda_2 \geq... \geq \lambda_d$ and $\mathbf{V} = (\mathbf{v}_1, ..., \mathbf{v}_d) \in \mathcal{O}_{d \times d}$. For a given $K \in [d]$, let $\mathbf{V}_K = \left(\mathbf{v}_1, \cdots, \mathbf{v}_K\right)$. Our goal is to estimate $\mathrm{Col}(\mathbf{V}_K)$, i.e., the linear space spanned by the top K eigenvectors of Σ. To ensure the identifiability of $\mathrm{Col}(\mathbf{V}_K)$, we assume $\Delta := \lambda_K - \lambda_{K+1} > 0$ and define $\kappa := \lambda_1/\Delta$ to be the condition number. Let $r = r(\Sigma) := \mathrm{Tr}(\Sigma)/\lambda_1$ be the effected rank of Σ.

The standard way of estimating $\mathrm{Col}(\mathbf{V}_K)$ is to use the top K eigenspace of the sample covariance $\widehat{\Sigma} = \frac{1}{N}\sum_{i=1}^N \mathbf{X}_i\mathbf{X}_i^T$. Let $\widehat{\Sigma} = \widehat{\mathbf{V}}\widehat{\Lambda}\widehat{\mathbf{V}}_T$ be spectral decomposition of $\widehat{\Sigma}$, where $\widehat{\Lambda} = \mathrm{diag}\left(\hat{\lambda}_1, \cdots, \hat{\lambda}_d\right)$ with $\hat{\lambda}_1 \geq \cdots \geq \hat{\lambda}_d$ and $\widehat{\mathbf{V}} = \left(\hat{\mathbf{v}}_1, \cdots, \hat{\mathbf{v}}_d\right)$. We use the empirical top K engenspace $\mathrm{Col}(\widehat{\mathbf{V}}_K)$, where $\widehat{\mathbf{V}}_K = \left(\hat{\mathbf{v}}_1, \cdots, \hat{\mathbf{v}}_K\right)$, to estimate the eigenspace $\mathrm{Col}(\mathbf{V}_K)$. To measure the statistical error, we adopt $\rho\left(\widehat{\mathbf{V}}_K, \mathbf{V}_K\right) := \left\|\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T - \mathbf{V}_K\mathbf{V}_K^T\right\|_F$, which is the Frobenius norm of the difference between projection matrices of two spaces and is a well-defined distance between linear subspaces. In fact, $\rho\left(\mathbf{V}_K, \widehat{\mathbf{V}}_K\right)$ is equivalent to the so-called $\sin\Theta$ distance. Denote the singular values of $\widehat{\mathbf{V}}_K^T\mathbf{V}_K$ by $\{\sigma_i\}_{i=1}^K$ in descending order. Recall that $\Theta\left(\widehat{\mathbf{V}}_K, \mathbf{V}_K\right) = \mathrm{diag}\left(\theta_1, \cdots, \theta_K\right)$, the principal angles between $\mathrm{Col}(\mathbf{V}_K)$ and $\mathrm{Col}(\widehat{\mathbf{V}}_K)$, are defined as $\mathrm{diag}\left(\cos^{-1}\sigma_1, \cdots, \cos^{-1}\sigma_K\right)$. Then we define $\sin\Theta\left(\widehat{\mathbf{V}}_K, \mathbf{V}_K\right)$ to be $\mathrm{diag}\left(\sin\theta_1, \cdots, \sin\theta_K\right)$. Note that

$$\rho^2\!\left(\mathbf{V}_K,\widehat{\mathbf{V}}_K\right)=\left\|\mathbf{V}_K\mathbf{V}_K^T\right\|_F^2+\left\|\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T\right\|_F^2-2\mathrm{Tr}\!\left(\mathbf{V}_K\mathbf{V}_K^T\widehat{\mathbf{V}}_K\widehat{\mathbf{V}}_K^T\right)=2K-2\left\|\widehat{\mathbf{V}}_K^T\mathbf{V}_K\right\|_F^2$$

$$=2\sum_{i=1}^{K}\left(1-\sigma_i^2\right)=2\sum_{i=1}^{K}\sin^2\theta_i=2\left\|\sin\Theta\left(\widehat{\mathbf{V}}_K,\mathbf{V}_K\right)\right\|_F^2. \tag{2.1}$$

Therefore, $\rho\!\left(\mathbf{V}_K,\widehat{\mathbf{V}}_K\right)$ and $\left\|\sin\Theta\left(\mathbf{V}_K,\widehat{\mathbf{V}}_K\right)\right\|_F$ are equivalent.

Now consider the estimation of top K eigenspace under the distributed data setting, where our $N=m\cdot n$ samples are scattered across m machines with each machine storing n samples*. Application of standard PCA here requires data or covariance aggregation, thus leads to huge communication cost for high-dimensional big data. In addition, for the areas such as genetic, biomedical studies and customer services, it is hard to communicate raw data because of privacy and ownership concerns. To address these problems, we need to avoid naive data aggregation and design a communication-efficient and privacy-preserving distributed algorithm for PCA. In addition, this new algorithm should be statistically accurate in the sense that it enjoys the same statistical error rate as the full sample PCA.

Throughout the paper, we assume that all the random samples $\left\{\mathbf{X}_i\right\}_{i=1}^{N}$ are i.i.d sub-Gaussian. We adopt the definition of sub-Gaussian random vectors in Koltchinskii and Lounici (2017) and Reiss and Wahl (2016) as specified below, where M is assumed to be a constant. It is not hard to show that the following definition is equivalent to the definition $\left\|\left(\mathbf{\Sigma}^{1/2}\right)^{\dagger}\mathbf{X}\right\|_{\psi_2}\le M$ used in Vershynin (2012), Wang and Fan (2017), and many other authors.

Definition 2.1. *We say the random vector* $\mathbf{X}\in\mathbb{R}^d$ *is sub-Gaussian if there exists $M>0$ such that* $\left\|\mathbf{u}^T\mathbf{X}\right\|_{\psi_2}\le M\sqrt{\mathbb{E}\left(\mathbf{u}^T\mathbf{X}\right)^2},\forall\mathbf{u}\in\mathbb{R}^d.$

We emphasize here that the global i.i.d assumption on $\left\{\mathbf{X}_i\right\}_{i=1}^{N}$ can be further relaxed. In fact, our statistical analysis only requires the following three conditions: (i) within each server ℓ, data are i.i.d.; (ii) across different servers, data are independent; (iii) the covariance matrices of the data in each server $\left\{\mathbf{\Sigma}^{(\ell)}\right\}_{\ell=1}^{m}$ share similar top K eigenspaces. We will further study this heterogeneous regime in Section 5. To avoid future confusion, unless specified, we always assume i.i.d. data across servers.

---

*Note that here for simplicity we assume the subsample sizes are homogeneous. We can easily extend our analysis to the case of heterogeneous sub-sample sizes with similar theoretical results.

3 Methodology

We now introduce our distributed PCA algorithm. For $\ell \in [m]$, let $\left\{ \mathbf{X}_i^{(\ell)} \right\}_{i=1}^n$ denote the samples stored on the ℓ-th machine. We specify the distributed in Algorithm 1.

Algorithm 1

Distributed PCA

---

1   On each server, compute locally the K leading eigenvectors $\widehat{\mathbf{V}}_K^{(\ell)} = \left(\widehat{\mathbf{v}}_1^{(\ell)}, \cdots, \widehat{\mathbf{v}}_K^{(\ell)} \right) \in \mathbb{R}^{d \times K}$ of the sample covariance matrix $\widehat{\boldsymbol{\Sigma}}^{(\ell)} = (1/n) \mathbf{X}_i^{(\ell)} \mathbf{X}_i^{(\ell)^T}$. Send $\widehat{\mathbf{V}}_K^{(\ell)}$ to the central processor.

2   On the central processor, compute $\widetilde{\boldsymbol{\Sigma}} = (1/m) \sum_{\ell=1}^m \widehat{\mathbf{V}}_K^{(\ell)} \widehat{\mathbf{V}}_K^{(\ell)^T}$, and its K leading eigenvectors $\left\{ \widetilde{\mathbf{V}}_j \right\}_{j=1}^K$. Output: $\widetilde{V}_K = \left(\widetilde{v}_1, \cdots, \widetilde{v}_K \right) \in \mathbb{R}^{d \times K}$.

In other words, each server frst calculates the top K eigenvectors of the local sample covariance matrix, and then transmits these eigenvectors $\left\{ \widehat{\mathbf{V}}_K^{(\ell)} \right\}_{\ell=1}^m$ to a central server, where the estimators get aggregated. This procedure has similar spirit as distributed estimation based on one-shot averaging in Zhang et al. (2013), Battey et al. (2015), Lee et al. (2017), among others. To see this, we recall the SDP formulation of the eigenvalue problem. Let $\widehat{\mathbf{V}}_K = \left(\widehat{\mathbf{v}}_1, \cdots, \widehat{\mathbf{v}}_K \right)$ contain the K leading eigenvectors of $\widehat{\boldsymbol{\Sigma}} = \frac{1}{m} \sum_{\ell=1}^m \widehat{\boldsymbol{\Sigma}}^{(\ell)}$. Lemma 5 in Section 8.2.2 asserts that $\widehat{\mathbf{P}}_K = \widehat{\mathbf{V}}_K \widehat{\mathbf{V}}_K^T$ solves the SDP:

$$\min_{\mathbf{P} \in S^{d \times d}} - \mathrm{Tr}\left(\mathbf{P}^T \widehat{\boldsymbol{\Sigma}} \right)$$

$$\text{s.t } \mathrm{Tr}(\mathbf{P}) \leq K, \|\mathbf{P}\|_2 \leq 1, \mathbf{P} \succcurlyeq 0. \tag{3.1}$$

Here $S^{d \times d}$ refers to the set of $d \times d$ symmetric matrices. In the traditional setting, we have access to all the data, and $\widehat{\mathbf{P}}_K$ is a natural estimator for $\mathbf{V}_K \mathbf{V}_K^T$. In the distributed setting, each machine can only access $\widehat{\boldsymbol{\Sigma}}^{(\ell)}$. Consequently, it solves a local version of (3.1):

$$\min_{\mathbf{P} \in S^{d \times d}} - \mathrm{Tr}\left(\mathbf{P}^T \widehat{\boldsymbol{\Sigma}}^{(\ell)} \right)$$

$$\text{s.t } \mathrm{Tr}(\mathbf{P}) \leq K, \ \|\mathbf{P}\|_2 \leq 1, \ \mathbf{P} \succcurlyeq 0. \tag{3.2}$$

The optimal solution is $\widehat{\mathbf{P}}_K^{(\ell)} = \widehat{\mathbf{V}}_K^{(\ell)} \widehat{\mathbf{V}}_K^{(\ell)T}$. Since the loss function in (3.1) is the average of local loss functions in (3.2), we can intuitively average the optimal solutions $\widehat{\mathbf{P}}_K^{(\ell)}$ to approximate $\widehat{\mathbf{P}}_K$. However, the average $\frac{1}{m}\sum_{\ell=1}^{m}\widehat{\mathbf{P}}_K^{(\ell)}$ may no longer be a rank-K projection matrix. Hence a rounding step is needed, extracting the leading eigenvectors of that average to get a projection matrix.

Here is another way of understanding the aggregation procedure. Given a collection of estimators $\left\{\widehat{\mathbf{V}}_K^{(\ell)}\right\}_{\ell=1}^{m} \subseteq \mathcal{O}_{d \times K}$ and the loss $\rho(\,\cdot\,,\,\cdot\,)$, we want to find the center $\mathbf{U} \in \mathcal{O}_{d \times K}$ that minimizes the sum of squared losses $\sum_{\ell=1}^{m}\rho^2\!\left(\mathbf{U}, \widehat{\mathbf{V}}_K^{(\ell)}\right)$. Lemma 6 in Section 8.2.2 indicates that $\mathbf{U} = \widetilde{V}_K$ is an optimal solution. Therefore, our distributed PCA estimator $\widetilde{\mathbf{V}}_K$ is a generalized "center" of individual estimators.

It is worth noting that in this algorithm, we do not really need to compute $\left\{\widehat{\boldsymbol{\Sigma}}^{(\ell)}\right\}_{\ell=1}^{m}$ and $\widetilde{\boldsymbol{\Sigma}}$. $\left\{\widehat{\mathbf{V}}_K^{(\ell)}\right\}_{\ell=1}^{m}$ and $\widetilde{\mathbf{V}}_K$ can be derived from top-K SVD of data matrices. This is far more expeditious than the entire SVD and highly scalable, by using, for example, the power method (Golub and Van Loan, 2012). As regard to the estimation of the top eigenvalues of $\boldsymbol{\Sigma}$, we can send the aggregated eigenvectors $\left\{\widetilde{\mathbf{v}}_j\right\}_{j=1}^{K}$ back to the m servers, where each one computes $\left\{\lambda_j^{(\ell)}\right\}_{j=1}^{K} = \left\{\widetilde{\mathbf{v}}_j^T\,\widehat{\boldsymbol{\Sigma}}^{(\ell)}\,\widetilde{\mathbf{v}}_j\right\}_{j=1}^{K}$. Then the central server collect all the eigenvalues and deliver the average eigenvalues $\left\{\widetilde{\lambda}_j\right\}_{j=1}^{K} = \left\{\frac{1}{m}\sum_{\ell=1}^{m}\lambda_j^{(\ell)}\right\}_{j=1}^{K}$ as the estimators of all eigenvalues.

As we can see, the communication cost of the proposed distributed PCA algorithm is of order $O(mKd)$. In contrast, to share all the data or entire covariance, the communication cost will be of order $O(md\min(n, d))$. Since in most cases $K = o(\min(n, d))$, our distributed PCA requires much less communication cost than naive data aggregation.

4   Statistical error analysis

Algorithm 1 delivers $\widetilde{V}_K$ to estimate the top K eigenspace of $\boldsymbol{\Sigma}$. In this section we analyze the statistical error of $\widetilde{V}_K$, i.e., $\rho(\widetilde{V}_K, \mathbf{V}_K)$. The main message is that $\widetilde{V}_K$ enjoys the same statistical error rate as the full sample counterpart $\widetilde{V}_K$ as long as the subsample size n is sufficiently large.

We first conduct a bias and variance decomposition of $\rho(\widetilde{V}_K, \mathbf{V}_K)$, which serves as the key step in establishing our theoretical results. Recall that $\widetilde{\boldsymbol{\Sigma}} = (1/m)\sum_{\ell=1}^{m}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}$ and $\widetilde{V}_K$ consists of the top K eigenvectors of $\widetilde{\boldsymbol{\Sigma}}$. Define $\boldsymbol{\Sigma}^* := \mathbb{E}\!\left(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}\right)$ and denote its top K eigenvectors by $\mathbf{V}_K^* = \left(\mathbf{v}_1^*, \cdots, \mathbf{v}_K^*\right) \in \mathbb{R}^{d \times K}$. When the number of machines goes to infinity, $\widetilde{\boldsymbol{\Sigma}}$

converges to Σ^*, and naturally we expect $\mathrm{Col}\big(\widetilde{V}_K\big)$ to converge to $\mathrm{Col}(\mathbf{V}_K^*)$ as well. This line of thinking inspires us to decompose the statistical error $\rho(\widetilde{V}_K, \mathbf{V}_K)$ into the following bias and sample variance terms:

$$\rho\big(\widetilde{V}_K, \mathbf{V}_K\big) \leq \underbrace{\rho\big(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*\big)}_{\text{sample variance term}} + \underbrace{\rho\big(\mathbf{V}_K^*, \mathbf{V}_K\big)}_{\text{bias term}}. \tag{4.1}$$

The first term is stochastic and the second term is deterministic. Here we elucidate on why we call $\rho\big(\widetilde{V}_K, \mathbf{V}_K^*\big)$ the sample variance term and $\rho\big(\mathbf{V}_K^*, \mathbf{V}_K\big)$ the bias term respectively.

1. Sample variance term $\rho\big(\widetilde{V}_K, \mathbf{V}_K^*\big)$:

   By Davis-Kahan's Theorem (Theorem 2 in Yu et al. (2015)) and (2.1), we have

   $$\rho\big(\widetilde{V}_K, \mathbf{V}_K^*\big) \lesssim \frac{\big\|\ \widetilde{\Sigma}\ -\ \Sigma^*\big\|_F}{\lambda_K(\Sigma^*) - \lambda_{K+1}(\Sigma^*)}. \tag{4.2}$$

   As we can see, $\rho\big(\widetilde{V}_K, \mathbf{V}_K^*\big)$ depends on how the average $\widetilde{\Sigma} = \frac{1}{m}\sum_{\ell=1}^{m}\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}$ concentrates to its mean Σ^*. This explains why we call $\rho\big(\widetilde{V}_K, \mathbf{V}_K^*\big)$ the sample variance term. We will show in the sequel that for sub-Gaussian random samples, $\big\{\big\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \Sigma^*\big\|_F\big\}_{\ell=1}^{m}$ and $\big\|\ \widetilde{\Sigma}\ -\Sigma^*\big\|_F$ are sub-exponential random variables and under appropriate regularity assumptions,

   $$\big\|\ \|\ \widetilde{\Sigma}\ -\ \Sigma^*\ \|_F\ \big\|_{\psi_1} \lesssim \frac{1}{\sqrt{m}}\big\|\ \|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)T} - \Sigma^*\|_F\ \big\|_{\psi_1}. \tag{4.3}$$

   If we regard ψ_1-norm as a proxy for standard deviation, this result is a counterpart to the formula for the standard deviation of the sample mean under the context of matrix concentration. By (4.3), the average of projection matrices $\widetilde{\Sigma}$ enjoys a similar square-root convergence, so does $\rho\big(\widetilde{V}_K, \mathbf{V}_K^*\big)$.

2. Bias term $\rho\big(\mathbf{V}_K^*, \mathbf{V}_K\big)$:

   The error $\rho\big(\mathbf{V}_K^*, \mathbf{V}_K\big)$ is deterministic and independent of how many machines we have, and is therefore called the bias term. We will show this bias term is exactly zero when the random sample has a symmetric innovation (to be defined later). In general, we will show that the bias term is negligible in comparison with the sample variance term when the number of nodes m is not unreasonably large.

   In the following subsections, we will analyze the sample variance term and bias term respectively and then combine these results to obtain the convergence rate for $\rho\big(\widetilde{V}_K, \mathbf{V}_K\big)$.

4.1 Analysis of the sample variance term

To analyze $\rho\left(\widetilde{V}_K, \mathbf{V}_K^*\right)$, as shown by (4.2), we need to derive the order of the numerator $\left\|\ \widetilde{\mathbf{\Sigma}}\ -\ \mathbf{\Sigma}^*\right\|_F$ and denominator $\lambda_K(\mathbf{\Sigma}^*) - \lambda_{K+1}(\mathbf{\Sigma}^*)$. We first focus on the matrix concentration term $\left\|\ \widetilde{\mathbf{\Sigma}}\ -\ \mathbf{\Sigma}^*\right\|_F = \left\|\frac{1}{m}\sum_{\ell=1}^{m}\left(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \mathbf{\Sigma}^*\right)\right\|_F$. Note that $\widetilde{\mathbf{\Sigma}}\ -\ \mathbf{\Sigma}^*$ is an average of m centered random matrices. To establish the correspondent concentration inequality, we first investigate each individual term in the average, i.e., $\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \mathbf{\Sigma}^*$ for $\ell \in [m]$. In the following lemma, we show that when random samples are sub-Gaussian, $\left\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \mathbf{\Sigma}^*\right\|_F$ is sub-exponential and we can give an explicit upper bound of its $\psi_1 -$ norm.

