

On the Error-Propagation of Inexact Hotelling’s Deflation for Principal Component Analysis

Fangshuo Liao¹ Junhyung Lyle Kim¹ Cruz Barnum² Anastasios Kyriillidis¹

Abstract

Principal Component Analysis (PCA) aims to find subspaces spanned by the so-called *principal components* that best represent the variance in the dataset. The deflation method is a popular meta-algorithm that sequentially finds individual principal components, starting from the most important ones and working towards the less important ones. However, as deflation proceeds, numerical errors from the imprecise estimation of principal components propagate due to its sequential nature. This paper mathematically characterizes the error propagation of the inexact Hotelling’s deflation method. We consider two scenarios: *i*) when the sub-routine for finding the leading eigenvector is abstract and can represent various algorithms; and *ii*) when power iteration is used as the sub-routine. In the latter case, the additional directional information from power iteration allows us to obtain a tighter error bound than the sub-routine agnostic case. For both scenarios, we explicitly characterize how the errors progress and affect subsequent principal component estimations.

1. Introduction

Principal Component Analysis (PCA) (Pearson, 1901; Hotelling, 1933) is a fundamental tool for data analysis with applications that range from statistics to machine learning and can be used for dimensionality reduction, classification, and clustering, to name a few (Majumdar, 2009; Wang et al., 2013; d’Aspremont et al., 2007; Jiang et al., 2011; Zou et al., 2006). Since its proposition Pearson (1901), many algorithms for PCA have been devised, even recently (Allen-Zhu & Li, 2017), to efficiently find approximations

^{*}Equal contribution ¹Department of Computer Science, Rice University, Houston, U.S.A. ²Department of Computer Science, UIUC. Correspondence to: Fangshuo Liao <Fangshuo.Liao@rice.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

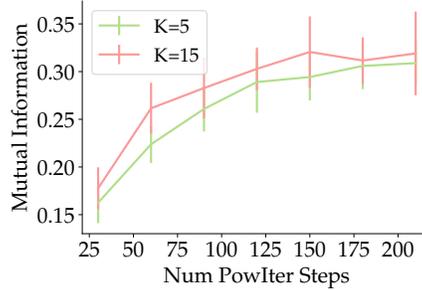


Figure 1. Spectral clustering of MNIST dataset using inexact deflation method in Algorithm 1. As the number of power iteration steps increases (*x*-axis), clustering performance, measured by the mutual information metric, also improves. A similar pattern is observed for recovering different numbers of eigenvectors.

of the first principal component, which is a vector that captures the most variance in the data. Mathematically, single-component PCA can be described as:

$$\mathbf{u}_1^* = \text{PCA}(\Sigma) = \arg \max_{\mathbf{v} \in \mathbb{R}^d: \|\mathbf{v}\|_2=1} \mathbf{v}^\top \Sigma \mathbf{v}, \quad (1)$$

where $\Sigma = \frac{1}{n} \cdot \mathbf{Y} \mathbf{Y}^\top \in \mathbb{R}^{d \times d}$ is the empirical covariance matrix given a set of n centered d -dimensional data-points $\mathbf{Y} \in \mathbb{R}^{d \times n}$.

It is sometimes helpful to know not only the first but also the top- K principal components, especially in the high-dimensional regime. This is becoming more relevant in the modern machine learning era, where the number of features reaches as large as billions (Fedus et al., 2022; Touvron et al., 2023). Mathematically, the multi-component PCA ($K \geq 1$) can be formulated as the following constrained optimization problem:

$$\mathbf{U}^* = \arg \max_{\mathbf{V} \in \{\mathbf{Q}_{:,1:K}: \mathbf{Q} \in \text{SO}(d)\}} \langle \Sigma \mathbf{V}, \mathbf{V} \rangle, \quad (2)$$

where $\text{SO}(d)$ denotes the group of rotations about a fixed point in d -dimensional Euclidean space. Since \mathbf{U}^* contains the first K columns of an orthogonal matrix, each principal component $\mathbf{U}_{:,k}$ must be a unit vector, and the pairwise principal components, say $\mathbf{U}_{:,k_1}^*, \mathbf{U}_{:,k_2}^*$ for $k_1 \neq k_2$, are orthogonal. A popular method for solving (2) is the *deflation method* (Hotelling, 1933) via sequentially solving (1), which is the focus of our work.

Algorithm 1 Deflation with inexact PCA

Require: $\Sigma \in \mathbb{R}^{d \times d}$, # of eigenvectors K ; sub-routine for top eigenvector $\text{PCA}(\cdot, \cdot)$; # of iterations t .

Ensure: Approximate eigenvectors $\mathbf{V} = \{\mathbf{v}_k\}_{k=1}^K$.

```

1:  $\mathbf{V} \leftarrow \emptyset$ 
2:  $\Sigma_1 \leftarrow \Sigma$ 
3: for  $k = 1, \dots, K$  do
4:    $\mathbf{v}_k \leftarrow \text{PCA}(\Sigma_k, t)$ 
5:    $\Sigma_{k+1} \leftarrow \Sigma_k - \mathbf{v}_k \mathbf{v}_k^\top \Sigma_k \mathbf{v}_k \mathbf{v}_k^\top$ 
6:    $\mathbf{V} \leftarrow \mathbf{V} \cup \{\mathbf{v}_k\}$ 
7: end for
8: return  $\mathbf{V}$ 
    
```

Deflation proceeds as follows. After approximating the top component \mathbf{u}_1^* , the matrix Σ is further processed to exist on the subspace that is orthogonal to the subspace which is spanned by the first component. This process is repeated by applying single-component PCA in (1) on the deflated matrix, which leads to an approximation of the second component \mathbf{u}_2^* , and so on, as below:

$$\begin{aligned} \Sigma_1 &= \Sigma; & \mathbf{v}_k &= \text{PCA}(\Sigma_k, t); \\ \Sigma_{k+1} &= \Sigma_k - \mathbf{v}_k \mathbf{v}_k^\top \Sigma_k \mathbf{v}_k \mathbf{v}_k^\top, \end{aligned} \quad (3)$$

where $\text{PCA}(\Sigma_k, \cdot)$ returns a normalized approximation of the top eigenvector of the deflated matrix Σ_k . One estimates the subsequent principal components by running the same single-component PCA algorithm repetitively.

In this work, we focus on the more realistic scenario where the sub-routine, $\text{PCA}(\Sigma_k, \cdot)$, incurs numerical errors due to limited computation budget and finite precision: even solving (1) for a single component, our estimate is only an approximation to the true principal component. The overall procedure with inexact $\text{PCA}(\Sigma_k, \cdot)$ sub-routine is described in Algorithm 1. Line 4 performs the sub-routine algorithm for t iterations, denoted by $\text{PCA}(\Sigma_k, t)$, which approximately computes the top eigenvector of Σ_k . In line 5, Σ_k is deflated using \mathbf{v}_k ; hence, the numerical error from approximately solving \mathbf{v}_k in line 4 affects the quality of the deflated matrix Σ_{k+1} in the next iteration.

To motivate why studying the accumulation of errors is critical, we apply spectral clustering (Pothen et al., 1990) on the MNIST dataset and measure its performance using the mutual information metric¹. Spectral clustering involves computing the top- K eigenvectors from a similarity matrix and uses the corresponding entries in the K eigenvectors as features for the following clustering procedure. In our experiments, eigenvectors are computed using the deflation in Algorithm 1 with different precision levels of the PCA

¹Higher mutual information indicates more accurate clustering. For more details about the experiment, please see Appendix D.

sub-routine, indicated by different power iteration steps. We can observe from Figure 1 that, as the number of power iteration steps increases, the mutual information increases, showing a higher clustering accuracy. This shows that the errors incurred from inexact solving each top eigenvector in the deflation process indeed have a non-negligible impact on the downstream machine learning tasks.

This interdependent accumulation of errors coming from approximately solving the sub-routine—which results in the less accurate deflation matrix, and further affects the quality of the approximation of the top eigenvector in the next iteration—is precisely what we characterize in this work. To the best of our knowledge, *this is the first work that analyzes the inexact setting, although this procedure is vastly used in practice*. Our contributions can be summarized as follows:

- In Theorem 3.1, we characterize how the errors from approximately solving the sub-routine (PCA) propagate into the overall error of the deflation procedure. Informally, the error between the approximate principal component and the actual one needs to be exponentially small to control the accuracy for the top- K eigenvectors for an increasing K , where the base of the exponential growth depends on the eigengap of the target matrix.
- In Theorem 4.1, we consider the case where PCA is fixed to the power iteration (Müntz et al., 1913). By leveraging the directional information of the numerical error vector produced by each power iteration, we improve the error bound in Theorem 3.1 (which is agnostic to the PCA sub-routine) to exponential growth with constant base.
- Along with the proof, we also present empirical results to justify the core ideas in our theoretical analysis.

Related work. Most works focus on the convergence proof of numerically solving the top eigenvector of matrices (Lanczos, 1950; Jain et al., 2016; Xu et al., 2018b). However, these results only concern the case of the first eigenvector, and their result does not readily apply to numerical algorithms, such as deflation, that solve for multiple top eigenvectors.

Iterative methods for multi-component PCA include *Schur complement deflation* (Zhang, 2006), *orthogonalized deflation* (Sriperumbudur et al., 2007), and *projection methods* (Saad, 1988). Among them, Hotelling’s deflation, our focus, was reported to have the highest accuracy in practice (Danisman et al., 2014). For the comparison of these methods, we refer the readers to Mackey (2008, Section 2), which extends some of the above deflation methods to the sparse PCA setting (Zou et al., 2006; Johnstone & Lu, 2009), which is out of scope of our work. Importantly, all the works mentioned above assume the PCA step in (3) is solved exactly, i.e., $t = \infty$. Works that analyze the resulting error when solving multiple top eigenvectors include Allen-Zhu & Li (2017);

Gemp et al. (2021). However, the theoretical analysis in Allen-Zhu & Li (2017) and Gemp et al. (2021) only applies to specific algorithms such as EigenGame and Oja++. At the same time, our work focuses on deflation, a widely used algorithmic framework for solving multiple eigenvectors.