Lemma 1. *Suppose that on the ℓ-th server we have n i.i.d. sub-Gaussian random samples $\{\mathbf{X}_i\}_{i=1}^{n}$ in $\mathbb{R}^d$ with covariance matrix $\mathbf{\Sigma}$. There exists a constant $C > 0$ such that when $n \geq r$,*

$$\left\|\ \|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \mathbf{\Sigma}^*\ \|_F\ \right\|_{\psi_1} \leq C\kappa\sqrt{\frac{Kr}{n}}.$$

Note that here we use the Frobenius norm to measure the distance between two matrices. Therefore, it is equivalent to treat $\left\{\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}\right\}_{\ell=1}^{K}$ and $\mathbf{\Sigma}^*$ as d^2—dimensional vectors and apply the concentration inequality for random vectors to bound $\left\|\ \widetilde{\mathbf{\Sigma}}\ -\ \mathbf{\Sigma}^*\right\|_F$. As we will demonstrate in the proof of Theorem 1, $\left\|\ \|\ \widetilde{\mathbf{\Sigma}}\ -\ \mathbf{\Sigma}^*\ \|_F\ \right\|_{\psi_1} \lesssim \frac{1}{\sqrt{m}}\left\|\ \|\widehat{\mathbf{V}}_K^{(l)}\widehat{\mathbf{V}}_K^{(l)T} - \mathbf{\Sigma}^*\ \|_F\ \right\|_{\psi_1}$.

With regard to $\lambda_K(\mathbf{\Sigma}^*) - \lambda_{K+1}(\mathbf{\Sigma}^*)$, when the individual node has enough samples, $\widehat{\mathbf{V}}_K^{(\ell)}$ and $\mathbf{V}_K$ will be close to each other and so will $\mathbf{\Sigma}^* = \mathbb{E}\left(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}\right)$ and $\mathbf{V}_K\mathbf{V}_K^T$. Given $\lambda_K\left(\mathbf{V}_K\mathbf{V}_K^T\right) = 1$ and $\lambda_{K+1}\left(\mathbf{V}_K\mathbf{V}_K^T\right) = 0$, we accordingly expect $\lambda_K(\mathbf{\Sigma}^*)$ and $\lambda_{K+1}(\mathbf{\Sigma}^*)$ be separated by a positive constant as well.

All the arguments above lead to the following theorem on $\rho\left(\widetilde{V}_K, \mathbf{V}_K^*\right)$.

Theorem 1. *Suppose $\mathbf{X}_1, \cdots, \mathbf{X}_N$ are i.i.d. sub-Gaussian random vectors in $\mathbb{R}^d$ with covariance matrix $\mathbf{\Sigma}$ and they are scattered across m machines. If $n \geq r$ and $\left\|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_2 \leq 1/4$, then*

$$\left\|\rho\left(\widetilde{V}_K, \mathbf{V}_K^*\right)\right\|_{\psi_1} \leq C\kappa\sqrt{\frac{Kr}{N}},$$

Where C is some universal constant.

4.2 Analysis of the bias term

In this section, we study the bias term $\rho\left(\mathbf{V}_K^*, \mathbf{V}_K\right)$ in (4.1). We first focus on a special case where the bias term is exactly zero. For a random vector $\mathbf{X}$ with covariance $\boldsymbol{\Sigma} = \mathbf{V}\,\boldsymbol{\Lambda}\,\mathbf{V}^T$, let $\mathbf{Z} = \boldsymbol{\Lambda}^{-\frac{1}{2}}\mathbf{V}^T\mathbf{X}$. We say $\mathbf{X}$ has symmetric innovation if $\mathbf{Z} \stackrel{d}{=} \left(\mathbf{I}_d - 2\mathbf{e}_j\mathbf{e}_j^T\right)\mathbf{Z}, \quad \forall j \in [d]$. In other words, flipping the sign of one component of $\mathbf{Z}$ will not change the distribution of $\mathbf{Z}$. Note that if $\mathbf{Z}$ has density, this is equivalent to say that its density function has the form $p\left(|z_1|, |z_2|, \cdots, |z_d|\right)$. All elliptical distributions centered at the origin belong to this family. In addition, if $\mathbf{Z}$ has symmetric and independent entries, $\mathbf{X}$ has also symmetric innovation. It turns out that when the random samples have symmetric innovation, $\boldsymbol{\Sigma}^* := \mathbb{E}\left(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}\right)$ and $\boldsymbol{\Sigma}$ share exactly the same set of eigenvectors. When we were finishing the paper, we noticed that Chen et al. (2016) had independently established a similar result for the Gaussian case.

Definition 4.1. *Let V be a K-dimensional linear subspace of $\mathbb{R}^d$. For a subspace estimator represented by $\widehat{\mathbf{V}} \in \mathcal{O}_{d \times K}$, we say it is **unbiased** for V if and only if the top K eigenspace of $\mathbb{E}\left(\widehat{\mathbf{V}}\widehat{\mathbf{V}}^T\right)$ is $\mathcal{V}$.*

If $\widehat{\mathbf{V}}_K^{(\ell)}$ is unbiased for $\mathrm{Col}(\mathbf{V}_K)$, then $\rho\left(\mathbf{V}_K^*, \mathbf{V}_K\right) = 0$ and we will only have the sample variance term in (4.1). In that case, aggregating $\left\{\widehat{\mathbf{V}}_K^{(\ell)}\right\}_{\ell=1}^m$ reduces variance and yields a better estimator $\widetilde{V}_K$. Theorem 2 shows that this is the case so long as the distribution has symmetric innovation and the sample size is large enough.

Theorem 2. *Suppose on the ℓ-th server we have n i.i.d. random samples $\left\{\mathbf{X}_i\right\}_{i=1}^n$ with covariance $\boldsymbol{\Sigma}$. If $\left\{\mathbf{X}_i\right\}_{i=1}^n$ have symmetric innovation, then $\mathbf{V}^T\boldsymbol{\Sigma}^*\mathbf{V}$ is diagonal, i.e., $\boldsymbol{\Sigma}^*$ and $\boldsymbol{\Sigma}$ share the same set of eigenvectors. Furthermore, if $\left\|\boldsymbol{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_2 < 1/2$, then $\left\{\widehat{\mathbf{V}}_K^{(\ell)}\right\}_{\ell=1}^m$ are unbiased for $\mathrm{Col}(\mathbf{V}_K)$ and $\rho\left(\mathbf{V}_K^*, \mathbf{V}_K\right) = 0$.*

It is worth pointing out that distributed PCA is closely related to aggregation of random sketches of a matrix (Halko et al., 2011; Tropp et al., 2016). To approximate the subspace spanned by the K leading left singular vectors of a large matrix $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$, we could construct a suitable random matrix $\mathbf{Y} \in \mathbb{R}^{d_2 \times n}$ with $n \geq K$, and use the left singular subspace of $\mathbf{AY} \in \mathbb{R}^{d_1 \times n}$ as an estimator. $\mathbf{AY}$ is called a random sketch of $\mathbf{A}$. It has been shown that to obtain reasonable statistical accuracy, n can be much smaller than $\min(d_1, d_2)$ as long as $\mathbf{A}$ is approximately low rank. Hence it is much cheaper to compute SVD on $\mathbf{AY}$ than on $\mathbf{A}$. When we want to aggregate a number of such subspace estimators, a smart choice of the random matrix ensemble for $\mathbf{Y}$ is always preferable. It follows from Theorem 2 that if we let $\mathbf{Y}$ have i.i.d. columns from a distribution with symmetric innovation (e.g., Gaussian distribution or

independent entries), then the subspace estimators are unbiased, which facilitates aggregation.

Here we explain why we need the condition $\left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_2 < 1/2$ to achieve zero bias. First of all, the condition is similar to a bound on the "variance" of the random matrix $\widehat{\mathbf{V}}_K^{(\ell)}$ whose covariance $\mathbf{\Sigma}^*$ is under investigation. As demonstrated above, with the symmetric innovation, $\mathbf{\Sigma}^*$ has the same set of eigenvectors as $\mathbf{\Sigma}$, but we still cannot guarantee that the top K eigenvectors of $\mathbf{\Sigma}^*$ match with those of $\mathbf{\Sigma}$. For example, the $(K+1)$-th eigenvector of $\mathbf{\Sigma}$ might be the K-th eigenvector of $\mathbf{\Sigma}^*$. In order to ensure the top K eigenspace of $\mathbf{\Sigma}^*$ is exactly the same as that of $\mathbf{\Sigma}$, we require $\widehat{\mathbf{V}}_K^{(\ell)}$ to not deviate too far from $\mathbf{V}_K$ so that $\mathbf{\Sigma}^*$ is close enough to $\mathbf{V}_K \mathbf{V}_K^T$. Both Theorems 1 and 2 require control of $\left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_2$, which will be studied shortly.

For general distributions, the bias term is not necessarily zero. However, it turns out that when the subsample size is large enough, the bias term $\rho\left(\mathbf{V}_K^*, \mathbf{V}_K\right)$ is of high-order compared with the statistical error of $\widehat{\mathbf{V}}_K^{(\ell)}$ on the individual subsample. By the decomposition (4.1) and Theorem 1, we can therefore expect the aggregated estimator $\widetilde{V}_K$ to enjoy sharper statistical error rate than PCA on the individual subsample. In other words, the aggregation does improve the statistical efficiency. A similar phenomenon also appears in statistical error analysis of the average of the debiased Lasso estimators in Battey et al. (2015) and Lee et al. (2017). Recall that in sparse linear regression, the Lasso estimator $\widehat{\beta}$ satisfies that $\left\| \widehat{\mathbf{\beta}} - \mathbf{\beta}^* \right\|_2 = O_P(\sqrt{s\log d/n})$, where β^* is the true regression vector, s is the number of nonzero coefficients of β^* and d is the dimension. The debiasing step reduces the bias of $\widehat{\beta}$ to the order $O_P(s \log d/n)$, which is negligible when m is not too large, compared with the statistical error of $\widehat{\beta}$ and thus enables the average of the debiased Lasso estimators to enhance the statistical efficiency.

Below we present Lemma 2, a high-order Davis-Kahan theorem that explicitly characterizes the linear term and high-order error on top K eigenspace due to matrix perturbation. This is a genuine generalization of the former high-order perturbation theorems on a single eigenvector, e.g., Lemma 1 in Kneip and Utikal (2001) and Theorem 2 in El Karoui and d'Aspremont (2010). An elegant result on eigenspace perturbation is Lemma 2 in Koltchinskii and Lounici (2016). Our error bound uses Frobenius norm while theirs uses spectral norm. Besides, when the top K eigenspace is of interest, the upper bound in Lemma 2 in Koltchinskii and Lounici (2016) contains an extra factor $1 + \left(\lambda_1 - \lambda_K\right)/\Delta$. Hence we have better dependence on problem parameters. Other related works in the literature consider asymptotic expansions of perturbation (Kato, 1966; Vaccaro, 1994; Xu, 2002), and singular space of a matrix contaminated by Gaussian noise (Wang, 2015). Our result is both non-asymptotic and deterministic. It serves as the core of bias analysis.

Lemma 2. *Let* $\mathbf{A}, \widehat{\mathbf{A}} \in \mathbb{R}^{d \times d}$ *be symmetric matrices with eigenvalues* $\lambda_1 \geq \cdots \geq \lambda_d$, *and* $\hat{\lambda}_1 \geq \cdots \geq \hat{\lambda}_d$, *respectively. Let* $\left\{ \mathbf{u}_j \right\}_{j=1}^d$, $\left\{ \hat{\mathbf{u}}_j \right\}_{j=1}^d$ *be two orthonormal bases of* $\mathbb{R}^d$ *such that* $\mathbf{A}\mathbf{u}_j = \lambda_j \mathbf{u}_j$ *and* $\widehat{\mathbf{A}}\hat{\mathbf{u}}_j = \hat{\lambda}_j \hat{\mathbf{u}}_j$ *for all* $j \in [d]$ *Fix* $s \in \{0, 1, \cdots, d - K\}$ *and assume that*

$\Delta = \min\left\{ \lambda_s - \lambda_{s+1}, \lambda_{s+K} - \lambda_{s+K+1} \right\} > 0$, *where* $\lambda_0 = +\infty$ *and* $\lambda_{d+1} = -\infty$. *Define* $\mathbf{U} = \left(\mathbf{u}_{s+1}, \cdots, \mathbf{u}_{s+K} \right), \widehat{\mathbf{U}} = \left(\hat{\mathbf{u}}_{s+1}, \cdots, \hat{\mathbf{u}}_{s+K} \right)$. *Define* $\mathbf{E} = \widehat{\mathbf{A}} - \mathbf{A}, S = \left\{ s+1, \cdots, s+K \right\}, \mathbf{G}_j = \sum_{i \notin S} \left(\lambda_i - \lambda_{s+j} \right)^{-1} \mathbf{u}_i \mathbf{u}_i^T$ *for* $j \in [K]$, *and*

$$ f : \mathbb{R}^{d \times K} \to \mathbb{R}^{d \times K}, \quad \left(\mathbf{w}_1, \cdots, \mathbf{w}_K \right) \mapsto \left(-\mathbf{G}_1 \mathbf{w}_1, \cdots, -\mathbf{G}_K \mathbf{w}_K \right). $$

When $\varepsilon = \|\mathbf{E}\|_2 / \Delta \leq 1/10$, we have

$$ \left\| \widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T - \left[f(\mathbf{E}\mathbf{U})\mathbf{U}^T + \mathbf{U}f(\mathbf{E}\mathbf{U})^T \right] \right\|_F \leq 24\sqrt{K}\varepsilon^2. $$

Similar to Taylor expansion, the difference is decomposed into the linear leading term and residual of higher order with respect to the perturbation. Here we only present a version that is directly applicable to bias analysis. Stronger results are summarized in Lemma 8 in Section 8.2.2, which may be of independent interest in perturbation analysis of spectral projectors.

Now we apply Lemma 2 to the context of principal eigenspace estimation. Let $\mathbf{A} = \Sigma, \widehat{\mathbf{A}} = \widehat{\Sigma}^{(1)}$ and $S = [K]$. It thus follows that $\mathbf{U} = \mathbf{V}_K, \widehat{\mathbf{U}} = \widehat{\mathbf{V}}_K^{(1)}$ and $\mathbf{E} = \widehat{\Sigma}^{(1)} - \Sigma$. From the second inequality in Lemma 2 we can conclude that the bias term $\rho(\mathbf{V}_K^*, \mathbf{V}_K)$ is a high-order term compared with the linear leading term. More specifically, the Davis-Kahan theorem helps us control the bias as follows:

$$ \rho\left(\mathbf{V}_K^*, \mathbf{V}_K \right) \lesssim \left\| \Sigma^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_F = \left\| \mathbb{E}[\widehat{\mathbf{V}}_K^{(1)} \widehat{\mathbf{V}}_K^{(1)T} - \mathbf{V}_K \mathbf{V}_K^T] \right\|_F. $$

By the facts that $\mathbb{E}(\mathbf{E}) = 0$ and f is linear, we have

$$ \rho\left(\mathbf{V}_K^*, \mathbf{V}_K \right) = \left\| \mathbb{E}\left[\widehat{\mathbf{V}}_K^{(1)} \widehat{\mathbf{V}}_K^{(1)T} - \left(\mathbf{V}_K \mathbf{V}_K^T + f(\mathbf{E}\mathbf{V}_K)\mathbf{V}_K^T + \mathbf{V}_K f(\mathbf{E}\mathbf{V}_K)^T \right) \right] \right\|_F. $$

By Jensen's inequality, the right hand side above is further bounded by

$$ \mathbb{E}\left\| \widehat{\mathbf{V}}_K^{(1)} \widehat{\mathbf{V}}_K^{(1)T} - \left(\mathbf{V}_K \mathbf{V}_K^T + f(\mathbf{E}\mathbf{V}_K)\mathbf{V}_K^T + \mathbf{V}_K f(\mathbf{E}\mathbf{V}_K)^T \right) \right\|_F. \tag{4.4} $$

When n is large enough, the typical size of $\varepsilon = \|\mathbf{E}\|_2 / \Delta$ is small, and Lemma 1 controls it tail and all of the moments. Together with Lemma 2, this fact implies that (4.4) has roughly

the same order as $\sqrt{K} \cdot \mathbb{E} \varepsilon^2$, which should be much smaller than the typical size of $\sqrt{K} \varepsilon$, i.e. the upper bound for $\rho\left(\widehat{\mathbf{V}}_K^{(1)}, \mathbf{V}_K\right)$ given by Davis-Kahan theorem. The following theorem makes our hand-waving analysis rigorous.

Theorem 3. *There are constants C_1 and C_2 such that when $n \geq r$,*

$$\rho\left(\mathbf{V}_K^*, \mathbf{V}_K\right) \leq C_1 \left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_F \leq C_2 \kappa^2 \sqrt{K} r/n.$$

As a by-product, we get $\left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_2 \lesssim \kappa^2 \sqrt{K} r/n$. Hence when $n \geq C\kappa^2 \sqrt{K} r$ for some large enough C, the assumptions in Theorems 1 and 2 on $\left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_2$ are guaranteed to hold.

4.3 Properties of distributed PCA

We now combine the results we obtained in the previous two subsections to derive the statistical error rate of $\widetilde{V}_K$. We first present a theorem under the setting of global i.i.d. data and discuss its optimality.

Theorem 4. *Suppose we have N i.i.d. sub-Gaussian random samples with covariance Σ. They are scattered across m servers, each of which stores n samples. There exist constants C, C_1, C_2, C_3 and C_4 such that the followings hold when $n \geq C\kappa^2 \sqrt{K} r$.*

1. *Symmetric innovation:*

$$\left\| \rho\left(\widetilde{V}_K, \mathbf{V}_K \right) \right\|_{\psi_1} \leq C_1 \kappa \sqrt{\frac{Kr}{N}}. \tag{4.5}$$

2. *General distribution:*

$$\left\| \rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K \right) \right\|_{\psi_1} \leq C_1 \kappa \sqrt{\frac{Kr}{N}} + C_2 \kappa^2 \frac{\sqrt{K} r}{n}. \tag{4.6}$$

Furthermore, if we further assume $m \leq C_3 n/\left(\kappa^2 r\right)$,

$$\left\| \rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K \right) \right\|_{\psi_1} \leq C_4 \kappa \sqrt{\frac{Kr}{N}}. \tag{4.7}$$

As we can see, with appropriate scaling conditions on n, m and d, $\widetilde{V}_K$ can achieve the statistical error rate of order $\kappa\sqrt{Kr/N}$. The result is applicable to the whole sample or traditional PCA, in which $m = 1$. Hence the distributed PCA and the traditional PCA share the same error bound as long as the technical conditions are satisfied.

In the second part of Theorem 4, the purpose of setting restrictions on n and m is to ensure that the distributed PCA algorithm delivers the same statistical rate as the centralized PCA

which uses all the data. In the boundary case where $n \asymp \kappa^2 \sqrt{K} r$, the bias of the local empirical eigenspace is of constant order. Since our aggregation cannot kill bias, there is no hope to achieve the centralized rate unless the number of machines is of constant order so that the centralized PCA has constant error too. Besides, our result says that when n is large, we can tolerate more data splits (larger m) for achieving the centralized statistical rate.

We now illustrate our result through a simple spiked covariance model introduced by Johnstone (2001). Assume that $\mathbf{\Lambda} = \operatorname{diag}(\lambda, \underbrace{1, \cdots, 1}_{d-1})$, where $\lambda > 1$, and we are interested in the first eigenvector of $\mathbf{\Sigma}$. Note that $K = 1$, $r = \operatorname{Tr}(\mathbf{\Sigma})/\|\mathbf{\Sigma}\|_2 = (\lambda + d - 1)/\lambda \asymp d/\lambda$ when $\lambda = O(d)$, and $\kappa = \lambda/(\lambda - 1) \asymp 1$. It is easy to see from (4.5) or (4.7) that

$$\left\| \rho(\widetilde{V}_1, \mathbf{V}_1) \right\|_{\psi_1} \lesssim \kappa \sqrt{\frac{r}{N}} \lesssim \sqrt{\frac{d}{N\lambda}}.$$

Without loss of generality, we could always assume that the direction of $\widetilde{V}_1$ is chosen such that $\widetilde{\mathbf{V}}_1^T \mathbf{V}_1 \geq 0$, i.e. $\widetilde{V}_1$ is aligned with $\mathbf{V}_1$. Note that

$$\rho^2(\widetilde{\mathbf{V}}_1, \mathbf{V}_1) = \left\| \widetilde{V}_1 \widetilde{\mathbf{V}}_1^T - \mathbf{V}_1 \mathbf{V}_1^T \right\|_F^2 = 2\left(1 - \widetilde{V}_1^T \mathbf{V}_1\right)\left(1 + \widetilde{V}_1^T \mathbf{V}_1\right) \geq 2\left(1 - \widetilde{V}_1^T \mathbf{V}_1\right) = \left\| \widetilde{V}_1 - \mathbf{V}_1 \right\|_2^2.$$

Hence

$$\mathbb{E}\left\| \widetilde{V}_1 - \mathbf{V}_1 \right\|_2^2 \lesssim \left\| \rho(\widetilde{V}_1, \mathbf{V}_1) \right\|_{\psi_1}^2 \lesssim \frac{d}{N\lambda}. \tag{4.8}$$

We now compare this rate with the previous results under the spiked model. In Paul and Johnstone (2012), the authors derived the ℓ_2 risk of the empirical eigenvectors when random samples are Gaussian. It is not hard to derive from Theorem 1 therein that given N i.i.d d-dimensional Gaussian samples, when N, d and λ go to infinity,

$$\mathbb{E}\left\| \widehat{\mathbf{V}}_1 - \mathbf{V}_1 \right\|_2^2 \asymp \frac{d}{N\lambda},$$

Where $\widehat{\mathbf{V}}_1$ is the empirical leading eigenvector with $\widehat{\mathbf{V}}_1^T \mathbf{V}_1 \geq 0$. We see from (4.8) that the aggregated estimator $\widetilde{V}_1$ performs as well as the full sample estimator $\widehat{\mathbf{V}}_1$ in terms of the mean squared error. See Wang and Fan (2017) for generalization of the results for spiked covariance.

In addition, our result is consistent with the minimax lower bound developed in Cai et al. (2013). For $\lambda > 0$ and fixed $c \geq 1$, define

$$\Theta = \left\{ \mathbf{\Sigma} \text{ is symmetric and } \mathbf{\Sigma} \succcurlyeq 0 : \lambda + 1 \leq \lambda_K \leq \lambda_1 \leq c\lambda + 1, \lambda_j = 1 \text{ for } K + 1 \leq j \leq d \right\}.$$

Assume that $K \leq d/2$ and $1 \lesssim d/\lambda \lesssim N$. Theorem 8 in Cai et al. (2013) shows that under the Gaussian distribution with $\Sigma \in \Theta$, the minimax lower bound of $\mathbb{E}\rho^2(\widehat{\mathbf{V}}, \mathbf{V}_K)$ satisfies

$$\inf_{\widehat{\mathbf{V}}} \sup_{\Sigma \in \Theta} \mathbb{E}\rho^2(\widehat{\mathbf{V}}, \mathbf{V}_K) \gtrsim \min\left\{K, (d-K), \frac{K(\lambda+1)(d-K)}{N\lambda^2}\right\} \gtrsim \frac{Kd}{N\lambda}. \tag{4.9}$$

Based on $r = Tr(\Sigma)/\|\Sigma\|_2 \leq (cK\lambda + d)/(c\lambda + 1) \lesssim Kd/\lambda$ and $\kappa \leq c \lesssim 1$, our (4.5) gives an upper bound

$$\mathbb{E}\rho^2(\widetilde{v}_1, \mathbf{V}_1) \lesssim \kappa^2 \frac{Kr}{n} \lesssim \frac{Kd}{N\lambda},$$

which matches the lower bound in (4.9).

Although the upper bound $\kappa\sqrt{Kr/N}$ established in Theorem 4 is optimal in the minimax sense as discussed above, the non-minimax risk of empirical eigenvectors can be improved when the condition number k is large. See Vu et al. (2013), Koltchinskii and Lounici (2016) and Reiss and Wahl (2016) for sharper results. We use (4.5) as a benchmark rate for the centralized PCA only for the sake of simplicity.

Notice that in Theorem 4, the prerequisite for $\widetilde{v}_K$ to enjoy the sharp statistical error rate is a lower bound on the subsample size n, i.e.,

$$n \gtrsim k^2\sqrt{K}r. \tag{4.10}$$

As in the remarks after Lemma 2, this is the condition we used to ensure closeness between Σ^* and $\mathbf{V}_K\mathbf{V}_K^T$. It is natural to ask whether this required sample complexity is sharp, or in other words, is it possible for $\widetilde{V}_K$ to achieve the same statistical error rate with a smaller sample size on each machine? The answer is no. The following theorem presents a distribution family under which $\mathrm{Col}(\widetilde{V}_K)$ is even perpendicular to $\mathrm{Col}(\mathbf{V}_K)$ with high probability when n is smaller than the threshold given in (4.10). This means that having a smaller sample size on each machine is too uninformative such that the aggregation step completely fails in improving estimation consistency.

Theorem 5. *Consider a Bernoulli random variable W with $P(W = 0) = P(W = 1) = 1/2$, a Rademacher random variable $P(Y = 1) = P(Y = -1) = 1/2$, and a random vector $Z \in \mathbb{R}^{d-1}$ that is uniformly distributed over the $(d-1)$-dimensional unit sphere. For $\lambda \geq 2$, we say a random vector $\mathbf{X} \in \mathbb{R}^d$ follows the distribution $\mathcal{D}(\lambda)$ if*

$$\mathbf{X} \stackrel{d}{=} \begin{pmatrix} 1_{\{W=0\}}\sqrt{2\lambda}Y \\ 1_{\{W=1\}}\sqrt{2(d-1)}\mathbf{Z} \end{pmatrix}.$$

Now suppose we have $\left\{\mathbf{X}_i\right\}_{i=1}^N$ as N i.i.d. random samples of $\mathbf{X}$. They are stored across m servers, each of which has n samples. When $32\log d \le n \le (d-1)/(3\lambda)$, we have

$$P\left(\widetilde{V}_1 \perp \mathbf{V}_1\right) \ge \begin{cases} 1 - d^{-1}, & \text{if } m \le d^3 \\ 1 - e^{-d/2}, & \text{if } m > d^3 \end{cases}.$$

It is easy to verify that $\mathscr{D}(\lambda)$ is symmetric, sub-Gaussian and satisfies $\mathbb{E}\mathbf{X} = 0$ and $\mathbb{E}\left(\mathbf{X}\mathbf{X}^T\right) = diag\,(\lambda, 1, \cdots, 1)$. Besides, $\kappa = \lambda/(\lambda-1) \asymp 1$ and $r = (\lambda+d-1)/\lambda = d/\lambda + 1 - \lambda^{-1} \asymp d/\lambda$ when $2 \le \lambda \lesssim d$.

According to (4.10), we require $n \gtrsim d/\lambda$ to achieve the rate as demonstrated in (4.5). Theorem 5 shows that if we have fewer samples than this threshold, the aggregated estimator $\widetilde{V}_1$ will be perpendicular to the true top eigenvector $\mathbf{V}_1$ with high probability. Therefore, our lower bound for the subsample size n is sharp.

5 Extension to heterogeneous samples

We now relax global *i.i.d.* assumptions in the previous section to the setting of heterogeneous covariance structures across servers. Suppose data on the server ℓ has covariance matrix $\mathbf{\Sigma}^{(\ell)}$, whose top K eigenvalues and eigenvectors are denoted by $\left\{\lambda_k^{(\ell)}\right\}_{k=1}^K$ and $\mathbf{V}_K^{(\ell)} = \left(\mathbf{v}_1^{(\ell)}, \cdots, \mathbf{v}_K^{(\ell)}\right)$ respectively. We will study two specific cases of heterogeneous covariances: one requires all covariances to share exactly the same principal eigenspaces, while the other considers the heterogeneous factor models with common factor eigen-structures.

5.1 Common principal eigenspaces

We assume that $\left\{\mathbf{\Sigma}^{(\ell)}\right\}_{\ell=1}^m$ share the same top K eigenspace, i.e. there exists some $\mathbf{V}_K \in \mathcal{O}_{d \times K}$ such that $\mathbf{V}_K^{(\ell)}\mathbf{V}_K^{(\ell)T} = \mathbf{V}_K\mathbf{V}_K^T$ for all $\ell \in [m]$. The following theorem can be viewed as a generalization of Theorem 4.

Theorem 6. *Suppose we have in total N sub-Gaussian samples scattered across m servers, each of which stores n i.i.d. samples with covariance $\mathbf{\Sigma}^{(\ell)}$. Assume that $\left\{\mathbf{\Sigma}^{(\ell)}\right\}_{\ell=1}^m$ share the same top K eigenspace. For each $\ell \in [m]$, let $S_\ell = \kappa_\ell\sqrt{\dfrac{Kr_e}{N}}$ and $B_\ell = \dfrac{\kappa_\ell^2\sqrt{K}r_\ell}{n}$, where $r_\ell := Tr\left(\mathbf{\Sigma}^{(\ell)}\right)/\lambda_1^{(\ell)}$ and $\kappa_\ell := \lambda_1^{(\ell)}/\left(\lambda_K^{(\ell)} - \lambda_{K+1}^{(\ell)}\right)$.*

    1. *Symmetric innovation: There exist some positive constants C and C_1 such that*

$$\left\|\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)\right\|_{\psi_1} \le C_1 \sqrt{\frac{1}{m} \sum_{\ell=1}^{m} S_\ell^2} \tag{5.1}$$

so long as $n \ge C\sqrt{K}\max_{\ell \in [m]}\left(\kappa_\ell^2 r_\ell\right)$.

2. *General distribution: There exist positive constant C_2 and C_3 such that when* $n \ge \max_{\ell \in [m]} r_\ell$,

$$\left\|\rho\left(\widetilde{V}_K, \mathbf{V}_K\right)\right\|_{\psi_1} \le C_2 \sqrt{\frac{1}{m} \sum_{\ell=1}^{m} S_\ell^2} + \frac{C_3}{m} \sum_{\ell=1}^{m} B_{(\ell)}. \tag{5.2}$$

5.2 Heterogeneous factor models

Suppose on the server ℓ, the data conform to a factor model as below.

$$\mathbf{X}_i^{(\ell)} = \mathbf{B}^{(\ell)}\mathbf{f}_i^{(\ell)} + \mathbf{u}_i^{(\ell)}, \quad i \in [n],$$

Where $\mathbf{B}^{(\ell)} \in \mathbb{R}^{d \times K}$ is the loading matrix, $\mathbf{f}_i^{(\ell)} \in \mathbb{R}^K$ is the factor that satisfies $\text{Cov}\left(\mathbf{f}_i^{(\ell)}\right) = \mathbf{I}$ and $\mathbf{u}_i^{(\ell)} \in \mathbb{R}^d$ is the residual vector. It is not hard to see that $\boldsymbol{\Sigma}^{(\ell)} = \text{Cov}\left(\mathbf{X}_i^{(\ell)}\right) = \mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)T} + \boldsymbol{\Sigma}_u^{(\ell)}$, where $\boldsymbol{\Sigma}_u^{(\ell)}$ is the covariance matrix of $\mathbf{u}_i^{(\ell)}$.

Let $\mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)T} = \mathbf{V}_K^{(\ell)}\boldsymbol{\Lambda}_K^{(\ell)}\mathbf{V}_K^{(\ell)T}$ be the spectral decomposition of $\mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)T}$. We assume that there exists a projection matrix $\mathbf{P}_K = \mathbf{V}_K\mathbf{V}_K^T$, where $\mathbf{V}_K \in \mathscr{O}_{d \times K}$, such that $\mathbf{V}_K^{(\ell)}\mathbf{V}_K^{(\ell)T} = \mathbf{P}_K$ for all $\ell \in [m]$ In other words, $\left\{\mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)T}\right\}_{\ell=1}^m$ share the same top K eigenspace. Given the context of factor models, this implies that the factors have similar impact on the variation of the data across servers. Our goal now is to recover $\text{Col}(\mathbf{V}_K)$ by the distributed PCA approach, namely Algorithm 1.

Recall that $\widehat{\boldsymbol{\Sigma}}^{(\ell)} = \frac{1}{n}\Sigma_{l=1}^n \mathbf{X}_i^{(\ell)}\mathbf{X}_i^{(\ell)T}$ is the sample covariance matrix on the ℓ-th machine, and $\widehat{\mathbf{V}}_K^{(\ell)} = \left(\widehat{\mathbf{v}}_1^{(\ell)}, \cdots, \widehat{\mathbf{v}}_K^{(\ell)}\right) \in \mathscr{O}_{d \times K}$ stores K leading eigenvectors of $\widehat{\boldsymbol{\Sigma}}^{(\ell)}$. Define $\widetilde{\boldsymbol{\Sigma}} = \frac{1}{m}\Sigma_{l=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}$, and let $\widetilde{V}_K \in \mathscr{O}_{d \times K}$ be the top K eigenvectors of $\widetilde{\boldsymbol{\Sigma}}$. Below we present a theorem that characterizes the statistical performance of the distributed PCA under the heterogeneous factor models.

Theorem 7. *For each $\ell \in [m]$, let $S_\ell = \kappa_\ell\sqrt{\frac{Kr_\ell}{N}}$ and $B_\ell = \frac{\kappa_\ell^2\sqrt{K}r_\ell}{n}$. There exist some positive constants C_1, C_2 and C_3 such that when $n\, n \ge \max_{\ell \in [m]} r_\ell$,*

$$\left\| \rho\!\left(\widetilde{V}_K, \mathbf{V}_K\right) \right\|_{\psi_1} \le C_1 \sqrt{\frac{1}{m} \sum_{\ell=1}^{m} S_\ell^2} + \frac{C_2}{m} \sum_{\ell=1}^{m} B_\ell + C_3 \frac{\sqrt{K}}{m} \sum_{\ell=1}^{m} \frac{\left\| \mathbf{\Sigma}_u^{(\ell)} \right\|_2}{\lambda_K\!\left(\mathbf{\Lambda}_K^{(\ell)}\right)} . \tag{5.3}$$

The first two terms in the RHS of (5.3) are similar to those in (5.2), while the third term characterizes the effect of heterogeneity in statistical efficiency of $\widetilde{V}_K$. When $\left\| \mathbf{\Sigma}_u^{(\ell)} \right\|_2$ is small compared with $\lambda_K\!\left(\mathbf{\Lambda}_K^{(\ell)}\right)$ as in spiky factor models, $\mathbf{\Sigma}_u^{(\ell)}$ can hardly distort the eigenspace $\mathrm{Col}\!\left(\mathbf{V}_K\right)$ and thus has little influence on the final statistical error of $\widetilde{\mathbf{V}}_K$.

6 Simulation study

In this section, we conduct Monte Carlo simulations to validate the statistical error rate of $\widetilde{\mathbf{V}}_K$ that is established in the previous section. We also compare the statistical accuracy of $\widetilde{\mathbf{V}}_K$ and its full sample counterpart $\widehat{\mathbf{V}}_K$, that is, the empirical top K eigenspace based on the full sample covariance. The main message is that our proposed distributed estimator performs equally well as the full sample estimator $\widehat{\mathbf{V}}_K$ when the subsample size n is large enough.

6.1 Verification of the statistical error rate

Consider $\{\mathbf{x}_i\}_{i=1}^{N}$ i.i.d. following $N(\mathbf{0}, \mathbf{\Sigma})$, where $\mathbf{\Sigma} = \mathrm{diag}(\lambda, \lambda/2, \lambda/4, 1, \cdots, 1)$. Here the number of spiky eigenvalues $K = 3$ and $\mathbf{V}_K = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$. We generate m subsamples, each of which has n samples, and run our proposed distributed PCA algorithm (Algorithm 1) to calculate $\widetilde{V}_K$. Since the centered multivariate Gaussian distribution is symmetric, according to Theorem 4, when $\lambda = O(d)$ we have

$$\left\| \rho\!\left(\widetilde{V}_K, \mathbf{V}_K\right) \right\|_{\psi_1} = O\!\left(\frac{C_1 \left\| \mathbf{\Sigma} \right\|_2}{\lambda_K - \lambda_{K+1}} \sqrt{\frac{Kr(\mathbf{\Sigma})}{N}} \right) = O\!\left(\sqrt{\frac{d}{mn\delta}} \right), \tag{6.1}$$

Where $\delta := \lambda_K - \lambda_{K+1} = \lambda/4 - 1$. Now we provide numerical verification of the order of the number of servers m, the eigengap δ, the subsample size n and dimension d in the statistical error.

Figure 1 presents four plots that demonstrate how $\rho\!\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)$ changes as d, m, n and δ increases respectively. Each data point on the plots is based on 100 independent Monte Carlo simulations. Figure 1(a) demonstrates how $\rho\!\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)$ increases with respect to the increasing dimension d when $\lambda = 50$ and $n = 2000$. Each line on the plot represents a fixed number of machines m. Figure 1(b) shows the decay rate of $\rho\!\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)$ as the number of servers m increases when $\lambda = 50$ and $n = 2000$. Different lines on the plot correspond to different dimensions d. Figure 1(c) demonstrates how $\rho\!\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)$ decays as the subsample size n increases when $= 50$ and $m = 50$. Figure 1(d) shows the relationship between $\rho\!\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)$ and

the eigengap δ when $d = 800$ and $n = 2000$. The results from Figures 1(a)-1(d) show that $\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)$ is proportion to $d^{\frac{1}{2}}, m^{-\frac{1}{2}}, n^{-\frac{1}{2}}$ and $\delta^{-\frac{1}{2}}$ respectively when the other three parameters are fixed. These empirical results are all consistent with (6.1).

Figure 1 demonstrates the marginal relationship between $\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)$ and the four parameters m, n, d and δ. Now we study their joint relationship. Inspired by (6.1), we consider a multiple regression model as follows:

$$\log\left(\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)\right) = \beta_0 + \beta_1 \log(d) + \beta_2 \log(m) + \beta_3 \log(n) + \beta_4 \log(\delta) + \varepsilon, \qquad (6.2)$$

where ε is the error term. We collect all the data points $\left(d, m, n, \delta, \rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)\right)$ from four plots in Figure 1 to fit the regression model (6.2). The fitting result is that $\hat{\beta}_1 = 0.5043, \hat{\beta}_2 = -0.4995, \hat{\beta}_3 = -0.5011$ and $\hat{\beta}_4 = -0.5120$ with the multiple $R^2 = 0.99997$. These estimates are quite consistent with the theoretical results in (6.1). Moreover, Figure 2 plots all the observed values of $\log\left(\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)\right)$ against its fitted values by the linear model (6.2). We can see that the observed and fitted values perfectly match. It indicates that the multiple regression model (6.2) well explains the joint relationship between the statistical error and the four parameters m, n, d and δ.