Subspace methods such as orthogonal iterations (Golub & Van Loan, 2013) are also widely used for discovering multiple eigenvectors. However, the deflation serves a different purpose than subspace methods. In particular, the sequential discovery of the eigenvectors in the deflation methods applies to scenarios where users have a dynamic necessity of the eigenvectors and require flexibility, especially in modern ML settings (Fawzi et al., 2022; Coulaud et al., 2013). For instance, compared with orthogonal iteration, which outputs all k eigenvectors at the same time, the deflation method will be more suitable for scenarios where the top principal component is known but more eigenvectors need to be computed depending on the need, and scenarios where eigenvectors are output in a consecutive way to facilitate downstream tasks.

In terms of the analysis technique, our work utilizes Weyl’s inequality (Weyl, 1912) and Davis-Kahan sin Θ Theorem (Davis & Kahan, 1970) to characterize the differences between eigenvalues of two matrices. In our fine-grained analysis in Section 4, we also utilize the Neumann expansion introduced by Eldridge et al. (2017) and characterize the subspace that the top eigenvector of a perturbed matrix lies in with a similar technique to Chen et al. (2020). However, the problem they study is entirely orthogonal to ours.

2. Problem Setup and Notations

Notation. For a vector \mathbf{a} , we use $\|\mathbf{a}\|_2$ to denote its ℓ_2 -norm. For a matrix \mathbf{A} , $\|\mathbf{A}\|_2$ to denote its spectral norm and $\|\mathbf{A}\|_F$ its Frobenius norm. $\lambda_j(\mathbf{A})$ denotes its j -th eigenvalue.

Problem setup. Given a symmetric matrix $\Sigma \in \mathbb{R}^{d \times d}$ with eigenvectors $\mathbf{u}_1^*, \dots, \mathbf{u}_d^*$ and eigenvalues $\lambda_1^*, \dots, \lambda_d^*$, sorted in descending order (c.f., Assumption 2.1), we want to estimate the top $K \leq d$ eigenvectors using the deflation method. Formally, let PCA be a sub-routine that returns an approximation of the top eigenvector of a given matrix. The deflation method repeatedly computes approximation \mathbf{v}_k of the principal eigenvector of the current deflation matrix Σ_k . It forms a new deflation matrix Σ_{k+1} by subtracting the component of \mathbf{v}_k from Σ_k . This process is made formal in Algorithm 1 and is mathematically defined in (3). If the sub-routine PCA is solved precisely, then the output satisfies $\mathbf{v}_k = \mathbf{u}_k^*$. In this ideal scenario, the sequence of deflation matrices has the following form:

$$\Sigma_{k+1}^* = \Sigma_k^* - \mathbf{u}_k^* \mathbf{u}_k^{*\top} \Sigma_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top}; \quad \Sigma_1^* = \Sigma. \quad (4)$$

We term the scenario in (4) as *ideal deflation* and the matrices $\{\Sigma_k^*\}_{k=1}^d$ as “ground-truth” deflation matrices. The

ideal deflation has the nice property that \mathbf{u}_k^* is the top eigenvector of Σ_k^* . Therefore, solving for the top- K eigenvectors can be exactly reduced to repeatedly solving for the top-1 eigenvector of the “ground-truth” deflation matrices.

Yet, solving for the top eigenvector exactly is challenging, which makes the ideal deflation almost impossible in practice. Thus, we are interested in a more general case, where PCA computes the top-eigenvector *inexactly*, and results in numerical error δ_k , as defined below:

$$\delta_k := \text{PCA}(\Sigma_k, t) - \mathbf{u}_k = \mathbf{v}_k - \mathbf{u}_k; \quad \|\delta_k\|_2 > 0. \quad (5)$$

Here, \mathbf{u}_k is the top eigenvector of Σ_k ; recall that Σ_k in (3) is constructed recursively using \mathbf{v}_k . When \mathbf{v}_k is solved inexactly, we cannot guarantee that $\Sigma_k = \Sigma_k^*$. Consequently, it is almost always the case that the top eigenvector of Σ_k does not equal to \mathbf{u}_k^* , namely $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2 > 0$. However, since PCA only knows Σ_k , its output \mathbf{v}_k converges to \mathbf{u}_k instead of \mathbf{u}_k^* when t is large. This difference between the top eigenvector of the matrices Σ_k and Σ_k^* further complicates the analysis of the deflation process in (3), since \mathbf{v}_k returned by $\text{PCA}(\Sigma_k, t)$ approximates \mathbf{u}_k instead of \mathbf{u}_k^* , which further builds up the difference between \mathbf{v}_k and \mathbf{u}_k^* and thus the subsequent $\|\Sigma_{k+1} - \Sigma_{k+1}^*\|_F$.

Our work aims to characterize this complicated error propagation through the deflation steps. To build up our analysis, we make the following assumption on the matrix Σ .

Assumption 2.1. The matrix $\Sigma \in \mathbb{R}^{d \times d}$ is a real symmetric matrix with eigenvalues and eigenvectors $\{\lambda_i^*, \mathbf{u}_i^*\}_{i=1}^d$, satisfying $1 = \lambda_1^* > \lambda_2^* > \dots > \lambda_d^* > 0$.

Assumption 2.1 states the real symmetric matrix Σ has strictly decaying eigenvalues. Observing such phenomenon is quite common in practice; see Papailiopoulos et al. (2014, Figure 1), Pappayan (2020, Figure 1), and Goujaud et al. (2022, Figure 1) for instance. Further, $\lambda_1^* = 1$ is not strictly required but assumed without loss of generality.

Assumption 2.1 guarantees that Σ is positive definite. By construction, all Σ_k^* ’s are guaranteed to be positive semi-definite. To enforce a similar condition on Σ_k ’s, we need to characterize the difference between the eigenvalues of Σ_k^* and Σ_k : for λ_k being the top eigenvalue of Σ_k , we need to control $|\lambda_k^* - \lambda_k|$. This is possible by Weyl’s inequality:

Lemma 2.2 (Weyl’s Inequality (Weyl, 1912)). *Let $\mathbf{M}, \mathbf{M}^* \in \mathbb{R}^{d \times d}$ be real symmetric matrices. Let σ_j, σ_j^* be the j -th eigenvalue of \mathbf{M} and \mathbf{M}^* , respectively. Then:*

$$|\sigma_j^* - \sigma_j| \leq \|\mathbf{M}^* - \mathbf{M}\|_2.$$

A direct consequence of Weyl’s inequality to our scenario is that $|\lambda_k^* - \lambda_k| \leq \|\Sigma_k - \Sigma_k^*\|_2$.

Another tool that we will utilize in the following sections of our paper is the Davis-Kahan sin Θ Theorem:

Lemma 2.3 (sin Θ Theorem (Davis & Kahan, 1970)). Let $\mathbf{M}^* \in \mathbb{R}^{d \times d}$ and let $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$. Let \mathbf{a}_1^* and \mathbf{a}_1 be the top eigenvectors of \mathbf{M}^* and \mathbf{M} , respectively. Then we have:

$$\sin \angle \{\mathbf{a}_1^*, \mathbf{a}_1\} \leq \frac{\|\mathbf{H}\|_2}{\min_{j \neq k} |\sigma_k^* - \sigma_j|}.$$

3. Sub-routine Agnostic Error Propagation

In this section, we aim to provide a characterization of the error propagation in the deflation methods in (3) that is agnostic to the detail of the sub-routine PCA. In particular, the only information known about PCA is the magnitude of the error of PCA, namely $\|\delta_k\|_2 = \|\mathbf{v}_k - \mathbf{u}_k\|_2 = \|\text{PCA}(\Sigma_k, t) - \mathbf{u}_k\|_2$.

Theorem 3.1. Consider the scenario of looking for the top- K eigenvectors of $\Sigma \in \mathbb{R}^{d \times d}$, and suppose Σ satisfies Assumption 2.1. Let $\mathcal{T}_j = \lambda_j^* - \lambda_{j+1}^*$ for $j \in [d-1]$, and $\mathcal{T}_d = \lambda_d^*$. Also, let $\mathcal{T}_{K,\min} = \min_{k \in [K]} \mathcal{T}_k$, and let δ_k and \mathbf{v}_k be defined in (5). If $\|\delta_k\|_2$'s are small enough such that

$$\sum_{k'=1}^{k-1} \lambda_{k'}^* \|\delta_{k'}\|_2 \prod_{j=k'+1}^{k-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right) \leq \frac{1}{20} \mathcal{T}_{K,\min}. \quad (6)$$

then the output of Algorithm 1 satisfies for all $k \in [K]$:

$$\begin{aligned} \|\mathbf{v}_k - \mathbf{u}_k^*\|_2 &\leq \|\mathbf{u}_k - \mathbf{u}_k^*\|_2 + \|\delta_k\|_2 \\ &\leq 5 \sum_{k'=1}^k \frac{\lambda_{k'}^*}{\lambda_k^*} \|\delta_{k'}\|_2 \prod_{j=k'+1}^k \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right). \end{aligned} \quad (7)$$