6.2 The effects of splitting

In this section we investigate how the number of data splits m affects the statistical performance of $\widetilde{\mathbf{V}}_K$ when the total sample size N is fixed. Since $N = mn$, it is easy to see that the larger m is, the smaller n will be, and hence the less computational load there will be on each individual server. In this way, to reduce the time consumption of the distributed algorithm, we prefer more splits of the data. However, per the assumptions of Theorem 4, the subsample size n should be large enough to achieve the optimal statistical performance of $\widetilde{\mathbf{V}}_K$ This motivates us to numerically illustrate how $\rho\left(\widetilde{V}_K, \mathbf{V}_K\right)$ changes as m increases with N fixed.

We adopt the same data generation process as described in the beginning of Section 6.1 with $\lambda = 50$ and $N = 6000$. We split the data into m subsamples where m is chosen to be all the factors of N that are less than or equal to 300. Figure 3 plots with respect to the number of machines m. Each point on the plot is based on 100 simulations. Each line corresponds to a different dimension d.

The results show that when the number of machines is not unreasonably large, or equivalently the number of subsample size n is not small, the statistical error does not depend on the number of machines when N is fixed. This is consistent with (6.1) where the statistical error rate only depends on the total sample size $N = mn$. When the number of machines m is large ($\log m \geq 5$), or the subsample size n is small, we observe slightly growing statistical error of the distributed PCA. This is aligned with the required lower bound of n in Theorem 4 to achieve the optimal statistical performance of $\widetilde{V}_K$. Note that

even when $m = 300$ ($\log(m) \approx 5.7$) and $n = 20$, our distributed PCA performs very well. This demonstrates that distributed PCA is statistically efficient as long as m is within a reasonable range.

6.3 Comparison between distributed and full sample PCA

In this subsection, we compare the statistical performance of the following three methods:

1. Distributed PCA (DP)

2. Full sample PCA (FP), i.e., the PCA based on the all the samples

3. Distributed PCA with communication of five additional largest eigenvectors (DP5).

Here we explain more on the third method DP5. The difference between DP5 and DP is that on each server, DP5 calculates $\widehat{\mathbf{V}}_{K+5}^{(\ell)}$, the top $K + 5$ eigenvectors of $\mathbf{\Sigma}^{(\ell)}$ and send them to the central server, and the central server computes the top K eigenvectors of $(1/m) \Sigma_{l=1}^{m} \widehat{\mathbf{V}}_{K+5}^{(\ell)} \widehat{\mathbf{V}}_{K+5}^{(\ell)T}$ as the final output. Intuitively, DP5 communicates more information of the covariance structure and is designed to guide the spill-over effects of the eigenspace spanned by the top K eigenvalues. In Figure 4, we compare the performance of all the three methods under various scenarios.

From Figures 4(a)-4(d), we can see that all the three methods have similar finite sample performance. This means that it suffices to communicate K eigenvectors to enjoy the same statistical accuracy as the full sample PCA. For more challenging situations with large $p/(mn\delta)$ ratios, small improvements using FP are visible.

7. Discussion

Our theoretical results are established under sub-Gaussian assumptions of the data. We believe that similar results will hold under distributions with heavier tails than sub-Gaussian tails, or more specifically, with only bounded fourth moment. Typical examples are Student t-distributions with more than four degrees of freedom, Pareto distribution, etc. The only difference is that with heavy-tailed distribution, if the local estimators are still the top eigenspaces of the sample covariance matrix, we will not be able to derive exponential deviation bounds. To establish statistical rate with exponential deviation, special treatments of data, including shrinkage (Fan et al., 2016; Minsker, 2016; Wei and Minsker, 2017), are needed, and the bias induced by such treatments should be carefully controlled. This will be an interesting future problem to study.

8 Proofs and technical lemmas

8.1 Proof of main results

8.1.1 Proof of Lemma 1—*Proof.* It follows from concentration of sample covariance matrix (Lemma 3) that $\left\| \left\| \widehat{\mathbf{\Sigma}}^{(1)} - \mathbf{\Sigma} \right\|_2 \right\|_{\psi_1} \lesssim \lambda_1 \sqrt{r/n}$. By the variant of Davis-Kahan theorem in Yu et al. (2015),

$$\rho\left(\widehat{\mathbf{V}}_K^{(1)}, \mathbf{V}_K\right) = \left\|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)T} - \mathbf{V}_K\mathbf{V}_K^T\right\|_F = \sqrt{2}\sin\Theta\left(\widehat{\mathbf{V}}_K^{(1)}, \mathbf{V}_K\right) \lesssim \sqrt{K}\|\widehat{\boldsymbol{\Sigma}}^{(1)} - \boldsymbol{\Sigma}\|_2/\Delta.$$

Hence

$$\left\|\rho\left(\widehat{\mathbf{V}}_K^{(1)}, \mathbf{V}_K\right)\right\|_{\psi_1} \lesssim \sqrt{K}\left\|\left\|\widehat{\boldsymbol{\Sigma}}^{(1)} - \boldsymbol{\Sigma}\right\|_2\right\|_{\psi_1}/\Delta \lesssim \kappa\sqrt{Kr/n}.$$

By Jensen's inequality,

$$\left\|\boldsymbol{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_F = \left\|\mathbb{E}\left(\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)T}\right) - \mathbf{V}_K\mathbf{V}_K^T\right\|_F \le \mathbb{E}\left\|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)T} - \mathbf{V}_K\mathbf{V}_K^T\right\|_F = \mathbb{E}\rho\left(\widehat{\mathbf{V}}_K^{(1)}, \mathbf{V}_K\right) \le$$
$$\left\|\rho\left(\widehat{\mathbf{V}}_K^{(1)}, \mathbf{V}_K\right)\right\|_{\psi_1}.$$

Therefore,

$$\left\|\|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)T} - \boldsymbol{\Sigma}^*\|_F\right\|_{\psi_1} \le \left\|\|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)T} - \mathbf{V}_K\mathbf{V}_K^T\|_F\right\|_{\psi_1} + \|\boldsymbol{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\|_F \le 2\left\|\rho\left(\widehat{\mathbf{V}}_K^{(1)}, \mathbf{V}_K\right)\right\|_{\psi_1} \lesssim \kappa\sqrt{\frac{Kr}{n}}.$$

□

8.1.2 Proof of Theorem 1—*Proof.* When $\left\|\boldsymbol{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_2 < 1/4$, the Weyl's inequality

forces $\lambda_K(\boldsymbol{\Sigma}^*) > \frac{3}{4}$ and $\lambda_{K+1}(\boldsymbol{\Sigma}^*) < \frac{1}{4}$. The Theorem 2 in Yu et al. (2015) yields

$$\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*\right) = \sqrt{2}\sin\Theta\left(\widetilde{V}_K, \mathbf{V}_K^*\right) \lesssim \frac{\left\|\widetilde{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}^*\right\|_F}{\lambda_K(\boldsymbol{\Sigma}^*) - \lambda_{K+1}(\boldsymbol{\Sigma}^*)} \lesssim \left\|\widetilde{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}^*\right\|_F. \qquad (8.1)$$

When $n \ge r$, Lemma 4 and Lemma 1 imply that

$$\left\|\|\widetilde{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}^*\|_F\right\|_{\psi_1} = \left\|\|\frac{1}{m}\sum_{\ell=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \boldsymbol{\Sigma}^*\|_F\right\|_{\psi_1} \lesssim \frac{1}{\sqrt{m}}\left\|\|\widehat{\mathbf{V}}_K^{(1)}\widehat{\mathbf{V}}_K^{(1)T} - \boldsymbol{\Sigma}^*\|_F\right\|_{\psi_1} \lesssim \kappa\sqrt{\frac{Kr}{N}}.$$

Combining the two inequalities above finishes the proof. □

8.1.3 Proof of Theorem 2—*Proof.* Choose $j \in [d]$ and let

$\mathbf{D}_j = \mathbf{I} - 2\mathbf{e}_j\mathbf{e}_j^T$. Let $\boldsymbol{\Sigma} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^T$ be the spectral decomposition of $\boldsymbol{\Sigma}$. Assume that $\widehat{\lambda}$ is an

eigenvalue of the sample covariance $\widehat{\boldsymbol{\Sigma}} = (1/n)\sum_{i=1}^n \mathbf{X}_i\mathbf{X}_i^T$ and $\widehat{\mathbf{v}} \in \mathbb{S}^{d-1}$ is the correspondent

eigenvector that satisfies $\widehat{\boldsymbol{\Sigma}}\widehat{\mathbf{v}} = \widehat{\lambda}\widehat{\mathbf{v}}$.

Define $\mathbf{Z}_i = \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{V}^T \mathbf{X}_i$ and $\widehat{\mathbf{S}} = (1/n) \sum_{i=1}^{n} \mathbf{Z}_i \mathbf{Z}_i^T$. Note that $\widehat{\mathbf{\Sigma}} = \mathbf{V} \mathbf{\Lambda}^{\frac{1}{2}} \widehat{\mathbf{S}} \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{V}^T$. Consider the

matrix $\breve{\mathbf{\Sigma}} = \mathbf{V} \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{D}_j \widehat{\mathbf{S}} \mathbf{D}_j \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{V}^T$. By the sign symmetry, $\widehat{\mathbf{\Sigma}}$ and $\breve{\mathbf{\Sigma}}$ are identically distributed. It is not hard to verify that $\breve{\mathbf{\Sigma}}$ also has an eigenvalue $\hat{\lambda}$ with the correspondent eigenvector being $\mathbf{V} \mathbf{D}_j \mathbf{V}^T \widehat{\mathbf{v}}$. Denote the top K eigenvectors of $\widehat{\mathbf{\Sigma}}$ by $\widehat{\mathbf{V}}_K = (\widehat{\mathbf{v}}_1, \cdots, \widehat{\mathbf{v}}_K)$ and the top K eigenvectors of $\breve{\mathbf{\Sigma}}$ by $\breve{\mathbf{V}}_K$. Therefore we have

$$\mathbf{V}^T \mathbb{E}\left(\widehat{\mathbf{V}}_K \widehat{\mathbf{V}}_K^T\right) \mathbf{V} = \mathbf{V}^T \mathbb{E}\left(\breve{\mathbf{V}}_K \breve{\mathbf{V}}_K^T\right) \mathbf{V} = \mathbf{V}^T \mathbf{V} \mathbf{D}_j \mathbf{V}^T \mathbb{E}\left(\widehat{\mathbf{V}}_K \widehat{\mathbf{V}}_K^T\right) \mathbf{V} \mathbf{D}_j \mathbf{V}^T \mathbf{V} = \mathbf{D}_j \mathbf{V}^T \mathbb{E}\left(\widehat{\mathbf{V}}_K \widehat{\mathbf{V}}_K^T\right) \mathbf{V} \mathbf{D}_j .$$

Since the equation above holds for all $j \in [d]$, we can reach the conclusion that $\mathbf{V}^T \mathbb{E}\left(\widehat{\mathbf{V}}_K \widehat{\mathbf{V}}_K^T\right) \mathbf{V}$ is diagonal, i.e, $\mathbf{\Sigma}^* \mathbb{E}\left(\widehat{\mathbf{V}}_K \widehat{\mathbf{V}}_K^T\right)$) and $\mathbf{\Sigma}$ share the same set of eigenvectors.

Suppose that $\left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_2 < 1/2$. As demonstrated above, for any $k \in [K]$, the kth column of $\mathbf{V}_K$, which we denote by $\mathbf{v}_k$, should be an eigenvector of $\mathbf{\Sigma}^*$. Note that

$$\left\| \mathbf{\Sigma}^* \mathbf{v}_k \right\|_2 = \left\| \left(\mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T + \mathbf{V}_K \mathbf{V}_K^T \right) \mathbf{v}_k \right\|_2 \geq 1 - \left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_2 > 1 - \frac{1}{2} = \frac{1}{2} .$$

With regard to $\mathbf{\Sigma}^*$, the correspondent eigenvalue of $\mathbf{v}_k$ must be greater than 1/2. Denote any eigenvector of $\mathbf{\Sigma}$ that is not in $\{\mathbf{v}_k\}_{k=1}^{K}$ by $\mathbf{u}$, then analogously,

$$\left\| \mathbf{\Sigma}^* \mathbf{u} \right\|_2 = \left\| \left(\mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T + \mathbf{V}_K \mathbf{V}_K^T \right) \mathbf{u} \right\|_2 \leq \left\| \mathbf{\Sigma}^* - \mathbf{V}_K \mathbf{V}_K^T \right\|_2 < \frac{1}{2} .$$

For $\mathbf{\Sigma}^*$, the correspondent eigenvalue of $\mathbf{u}$ is smaller than 1/2. Therefore, the top K eigenspace of $\mathbf{\Sigma}^*$ is exactly $\mathrm{Col}(\mathbf{V}_K)$, and $\rho\left(\mathbf{V}_K^*, \mathbf{V}_K\right) = 0$. □

8.1.4 Proof of Lemma 2—*Proof.* Note that $\left\| f(\,\cdot\,) \right\|_F \leq \Delta^{-1} \left\| \,\cdot\, \right\|_F$ and

$$\left\| f(\mathbf{E}\mathbf{U}) \right\|_F \leq \Delta^{-1} \left\| \mathbf{E}\mathbf{U} \right\|_F \leq \Delta^{-1} \sqrt{K} \left\| \mathbf{E}\mathbf{U} \right\|_2 \leq \Delta^{-1} \sqrt{K} \left\| \mathbf{E} \right\|_2 = \sqrt{K} \varepsilon .$$

Hence Lemma 2 is a direct corollary of Lemma 8. □

8.1.5 Proof of Theorem 3—*Proof.* Define
$\mathbf{E} = \widehat{\mathbf{\Sigma}}^{(1)} - \mathbf{\Sigma}, \mathbf{P} = \mathbf{V}_K \mathbf{V}_K^T, \widehat{\mathbf{P}} = \widehat{\mathbf{V}}_K^{(1)} \widehat{\mathbf{V}}_K^{(1)T}, \mathbf{Q} = f\left(\mathbf{E}\mathbf{V}_K\right) \mathbf{V}_K^T + \mathbf{V}_K f\left(\mathbf{E}\mathbf{V}_K\right)^T, \mathbf{W} = \widehat{\mathbf{P}} - \mathbf{P} - \mathbf{Q}$ and $\varepsilon = \left\| \mathbf{E} \right\|_2 / \Delta$. From $\mathbb{E}\mathbf{Q} = 0$
and

$$\widehat{P} - P - Q = W = W1_{\{\varepsilon \le 1/10\}} + (W + Q)1_{\{\varepsilon > 1/10\}} - Q1_{\{\varepsilon > 1/10\}} = W1_{\{\varepsilon \le 1/10\}} + (\widehat{P} - P)1_{\{\varepsilon > 1/10\}} - Q1_{\{\varepsilon > 1/10\}},$$

we derive that

$$\mathbb{E}\widehat{P} - P = \mathbb{E}\big(W1_{\{\varepsilon \le 1/10\}}\big) + \mathbb{E}\big[(\widehat{P} - P)1_{\{\varepsilon > 1/10\}}\big] - \mathbb{E}\big(Q1_{\{\varepsilon > 1/10\}}\big), \left\|\mathbb{E}\widehat{P} - P\right\|_F \le \mathbb{E}\big(\left\|W\right\|_F 1_{\{\varepsilon \le 1/10\}}\big) + \mathbb{E}\big(\left\|\widehat{P} - P\right\|_F 1_{\{\varepsilon > 1/10\}}\big) + \mathbb{E}\big(\left\|Q\right\|_F 1_{\{\varepsilon > 1/10\}}\big). \tag{8.2}$$

We are going to bound the three terms separately. On the one hand, Lemma 2 implies that $\left\|W\right\|_F \le 24\sqrt{K}\varepsilon^2$ when $\varepsilon \le 1/10$. Hence

$$\mathbb{E}\big(\left\|W\right\|_F 1_{\{\varepsilon \le 1/10\}}\big) \le \mathbb{E}\big(24\sqrt{K}\varepsilon^2 1_{\{\varepsilon \le 1/10\}}\big) \lesssim \sqrt{K}\mathbb{E}\varepsilon^2. \tag{8.3}$$

On the other hand, the Davis-Kahan theorem shows that $\left\|\widehat{P} - P\right\|_F \lesssim \sqrt{K}\varepsilon$. Besides, it is easily seen that $\left\|Q\right\|_F \lesssim \left\|f(\mathbf{E}\mathbf{V}_K)\right\|_F \le \sqrt{K}\left\|\mathbf{E}\right\|_2/\Delta = \sqrt{K}\varepsilon$. Hence

$$\mathbb{E}\big(\left\|\widehat{P} - P\right\|_F 1_{\{\varepsilon > 1/10\}}\big) + \mathbb{E}\big(\left\|Q\right\|_F 1_{\{\varepsilon > 1/10\}}\big) \lesssim \sqrt{K}\mathbb{E}\big(\varepsilon 1_{\{\varepsilon > 1/10\}}\big)$$
$$\le 10\sqrt{K}\mathbb{E}\big(\varepsilon^2 1_{\{\varepsilon > 1/10\}}\big) \lesssim \sqrt{K}\mathbb{E}\varepsilon^2. \tag{8.4}$$

By (8.2),(8.3), (8.4) and Lemma 3 we have

$$\left\|\mathbb{E}\widehat{P} - P\right\|_F \lesssim \sqrt{K}\mathbb{E}\varepsilon^2 = \sqrt{K}\Delta^{-2}\mathbb{E}\left\|\mathbf{E}\right\|_2^2 \lesssim \sqrt{K}\Delta^{-2}\left\|\left\|\mathbf{E}\right\|_2\right\|_{\psi_1}^2 \lesssim \frac{\kappa^2\sqrt{K}r}{n}. \tag{8.5}$$

□

8.1.6 Proof of Theorem 4—*Proof.* According to Theorem 3, there exists a constant C such that $\left\|\Sigma^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_2 \le 1/4$ as long as $n \ge C\kappa^2\sqrt{K}r \ge r$. Then Theorem 1 implies that $\left\|\rho\big(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*\big)\right\|_{\psi_1} \le C_1\kappa\sqrt{\frac{Kr}{N}}$ for some constant C_1.

When random samples have symmetric innovation, we have $\rho\big(\mathbf{V}_K^*, \mathbf{V}_K\big) = 0$ and

$$\left\|\rho\big(\widetilde{V}_K, \mathbf{V}_K\big)\right\|_{\psi_1} = \left\|\rho\big(\widetilde{V}_K, \mathbf{V}_K^*\big)\right\|_{\psi_1} \le C_1\kappa\sqrt{\frac{Kr}{N}}.$$

For general distribution, Theorem 3 implies that $\rho\big(\mathbf{V}_K^*, \mathbf{V}_K\big) \le C_2\kappa^2\sqrt{K}r/n$ for some constant C_2 and

$$\left\|\rho\big(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\big)\right\|_{\psi_1} \le \left\|\rho\big(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*\big)\right\|_{\psi_1} + \rho\big(\mathbf{V}_K^*, \mathbf{V}_K\big) \le C_1 \kappa \sqrt{\frac{Kr}{N}} + C_2 \kappa^2 \frac{\sqrt{Kr}}{n}. \qquad (8.6)$$

When $m < C_3 n/(K^2 r)$ for some constant C_3, we have

$$\kappa \sqrt{\frac{Kr}{N}} = \sqrt{\frac{\kappa^2 Kr}{nm}} \ge \sqrt{\frac{\kappa^2 Kr}{n \cdot C_3 n/(\kappa^2 r)}} = \frac{1}{\sqrt{C_3}} \cdot \frac{\kappa^2 \sqrt{Kr}}{n},$$

and (8.6) forces

$$\left\|\rho\big(\widetilde{V}_K, \mathbf{V}_K\big)\right\|_{\psi_1} \le \big(C_1 + C_2\sqrt{C_3}\big)\kappa\sqrt{\frac{Kr}{N}}.$$

$\square$

8.1.7 Proof of Theorem 5

—*Proof.* We first focus on the first subsample $\big\{\mathbf{X}_i^{(1)}\big\}_{i=1}^n$ and the associated top eigenvector $\widehat{\mathbf{V}}_1^{(1)}$. For ease of notation, we temporarily drop the superscript. Let

$S = \Sigma_{i=1}^n W_i$ and $\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}} = \frac{d-1}{n} \Sigma_{i=1}^n \mathbf{1}_{\{W-1\}} \mathbf{Z}_i \mathbf{Z}_i^T$. From $\widehat{\Sigma} = \begin{pmatrix} \frac{2\lambda}{n}(n-S) & 0_{1\times(d-1)} \\ 0_{(d-1)\times 1} & 2\,\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}} \end{pmatrix}$ lead to

we know that $\left\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\right\|_2 > (\lambda/n)(n-S)$ and $\left\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\right\|_2 < (\lambda/n)(n-S)$

$\widehat{\mathbf{V}}_1 \perp \mathbf{V}_1$ and $\widehat{\mathbf{V}}_1 /\!/ \mathbf{V}_1 \big(\text{i.e. } \widehat{\mathbf{V}}_1 = \pm\mathbf{V}_1\big)$, respectively. Besides, $\left\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\right\|_2$ is a continuous random variable. Hence $\mathbb{P}\big(\widehat{\mathbf{V}}_1 \perp \mathbf{V}_1\big) + \mathbb{P}\big(\widehat{\mathbf{V}}_1 /\!/ \mathbf{V}_1\big) = 1$. Note that

$$Tr\big(\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\big) = \frac{d-1}{n}\sum_{i=1}^n \mathbf{1}_{\{W_i=1\}} = \frac{(d-1)S}{n}, \left\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\right\|_2 \ge \frac{Tr\big(\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\big)}{rank\big(\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\big)} \ge \frac{Tr\big(\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\big)}{\min\{n,d-1\}} \ge \frac{(d-1)S}{n^2}.$$

Then

$$\mathbb{P}\big(\widehat{\mathbf{V}}_1 /\!/ \mathbf{V}_1\big) \le \mathbb{P}\Big(\left\|\widehat{\boldsymbol{\Sigma}}_{\mathbf{Z}}\right\|_2 \le \frac{\lambda}{n}(n-S)\Big) \le \mathbb{P}\Big(\frac{(d-1)S}{n^2} \le \frac{\lambda}{n}(n-S)\Big) = \mathbb{P}\Bigg(\frac{S}{n} \le \frac{1}{1+\frac{d-1}{n\lambda}}\Bigg) \le \mathbb{P}\Big(\frac{S}{n} \le \frac{1}{4}\Big)$$

$$= \mathbb{P}\Big(\frac{S}{n} - \frac{1}{2} \le -\frac{1}{4}\Big) \le e^{-n/8}.$$

Above we used the assumption $d \ge 3n\lambda + 1$ and Hoeffding's inequality. Now we finish the analysis of $\widehat{\mathbf{V}}_1^{(1)}$ and collect back the superscript.

From now on we define $S = \Sigma_{l=1}^{m} 1_{\left\{\widehat{\mathbf{V}}_1^{(\ell)}//\mathbf{V}_1\right\}}$. For $\widehat{\mathbf{V}}_1^{(\ell)}$, let a_ℓ be its first entry and $\mathbf{b}_\ell$ be the vector of its last $(d-1)$ entries. The dichotomy $\mathbb{P}\left(\widehat{\mathbf{V}}_1^{(\ell)}//\mathbf{V}_1\right) + \mathbb{P}\left(\widehat{\mathbf{V}}_1^{(\ell)} \perp \mathbf{V}_1\right) = 1$ mentioned above forces

$$|a_\ell| = 1_{\left\{\widehat{\mathbf{V}}_1^{(\ell)}/\mathbf{V}_1\right\}}, \|\mathbf{b}_\ell\|_2 = 1_{\left\{\widehat{\mathbf{V}}_1^{(\ell)} \perp \mathbf{V}_1\right\}}, \widehat{\mathbf{V}}_1^{(\ell)}\widehat{\mathbf{V}}_1^{(\ell)T} = \begin{pmatrix} 1_{\left\{\widehat{\mathbf{V}}_1^{(\ell)}/\mathbf{V}_1\right\}} & \mathbf{0}_{1\times(d-1)} \\ \mathbf{0}_{(d-1)\times 1} & \mathbf{b}_\ell\mathbf{b}_\ell^T \end{pmatrix}, \text{ and}$$

$$\widetilde{\mathbf{\Sigma}} = \frac{1}{m}\sum_{\ell=1}^{m}\widehat{\mathbf{V}}_1^{(\ell)}\widehat{\mathbf{V}}_1^{(\ell)T} = \begin{pmatrix} \frac{1}{m}S & \mathbf{0}_{1\times(d-1)} \\ \mathbf{0}_{(d-1)\times 1} & \frac{1}{m}\sum_{\ell=1}^{m}\mathbf{b}_\ell\mathbf{b}_\ell^T \end{pmatrix}.$$

Note that $n \geq 32\log d$ forces $n \geq 32\log d$ forces $\mathbb{P}\left(\widehat{\mathbf{V}}_1^{(\ell)}//\mathbf{V}_1\right) \leq e^{-n/8} \leq d^{-4}$.

Case 1: $m \leq d^3$

In this case, $\mathbb{P}(S=0) = \left[1 - \mathbb{P}\left(\widehat{\mathbf{V}}_1^{(1)}//\mathbf{V}_1\right)\right]^m \geq 1 - m\mathbb{P}\left(\widehat{\mathbf{V}}_1^{(1)}//\mathbf{V}_1\right) \geq 1 - d^{-1}$. When $S = 0$, we have $\|\mathbf{b}_\ell\|_2 = 1$ for all $\ell \in [m]$ and $\left\|(1/m)\Sigma_{l=1}^{m}\mathbf{b}_\ell\mathbf{b}_\ell^T\right\|_2 > 0$, leading to $\widetilde{\mathbf{V}}_1 \perp \mathbf{V}_1$.

Case 2: $m > d^3$

On the one hand, by Hoeffding's inequality we obtain

$$\mathbb{P}\left(\frac{S}{m} \geq \frac{1}{d}\right) \leq \mathbb{P}\left(\frac{1}{m}(S - \mathbb{E}S) \geq \frac{1}{2d}\right) \leq e^{-2m\left(\frac{1}{2d}\right)^2} = e^{-\frac{m}{2d^2}} < e^{-d/2}.$$

On the other hand, note that

$$\left\|\frac{1}{m}\sum_{k=1}^{m}\mathbf{b}_\ell\mathbf{b}_{\ell\ell}^T\right\|_2 \geq \frac{Tr\left(\frac{1}{m}\Sigma_{k=1}^{m}\mathbf{b}_\ell\mathbf{b}_\ell^T\right)}{d-1} = \frac{\frac{1}{m}\Sigma_{k=1}^{m}\|\mathbf{b}_\ell\|_2^2}{d-1} = \frac{1}{d-1}\left(1 - \frac{S}{m}\right).$$

Hence

$$\mathbb{P}\left(\widetilde{V}_1 \perp \mathbf{V}_1\right) \geq \mathbb{P}\left(\left\|\frac{1}{m}\sum_{k=1}^{m}\mathbf{b}_\ell\mathbf{b}_\ell^T\right\|_2 > \frac{S}{m}\right) \geq \mathbb{P}\left[\frac{1}{d-1}\left(1 - \frac{S}{m}\right) > \frac{S}{m}\right] = \mathbb{P}\left(\frac{S}{m} < \frac{1}{d}\right) \geq 1 - e^{-d/2}.$$

$\square$

8.1.8 Proof of Theorem 6—*Proof.* With slight abuse of notations, here we define $\mathbf{\Sigma}_\ell^* = \mathbb{E}\left(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}\right), \mathbf{\Sigma}^* = \frac{1}{m}\Sigma_{\ell=1}^{m}\mathbf{\Sigma}_\ell^*$, and $\mathbf{V}_K^* \in \mathbb{R}^{d\times K}$ to be the top K eigenvectors of Σ^*.

First we consider the general case. Note that $\lambda_K\left(\mathbf{V}_K\mathbf{V}_K^T\right) = 1$ and $\lambda_K\left(\mathbf{V}_K\mathbf{V}_K^T\right) = 0$. By the Davis-Kahan theorem, we have

$$\rho\left(\widetilde{V}_K, \mathbf{V}_K\right) \lesssim \left\|\ \widetilde{\boldsymbol{\Sigma}}\ - \mathbf{V}_K\mathbf{V}_K^T\right\|_F \leq \left\|\ \widetilde{\boldsymbol{\Sigma}}\ - \boldsymbol{\Sigma}^*\right\|_F + \left\|\ \boldsymbol{\Sigma}^*\ - \mathbf{V}_K\mathbf{V}_K^T\right\|_F. \tag{8.7}$$

Note that $\boldsymbol{\Sigma}^* = \frac{1}{m}\Sigma_{\ell=1}^m \boldsymbol{\Sigma}_\ell^*$. The first term in (8.7) is the norm of independent sums

$$\left\|\ \widetilde{\boldsymbol{\Sigma}}\ - \boldsymbol{\Sigma}^*\right\|_F = \left\|\frac{1}{m}\sum_{\ell=1}^m \left(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \boldsymbol{\Sigma}_\ell^*\right)\right\|_F$$

It follows from Lemma 1 that $\left\|\left\|\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T} - \boldsymbol{\Sigma}_\ell^*\ \right\|_F\right\|_{\psi_1} \lesssim \kappa_\ell\sqrt{\dfrac{Kr_\ell}{n}} = \sqrt{m}S_\ell$, from which Lemma 4 leads to

$$\left\|\left\|\ \widetilde{\boldsymbol{\Sigma}}\ - \boldsymbol{\Sigma}^*\ \right\|_F\right\|_{\psi_1} \lesssim \frac{1}{m}\sqrt{\sum_{\ell=1}^m \left(\sqrt{m}S_\ell\right)^2} = \sqrt{\frac{1}{m}\sum_{\ell=1}^m S_\ell^2}. \tag{8.8}$$

The second term in (8.7) is bounded by

$$\left\|\ \boldsymbol{\Sigma}^*\ - \mathbf{V}_K\mathbf{V}_K^T\right\|_F = \left\|\frac{1}{m}\sum_{\ell=1}^m \left(\boldsymbol{\Sigma}_\ell^* - \mathbf{V}_K\mathbf{V}_K^T\right)\right\|_F \leq \frac{1}{m}\sum_{\ell=1}^m \left\|\ \boldsymbol{\Sigma}_\ell^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_F.$$

Theorem 3 implies that when $n \geq r_\ell$,

$$\left\|\boldsymbol{\Sigma}_\ell^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_F \lesssim \kappa_\ell^2\sqrt{K}r_\ell/n = B_\ell. \tag{8.9}$$

Hence

$$\left\|\ \boldsymbol{\Sigma}^*\ - \mathbf{V}_K\mathbf{V}_K^T\right\|_F \lesssim \frac{1}{m}\sum_{\ell=1}^m B_\ell. \tag{8.10}$$

The claim under general case follows from (8.7), (8.8) and (8.10).

Now we come to the symmetric case. If $\left\|\boldsymbol{\Sigma}_\ell^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_2 < 1/2$ for all $\ell \in [m]$, then Theorem 2 implies that the top K eigenspace of $\boldsymbol{\Sigma}_\ell^*$ is $Col\left(\mathbf{V}_K\right)$. Therefore, the top K eigenspace of $\boldsymbol{\Sigma}^*$ is still $Col(\mathbf{V}_K)$ and $\rho\left(\mathbf{V}_K, \mathbf{V}_K^*\right) = 0$.

When $n \geq C\sqrt{K}\max_{\ell \in [m]}\left(\kappa_\ell^2 r_\ell\right)$ for large C, (8.9) ensures $\max_{\ell \in [m]}\left\|\mathbf{\Sigma}_\ell^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_2 \leq 1/4$, $\left\|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_2 \leq 1/4$ and $\rho(\mathbf{V}_K, \mathbf{V}_K^*) = 0$. Weyl's inequality forces $\lambda_K(\mathbf{\Sigma}^*) \geq 3/4$ and $\lambda_{K+1}(\mathbf{\Sigma}^*) \leq 1/4$. By the Davis-Kahan theorem and (8.8),

$$\left\|\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K\right)\right\|_{\psi_1} = \left\|\rho\left(\widetilde{\mathbf{V}}_K, \mathbf{V}_K^*\right)\right\|_{\psi_1} \left\| \left\| \widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\right\|_F \right\|_{\psi_1} \lesssim \sqrt{\frac{1}{m}\sum_{\ell=1}^m s_\ell^2}.$$

$\square$

8.1.9 Proof of Theorem 7—*Proof.* We define

$\mathbf{\Sigma}_\ell^* = \mathbb{E}\left(\widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}\right)$ and $\mathbf{\Sigma}^* = \frac{1}{m}\Sigma_{\ell=1}^m \mathbf{\Sigma}_\ell^*$. Let $\mathbf{V}_K^*, \overline{\mathbf{V}}_K^{(\ell)} \in \mathcal{O}_{d \times K}$ be the top K eigenvectors of Σ^* and $\mathbf{\Sigma}^*$ and $\mathbf{\Sigma}^{(\ell)}$, respectively. By the Davis-Kahan theorem,

$$\rho\left(\widetilde{V}_K, \mathbf{V}_K\right) \lesssim \left\| \widetilde{\mathbf{\Sigma}} - \mathbf{V}_K\mathbf{V}_K^T\right\|_F \leq \left\| \widetilde{\mathbf{\Sigma}} - \mathbf{\Sigma}^*\right\|_F + \left\| \mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_F. \tag{8.11}$$

The first term in (8.11) is controlled in exactly the same way as (8.8). The second term is further decomposed as

$$\left\|\mathbf{\Sigma}^* - \mathbf{V}_K\mathbf{V}_K^T\right\|_F = \left\|\frac{1}{m}\sum_{\ell=1}^m \left(\mathbf{\Sigma}_\ell^* - \mathbf{V}_K\mathbf{V}_K^T\right)\right\|_F \leq \left\|\frac{1}{m}\sum_{\ell=1}^m \left(\mathbf{\Sigma}_\ell^* - \overline{\mathbf{V}}_K^{(\ell)}\overline{\mathbf{V}}_K^{(\ell)T}\right)\right\|_F$$
$$+ \left\|\frac{1}{m}\sum_{\ell=1}^m \left(\overline{\mathbf{V}}_K^{(\ell)}\overline{\mathbf{V}}_K^{(\ell)T} - \mathbf{V}_K\mathbf{V}_K^T\right)\right\|_F. \tag{8.12}$$

Similar to (8.9) and (8.10), with $n \geq r_\ell$ we have $\left\| \mathbf{\Sigma}_\ell^* - \overline{\mathbf{V}}_K^{(\ell)}\overline{\mathbf{V}}_K^{(\ell)T}\right\|_F \lesssim B_\ell$ and

$$\left\|\frac{1}{m}\sum_{\ell=1}^m \left(\mathbf{\Sigma}_\ell^* - \overline{\mathbf{V}}_K^{(\ell)}\overline{\mathbf{V}}_K^{(\ell)T}\right)\right\|_F \leq \frac{1}{m}\sum_{\ell=1}^m \left\|\mathbf{\Sigma}_\ell^* - \overline{\mathbf{V}}_K^{(\ell)}\overline{\mathbf{V}}_K^{(\ell)T}\right\|_F \lesssim \frac{1}{m}\sum_{\ell=1}^m B_\ell. \tag{8.13}$$

For the last part in (8.12), note that $\overline{\mathbf{V}}_K^{(\ell)}$ and $\mathbf{V}_K$ and V_K contain eigenvectors of $\mathbf{\Sigma}^{(\ell)}$ and $\left|\mathbf{B}^{(\ell)}\mathbf{B}^{(\ell)T}\right|$. Hence the Davis-Kahan theorem forces

$$\left\|\overline{\mathbf{V}}_K^{(\ell)}\overline{\mathbf{V}}_K^{(\ell)T} - \mathbf{V}_K\mathbf{V}_K^T\right\|_F \lesssim \frac{\sqrt{K}\left\|\mathbf{\Sigma}_u^{(\ell)}\right\|_2}{\lambda_K\left(\mathbf{\Lambda}_K^{(\ell)}\right)}.$$

and

$$\left\| \frac{1}{m} \sum_{\ell=1}^{m} \left(\overline{\mathbf{V}}_K^{(\ell)} \overline{\mathbf{V}}_K^{(\ell)T} - \mathbf{V}_K \mathbf{V}_K^T \right) \right\|_F \lesssim \frac{\sqrt{K}}{m} \sum_{\ell=1}^{m} \frac{\left\| \mathbf{\Sigma}_u^{(\ell)} \right\|_2}{\lambda_K \left(\mathbf{\Lambda}_K^{(\ell)} \right)}. \tag{8.14}$$

The proof is completed by collecting (8.11), (8.12), (8.13) and (8.14). □

8.2 Technical lemmas

8.2.1 Tail bounds—Lemma 3. *Suppose* $\mathbf{X}$ *and* $\{\mathbf{X}_i\}_{i=1}^{n}$ *are i.i.d. sub-Gaussian random vectors in* $\mathbf{R}^d$ *with zero mean and covariance matrix* $\mathbf{\Sigma} \succcurlyeq 0$. *Let* $\widehat{\mathbf{\Sigma}} = \frac{1}{n} \Sigma_{i=1}^{n} \mathbf{X}_i \mathbf{X}_i^T$ *be the sample covariance matrix,* $\{\lambda_j\}_{j=1}^{d}$ *be the eigenvalues of* Σ *sorted in descending order, and* $r = Tr(\mathbf{\Sigma})/\|\mathbf{\Sigma}\|_2$. *There exist constants* $c \geq 1$ *and* $C \geq 0$ *such that when* $n > r$, *we have*

$$\mathbb{P}\left(\left\| \widehat{\mathbf{\Sigma}} - \mathbf{\Sigma} \right\|_2 \geq s \right) \leq \exp\left(-\frac{s}{c\lambda_1 \sqrt{r/n}} \right), \quad \forall s \geq 0,$$

and $\left\| \left\| \widehat{\mathbf{\Sigma}} - \mathbf{\Sigma} \right\|_2 \right\|_{\psi_1} \leq C\lambda_1 \sqrt{r/n}$.