Remark 1. Theorem 3.1 characterizes how the errors from approximately solving the sub-routine (PCA) propagate into the overall error of the deflation procedure. The condition in (6) is to make sure that the difference between the ‘‘ground-truth’’ deflated matrices in (4) and the empirical ones in (3) is controlled and is sufficiently small, i.e., $\|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{4} \mathcal{T}_{k,\min}$ for all $k \in [K]$. As a simplification, this condition can be guaranteed as long as

$$\|\delta_k\|_2 \leq \frac{\mathcal{T}_{K,\min}}{20K} \prod_{j=k+1}^{K-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)^{-1}.$$

Remark 2. Notice that the error upper bound in (7) takes the form of a summation over components that depend on the error of the sub-routine $\{\|\delta_{k'}\|_2\}_{k'=1}^k$. In particular, not only do the number of summands grow as k grows, but each summand also has a multiplicative factor of $\prod_{j=k'+1}^k \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)$ which grows as k becomes larger. When the eigen spectrum of Σ^* decays slowly, this multiplicative factor can grow near factorially. For instance, in the case of a power-law decay spectrum where $\lambda_j = \frac{1}{j}$, the multiplicative factor becomes $\prod_{j=k'+1}^k (3 + 2j)$, since

$\mathcal{T}_j = \frac{1}{j(j+1)}$. In this case, the errors $\{\|\delta_{k'}\|_2\}_{k'=1}^k$ need to be nearly factorially small as we attempt to solve for more eigenvectors. To be more specific, based on Theorem 3.1, to guarantee that $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \varepsilon$, it is necessary to have:

$$\|\delta_{k'}\|_2 \leq \frac{\varepsilon \lambda_k^*}{5 \lambda_{k'}^*} \prod_{j=k'+1}^k \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)^{-1}.$$

The corollary below states a sufficient condition to guarantee that $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \varepsilon$ for all $k \in [K]$.

Corollary 3.2. Consider the scenario of solving the top- K eigenvectors of $\Sigma \in \mathbb{R}^{d \times d}$, and suppose Σ satisfies Assumption 2.1. Let $\mathcal{T}_j = \lambda_j^* - \lambda_{j+1}^*$ for $j \in [d-1]$, and $\mathcal{T}_d = \lambda_d^*$. Also let $\mathcal{T}_{K,\min} = \min_{k \in [K]} \mathcal{T}_k$. Finally, let δ_i and \mathbf{v}_i be defined as in (5). If for all $k \in [K]$ it holds that:

$$\|\delta_k\|_2 \leq \frac{\min\{\varepsilon \lambda_k^*, \mathcal{T}_{K,\min}\}}{20K} \prod_{j=k+1}^K \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)^{-1}, \quad (8)$$

then the output of Algorithm 1 satisfies $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \varepsilon$ for all $k \in [K]$.

Remark 3. Corollary 3.2 characterizes how accurately the sub-routine (PCA) has to be solved to achieve $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \varepsilon$, $\forall k \in [K]$, i.e., the desired accuracy of the overall deflation procedure.

Below, we provide lower bounds on the number of sub-routine iterations to satisfy (8), assuming a linearly-converging sub-routine with convergence rate α_k exists. In Section 4, we improve this rate by specifically setting the sub-routine to be the power iteration (Müntz et al., 1913).

Corollary 3.3. Under the same assumptions as Corollary 3.2, we further assume a sub-routine exists for computing the top eigenvalue with linear convergence rate $\alpha_k \in (0, 1)$. If $\|\delta_k\|_2 \leq c_0 \cdot \alpha_k^t$ for some constant $c_0 > 0$ for all $k \in [K]$, where t is the number of steps in the sub-routine satisfying:

$$t \geq \Omega \left(\frac{\log \frac{c_0 K}{\min\{\varepsilon \lambda_k^*, \mathcal{T}_{K,\min}\}} + \sum_{j=k+1}^K \log \left(\frac{\lambda_j^*}{\mathcal{T}_j} + 1\right)}{\log \alpha_k^{-1}} \right), \quad (9)$$

then the output of Algorithm 1 satisfies $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \varepsilon$ for all $k \in [K]$.

Remark 4. Corollary 3.3 is constructed by further assuming that $\|\delta_k\|_2 \leq c_0 \cdot \alpha_k^t$ based on Corollary 3.2 and derive the lower bound on t . The linear convergence property of solving the top eigenvector is shared among multiple existing algorithms (Xu et al., 2018a; Golub & Van Loan, 1996). In particular, when PCA is chosen to be the power iteration, we can show that in our setting $\|\delta_k\|_2 \leq c_0 \left(\frac{\lambda_2(\Sigma_k)}{\lambda_k}\right)^t$. In

the case of a power-law decay spectrum, when we treat c_0 as constants, the bound in (9) can be simplified to

$$t \geq \Omega \left(\log \frac{K}{\varepsilon} + \sum_{j=k+1}^K \log(j+1) \right).$$

Applying a slightly loose bound over the summation, we can see that based on the analysis of Theorem 3.1, one needs $t \geq \Omega \left(\log \frac{K}{\varepsilon} + K \log K \right)$ to achieve ε -accuracy. Section 4 compares this rate with the lower bound we obtain when PCA is fixed to power iteration.

3.1. Proof Overview

The proof involves bounding $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2$ and $\|\Sigma_k - \Sigma_k^*\|_F$. The former quantifies the difference between the true k -th eigenvector of Σ and the leading eigenvector of the k -th deflation step Σ_k . The latter quantifies the difference between the ‘‘ground-truth’’ deflation matrices defined in (4) and the ones in practice. Recall that \mathbf{u}_k^* is the top eigenvector of Σ_k^* (c.f., Lemma C.1). Let $\text{nev}_1(\Sigma_k)$ denote a normalized top eigenvector of Σ_k . Since both $\text{nev}_1(\Sigma_k)$ and $-\text{nev}_1(\Sigma_k)$ are unit-norm eigenvectors of Σ_k , we will choose \mathbf{u}_k to be the one such that $\mathbf{u}_k^\top \mathbf{u}_k^* \geq 0$:

$$\mathbf{u}_k := \arg \min_{s \in \{\pm 1\}} \|s \cdot \text{nev}_1(\Sigma_k) - \mathbf{u}_k^*\|_2 \cdot \text{nev}_1(\Sigma_k). \quad (10)$$

The difference between \mathbf{v}_k , the k -th approximate eigenvector returned by Algorithm 1, and \mathbf{u}_k^* , the k -th ground-truth eigenvector, consists of two components: *i*) $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2$, the difference between the top eigenvectors of the empirical deflated matrix Σ_k and the ‘‘ground-truth’’ deflated matrix Σ_k^* ; and *ii*) $\|\mathbf{v}_k - \mathbf{u}_k\|_2$, the numerical error from not solving the sub-routine $\text{PCA}(\cdot)$ exactly. We can decompose the difference $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2$ in the following sense:

$$\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \underbrace{\|\mathbf{v}_k - \mathbf{u}_k\|_2}_{:= \|\delta_k\|_2} + \|\mathbf{u}_k - \mathbf{u}_k^*\|_2. \quad (11)$$

where the last equality follows from the definition in (5). The norm of $\|\delta_k\|_2$ depends largely on the sub-routine PCA and can be controlled as long as t , the number of sub-routine iterations, is large enough. Therefore, our analysis will mainly focus on upper-bounding $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2$. Intuitively, $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2$ depends on the difference between Σ_k and Σ_k^* . As a building block of our proof, we will first provide a characterization of $\|\Sigma_k - \Sigma_k^*\|_F$.

Lemma 3.4. *Suppose that $\|\delta_k\|_2 \leq \frac{1}{6}$ for all $k \in [K-1]$, then we have that for all $k \in [K-1]$,*

$$\|\Sigma_{k+1} - \Sigma_{k+1}^*\|_F \leq 3 \|\Sigma_k - \Sigma_k^*\|_F + 5\lambda_k^* \|\delta_k\|_2 + 2\lambda_k^* \|\mathbf{u}_k - \mathbf{u}_k^*\|_2. \quad (12)$$

Lemma 3.4 upper-bounds $\|\Sigma_{k+1} - \Sigma_{k+1}^*\|_F$ in terms of $\|\Sigma_k - \Sigma_k^*\|_F$, $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2$, and $\|\delta_k\|_2$. To establish a recursive characterization of $\|\Sigma_k - \Sigma_k^*\|_F$, we need to obtain an upper bound for $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2$. Indeed, this can be easily obtained by applying the Davis-Kahan sin Θ theorem (Davis & Kahan, 1970; Eldridge et al., 2017).

Lemma 3.5. *Let $\mathcal{T}_k := \min_{j \neq k} |\lambda_k^* - \lambda_j^*|$. If $\|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{4} \mathcal{T}_k$, then we shall have that*

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2 \leq \frac{2}{\mathcal{T}_k} \|\Sigma_k - \Sigma_k^*\|_F. \quad (13)$$

Plugging (13) into (12) gives a recurrence that depends purely on $\|\delta_k\|_2$ and the spectrum of Σ :

$$\|\Sigma_{k+1} - \Sigma_{k+1}^*\|_F \leq \left(3 + \frac{2}{\mathcal{T}_k} \right) \|\Sigma_k - \Sigma_k^*\|_F + 5\lambda_k^* \|\delta_k\|_2.$$

Unrolling this recurrence gives a closed-form upper bound for $\|\Sigma_k - \Sigma_k^*\|_F$. Combining this upper bound with (13) and plugging the result into (11) gives us an upper bound for $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2$.

4. Error Propagation When Using Power Iteration

In this section, we turn our focus to a specific sub-routine algorithm (PCA in line 5 of Algorithm 1) for finding the top eigenvectors: the power iteration method (Müntz et al., 1913). Given a symmetric matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$ with rank r and an initialization vector $\mathbf{x}_0 \in \mathbb{R}^d$, the power iteration executes:

$$\hat{\mathbf{x}}_{t+1} = \mathbf{M}\mathbf{x}_t; \quad \mathbf{x}_{t+1} = \frac{\hat{\mathbf{x}}_{t+1}}{\|\hat{\mathbf{x}}_{t+1}\|_2}. \quad (14)$$

Assuming that PCA is the power iteration not only gives us the convergence rate of $\|\delta_k\|_2$ in terms of the number of steps in each power iteration, but also provides the directional information δ_k . In particular, one could see that $|\delta_k^\top \mathbf{u}_{k+j-1}| \leq c_0 \left(\frac{\lambda_j(\Sigma)_k}{\lambda_k^*} \right)^t$ for some constant c_0^2 . Namely, the component of δ_k along \mathbf{u}_j has an exponentially smaller magnitude if j is larger, depending on the eigen spectrum of Σ^* . By leveraging this additional information, our theorem below shows that we can improve the bound on $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2$, when PCA is specified to the power iteration, compared to the sub-routine agnostic result in Theorem 3.1.

Theorem 4.1. *Consider the problem of looking for the top- K eigenvectors of $\Sigma \in \mathbb{R}^{d \times d}$ under Assumption 2.1. Let $\mathcal{T}_i = \lambda_i^* - \lambda_{i+1}^*$ for $i \in [d-1]$ and $\mathcal{T}_d = \lambda_d^*$, and let $\mathcal{T}_{k,\min} = \min_{i \in [k]} \mathcal{T}_i$. Assume that there exists a constant $c_0 \in (0, \infty)$ such that the initialized vector $\mathbf{x}_{0,k}$ in the k -th*

²For the detail of proving this statement please see Lemma C.6

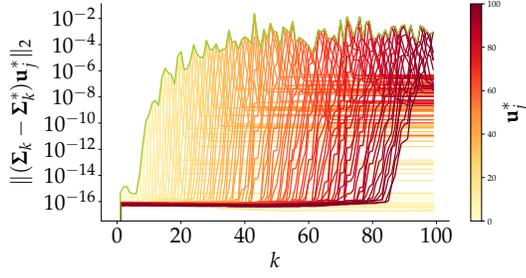


Figure 2. Dynamics of $\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2$ with respect to the change of k . Each \mathbf{u}_j^* is represented by a different color, with light color for small j and dark color for large j . Experiments done for $\Sigma \in \mathbb{R}^{100 \times 100}$, $\lambda_k^* = \frac{1}{k}$, $\{\mathbf{u}_k^*\}_{k=1}^d$ being randomly generated orthogonal basis, and $t = 200$. The orthogonal basis $\{\mathbf{u}_k^*\}_{k=1}^d$ is generated by randomly sampling a matrix with I.I.D. Gaussian entries, and computing its left singular vectors.

power iteration procedure satisfies $|\mathbf{x}_{0,k}^\top \mathbf{u}_k| > c_0^{-1}$ for all $k \in [K]$.³ Let $\mathcal{G} = \max_{k \in [K]} \left(1 + \frac{c_0 \lambda_k^* \lambda_{k+1}^*}{\lambda_k^* - \lambda_{k+1}^*}\right)$. If we use t steps of power iteration to solve for the top eigenvectors, where t satisfies:

$$t \geq \begin{cases} \log 2\mathcal{G}_k \left(\log \frac{\lambda_{k'+1}^* + 7\lambda_{k'}^*}{\lambda_{k'}^* + 7\lambda_{k'+1}^*}\right)^{-1}, & \forall k \leq K \\ \frac{1}{\log \lambda_k^* - \log \lambda_{k+1}^*}, & \forall k \leq d \end{cases} \quad (15)$$

Moreover, if t is also large enough to guarantee that

$$\sum_{k'=1}^{K-1} 8^{K-k'} \frac{\lambda_{k'}^*}{\mathcal{T}_{k'}} \left(\frac{7\lambda_{k+1}^* + \lambda_k^*}{7\lambda_k^* + \lambda_{k+1}^*}\right)^t \leq \frac{\mathcal{T}_{K,\min}}{140c_0}, \quad (16)$$

Then we can guarantee that for all $k \in [K]$

$$\begin{aligned} & \|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \\ & \leq 3 \sum_{k'=1}^k 8^{k-k'} \frac{\lambda_{k'}^*}{\lambda_k^*} \left(5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*}\right)^t\right). \end{aligned} \quad (17)$$

Remark 5. Similar to Theorem 3.1, the requirement of t in (16) is to guarantee that $\|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{8} \mathcal{T}_{K,\min}$. In particular, (16) combined with (15) can be satisfied by the following simplified condition on t

$$t \geq \begin{cases} \left(\frac{\max\left\{\log 2\mathcal{G}, 3(K-k) + \log \frac{140c_0K}{\mathcal{T}_{K,\min}^2}\right\}}{\log(7\lambda_{k+1}^* + \lambda_k^*) - \log(7\lambda_k^* + \lambda_{k+1}^*)}\right), & \forall k \leq K \\ \frac{1}{\log \lambda_k^* - \log \lambda_{k+1}^*}, & \forall k \leq d \end{cases} \quad (18)$$

³This is a standard assumption to guarantee the convergence of the power iteration, as in Xu et al. (2018a); Golub & Van Loan (1996)

The upper bound on $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2$ in (17) also takes a similar form to (7). Recall that when PCA is chosen to be the power iteration, we have $\|\delta_k\|_2 \leq c_0 \left(\frac{\lambda_2(\Sigma_k)}{\lambda_k}\right)^t \approx c_0 \left(\frac{\lambda_{k+1}^*}{\lambda_k^*}\right)^t$.

Therefore, one can treat $\left(5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*}\right)^t\right)$ in (17) as roughly $\frac{C}{\mathcal{T}_{k'}} \|\delta_{k'}\|_2$ for some constant C . Thus, the $\prod_{j=k+1}^K \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)$ scaling of each summand in (7) that depends on \mathcal{T}_j is improved to the exponential scaling of 8^{K-k} in (17). This improvement is further made clear in the corollary below.

Corollary 4.2. Under the same assumptions as in Theorem 4.1, if we use t steps of power iteration to solve for the top eigenvectors such that:

$$t \geq \begin{cases} \Omega \left(\frac{\max\left\{\log \mathcal{G}, K-k + \log \frac{c_0K}{\epsilon \mathcal{T}_{K,\min}}\right\}}{\log(7\lambda_k^* + \lambda_{k+1}^*) - \log(7\lambda_{k+1}^* + \lambda_k^*)}\right), & \forall k \leq K, \text{ and} \\ \Omega \left(\frac{1}{\log \lambda_k^* - \log \lambda_{k+1}^*}\right), & \forall k \leq d, \end{cases} \quad (19)$$

then we can guarantee $\|\mathbf{v}_k - \mathbf{u}_k\|_2 \leq \epsilon$ for all $k \in [K]$.

Remark 6. Corollary 4.2 is obtained by setting $\|\delta_k\|_2 \leq c_0 \left(\frac{\lambda_2(\Sigma_k)}{\lambda_k}\right)^t$ in power iteration and notice that $\frac{\lambda_2(\Sigma_k)}{\lambda_k} \leq \frac{7\lambda_{k+1}^* + \lambda_k^*}{7\lambda_k^* + \lambda_{k+1}^*}$ when $\|\Sigma_k - \Sigma_k^*\|_2 \leq \frac{1}{8} \mathcal{T}_k$. When we treat c_0 and $\{\lambda_k^*\}_{k=1}^d$ as constants, \mathcal{G} also becomes a constant. In this case, the lower bound on t in (19) becomes $t \geq \Omega \left(\log \frac{K}{\epsilon} + K\right)$. Compared with the $t \geq \Omega \left(\log \frac{K}{\epsilon} + K \log K\right)$ requirement we derived in the discussion of Corollary 3.3. Thus, our lower bound in Corollary 4.2 saves a factor of $\log K$.

Remark 7. In the discussion about Theorem 4.1 and Corollary 4.2 above, we deliberately ignored the effect of c_0 and \mathcal{G} . For \mathcal{G} , we observe that \mathcal{G} can also be written as $\mathcal{G} = \max_{k \in [K]} \left(1 + c_0 \left(\frac{1}{\lambda_{k+1}^*} - \frac{1}{\lambda_k^*}\right)^{-1}\right)$. If the eigenvalues of Σ follow a power law decay as $\lambda_k^* = \frac{1}{k^\gamma}$ for some $\gamma \geq 1$, then we have $\frac{1}{\lambda_{k+1}^*} - \frac{1}{\lambda_k^*} \geq 1$ and thus $\mathcal{G} \leq c_0 + 1$. Next, we focus on c_0 . If in each power iteration sub-routine, the initialized vector $\mathbf{x}_{0,k}$ is generated randomly, then $|\mathbf{x}_{0,k}^\top \mathbf{u}_k|$ is independent of \mathbf{u}_k , and $\mathbb{P}(|\mathbf{x}_{0,k}^\top \mathbf{u}_k| > 0) = 1$. This idea can be made concrete in the following lemma

Lemma 4.3. Let $\mathbf{x}_{0,k}$ be sampled uniformly from a unit sphere in \mathbb{R}^d . Then with probability at least $1 - \frac{2K}{3d}$ we have that $|\mathbf{x}_{0,k}^\top \mathbf{u}_k| \geq \frac{1}{\sqrt{1+2d^3}}$ for all $k \in [K]$.