Proof. By the Theorem 9 in Koltchinskii and Lounici (2017) and the simple fact

$$\left(\mathbb{E}\|\mathbf{X}\|_2 \right)^2 / \left\| \mathbf{\Sigma} \right\|_2 \leq \mathbb{E}\|\mathbf{X}\|_2^2 / \left\| \mathbf{\Sigma} \right\|_2 = Tr(\mathbf{\Sigma}) / \left\| \mathbf{\Sigma} \right\|_2 = r(\mathbf{\Sigma}),$$

we know the existence of a constant $c > 1$ such that

$$\mathbb{P}\left(\left\| \widehat{\mathbf{\Sigma}} - \mathbf{\Sigma} \right\|_2 \geq c\lambda_1 \max\left\{ \sqrt{\frac{r}{n}}, \frac{r}{n}, \sqrt{\frac{t}{n}}, \frac{t}{n} \right\} \right) \leq e^{-t}, \quad \forall t \geq 1. \tag{8.15}$$

Since $1 \leq r \leq n$, (8.15) yields

$$\mathbb{P}\left(\left\| \widehat{\mathbf{\Sigma}} - \mathbf{\Sigma} \right\|_2 \geq c\lambda_1 \sqrt{\frac{t}{n}} \right) \leq e^{-t}, \quad r \leq t \leq n, \tag{8.16}$$

$$\mathbb{P}\left(\left\| \widehat{\mathbf{\Sigma}} - \mathbf{\Sigma} \right\|_2 \geq c\lambda_1 \frac{t}{n} \right) \leq e^{-t}, \quad t \geq n. \tag{8.17}$$

When $r \leq t \leq n$, we have $\sqrt{\frac{t}{n}} \leq \frac{t}{n}\sqrt{\frac{n}{r}}$. By letting $s = c\lambda_1 \frac{t}{n}\sqrt{\frac{n}{r}}$ we derive from (8.16) that for $c\lambda_1 \sqrt{\frac{r}{n}} \leq s \leq c\lambda_1 \sqrt{\frac{n}{r}}$,

$$\mathbb{P}\Big(\big\| \, \widehat{\boldsymbol{\Sigma}} \, - \, \boldsymbol{\Sigma} \, \big\|_2 \ge s \Big) \le \mathbb{P}\Bigg(\bigg\| \, \widehat{\boldsymbol{\Sigma}} \, - \, \boldsymbol{\Sigma} \, \bigg\|_2 \ge c\lambda_1 \sqrt{\frac{t}{n}} \Bigg) \le e^{-t} = \exp\Big(-\frac{s\sqrt{nr}}{c\lambda_1} \Big). \qquad (8.18)$$

When $t \ge n$, we let $s = c\lambda_1 \frac{t}{n}$ and derive from (8.17) that for $s \ge c\lambda_1$,

$$\mathbb{P}\Big(\big\| \, \widehat{\boldsymbol{\Sigma}} \, - \, \boldsymbol{\Sigma} \, \big\|_2 \ge s \Big) = \mathbb{P}\Bigg(\bigg\| \, \widehat{\boldsymbol{\Sigma}} \, - \, \boldsymbol{\Sigma} \, \bigg\|_2 \ge c\lambda_1 \frac{t}{n} \Bigg) \le e^{-t} = \exp\Big(-\frac{ns}{c\lambda_1} \Big). \qquad (8.19)$$

(8.18), (8.19) and $n \ge r$ lead to

$$\mathbb{P}\Big(\big\| \, \widehat{\boldsymbol{\Sigma}} \, - \, \boldsymbol{\Sigma} \, \big\|_2 \ge s \Big) \le \exp\Big(-\frac{s\sqrt{nr}}{c\lambda_1} \Big), \qquad \forall s \ge c\lambda_1 \sqrt{r/n}.$$

and thus

$$\mathbb{P}\Big(\big\| \, \widehat{\boldsymbol{\Sigma}} \, - \, \boldsymbol{\Sigma} \, \big\|_2 \ge s \Big) \le \exp\Big(1 - \frac{s}{c\lambda_1 \sqrt{r/n}} \Big), \qquad \forall s \ge 0.$$

According to the Definition 5.13 in Vershynin (2012), we get $\Big\| \big\| \, \widehat{\boldsymbol{\Sigma}} \, - \, \boldsymbol{\Sigma} \, \big\|_2 \Big\|_{\psi_1} \le C\lambda_1 \sqrt{r/n}$ for some constant C. $\square$

The next lemma investigates the sum of independent random vectors in a Hilbert space whose norms are sub-exponential, which directly follows from Theorem 2.5 in Bosq (2000).

Lemma 4. *If* $\{\mathbf{X}_i\}_{i=1}^{n}$ *are independent random vectors in a separable Hilbert space (where the norm is denoted by* $\| \cdot \|$*) with* $\mathbb{E}\mathbf{X}_i = 0$ *and* $\big\| \|\mathbf{X}_i\| \big\|_{\psi_1} \le L_i < \infty$*. We have*

$$\Big\| \, \| \sum_{i=1}^{n} \mathbf{X}_i \| \, \Big\|_{\psi_1} \lesssim \sqrt{\sum_{i=1}^{n} L_i^2}.$$

Proof. We are going to apply Theorem 2.5 in Bosq (2000). By definition $k^{-1}\mathbb{E}^{1/k}\|\mathbf{X}_i\|^k \le \big\| \|\mathbf{X}_i\| \big\|_{\psi_1} \le L_i$ for all $k \ge 1$, and

$$\mathbb{E}\|\mathbf{X}_i\|^k \le (kL_i)^k \le \sqrt{2\pi k}(k/e)^k (eL_i)^k \lesssim k!(eL_i)^k.$$

Hence there exists some constant c such that $\mathbb{E}\|\mathbf{X}_i\|^k \le \frac{k!}{2}(cL_i)^k$ for $k \ge 2$. Let $\ell = \sqrt{c^2 \sum_{i=1}^{n} L_i^2}$ and $b = c \cdot \max_{i \in [n]} L_i$. We have

$$\sum_{i=1}^{n} \mathbb{E} \left\| \mathbf{X} \right\|^{k} \leq \frac{k!}{2} \sum_{i=1}^{n} \left(cL_i\right)^k \leq \frac{k!}{2} \left(\sum_{i=1}^{n} c^2 L_i^2 \right) \left(c \cdot \max_{i \in [n]} L_i\right)^{k-2} = \frac{k!}{2} \ell^2 b^{k-2}, \quad \forall k \geq 2.$$

Let $\mathbf{S}_n = \Sigma_{i=1}^{n} \mathbf{X}_i$. Theorem 2.5 in Bosq (2000) implies that

$$\mathbb{P}\left(\left\| \mathbf{S}_n \right\| \geq t \right) \leq 2\exp\left(-\frac{t^2}{2\ell^2 + 2bt}\right), \quad \forall t > 0.$$

When $4\ell \leq t \leq \ell^2/b$ (this cannot happen if $4b > \ell$), we have $2\ell^2 \geq 2bt$ and

$$\mathbb{P}\left(\left\| \mathbf{S}_n \right\| \geq t \right) \leq 2\exp\left(-\frac{t^2}{\sqrt{2\ell^2 + 2\ell^2}}\right) \leq 2\exp\left(-\frac{4\ell \cdot t}{4\ell^2}\right) = 2\exp\left(-\frac{t}{\ell}\right) \leq \exp\left(1 - \frac{t}{4\ell}\right).$$

When $t \geq \ell^2/b$, we have $2bt \geq 2\ell^2$ and

$$\mathbb{P}\left(\left\| \mathbf{S}_n \right\| \geq t \right) \leq 2\exp\left(-\frac{t^2}{2bt + 2bt}\right) = 2\exp\left(-\frac{t}{4b}\right) \leq \exp\left(1 - \frac{t}{4\ell}\right),$$

where the last inequality follows from $2 \leq e$ and $b \leq \ell$. It is then easily seen that

$$\mathbb{P}\left(\left\| \mathbf{S}_n \right\| \geq t \right) \leq \exp\left(1 - \frac{t}{4\ell}\right), \quad \forall t \geq 0.$$

With the help of Definition 5.13 in Vershynin (2012), we can conclude that

$$\left\| \left\| \mathbf{S}_n \right\| \right\|_{\psi_1} \lesssim \ell \lesssim \sqrt{\sum_{i=1}^{n} L_i^2}.$$

□

8.2.2 Matrix analysis—Lemma 5. *Suppose that* $\mathbf{A} \in \mathbb{R}^{d \times d}$ *is a symmetric matrix with eigenvalues* $\left\{\lambda_j\right\}_{j=1}^{d}$ *(in descending order) and corresponding eigenvectors* $\left\{\mathbf{u}_j\right\}_{j=1}^{d}$. *When* $K \in [d]$, $\mathbf{P}_K = \sum_{j=1}^{K} \mathbf{u}_j \mathbf{u}_j^T$ *an optimal solution to the SDP:*

$$\begin{aligned} \min_{\mathbf{P} \in S^{d \times d}} \quad & -\operatorname{Tr}\left(\mathbf{P}^T \mathbf{A}\right) \\ \text{s.t.} \quad & \operatorname{Tr}(\mathbf{P}) \leq K, \left\|\mathbf{P}\right\|_2 \leq 1, \mathbf{P} \succeq 0 \end{aligned} \tag{8.20}$$

Proof. By orthonormal invariance of the problem formulation, we assume without loss of generality that $\left\{\mathbf{u}_j\right\}_{j=1}^d$ are the canonical bases $\left\{\mathbf{e}_j\right\}_{j=1}^d$. Then $\mathbf{A} = \operatorname{diag}(\lambda_1, \cdots \lambda_d)$ and $\operatorname{Tr}\left(\mathbf{P}^T\mathbf{A}\right) = \sum_{j=1}^d \lambda_j P_{jj}$. The constraints on $\mathbf{P}$ force $0 \le P_{jj} \le$ and $\sum_{j=1}^d P_{jj} \le K$. Hence $-\operatorname{Tr}(\mathbf{P}^T\mathbf{A}) \ge -\sum_{j=1}^K \lambda_j$ always holds, and $\mathbf{P}_{\mathrm{K}} = \sum_{j=1}^K \mathbf{e}_j \mathbf{e}_j^T$ is a feasible solution that attains this minimum. $\square$

Lemma 6. *Suppose*

$\widehat{\mathbf{V}}_K^{(\ell)} \in \mathcal{O}_{d \times K}, \forall \ell \in [m]$, and define $\widetilde{\mathbf{\Sigma}} = \frac{1}{m}\sum_{\ell=1}^m \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}$. Let $\widetilde{\mathbf{\Sigma}} = \sum_{j=1}^d \widetilde{\lambda}_j \widetilde{\mathbf{v}}_j \widetilde{\mathbf{v}}_j^T$ *be its eigen-decomposition, where* $\widetilde{\lambda}_1 \ge \cdots \ge \widetilde{\lambda}_d$. *Then*

$$\widetilde{\mathbf{V}}_K = \left(\widetilde{\mathbf{v}}_1, \cdots, \widetilde{\mathbf{v}}_K\right) \in \operatorname{argmin}_{\mathbf{U} \in \mathcal{O}_{d \times K}} \sum_{\ell=1}^m \rho^2\left(\mathbf{U}, \widehat{\mathbf{V}}_K^{(\ell)}\right).$$

Proof. Let $\widehat{\mathbf{P}}^{(\ell)} = \widehat{\mathbf{V}}_K^{(\ell)}\widehat{\mathbf{V}}_K^{(\ell)T}$ and $R(\mathbf{U}) = \frac{1}{m}\sum_{\ell=1}^m \rho^2\left(\mathbf{U}, \widehat{\mathbf{V}}_K^{(\ell)}\right)$. Then $\widetilde{\mathbf{\Sigma}} = \frac{1}{m}\sum_{\ell=1}^m \widehat{\mathbf{P}}^{(\ell)}$ and

$$R(\mathbf{U}) = \frac{1}{m}\sum_{\ell=1}^m \rho^2\left(\mathbf{U}, \widehat{\mathbf{V}}_K^{(\ell)}\right) = \frac{1}{m}\sum_{\ell=1}^m \left\|\mathbf{U}\mathbf{U}^T - \widehat{\mathbf{P}}^{(\ell)}\right\|_F^2 = \left\|\mathbf{U}\mathbf{U}^T - \widetilde{\mathbf{\Sigma}}\right\|_F^2 + \frac{1}{m}\sum_{\ell=1}^m \left\|\widetilde{\mathbf{\Sigma}} - \widehat{\mathbf{P}}^{(\ell)}\right\|_F^2 = \left\|\mathbf{U}\mathbf{U}^T\right\|_F^2$$

$$+ \left\|\widetilde{\mathbf{\Sigma}}\right\|_F^2 - 2\operatorname{Tr}\left(\mathbf{U}\mathbf{U}^T\widetilde{\mathbf{\Sigma}}\right) + \frac{1}{m}\sum_{\ell=1}^m \left\|\widetilde{\mathbf{\Sigma}} - \widehat{\mathbf{P}}^{(\ell)}\right\|_F^2.$$

The fact $\mathbf{U} \in \mathcal{O}_{d \times K}$ forces $\left\|\mathbf{U}\mathbf{U}^T\right\|_F^2 = K$. Hence

$$\operatorname{argmin}_{\mathbf{U} \in \mathcal{O}_{d \times K}} R(\mathbf{U}) = \operatorname{argmax}_{\mathbf{U} \in \mathcal{O}_{d \times K}} \operatorname{Tr}\left(\mathbf{U}\mathbf{U}^T \widetilde{\mathbf{\Sigma}}\right).$$

By slightly modifying the proof for Lemma 5 we get the desired result. $\square$

Suppose that $\mathbf{U}, \mathbf{V} \in \mathcal{O}_{d \times K}$. Let $\mathbf{P}_{\mathbf{U}} = \mathbf{U}\mathbf{U}^T, \mathbf{P}_{\mathbf{V}} = \mathbf{V}\mathbf{V}^T$, $\mathbf{H} = \mathbf{V}^T\mathbf{U}$, and $\left\{\sigma_j\right\}_{j=1}^K$ be the singular values (sorted in descending order) of $\mathbf{H}$. By the Corollary 5.4 in Chapter I, Stewart and Sun (1990), $\left\{\sigma_j\right\}_{j=1}^K$ are cosines of the canonical angles $\left\{\theta_j\right\}_{j=1}^K \subseteq [0, \pi/2]$ between $\operatorname{Col}(\mathbf{U})$ and $\operatorname{Col}(\mathbf{V})$. Let $\sin \Theta\,(\mathbf{U}, \mathbf{V}) = \operatorname{diag}(\sin \theta_1, \sin \theta_K)$.

Define $\widehat{\mathbf{H}} = \operatorname{sgn}(\mathbf{H})$. Here $\operatorname{sgn}(\cdot)$ is the matrix sign function (see Gross (2011)) defined as follows: let $\mathbf{H} = \sum_{j=1}^K \sigma_j \widetilde{\mathbf{u}}_j \widetilde{\mathbf{v}}_j^T$ be the singular value decomposition, where $\left\{\widetilde{\mathbf{u}}_j\right\}_{j=1}^K, \left\{\widetilde{\mathbf{v}}_j\right\}_{j=1}^K$ are two orthonormal bases in $\mathbb{R}^K$ and $\left\{\sigma_j\right\}_{j=1}^K \subseteq [0, +\infty)$, then $\widehat{\mathbf{H}} = \sum_{j=1}^K \operatorname{sgn}\left(\sigma_j\right)\widetilde{\mathbf{u}}_j \widetilde{\mathbf{v}}_j^T = \sum_{\sigma_j > 0} \widetilde{\mathbf{u}}_j \widetilde{\mathbf{v}}_j^T$.

Lemma 7. *We have*

$$\left\| \mathbf{P_U} - \mathbf{P_V} \right\|_2 = \left\| \sin \Theta \, (\mathbf{U}, \mathbf{V}) \right\|_2 \text{ and } \left\| \mathbf{P_U} - \mathbf{P_V} \right\|_F = \sqrt{2} \left\| \sin \Theta \, (\mathbf{U}, \mathbf{V}) \right\|_F. \text{ If } \left\| \mathbf{P_U} - \mathbf{P_V} \right\|_2 < 1, \text{ then } \widehat{\mathbf{H}}$$

is orthonormal,

$$\left\| \mathbf{H} - \widehat{\mathbf{H}} \right\|_2 \le \frac{\left\| \mathbf{P_U} - \mathbf{P_V} \right\|_2^2}{2 - \left\| \mathbf{P_U} - \mathbf{P_V} \right\|_2^2},$$

$$\| \mathbf{V}H - \mathbf{U} \|_F \le \| \mathbf{V}\widehat{\mathbf{H}} - \mathbf{U} \|_F = \sqrt{2} \| \mathbf{H} - \widehat{\mathbf{H}} \|_*^{1/2}, \frac{1}{\sqrt{2}} \| \mathbf{P_U} - \mathbf{P_V} \|_F \le \| \mathbf{V}\widehat{\mathbf{H}} - \mathbf{U} \|_F \le \frac{\| \mathbf{P_U} - \mathbf{P_V} \|_F}{\sqrt{2 - \| \sin \Theta \, (\mathbf{U}, \mathbf{V}) \|_2^2}}.$$

Proof. By the Theorem 5.5 in Chapter I, Stewart and Sun (1990), the singular values of $\mathbf{P_U} - \mathbf{P_V}$ are $\sin \theta_K$, $\sin \theta_K$, $\sin \theta_{K-1}$, $\sin \theta_{K-1}$,..., $\sin \theta_1$, $\sin \theta_1$, 0,..., 0. This immediately leads to $\left\| \mathbf{P_U} - \mathbf{P_V} \right\|_2 = \left\| \sin \Theta \, (\mathbf{U}, \mathbf{V}) \right\|_2$ and $\left\| \mathbf{P_U} - \mathbf{P_V} \right\|_F = \sqrt{2} \left\| \sin \Theta \, (\mathbf{U}, \mathbf{V}) \right\|_F$..

When $\left\| \sin \Theta \, (\mathbf{U}, \mathbf{V}) \right\|_2 = \left\| \mathbf{P_U} - \mathbf{P_V} \right\|_2 < 1$, we have $\theta_K < \pi/2$. Thus the smallest singular value of $\mathbf{H}$ is $\sigma_K = 1 - \cos \theta_K > 0$, and $\widehat{\mathbf{H}}$ is orthonormal. Observe that

$$\| \mathbf{V}\widehat{\mathbf{H}} - \mathbf{U} \|_F^2 = \| \mathbf{V}\widehat{\mathbf{H}} \|_F^2 + \| \mathbf{U} \|_F^2 - 2\mathrm{Tr}\left(\widehat{\mathbf{H}}^T \mathbf{V}^T \mathbf{U} \right) = 2K - 2\mathrm{Tr}\left(\widehat{\mathbf{H}}^T \mathbf{H} \right)$$

$$= 2 \sum_{j=1}^{K} \left(1 - \sigma_j \right) = 2\| \mathbf{H} - \widehat{\mathbf{H}} \|_*. \tag{8.21}$$

Hence $\left\| \mathbf{V}H - \mathbf{U} \right\|_F \le \left\| \mathbf{V}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F$ follows from

$$\| \mathbf{V}H - \mathbf{U} \|_F^2 = \| \mathbf{V}H \|_F^2 + \| \mathbf{U} \|_F^2 - 2\mathrm{Tr}\left(\mathbf{H}^T \mathbf{V}^T \mathbf{U} \right) = K - \| \mathbf{H} \|_F^2 = \sum_{j=1}^{K} \left(1 - \sigma_j^2 \right) = \sum_{j=1}^{K} \left(1 - \sigma_j \right)\left(1 + \sigma_j \right)$$

$$\le 2 \sum_{j=1}^{K} \left(1 - \sigma_j \right) = \| \mathbf{V}\widehat{\mathbf{H}} - \mathbf{U} \|_F^2.$$

For any $\theta \in [0, \pi/2)$, we have

$1 - \cos\theta = \frac{1 - \cos^2\theta}{1 + \cos\theta} = \frac{\sin^2\theta}{2 - (1 - \cos\theta)}$, which leads to $\frac{1}{2}\sin^2\theta \le 1 - \cos\theta \le \sin^2\theta$ and furthermore, $1 - \cos\theta = \frac{\sin^2\theta}{2 - (1 - \cos\theta)} \le \frac{\sin^2\theta}{2 - \sin^2\theta}$.