The proof of Lemma 4.3 follows the idea of (Wang et al., 2020), and is deferred to Appendix B.1. Next, we will first explore the property of power iteration crucial to our proof

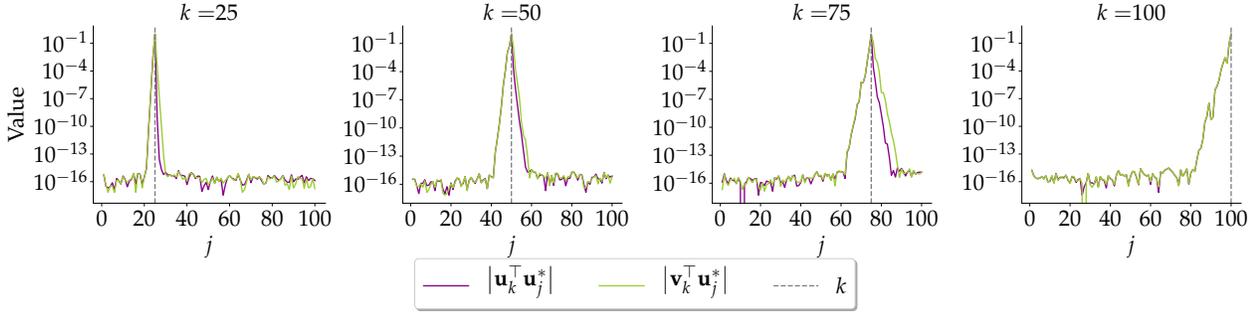


Figure 3. The comparison among the dynamics of $\mathbf{u}_k^\top \mathbf{u}_j^*$ and $\mathbf{v}_k^\top \mathbf{u}_j^*$ with respect to the change of k for $j \in \{25, 50, 75, 100\}$. Experiments are performed for $\Sigma \in \mathbb{R}^{100 \times 100}$, $\lambda_k^* = \frac{1}{k}$, $\{\mathbf{u}_k^*\}_{k=1}^d$ being randomly generated orthogonal basis, with $t = 200$. The results show that both $|\mathbf{u}_k^\top \mathbf{u}_j^*|$ and $|\mathbf{v}_k^\top \mathbf{u}_j^*|$ are small only when k is near j . The orthogonal basis $\{\mathbf{u}_k^*\}_{k=1}^d$ is generated by randomly sampling a matrix with I.I.D. Gaussian entries, and computing its left singular vectors.

of Theorem 4.1 in Section 4.1, and then provide an overview of the proof in Section 4.2.

4.1. What Does Power Iteration Tell Us?

Recall our previous observation that the error vector returned by power iteration has different magnitudes along different eigenvector directions of the matrix of interest. In this section, we perform the same analysis, assuming that the matrix passed into the power iteration sub-routine is a perturbation of some “ground-truth” matrix. Let \mathbf{a}_j^* be the j -th eigenvector of \mathbf{M}^* and $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$. Given the definition of \mathbf{x}_t as the result of a t -step power iteration on \mathbf{M} , defined in (14), we will show that $\mathbf{x}_t^\top \mathbf{a}_j^*$ ’s take different forms depending on the index of the eigenvector j .

Lemma 4.4. *Let \mathbf{x}_t be the result of running power iteration starting from \mathbf{x}_0 for t iterations, as defined in (14). Let $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$. Moreover, let σ_j^* , \mathbf{a}_j^* be the j -th eigenvalue and eigenvector of \mathbf{M}^* , and let σ_j be the j -th eigenvalue of \mathbf{M} . Assume that there exists $c_0 \in (0, \infty)$ such that $|\mathbf{x}_0^\top \mathbf{a}_1| \geq c_0^{-1}$. Then, for all $j = 2, \dots, r$, we have:*

$$|\mathbf{x}_t^\top \mathbf{a}_j^*| \leq c_0 \left(\left(\frac{\sigma_j^*}{\sigma_1} \right)^t + \frac{\sigma_j^*}{\sigma_1 - \sigma_j^*} \|\mathbf{H} \mathbf{a}_j^*\|_2 \right). \quad (20)$$

Lemma 4.4 upper bounds the alignment of the vector in the t th step of the power iteration with any incorrect eigenvector \mathbf{a}_j^* . The magnitude of $\mathbf{x}_t^\top \mathbf{a}_j^*$ depends both on $\left(\frac{\sigma_j^*}{\sigma_1} \right)^t$ and how the perturbation matrix aligns with the j -th eigenvector. When \mathbf{M}^* has a decaying spectrum, the first term in the upper bound of $|\mathbf{x}_t^\top \mathbf{a}_j^*|$ in (20) decreases as j becomes larger. Overloading Lemma 4.4 in the notations of our deflation procedure, we have

$$|\mathbf{v}_k^\top \mathbf{u}_j^*| \leq c_0 \left(\left(\frac{\lambda_j^*}{\lambda_1} \right)^t + \frac{\lambda_j^*}{\lambda_1 - \lambda_j^*} \|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2 \right)$$

In short, the component of \mathbf{v}_k along an incorrect director, $|\mathbf{v}_k^\top \mathbf{u}_j^*|$, is determined by both the convergence along the j th eigenvector, and by the difference between the two deflated matrix along the direction of \mathbf{u}_j^* . This calls for analysis on $\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2$. Indeed, since the matrix residue $\Sigma_k - \Sigma_k^*$ is resulted from $\mathbf{v}_k - \mathbf{u}_k^*$ ’s, it is natural to believe that $\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2$ can also be upper-bounded by quantities depending on j . Our next lemma characterizes this behavior

Lemma 4.5. *Let \mathbf{v}_k be the output of the k -th power iteration. For all $k \in [d]$ and $j \geq k$, we have:*

$$\|(\Sigma_k^* - \Sigma_k) \mathbf{u}_j^*\|_2 \leq \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*|. \quad (21)$$

The proof of Lemma 4.5 is deferred to Appendix B.3. In particular, (21) in Lemma 4.5 upper-bounds $\|(\Sigma_k^* - \Sigma_k) \mathbf{u}_j^*\|_2$ using a linear combination of $\{|\mathbf{v}_{k'}^\top \mathbf{u}_j^*|\}_{k'=1}^{k-1}$. Indeed, recall that \mathbf{v}_k is the output of the k -th power iteration procedure. This means that we can invoke Lemma 4.4 with $\mathbf{M} = \Sigma_k$, $\mathbf{M}^* = \Sigma_k^*$, and correspondingly $\mathbf{H} = \Sigma_k^* - \Sigma_k$. Plugging the bound of $\|(\Sigma_k^* - \Sigma_k) \mathbf{u}_j^*\|_2$ from (21) into then gives the recursively defined $|\mathbf{v}_k^\top \mathbf{u}_j^*|$:

$$|\mathbf{v}_k^\top \mathbf{u}_j^*| \leq c_0 \left(\left(\frac{\lambda_j^*}{\lambda_k} \right)^t + \frac{\lambda_j^*}{\lambda_k - \lambda_j^*} \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \right). \quad (22)$$

(22) shows the independence of each $\{|\mathbf{v}_k^\top \mathbf{u}_j^*|\}_{k=1}^d$ for different j , which allows each $\{|\mathbf{v}_k^\top \mathbf{u}_j^*|\}_{k=1}^d$ to grow at different speeds. To see whether this characterization of independent growth is crucial, observe that in Figure 3, each $|\mathbf{v}_k^\top \mathbf{u}_j^*|$ only goes through a fast growth when j is near k . This behavior results in the phenomena that at each step k , there are only a few j ’s where $\|(\Sigma_k^* - \Sigma_k) \mathbf{u}_j^*\|_2$ is large.

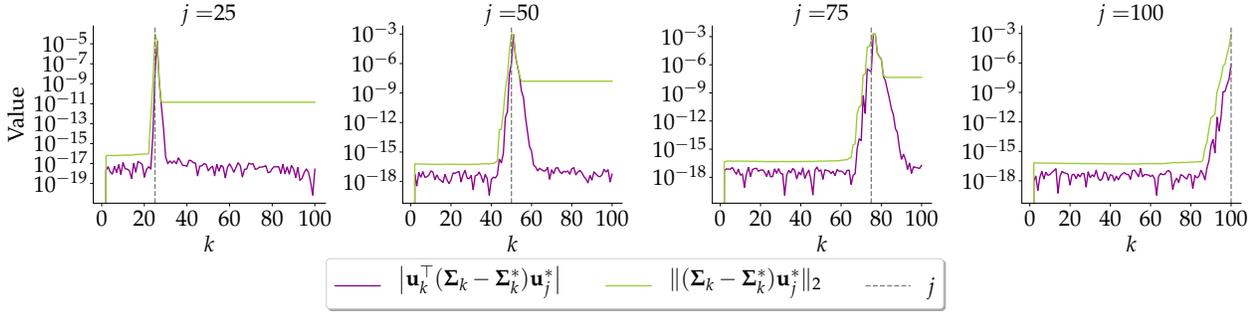


Figure 4. The comparison of between the dynamics of $|\mathbf{u}_k^\top (\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*|$ and $\|(\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*\|_2$ with respect to the change of k for $j \in \{25, 50, 75, 100\}$. Experiments are performed for $\boldsymbol{\Sigma} \in \mathbb{R}^{100 \times 100}$, $\lambda_k^* = \frac{1}{k}$, $\{\mathbf{u}_k^*\}_{k=1}^d$ as randomly generated orthogonal basis, with $t = 200$. The results show that $\|(\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*\|_2$ is a good approximation of $|\mathbf{u}_k^\top (\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*|$ when the latter is large. The orthogonal basis $\{\mathbf{u}_k^*\}_{k=1}^d$ is generated by randomly sampling a matrix with I.I.D. Gaussian entries, and computing its left singular vectors.