Hence

$$\left\| \mathbf{V}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F^2 = 2 \sum_{j=1}^{K} \left(1 - \sigma_j \right) = 2 \sum_{j=1}^{K} \left(1 - \cos\theta_j \right) \ge \sum_{j=1}^{K} \sin^2\theta_j = \frac{1}{2} \left\| \mathbf{P_U} - \mathbf{P_V} \right\|_F^2,$$

$$\left\|\mathbf{V}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F^2 \leq 2 \sum_{j=1}^K \frac{\sin^2\theta_j}{2 - \sin^2\theta_j} \leq \frac{2\sum_{j=1}^K \sin^2\theta_j}{2 - \sin^2\theta_K} = \frac{\left\|\mathbf{P_U} - \mathbf{P_V}\right\|_F^2}{2 - \left\|\sin\Theta\left(\mathbf{U},\mathbf{V}\right)\right\|_2^2},$$

$$\left\|\mathbf{H} - \widehat{\mathbf{H}}\right\|_2 = 1 - \sigma_K = 1 - \cos\theta_K \leq \frac{\sin^2\theta_K}{2 - \sin^2\theta_K} = \frac{\left\|\sin\Theta\left(\mathbf{U},\mathbf{V}\right)\right\|_2^2}{2 - \left\|\sin\Theta\left(\mathbf{U},\mathbf{V}\right)\right\|_2^2}.$$

□

Lemma 8. *Consider the settings for Lemma 2 and define*
$\mathbf{H} = \widehat{\mathbf{U}}^T\mathbf{U}$ *and* $\widehat{\mathbf{H}} = sgn\left(\mathbf{H}\right)$. *When* $\varepsilon = \left\|\mathbf{E}\right\|_2 / \Delta \leq 1/10$, *we have* $\left\|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T\right\|_2 \leq \varepsilon/(1 - \varepsilon)$,

$$\frac{\left\|f(\mathbf{EU})\right\|_F}{1 + 5\varepsilon} \leq \left\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F \leq \frac{\left\|f(\mathbf{EU})\right\|_F}{1 - 5\varepsilon}, \tag{8.22}$$

$$\left\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{EU})\right\|_F \leq 9\varepsilon\left\|f(\mathbf{EU})\right\|_F, \tag{8.23}$$

$$\frac{\sqrt{2}\left\|f(\mathbf{EU})\right\|_F}{1 + 7\varepsilon} \leq \left\|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T\right\|_F \leq \frac{\sqrt{2}\left\|f(\mathbf{EU})\right\|_F}{1 - 7\varepsilon}, \tag{8.24}$$

$$\left\|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T - \left[f(\mathbf{EU})\mathbf{U}^T + \left|\mathbf{U}f(\mathbf{EU})^T\right|\right]\right\|_F \leq 24\varepsilon\left\|f(\mathbf{EU})\right\|_F. \tag{8.25}$$

Besides, $\left\|f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right\|_F = \sqrt{2}\left\|f(\mathbf{EU})\right\|_F$.

Proof. Define $\widehat{\mathbf{P}} = \widehat{\mathbf{U}}\widehat{\mathbf{U}}^T$, $\mathbf{P} = \mathbf{U}\mathbf{U}^T$ and $\mathbf{P}_\perp = \mathbf{I} - \mathbf{P}$. The Davis-Kahan sin Θ theorem (Davis and Kahan, 1970) and Lemma 7 force that
$\delta\left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_2 \leq \left\|\mathbf{EP}\right\|_2 \leq \left\|\mathbf{E}\right\|_2$, where $\delta = \min\left\{\left(\hat{\lambda}_s - \lambda_{s+1}\right)_+, \left(\lambda_{s+r} - \hat{\lambda}_{s+r+1}\right)_+\right\}$ and we define x_+
$= \max\left\{x, 0\right\}$ for $x \in \mathbb{R}$.

Since the Weyl's inequality (Stewart and Sun, 1990, Corollary IV.4.9) leads to
$\delta \geq \Delta - \left\|\mathbf{E}\right\|_2 = (1 - \varepsilon)\Delta$, we get $\left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_2 \leq \varepsilon/(1 - \varepsilon)$.

To attack (8.22) and (8.23), we divide the difference

$$\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{EU}) = \left[\mathbf{P}_\perp\widehat{\mathbf{U}}\widehat{\mathbf{H}} - f(\mathbf{EU})\right] + \mathbf{P}_\perp\widehat{\mathbf{U}}(\widehat{\mathbf{H}} - \mathbf{H}) + (\mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}) \tag{8.26}$$

and conquer the terms separately. Since $\varepsilon < 1/2$, the first claim in Lemma 8 yields $\left\| \widehat{\mathbf{U}}\widehat{\mathbf{U}}^T - \mathbf{U}\mathbf{U}^T \right\|_2 < 1$. Then according to Lemma 7, $\widehat{\mathbf{H}}$ is orthonormal,

$$\left\| \mathbf{P}_\perp \widehat{\mathbf{U}}(\widehat{\mathbf{H}} - \mathbf{H}) \right\|_F = \left\| \mathbf{P}_\perp (\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})\widehat{\mathbf{H}}^T(\widehat{\mathbf{H}} - \mathbf{H}) \right\|_F \le \left\| \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F \left\| \widehat{\mathbf{H}} - \mathbf{H} \right\|_2,$$

and

$$\left\| \mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F = \left\| \mathbf{U}\left(\mathbf{H}^T\widehat{\mathbf{H}} - \mathbf{I}\right) \right\|_F \le \left\| \mathbf{H}^T\widehat{\mathbf{H}} - \mathbf{I} \right\|_F = \left\| (\mathbf{H} - \widehat{\mathbf{H}})^T\widehat{\mathbf{H}} \right\|_F = \left\| \mathbf{H} - \widehat{\mathbf{H}} \right\|_F \le \left\| \mathbf{H} - \widehat{\mathbf{H}} \right\|_2^{1/2} \left\| \mathbf{H} - \widehat{\mathbf{H}} \right\|_*^{1/2} = \left\| \mathbf{H} - \widehat{\mathbf{H}} \right\|_2^{1/2} \left\| \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F / \sqrt{2}. \tag{8.27}$$

Observe that when $\varepsilon < 1/10$, Lemma 7 forces that

$$\|\mathbf{H} - \widehat{\mathbf{H}}\|_2 \le \frac{\|\widehat{\mathbf{P}} - \mathbf{P}\|_2^2}{2 - \|\widehat{\mathbf{P}} - \mathbf{P}\|_2^2} \le \frac{\left(\frac{\varepsilon}{1-\varepsilon}\right)^2}{2 - \left(\frac{\varepsilon}{1-\varepsilon}\right)^2} = \frac{\varepsilon^2}{2(1-\varepsilon)^2 - \varepsilon^2} \le \frac{5}{8}\varepsilon^2 \le \frac{1}{16}\varepsilon. \tag{8.28}$$

Combining the estimates above yields

$$\|\mathbf{P}_\perp \widehat{\mathbf{U}}(\widehat{\mathbf{H}} - \mathbf{H}) + (\mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})\|_F \le \left(\frac{1}{16} + \sqrt{\frac{5}{16}}\right)\varepsilon\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F \le \frac{16}{25}\varepsilon\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F \tag{8.29}$$

.

We start to work on $\mathbf{P}_\perp \widehat{\mathbf{U}}\mathbf{H} - f(\mathbf{E}\mathbf{U})$. Define $\mathbf{\Lambda} = \operatorname{diag}\left(\lambda_{s+1}, \cdots, \lambda_{s+K}\right)$, and $L(\mathbf{V}) = \mathbf{A}\mathbf{V} - \mathbf{V}\mathbf{\Lambda}$ for $\mathbf{V} \in \mathbb{R}^{d \times K}$. Note that $L\left(\mathbf{v}_1, \cdots, \mathbf{v}_K\right) = \left(\left(\mathbf{A} - \lambda_{s+1}\mathbf{I}\right)\mathbf{v}_1, \cdots, \left(\mathbf{A} - \lambda_{s+K}\mathbf{I}\right)\mathbf{v}_K\right)$, and $\mathbf{G}_j\left(\mathbf{A} - \lambda_{s+j}\mathbf{I}\right) = \mathbf{P}_\perp$ holds for all $j \in [K]$. As a result, $f(L(\mathbf{V})) = -\mathbf{P}_\perp\mathbf{V}$ for any $\mathbf{V} \in \mathbb{R}^{d \times K}$. This motivates us to work on $L(\widehat{\mathbf{U}}\mathbf{H})$ in order to study $\mathbf{P}_\perp\widehat{\mathbf{U}}\mathbf{H}$.

Let $\widehat{\mathbf{\Lambda}} = \operatorname{diag}\left(\widehat{\lambda}_{s+1}, \cdots, \widehat{\lambda}_{s+K}\right)$. By definition, $\widehat{\mathbf{A}}\widehat{\mathbf{U}} = \widehat{\mathbf{U}}\widehat{\mathbf{\Lambda}}$ and

$$L(\widehat{\mathbf{U}}\mathbf{H}) = \mathbf{A}\widehat{\mathbf{U}}\mathbf{H} - \widehat{\mathbf{U}}\mathbf{H}\mathbf{\Lambda} = (\mathbf{A} - \widehat{\mathbf{A}})\widehat{\mathbf{U}}\mathbf{H} + (\widehat{\mathbf{A}}\widehat{\mathbf{U}} - \widehat{\mathbf{U}}\widehat{\mathbf{\Lambda}})\mathbf{H} + \widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H} + \widehat{\mathbf{U}}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}) = -\mathbf{E}\widehat{\mathbf{U}}\mathbf{H} + \widehat{\mathbf{U}}(\widehat{\mathbf{\Lambda}} - \mathbf{\Lambda})\mathbf{H} + \widehat{\mathbf{U}}(\mathbf{\Lambda}\mathbf{H} - \mathbf{H}\mathbf{\Lambda}). \tag{8.30}$$

Now we study the images of these three terms under the linear mapping f. First, the facts $\left\| f(\cdot) \right\|_F \le \Delta^{-1} \left\| \cdot \right\|_F$ and $\left\| \widehat{\mathbf{U}}\mathbf{H} - \mathbf{U} \right\|_F \le \left\| \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F$ (by Lemma 7) imply that

$$\left\|f(\mathbf{E}\widehat{\mathbf{U}}\mathbf{H}) - f(\mathbf{E}\mathbf{U})\right\|_F = \left\|f[\mathbf{E}(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})]\right\|_F \leq \ \Delta^{-1}\ \left\|\mathbf{E}(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\right\|_F \leq \ \Delta^{-1}$$
$$\left\|\mathbf{E}\right\|_2\left\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\right\|_F \leq \varepsilon\left\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\right\|_F. \tag{8.31}$$

Second, the definition of f forces $f(\mathbf{U}M) = 0$ for all $\mathbf{M} \in \mathbb{R}^{K \times K}$.

$$\begin{aligned}\|f[\widehat{\mathbf{U}}(\widehat{\boldsymbol{\Lambda}} \ - \ \boldsymbol{\Lambda})\mathbf{H}]\|_F &= \left\|f\left[(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\widehat{\mathbf{H}}^T(\widehat{\boldsymbol{\Lambda}} \ - \ \boldsymbol{\Lambda})\mathbf{H}\right]\right\|_F \\ &\leq \ \Delta^{-1}\ \|(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\widehat{\mathbf{H}}^T(\widehat{\boldsymbol{\Lambda}} \ - \ \boldsymbol{\Lambda})\mathbf{H}\|_F \leq \ \Delta^{-1}\ \|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F\|\widehat{\mathbf{H}}^T(\widehat{\boldsymbol{\Lambda}} \ - \ \boldsymbol{\Lambda})\mathbf{H}\|_2 \\ &\leq \ \Delta^{-1}\ \|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F\|\mathbf{E}\|_2\|\mathbf{H}\|_2 \leq \varepsilon\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_F. \end{aligned} \tag{8.32}$$

Here we applied Weyl's inequality $\| \widehat{\boldsymbol{\Lambda}} \ - \ \boldsymbol{\Lambda} \|_2 \leq \|\mathbf{E}\|_2$ and used the fact that $\left\|\mathbf{H}\right\|_2 = \left\|\widehat{\mathbf{U}}^T\mathbf{U}\right\|_2 \leq 1$. Third, by similar tricks we work on the third term

$$\begin{aligned}\left\|f[\widehat{\mathbf{U}}(\boldsymbol{\Lambda}\ \mathbf{H} - \mathbf{H}\ \boldsymbol{\Lambda})]\right\|_F &= \left\|f\left[(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\widehat{\mathbf{H}}^T(\boldsymbol{\Lambda}\ \mathbf{H} - \mathbf{H}\ \boldsymbol{\Lambda})\right]\right\|_F \\ &\leq \ \Delta^{-1}\ \left\|(\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U})\widehat{\mathbf{H}}^T(\boldsymbol{\Lambda}H - \mathbf{H}\ \boldsymbol{\Lambda})\right\|_F \leq \ \Delta^{-1}\ \left\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\right\|_F\left\|\boldsymbol{\Lambda}\ \mathbf{H} - \mathbf{H}\ \boldsymbol{\Lambda}\right\|_2. \end{aligned} \tag{8.33}$$

As an intermediate step, we are going to show that $\|\boldsymbol{\Lambda}H - \mathbf{H}\ \boldsymbol{\Lambda}\|_2 \leq 2\|\mathbf{E}\|_2$. On the one hand, $\widehat{\mathbf{A}}\widehat{\mathbf{U}} = \widehat{\mathbf{U}}\ \widehat{\boldsymbol{\Lambda}}$ yields

$$L(\widehat{\mathbf{U}}) = (\mathbf{A} - \widehat{\mathbf{A}})\widehat{\mathbf{U}} + (\widehat{\mathbf{A}}\widehat{\mathbf{U}} - \widehat{\mathbf{U}}\ \widehat{\boldsymbol{\Lambda}}) + \widehat{\mathbf{U}}(\widehat{\boldsymbol{\Lambda}} \ - \ \boldsymbol{\Lambda}) = \ - \mathbf{E}\widehat{\mathbf{U}} + \widehat{\mathbf{U}}(\widehat{\boldsymbol{\Lambda}} \ - \ \boldsymbol{\Lambda}). \tag{8.34}$$

On the other hand, let
$\mathbf{U}_1 = (\mathbf{u}_1, \cdots, \mathbf{u}_s, \mathbf{u}_{s+K+1}, \cdots, \mathbf{u}_d), \widehat{\mathbf{U}}_1 = (\widehat{\mathbf{u}}_1, \cdots, \widehat{\mathbf{u}}_s, \widehat{\mathbf{u}}_{s+K+1}, \cdots, \widehat{\mathbf{u}}_d)$, and $\boldsymbol{\Lambda}_1 = $ We have $\text{diag}(\lambda_1, \cdots, \lambda_s, \lambda_{s+K+1}, \cdots, \lambda_d)$.

$$\mathbf{A}\widehat{\mathbf{U}} = \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\boldsymbol{\Lambda} & 0 \\ 0 & \boldsymbol{\Lambda}_1\end{pmatrix}\begin{pmatrix}\mathbf{U}^T \\ \mathbf{U}_1^T\end{pmatrix}\widehat{\mathbf{U}} = \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\boldsymbol{\Lambda}\mathbf{H}^T \\ \boldsymbol{\Lambda}_1 \mathbf{U}_1^T\widehat{\mathbf{U}}\end{pmatrix},$$

$$\widehat{\mathbf{U}}_{\boldsymbol{\Lambda}} = \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\mathbf{U}^T \\ \mathbf{U}_1^T\end{pmatrix}\widehat{\mathbf{U}}\boldsymbol{\Lambda} = \begin{pmatrix}\mathbf{U} & \mathbf{U}_1\end{pmatrix}\begin{pmatrix}\mathbf{H}^T\boldsymbol{\Lambda} \\ \mathbf{U}_1^T\widehat{\mathbf{U}}\boldsymbol{\Lambda}\end{pmatrix}.$$

As a result, (8.34) yields that

$$\left\|\boldsymbol{\Lambda}\ \mathbf{H} - \mathbf{H}\ \boldsymbol{\Lambda}\right\|_2 = \left\|\mathbf{H}^T\boldsymbol{\Lambda} \ - \ \boldsymbol{\Lambda}\ \mathbf{H}^T\right\|_2 \leq \left\|L(\widehat{\mathbf{U}})\right\|_2 = \left\| \ - \mathbf{E}\widehat{\mathbf{U}} + \widehat{\mathbf{U}}(\widehat{\boldsymbol{\Lambda}} \ - \boldsymbol{\Lambda})\right\|_2$$
$$\leq 2\left\|\mathbf{E}\right\|_2. \tag{8.35}$$

By combining (8.30), (8.31), (8.32), (8.33) and (8.35), we obtain that

$$\left\| \mathbf{P}_{\perp} \widehat{\mathbf{U}} \mathbf{H} - f(\mathbf{E}\mathbf{U}) \right\|_F = \left\| - f[L(\widehat{\mathbf{U}}\mathbf{H})] - f(\mathbf{E}\mathbf{U}) \right\|_F \leq 4\varepsilon \left\| \widehat{\mathbf{U}}\widehat{\mathbf{U}} - \mathbf{U} \right\|_F. \tag{8.36}$$

Based on (8.26), (8.29) and (8.36), we obtain that

$$\left\| \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{E}\mathbf{U}) \right\|_F \leq \frac{116}{25}\varepsilon \left\| \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F. \tag{8.37}$$

It follows from the triangle's inequality that

$$\frac{\|f(\mathbf{E}\mathbf{U})\|_F}{1+5\varepsilon} \leq \frac{\|f(\mathbf{E}\mathbf{U})\|_F}{1+116\varepsilon/25} \leq \left\| \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} \right\|_F \leq \frac{\|f(\mathbf{E}\mathbf{U})\|_F}{1-116\varepsilon/25} \leq \frac{\|f(\mathbf{E}\mathbf{U})\|_F}{1-5\varepsilon},$$

$$\left\| \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{E}\mathbf{U}) \right\|_F \leq \frac{116}{25}\varepsilon \frac{\|f(\mathbf{E}\mathbf{U})\|_F}{1 - \frac{116}{25} \cdot \frac{1}{10}} \leq 8.66\varepsilon \left\| f(\mathbf{E}\mathbf{U}) \right\|_F \leq 9\varepsilon \left\| f(\mathbf{E}\mathbf{U}) \right\|_F.$$

Hence we have proved (8.22) and (8.23). Now we move on to (8.24) and (8.25). Note that

$$\widehat{\mathbf{P}} - \mathbf{P} = \widehat{\mathbf{U}}\widehat{\mathbf{H}}(\widehat{\mathbf{U}}\widehat{\mathbf{H}})^T - \mathbf{U}\mathbf{U}^T = (\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})(\widehat{\mathbf{U}}\widehat{\mathbf{H}})^T + \mathbf{U}(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})^T = (\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})^T + (\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})\mathbf{U}^T + \mathbf{U}(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})^T. \tag{8.38}$$

The first term is controlled by

$$\|(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})^T\|_F \leq \|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_2 \|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F$$

$$\leq \left(\|\widehat{\mathbf{U}}\mathbf{H} - \mathbf{U}\|_2 + \|\widehat{\mathbf{U}}(\widehat{\mathbf{H}} - \mathbf{H})\|_2 \right) \|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F = \left(\|(\widehat{\mathbf{P}} - \mathbf{P})\mathbf{U}\|_2 + \|\widehat{\mathbf{U}}(\widehat{\mathbf{H}} - \mathbf{H})\|_2 \right)$$

$$\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F \leq \left(\|\widehat{\mathbf{P}} - \mathbf{P}\|_2 + \|\widehat{\mathbf{H}} - \mathbf{H}\|_2 \right) \|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F \leq \left(\frac{1}{1-\varepsilon} + \frac{1}{16} \right)\varepsilon\|\widehat{\mathbf{U}}\widehat{\mathbf{H}}$$

$$- \mathbf{U}\|_F \leq 1.18\varepsilon\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\|_F, \tag{8.39}$$

where the penultimate inequality uses $\left\| \widehat{\mathbf{P}} - \mathbf{P} \right\|_2 \leq \varepsilon/(1-\varepsilon)$ and (8.28). By defining $\mathbf{W} = \widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{E}\mathbf{U})$ we can write

$$(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})\mathbf{U}^T + \mathbf{U}(\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U})^T = \left[f(\mathbf{E}\mathbf{U})\mathbf{U}^T + \mathbf{U}f(\mathbf{E}\mathbf{U})^T \right] + \left(\mathbf{W}\mathbf{U}^T + \mathbf{U}\mathbf{W}^T \right). \tag{8.40}$$

It is easily seen that

$$\left\|\mathbf{WU}^T + \mathbf{UW}^T\right\|_F = \left\|\left(\mathbf{P}_\perp \mathbf{WU}^T + \mathbf{P}w\mathbf{U}^T\right) + \left(\mathbf{UW}^T\mathbf{P}_\perp + \mathbf{U}^T\mathbf{W}P\right)\right\|_F \leq \left\|\mathbf{P}_\perp \mathbf{WU}^T + \mathbf{UW}^T\mathbf{P}_\perp\right\|_F$$

$$+ \left\|\mathbf{P}w\mathbf{U}^T + \mathbf{U}^T\mathbf{W}P\right\|_F = \left(\left\|\mathbf{P}_\perp \mathbf{WU}^T\right\|_F^2 + \left\|\mathbf{UW}^T\mathbf{P}_\perp\right\|_F^2\right)^{1/2} + \left\|\mathbf{P}w\mathbf{U}^T + \mathbf{U}^T\mathbf{W}P\right\|_F \leq \sqrt{2}\left\|\mathbf{P}_\perp \mathbf{W}\right\|_F + 2\left\|\mathbf{P}w\right\|_F.$$

On the one hand, (8.37) forces that $\left\|\mathbf{P}_\perp \mathbf{W}\right\|_F \leq \left\|\mathbf{W}\right\|_F \leq \frac{116}{25}\varepsilon\left\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F$. On the other hand, the fact $\mathbf{P}f(\mathbf{EU}) = \mathbf{0}$, (8.27) and (8.28) yield

$$\left\|\mathbf{P}\mathbf{W}\right\|_F = \left\|\mathbf{P}[\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U} - f(\mathbf{EU})]\right\|_F = \left\|\mathbf{P}\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F \leq \sqrt{\frac{5}{16}}\varepsilon\left\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F.$$

Hence

$$\left\|\mathbf{WU}^T + \mathbf{UW}^T\right\|_F \leq \left(\frac{116}{25}\sqrt{2} + 2\sqrt{\frac{5}{16}}\right)\varepsilon\left\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F \leq 7.68\varepsilon\left\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F. \tag{8.41}$$

By collecting (8.38), (8.39), (8.40) and (8.41) we derive that

$$\left\|\widehat{\mathbf{P}} - \mathbf{P} - \left[f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right]\right\|_F \leq 8.86\varepsilon\left\|\widehat{\mathbf{U}}\widehat{\mathbf{H}} - \mathbf{U}\right\|_F \leq 8.86\varepsilon\frac{\left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_F}{\sqrt{2 - \left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_2^2}} \leq 8.86\varepsilon\frac{\left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_F}{\sqrt{2 - 1/9^2}} \leq 6.29\varepsilon\left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_F,$$

where we also used Lemma 7 and $\left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_2 \leq \frac{\varepsilon}{1-\varepsilon} \leq 1/9$. Therefore,

$$\frac{\left\|f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right\|_F}{1 + 6.29\varepsilon} \leq \left\|\widehat{\mathbf{P}} - \mathbf{P}\right\|_F \leq \frac{\left\|f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right\|_F}{1 - 6.29\varepsilon}, \left\|\widehat{\mathbf{P}} - \mathbf{P} - \left[f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right]\right\|_F$$

$$\leq \frac{6.29\varepsilon}{1 - 6.29\varepsilon}\left\|f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right\|_F \leq 16.96\varepsilon\left\|f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right\|_F.$$

We finish the proof by

$$\left\|f(\mathbf{EU})\mathbf{U}^T + \mathbf{U}f(\mathbf{EU})^T\right\|_F^2 = \left\|f(\mathbf{EU})\mathbf{U}^T\right\|_F^2 + \left\|\mathbf{U}f(\mathbf{EU})^T\right\|_F^2 + 2\text{Tr}\left(\left[f(\mathbf{EU})\mathbf{U}^T\right]^T\mathbf{U}f(\mathbf{EU})^T\right) = \left\|f(\mathbf{EU})\mathbf{U}^T\right\|_F^2$$

$$+ \left\|\mathbf{U}f(\mathbf{EU})^T\right\|_F^2 + 0 = 2\left\|f(\mathbf{EU})\right\|_F^2.$$

□

Acknowledgments

References

Anderson TW (1963). Asymptotic theory for principal component analysis. The Annals of Mathematical Statistics 34 122–148.

Baik J, Ben Arous G and PEche S (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. The Annals of Probability 33 1643–1697.

Battey H, Fan J, Liu H, Lu J and Zhu Z (2015). Distributed estimation and inference with statistical guarantees. arXiv preprint arXiv:150905457.

Bertrand A and Moonen M (2014). Distributed adaptive estimation of covariance matrix eigenvectors in wireless sensor networks with application to distributed PCA. Signal Processing 104 120–135.

Blanchard G and Mucke N (2017). Parallelizing spectral algorithms for kernel learning. arXiv preprint arXiv:161007497.

Bosq D (2000). Stochastic processes and random variables in function spaces In Linear Processes in Function Spaces. Springer, 15–42.

Boutsidis C, Woodruff DP and Zhong P (2016). Optimal principal component analysis in distributed and streaming models. In Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing ACM.

Cai TT, Ma Z and Wu Y (2013). Sparse PCA: Optimal rates and adaptive estimation. The Annals of Statistics 41 3074–3110.

Chen T-L, Chang DD, Huang S-Y, Chen H, Lin C and Wang W (2016). Integrating multiple random sketches for singular value decomposition. arXiv preprint arXiv:160808285.

Chen X and Xie M-G (2014). A split-and-conquer approach for analysis of extraordinarily large data. Statistica Sinica 24 1655–1684.

Davis C and Kahan WM (1970). The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis* 7 1–46.

El Karoui N and d'Aspremont A (2010). Second order accurate distributed eigenvector computation for extremely large matrices. Electronic Journal of Statistics 4 1345–1385.

Fan J, Wang W and Zhu Z (2016). Robust low-rank matrix recovery. arXiv preprint arXiv:160308315.

Garber D, Shamir O and Srebro N (2017). Communication-efficient algorithms for distributed stochastic principal component analysis. arXiv preprint arXiv:170208169

Golub GH and Van Loan CF (2012). Matrix computations, vol. 3 JHU Press.

Gross D (2011). Recovering low-rank matrices from few coefficients in any basis. IEEE Transactions on Information Theory 57 1548–1566.

Guo Z-C, Lin S-B and Zhou D-X (2017). Learning theory of distributed spectral algorithms. Inverse Problems.

Halko N, Martinsson P-G and Tropp JA (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. SIAM review 53 217–288.

Hotelling H (1933). Analysis of a complex of statistical variables into principal components. Journal of Educational Psychology 24 417–441.

Johnstone IM (2001). On the distribution of the largest eigenvalue in principal components analysis. The Annals of Statistics 29 295–327.

Johnstone IM and Lu AY (2009). On consistency and sparsity for principal components analysis in high dimensions. Journal of the American Statistical Association 104 682–693. [PubMed: 20617121]

Jung S and Marron JS (2009). PCA consistency in high dimension, low sample size context. The Annals of Statistics 37 4104–4130.

Kannan R, Vempala S and Woodruff D (2014). Principal component analysis and higher correlations for distributed data. In Conference on Learning Theory.

Kargupta H, Huang W, Sivakumar K and Johnson E (2001). Distributed clustering using collective principal component analysis. Knowledge and Information Systems 3 422–448.

Kato T (1966). Perturbation theory for linear operators. Springer.

Kneip A and Utikal KJ (2001). Inference for density families using functional principal component analysis. Journal of the American Statistical Association 96 519–542.

Koltchinskii V and Lounici K (2016). Asymptotics and concentration bounds for bilinear forms of spectral projectors of sample covariance In Annales de l'Institut Henri Poincaré, Probabilités et Statistiques, vol. 52 Institut Henri Poincare.

Koltchinskii V and Lounici K (2017). Concentration inequalities and moment bounds for sample covariance operators. Bernoulli 23 110–133.

Lee JD, Liu Q, Sun Y and Taylor JE (2017). Communication-efficient sparse regression. Journal of Machine Learning Research 18 1–30.

Li L, Scaglione A and Manton JH (2011). Distributed principal subspace estimation in wireless sensor networks. IEEE Journal of Selected Topics in Signal Processing 5 725–738.

Liang Y, Balcan M-FF, Kanchanapally V and Woodruff D (2014). Improved distributed principal component analysis. In Advances in Neural Information Processing Systems.

Minsker S (2016). Sub-gaussian estimators of the mean of a random matrix with heavytailed entries. arXiv preprint arXiv:160507129.

Nadler B (2008). Finite sample approximation results for principal component analysis: A matrix perturbation approach. The Annals of Statistics 36 2791–2817.

Onatski A (2012). Asymptotics of the principal components estimator of large factor models with weakly influential factors. Journal of Econometrics 168 244–258.

Paul D (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. Statistica Sinica 17 1617–1642.

Paul D and Johnstone IM (2012). Augmented sparse principal component analysis for high dimensional data. arXiv preprint arXiv:12021242.

Pearson K (1901). On lines and planes of closest fit to systems of point in space. Philosophical Magazine Series 6 2 559–572.

Qu Y, Ostrouchov G, Samatova N and Geist A (2002). Principal component analysis for dimension reduction in massive distributed data sets. In IEEE International Conference on Data Mining (ICDM).

Reiss M and Wahl M (2016). Non-asymptotic upper bounds for the reconstruction error of PCA. arXiv preprint arXiv:160903779.

Schizas ID and Aduroja A (2015). A distributed framework for dimensionality reduction and denoising. IEEE Transactions on Signal Processing 63 6379–6394.

Shen D, Shen H and Marron JS (2013). Consistency of sparse PCA in high dimension, low sample size contexts. Journal of Multivariate Analysis 115 317–333.

Shen D, Shen H, Zhu H and Marron J (2016). The statistics and mathematics of high dimension low sample size asymptotics. Statistica Sinica 26 1747–1770. [PubMed: 28018116]

Stewart GW and Sun J (1990). Matrix perturbation theory. Academic Press.

Tropp JA, Yurtsever A, Udell M and Cevher V (2016). Randomized singleview algorithms for low-rank matrix approximation. arXiv preprint arXiv:160900048.

Vaccaro RJ (1994). A second-order perturbation expansion for the svd. SIAM Journal on Matrix Analysis and Applications 15 661–671.

Vershynin R (2012). Introduction to the non-asymptotic analysis of random matrices. Compressed Sensing, Theory and Applications 210–268.

Vu VQ, Lei J et al. (2013). Minimax sparse principal subspace estimation in high dimensions. The Annals of Statistics 41 2905–2947.

Wang R (2015). Singular vector perturbation under gaussian noise. SIAM Journal on Matrix Analysis and Applications 36 158–177.

Wang W and Fan J (2017). Asymptotics of empirical eigen-structure for ultra-high dimensional spiked covariance model. The Annals of Statistics.

Wei X and Minsker S (2017). Estimation of the covariance structure of heavy-tailed distributions. In Advances in Neural Information Processing Systems.

Xu Z (2002). Perturbation analysis for subspace decomposition with applications in subspace-based algorithms. IEEE Transactions on Signal Processing 50 2820–2830.

Yu Y, Wang T and Samworth R (2015). A useful variant of the Davis-Kahan theorem for statisticians. Biometrika 102 315–323.

Zhang Y, Duchi JC and Wainwright MJ (2013). Divide and conquer kernel ridge regression. In COLT.
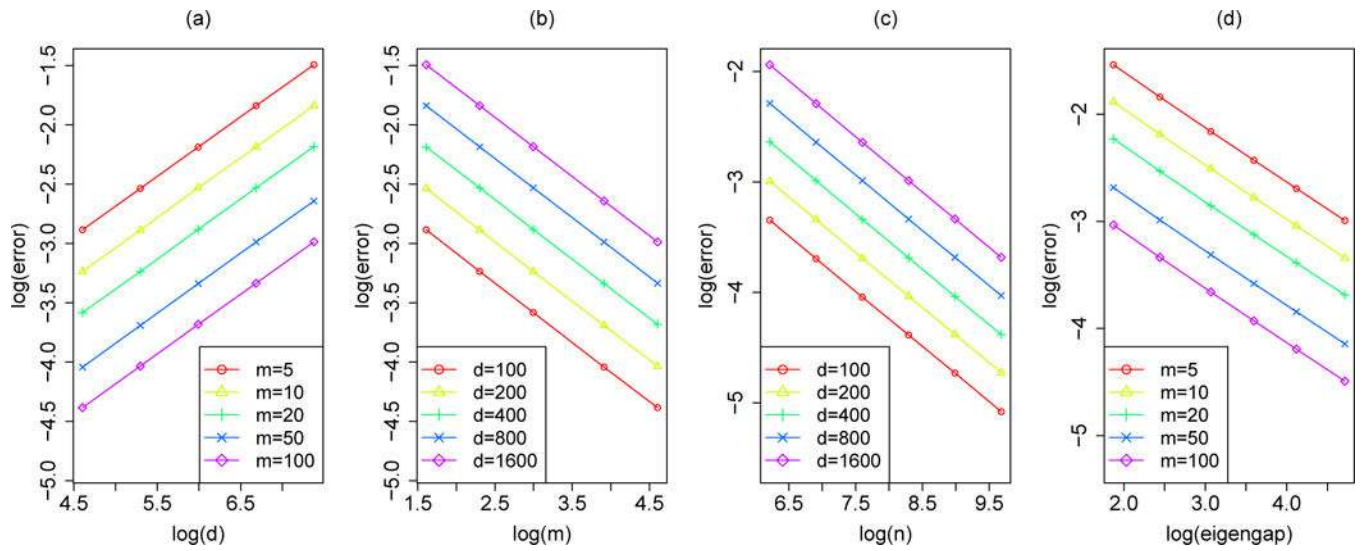
Figure 1:

Statistical error rate with respect to: (a) the dimension d when $\lambda = 50$ and $n = 2000$; (b) the number of servers m when $\lambda = 50$ and $n = 2000$; (c) the subsample size n when $\lambda = 50$ and $m = 50$; (d) the eigengap δ when $d = 800$ and $n = 2000$.

Figure 2:
Observed and fitted values of $\log\left(\rho\left(\tilde{V}_K, \mathbf{V}_K\right)\right)$.

Figure 3:
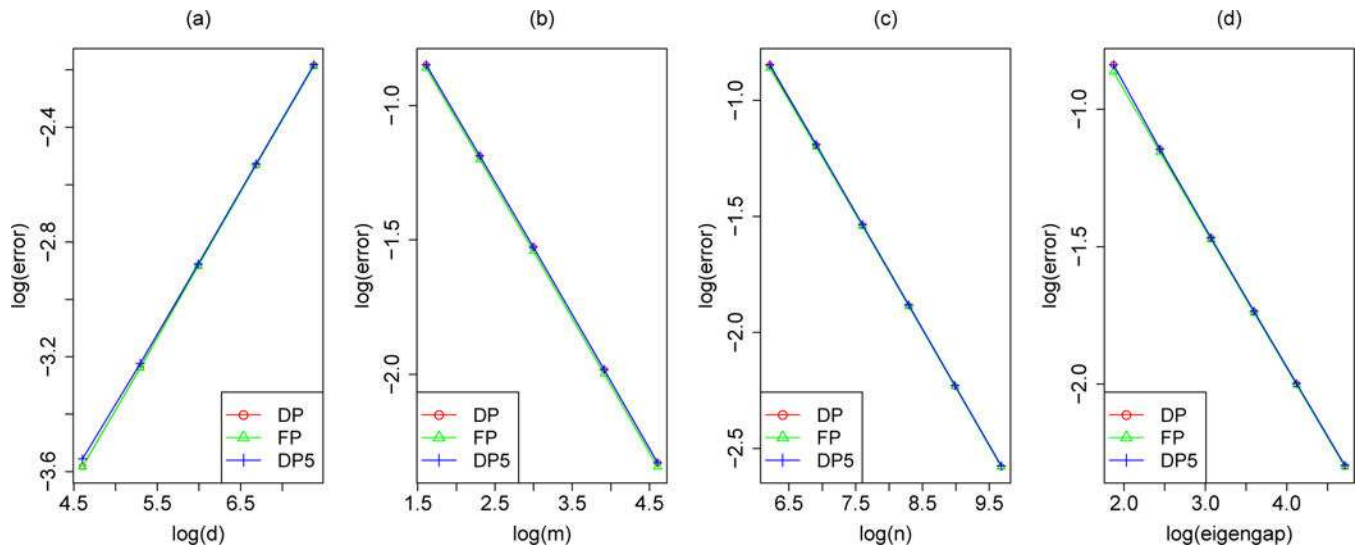Statistical error with respect to the number of machines when the total sample size $N = 6000$ is fixed.

Figure 4:
Comparison between DP, FP and DP5: (a) $m = 20$, $n = 2000$ and $\lambda = 50$; (b) $d = 1600$, $n = 1000$ and $\lambda = 30$; (c) $d = 800$, $m = 5$ and $\lambda = 30$; (d) $d = 1600$, $m = 10$ and $n = 500$.