We can see this as we unroll this recursive form and obtain

$$\begin{aligned} \|(\boldsymbol{\Sigma}_k^* - \boldsymbol{\Sigma}_k) \mathbf{u}_j^*\|_2 &\leq \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \\ &\leq c_0 \sum_{k'=1}^{k-1} \mathcal{G}_k^{k-k'-1} \left(\frac{\lambda_j^*}{\lambda_k}\right)^{t-1} \end{aligned} \quad (23)$$

where $\mathcal{G}_k = 1 + \frac{c_0 \lambda_k^* \lambda_{k+1}^*}{\lambda_k^* - \lambda_{k+1}^*}$. (23) upper bounds $\|(\boldsymbol{\Sigma}_k^* - \boldsymbol{\Sigma}_k) \mathbf{u}_j^*\|_2$ with a summation of the linear decay terms resulted from the power iteration. For the detailed computation of unrolling the recursive definition, we refer the reader to Appendix B.7 and Appendix C (see Lemma C.3). Noticeably, (23) provides a closed-form upper bound of $\|(\boldsymbol{\Sigma}_k^* - \boldsymbol{\Sigma}_k) \mathbf{u}_j^*\|_2$, which further demonstrates the independent growth of $\|(\boldsymbol{\Sigma}_k^* - \boldsymbol{\Sigma}_k) \mathbf{u}_j^*\|_2$ as k increases for different j . This behavior corresponds to what we observed in experiments in Figure 2, where each $\|(\boldsymbol{\Sigma}_k^* - \boldsymbol{\Sigma}_k) \mathbf{u}_j^*\|_2$ grows sequentially.

4.2. Proof Overview of Theorem 4.1

In this section, we will demonstrate how the properties of power iteration, in particular (23) is used to construct a proof for Theorem 4.1. We start with the following lemma.

Lemma 4.6. *Let $\mathbf{M}^* \in \mathbb{R}^{d \times d}$ be a rank- r matrix and let $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$ for some $\mathbf{H} \in \mathbb{R}^{d \times d}$. Let σ_i^* , \mathbf{a}_i^* be the i -th eigenvalue and eigenvector of \mathbf{M}^* , and σ_i , \mathbf{a}_i be the i -th eigenvalue and eigenvector of \mathbf{M} . Then, for all $i, j \in [r]$ with $\sigma_i \neq \sigma_j^*$, we have:*

$$\mathbf{a}_i^\top \mathbf{a}_j^* = \frac{\mathbf{a}_i^\top \mathbf{H} \mathbf{a}_j^*}{\sigma_i - \sigma_j^*}.$$

Unfortunately, we cannot directly characterize $\mathbf{u}_k^\top \mathbf{u}_k^*$ using Lemma 4.6 by plugging in $i = j = 1$ and $\mathbf{M} = \boldsymbol{\Sigma}_k$, $\mathbf{M}^* = \boldsymbol{\Sigma}_k^*$, since, as far as we know, it is impossible to show that

$\sigma_i = \lambda_k \neq \lambda_k^* = \sigma_j^*$. However, Lemma 4.6 will be useful to upper-bound $(\mathbf{u}_k^\top \mathbf{u}_j^*)^2$ when $j > k$, since in this case, we can guarantee that $\lambda_k \geq \lambda_k^* - \|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_2 > \lambda_j^*$ as long as $\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_2 < \mathcal{T}_k$, where \mathcal{T}_k is the k -th eigengap of $\boldsymbol{\Sigma}$. Under such conditions, we have:

$$|\mathbf{u}_k^\top \mathbf{u}_j^*| \leq \frac{\|(\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*\|_2}{\lambda_k - \lambda_j^*}. \quad (24)$$

Indeed, going from Lemma 4.6 to (24) loosens the bound by ignoring the directional information of \mathbf{u}_k . However, through experiments, we observe that (24) is a good enough upper bound since $\|(\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*\|_2$ already stays at a much smaller magnitude than $\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_2$. This is reflected in Figure 4, where we can see that the curve $\|(\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*\|_2$ enjoys the same decrease as $|\mathbf{u}_k^\top (\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*) \mathbf{u}_j^*|$ when j first surpasses k , and stops decreasing after reaching a small magnitude (lower than 10^{-5}). (24) successfully applied the property of the power iteration in (23) to bound each $|\mathbf{u}_k^\top \mathbf{u}_j^*|$. Our next lemma connects the upper bound of each $|\mathbf{u}_k^\top \mathbf{u}_j^*|$ to a lower bound of $|\mathbf{u}_k^\top \mathbf{u}_k^*|$.

Lemma 4.7. *Suppose that $\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_F \leq \frac{1}{8} \lambda_k^*$, then we have*

$$(\mathbf{u}_k^\top \mathbf{u}_k^*)^2 \geq 1 - \frac{2.4}{\lambda_k^{*2}} \|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_F^2 - \sum_{j=k+1}^d (\mathbf{u}_k^\top \mathbf{u}_j^*)^2 \quad (25)$$

Our proof of Lemma 4.7 is based on Neumann expansion (Eldridge et al., 2017; Chen et al., 2020), a powerful technique for analyses that involve the product between perturbation matrix and eigenvectors. Indeed, the summation of the last term of (25) may seem to introduce an additional factor of d . However, since the sequence $\{|\mathbf{u}_k^\top \mathbf{u}_j^*|\}_{j=k+1}^d$ decays fast, as observed in Figure 3, the summation will eventually be independent of d . Combining (23), (24), and (25) gives the following lemma that upper bounds $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2$.

Lemma 4.8. Assume that there exists a constant $c_0 \in (0, \infty)$ such that the initialized vector $\mathbf{x}_{0,k'}$ in the k' -th power iteration procedure satisfies $|\mathbf{x}_{0,k'}^\top \mathbf{u}_{k'}| > c_0^{-1}$ for all $k' \in [k]$. Let $\mathcal{T}_k = \lambda_k^* - \lambda_{k+1}^*$, $\mathcal{T}_{k,\min} = \min_{k' \in [k]} \mathcal{T}_{k'}$ and $\mathcal{G}_k = 1 + \frac{c_0 \lambda_k^* \lambda_{k+1}^*}{\lambda_k^* - \lambda_{k+1}^*}$. Suppose $\|\Sigma_{k'} - \Sigma_{k'}^*\|_F \leq \frac{1}{8} \mathcal{T}_{k,\min}$, and $t \geq \log 2\mathcal{G}_k \left(\log \frac{\lambda_{k'+1}^* + 7\lambda_{k'}^*}{\lambda_{k'}^* + 7\lambda_{k'+1}^*} \right)^{-1}$ for all $k' \in [k-1]$. Then, we have that:

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2^2 \leq \frac{4.8}{\lambda_k^*} \|\Sigma_k - \Sigma_k^*\|_2^2 + \frac{11c_0^2}{2\mathcal{T}_k^2} \sum_{j=k+1}^d \left(\frac{\lambda_j^*}{\lambda_k^*} \right)^{2t}.$$

Lemma 4.8 shows a similar result to Lemma 3.5. In particular, the first term in the upper bound of Lemma 4.8 also scales with $\|\Sigma_k - \Sigma_k^*\|_2$, and the second term involves a summation over $\left(\frac{\lambda_j^*}{\lambda_k^*} \right)^{2t}$ for $j > k$, which can be made small as long as t is large enough. The upper bound in Lemma 4.8 shows a significant improvement in that the scaling factor in front of $\|\Sigma_k - \Sigma_k^*\|_2$ depends only on λ_k^* compared with the $O\left(\frac{\sqrt{d-k}}{\mathcal{T}_k}\right)$ scaling in Lemma 3.5. This allows us to improve over the nearly factorial bound in Theorem 3.1. After obtaining Lemma 4.8, the rest of the proof for Theorem 4.1 becomes similar to what we have discussed in Section 3.1. We can plug the upper bound in Lemma 4.8 into Lemma 3.4 and unroll the recurrence to obtain a closed-form upper bound for $\|\Sigma_k - \Sigma_k^*\|_F$. After that, we can finish the proof by plugging the upper bound for $\|\Sigma_k - \Sigma_k^*\|_F$ back into Lemma 4.8 and utilize (11) to arrive at the desired upper bound for $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2$. The detailed proof of Theorem 4.1 is deferred to Appendix B.7.

5. Conclusion and Future Work

Numerical algorithms for solving the top eigenvector of a given matrix often produce non-negligible errors, which will propagate and accumulate through multiple deflation steps. In this paper, we mathematically analyze this error propagation for a scenario agnostic to the sub-routine for solving the top eigenvectors and the case where this sub-routine is fixed to power iteration. In particular, for the sub-routine agnostic case, our analysis gives an exponential growth of the errors as one aims to solve more eigenvectors, where the base of the exponential growth depends on the eigengap. This upper bound on the errors is improved to exponential growth with a constant base by utilizing the directional information of the error vector produced by power iteration. Our result implies a lower bound on the number of power iteration steps required to solve for multiple leading eigenvectors using deflation up to a certain accuracy.

This paper focuses on the case where Σ is positive definite and has distinct eigenvalues. However, the analysis can also be extended to accommodate repeated and zero eigenvalues. When eigenvalues are repeated, the corresponding

eigenvector and the method for measuring the error of the approximate eigenvector need to be redefined. Instead of considering a single eigenvector for each eigenvalue, we can generalize the approach by considering the subspace spanned by the eigenvectors corresponding to each repeated eigenvalue. The core ideas of our analysis can be applied to account for the error propagation between subspaces, but one must account for additional error propagation within subspaces. To extend the analysis to positive semi-definite matrices, Σ is of rank r and we are interested in finding top- K eigenvectors with $K \leq r$. In this case, our Theorem 3.1 still holds, as it does not depend on eigenvalues beyond the K th. Theorem 4.1 may require slight modification, but we hypothesize that the case where Σ is low rank will relax the condition on t in (16) to $k \leq r$ instead of $k \leq d$.

While this paper focuses on the sub-routine agnostic scenario and the power iteration, future work can consider extending the analysis to the case where other algorithms, such as Lanczos method (Golub & Van Loan, 1996) and power iteration with momentum (Xu et al., 2018a), are used to approximate the top eigenvectors in the deflation process.

Acknowledgements

This work is supported by NSF CMMI no. 2037545, NSF CAREER award no. 2145629, Welch Foundation Grant #A22-0307, a AWS Research Grant, a Microsoft Research Award, and a Rice Ken Kennedy Institute Seed Funding Award.

Impact Statement

This paper aims to deepen the theoretical understanding of principal component analysis, a long-standing method in machine learning. The theoretical results presented herein may change the computation power used for principal component analysis. While the direct impact of this work may be limited to computational efficiency gains, the broader implications extend beyond this, potentially catalyzing advancements in various domains where principal component analysis finds application. However, no specific societal consequences can be directly attributed to the work itself.

References

- Allen-Zhu, Z. and Li, Y. First efficient convergence for streaming k-pca: a global, gap-free, and near-optimal rate, 2017.
- Belkin, M. and Niyogi, P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003. doi: 10.1162/089976603321780317.

- Chen, Y., Cheng, C., and Fan, J. Asymmetry helps: Eigenvalue and eigenvector analyses of asymmetrically perturbed low-rank matrices, 2020.
- Coulaud, O., Giraud, L., Ramet, P., and Vasseur, X. Deflation and augmentation techniques in krylov subspace methods for the solution of linear systems, 2013.
- Danisman, Y., Yilmaz, M., Ozkaya, A., and Comlekçiler, I. A comparison of eigenvalue methods for principal component analysis. *Appl. and Comput. Math*, 13:316–331, 2014.
- d’Aspremont, A., El Ghaoui, L., Jordan, M., and Lanckriet, G. A direct formulation for sparse pca using semidefinite programming. *SIAM review*, 49(3):434–448, 2007.
- Davis, C. and Kahan, W. M. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970. ISSN 00361429. URL <http://www.jstor.org/stable/2949580>.
- Deng, L. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- Eldridge, J., Belkin, M., and Wang, Y. Unperturbed: spectral analysis beyond davis-kahan, 2017.
- Fawzi, A., Balog, M., Huang, A., and et al. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature*, 610:47–53, 2022. doi: 10.1038/s41586-022-05172-4. URL <https://doi.org/10.1038/s41586-022-05172-4>.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022.
- Gemp, I., McWilliams, B., Vernade, C., and Graepel, T. Eigengame: {PCA} as a nash equilibrium. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=NzTU59SYbNq>.
- Golub, G. H. and Van Loan, C. F. *Matrix Computations*. The Johns Hopkins University Press, third edition, 1996.
- Golub, G. H. and Van Loan, C. F. *Matrix computations*. JHU press, 2013.
- Goujoud, B., Scieur, D., Dieuleveut, A., Taylor, A. B., and Pedregosa, F. Super-acceleration with cyclical step-sizes. In *International Conference on Artificial Intelligence and Statistics*, pp. 3028–3065. PMLR, 2022.
- Hotelling, H. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- Hubert, L. J. and Arabie, P. Comparing partitions. *Journal of Classification*, 2:193–218, 1985. URL <https://api.semanticscholar.org/CorpusID:189915041>.
- Jain, P., Jin, C., Kakade, S. M., Netrapalli, P., and Sidford, A. Streaming pca: Matching matrix bernstein and near-optimal finite sample guarantees for oja’s algorithm, 2016.
- Jiang, R., Fei, H., and Huan, J. Anomaly localization for network data streams with graph joint sparse pca. In *Proceedings of the 17th ACM SIGKDD*, pp. 886–894. ACM, 2011.
- Johnstone, I. M. and Lu, A. Y. Sparse principal components analysis. *arXiv preprint arXiv:0901.4392*, 2009.
- Kuczynski, J. and Woźniakowski, H. Estimating the largest eigenvalue by the power and lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, 13(4):1094–1122, 1992.
- Lanczos, C. An iterative method for the solution of the eigenvalue problem of linear differential and integral. 1950. URL <https://api.semanticscholar.org/CorpusID:478182>.
- Mackey, L. Deflation methods for sparse pca. *Advances in neural information processing systems*, 21, 2008.
- Majumdar, A. Image compression by sparse pca coding in curvelet domain. *Signal, image and video processing*, 3(1):27–34, 2009.
- Müntz, H. et al. Direct solution of the ésecular equation and of some transcendent analogous problems. *CR Acad. Sci. Paris*, 156:43–46, 1913.
- Papailiopoulos, D., Kyriillidis, A., and Boutsidis, C. Provable deterministic leverage score sampling. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 997–1006, 2014.
- Papayan, V. Traces of class/cross-class structure pervade deep learning spectra. *The Journal of Machine Learning Research*, 21(1):10197–10260, 2020.
- Pearson, K. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11): 559–572, 1901.

- Pothen, A., Simon, H. D., and Liou, K.-P. Partitioning sparse matrices with eigenvectors of graphs. *SIAM journal on matrix analysis and applications*, 11(3):430–452, 1990.
- Saad, Y. Projection and deflation method for partial pole assignment in linear state feedback. *IEEE Transactions on Automatic Control*, 33(3):290–297, 1988.
- Sriperumbudur, B. K., Torres, D. A., and Lanckriet, G. R. Sparse eigen methods by DC programming. In *Proceedings of the 24th international conference on Machine learning*, pp. 831–838, 2007.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Wang, P., Zhou, Z., and So, A. M.-C. A nearly-linear time algorithm for exact community recovery in stochastic block model. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 10126–10135. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/wang20ac.html>.
- Wang, Z., Han, F., and Liu, H. Sparse principal component analysis for high dimensional multivariate time series. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, pp. 48–56, 2013.
- Weyl, H. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, 1912.
- Xu, P., He, B., De Sa, C., Mitliagkas, I., and Re, C. Accelerated stochastic power iteration. In *International Conference on Artificial Intelligence and Statistics*, pp. 58–67. PMLR, 2018a.
- Xu, P., He, B., De Sa, C., Mitliagkas, I., and Re, C. Accelerated stochastic power iteration. In Storkey, A. and Perez-Cruz, F. (eds.), *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 58–67. PMLR, 09–11 Apr 2018b. URL <https://proceedings.mlr.press/v84/xu18a.html>.
- Zhang, F. *The Schur complement and its applications*, volume 4. Springer Science & Business Media, 2006.
- Zou, H., Hastie, T., and Tibshirani, R. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.

A. Missing Proofs from Section 3

A.1. Proof of Lemma 3.4

Recall that $\Sigma_{k+1} = \Sigma_k - \mathbf{v}_k \mathbf{v}_k^\top \Sigma_k \mathbf{v}_k \mathbf{v}_k^\top$. Since $\mathbf{v}_k = \mathbf{u}_k + \boldsymbol{\delta}_k$, we can write Σ_{k+1} as

$$\begin{aligned} \Sigma_{k+1} &= \Sigma_k - (\mathbf{u}_k + \boldsymbol{\delta}_k) (\mathbf{u}_k + \boldsymbol{\delta}_k)^\top \Sigma_k (\mathbf{u}_k + \boldsymbol{\delta}_k) (\mathbf{u}_k + \boldsymbol{\delta}_k)^\top \\ &= \Sigma_k - (\mathbf{u}_k + \boldsymbol{\delta}_k) (\mathbf{u}_k^\top \Sigma_k \mathbf{u}_k + \boldsymbol{\delta}_k^\top \Sigma_k \mathbf{u}_k + \mathbf{u}_k^\top \Sigma_k \boldsymbol{\delta}_k + \boldsymbol{\delta}_k^\top \Sigma_k \boldsymbol{\delta}_k) (\mathbf{u}_k + \boldsymbol{\delta}_k)^\top \\ &= \Sigma_k - (\lambda_k + 2\lambda_k \mathbf{u}_k^\top \boldsymbol{\delta}_k + \boldsymbol{\delta}_k^\top \Sigma_k \boldsymbol{\delta}_k) (\mathbf{u}_k + \boldsymbol{\delta}_k) (\mathbf{u}_k + \boldsymbol{\delta}_k)^\top \\ &= \underbrace{\Sigma_k - \lambda_k \mathbf{u}_k \mathbf{u}_k^\top}_{\tilde{\Sigma}_{k+1}} - \underbrace{\lambda_k (\boldsymbol{\delta}_k \mathbf{u}_k^\top + \mathbf{u}_k \boldsymbol{\delta}_k^\top + \boldsymbol{\delta}_k \boldsymbol{\delta}_k^\top)}_{\mathcal{E}_{1,k}} \\ &\quad - \underbrace{(2\lambda_k \mathbf{u}_k^\top \boldsymbol{\delta}_k + \boldsymbol{\delta}_k^\top \Sigma_k \boldsymbol{\delta}_k) (\mathbf{u}_k + \boldsymbol{\delta}_k) (\mathbf{u}_k + \boldsymbol{\delta}_k)^\top}_{\mathcal{E}_{2,k}} \end{aligned}$$

For the convenience of the analysis, we let $\Delta_k = \|\Sigma_k - \Sigma_k^*\|_F$. In this way, we can write

$$\Delta_{k+1} = \left\| \tilde{\Sigma}_{k+1} - \Sigma_{k+1}^* - \mathcal{E}_{1,k} - \mathcal{E}_{2,k} \right\|_F \leq \left\| \tilde{\Sigma}_{k+1} - \Sigma_{k+1}^* \right\|_F + \|\mathcal{E}_{1,k}\|_F + \|\mathcal{E}_{2,k}\|_F \quad (26)$$

Recall the definition of Σ_{k+1}^* from (4). By Lemma C.1, we have $\Sigma_{k+1}^* = \Sigma_k^* - \lambda_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top}$. Thus, the first term in (26) can be bounded as

$$\begin{aligned} \left\| \tilde{\Sigma}_{k+1} - \Sigma_{k+1}^* \right\|_F &= \left\| \Sigma_k - \lambda_k \mathbf{u}_k \mathbf{u}_k^\top - \Sigma_k^* + \lambda_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top} \right\|_F \\ &\leq \left\| \Sigma_k - \Sigma_k^* \right\|_F + \left\| \lambda_k \mathbf{u}_k \mathbf{u}_k^\top - \lambda_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top} \right\|_F \\ &\leq \Delta_k + \left\| (\lambda_k - \lambda_k^*) \mathbf{u}_k \mathbf{u}_k^\top - \lambda_k^* (\mathbf{u}_k^* \mathbf{u}_k^{*\top} - \mathbf{u}_k \mathbf{u}_k^\top) \right\|_F \\ &\leq 2\Delta_k + \lambda_k^* \left\| \mathbf{u}_k \mathbf{u}_k^\top - \mathbf{u}_k^* \mathbf{u}_k^{*\top} \right\|_F \end{aligned}$$

We use Weyl's inequality in the last inequality to obtain that $|\lambda_k - \lambda_k^*| \leq \Delta_k$. The last term above can be further bounded as

$$\left\| \mathbf{u}_k \mathbf{u}_k^\top - \mathbf{u}_k^* \mathbf{u}_k^{*\top} \right\|_F = \left\| \mathbf{u}_k (\mathbf{u}_k - \mathbf{u}_k^*)^\top + (\mathbf{u}_k - \mathbf{u}_k^*) \mathbf{u}_k^{*\top} \right\|_F \leq 2 \|\mathbf{u}_k - \mathbf{u}_k^*\|_2 = 2\theta_k$$

Therefore, we have

$$\left\| \tilde{\Sigma}_{k+1} - \Sigma_{k+1}^* \right\|_F \leq 2\Delta_k + 2\lambda_k^* \theta_k \quad (27)$$

Now, we focus on $\mathcal{E}_{1,k}$. It can be bounded as

$$\|\mathcal{E}_{1,k}\|_F = \lambda_k \left\| \boldsymbol{\delta}_k \mathbf{u}_k^\top + \mathbf{u}_k \boldsymbol{\delta}_k^\top + \boldsymbol{\delta}_k \boldsymbol{\delta}_k^\top \right\|_F \leq \lambda_k \left(2\|\boldsymbol{\delta}_k\|_2 + \|\boldsymbol{\delta}_k\|_2^2 \right)$$

Given that $\lambda_k \leq \lambda_k^* + |\lambda_k - \lambda_k^*| \leq \lambda_k + \Delta_k$, we have

$$\|\mathcal{E}_{1,k}\|_F \leq (\lambda_k^* + \Delta_k) \left(2\|\boldsymbol{\delta}_k\|_2 + \|\boldsymbol{\delta}_k\|_2^2 \right) \quad (28)$$

Similarly, for $\mathcal{E}_{2,k}$ we have

$$\begin{aligned} \|\mathcal{E}_{2,k}\|_F &\leq \left| 2\lambda_k \mathbf{u}_k^\top \boldsymbol{\delta}_k + \boldsymbol{\delta}_k^\top \Sigma_k \boldsymbol{\delta}_k \right| \|\mathbf{u}_k + \boldsymbol{\delta}_k\|_2^2 \\ &\leq \left(2(\lambda_k^* + \Delta_k) \|\boldsymbol{\delta}_k\|_2 + (\lambda_k^* + \Delta_k) \|\boldsymbol{\delta}_k\|_2^2 \right) (1 + \|\boldsymbol{\delta}_k\|_2)^2 \\ &= (\lambda_k^* + \Delta_k) \left(2\|\boldsymbol{\delta}_k\|_2 + \|\boldsymbol{\delta}_k\|_2^2 \right) (1 + \|\boldsymbol{\delta}_k\|_2)^2 \end{aligned} \quad (29)$$

where the first inequality follows from $\|\Sigma_k\|_2 \leq \lambda_k$ since, by definition, λ_k is the largest eigenvalue of Σ_k . Plugging (27), (28), and (29) into (26), we have

$$\begin{aligned} \Delta_{k+1} &\leq 2\Delta_k + 2\lambda_k^* \theta_k + (\lambda_k^* + \Delta_k) \left(2\|\delta_k\|_2 + \|\delta_k\|_2^2 \right) \\ &\quad + (\lambda_k^* + \Delta_k) \left(2\|\delta_k\|_2 + \|\delta_k\|_2^2 \right) (1 + \|\delta_k\|_2)^2 \\ &\leq 2\Delta_k + 2\lambda_k^* \theta_k + (\lambda_k^* + \Delta_k) \left(2\|\delta_k\|_2 + \|\delta_k\|_2^2 \right) \left(1 + (1 + \|\delta_k\|_2)^2 \right) \\ &= 2\Delta_k + 2\lambda_k^* \theta_k + (\lambda_k^* + \Delta_k) \|\delta_k\|_2 (2 + \|\delta_k\|_2) \left(1 + (1 + \|\delta_k\|_2)^2 \right) \end{aligned}$$

When $\|\delta_k\|_2 \leq \frac{1}{6}$, we will have that

$$(2 + \|\delta_k\|_2) \left(1 + (1 + \|\delta_k\|_2)^2 \right) = \frac{13}{6} \cdot \left(1 + \left(\frac{7}{6} \right)^2 \right) = \frac{1014}{216} \leq 5$$

Thus, we have

$$\Delta_{k+1} \leq 2\Delta_k + 2\lambda_k^* \theta_k + 5(\lambda_k^* + \Delta_k) \|\delta_k\|_2 \leq 3\Delta_k + 2\lambda_k^* \theta_k + 5\lambda_k^* \|\delta_k\|_2$$

A.2. Proof of Lemma 3.5

To prove Lemma 3.5, we first recall the Davis-Kahan sin Θ theorem (Davis & Kahan, 1970; Eldridge et al., 2017).

Theorem A.1 (Davis-Kahan sin Θ theorem). *Let $\mathbf{M}^* \in \mathbb{R}^{d \times d}$ and let $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$. Let \mathbf{a}_1^* be the top eigenvector of \mathbf{M}^* and \mathbf{a}_1 be the top eigenvector of \mathbf{M} . Then we have*

$$\sin \angle \{ \mathbf{a}_1^*, \mathbf{a}_1 \} \leq \frac{\|\mathbf{H}\|_2}{\min_{j \neq k} |\sigma_k^* - \sigma_j|}.$$

To start, we notice that, under the assumption that $\|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{4}\mathcal{T}_k$, by the Weyl's inequality, we have

$$\min_{j \neq k} |\sigma_k^* - \sigma_j| \geq \min_{j \neq k} |\sigma_k^* - \sigma_j^*| - |\sigma_j - \sigma_j^*| \geq \mathcal{T}_k - \|\Sigma_k - \Sigma_k^*\|_F \geq \frac{3}{4}\mathcal{T}_k$$

By Lemma C.1, we have that \mathbf{u}_k^* is the top eigenvector of Σ_k^* with corresponding eigenvalue λ_k^* . Moreover, as defined in (10), \mathbf{u}_k is the top eigenvector of Σ_k . Therefore, a direct application of Theorem A.1 gives

$$\sin \angle \{ \mathbf{u}_k, \mathbf{u}_k^* \} \leq \frac{4\|\Sigma_k - \Sigma_k^*\|_F}{3\mathcal{T}_k} \quad (30)$$

Moreover, notice that

$$\sin \angle \{ \mathbf{u}_k, \mathbf{u}_k^* \}^2 = 1 - \cos \angle \{ \mathbf{u}_k, \mathbf{u}_k^* \}^2 = 1 - (\mathbf{u}_k^\top \mathbf{u}_k^*)^2$$

which gives that $(\mathbf{u}_k^\top \mathbf{u}_k^*)^2 = 1 - \sin \angle \{ \mathbf{u}_k, \mathbf{u}_k^* \}^2$. Therefore

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2^2 = 2 - 2\mathbf{u}_k^\top \mathbf{u}_k^* \leq 2 - 2(\mathbf{u}_k^\top \mathbf{u}_k^*)^2 = 2 \sin \angle \{ \mathbf{u}_k, \mathbf{u}_k^* \}^2$$

Here the first inequality follows from the fact that $|\mathbf{u}_k^\top \mathbf{u}_k^*| \leq 1$. Applying (30) gives

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2 \leq \sqrt{2} \sin \angle \{ \mathbf{u}_k, \mathbf{u}_k^* \} \leq \frac{4\sqrt{2}}{3\mathcal{T}_k} \|\Sigma_k - \Sigma_k^*\|_F \leq \frac{2}{\mathcal{T}_k} \|\Sigma_k - \Sigma_k^*\|_F$$

A.3. Proof of Theorem 3.1

Let \hat{K} be the smallest integer such that $\|\Sigma_{\hat{K}} - \Sigma_{\hat{K}}^*\|_2 > \frac{1}{4}\mathcal{T}_{K, \min}$, and let $K' = \min \{ \hat{K}, K + 1 \}$. By definition, $\mathcal{T}_{K, \min} = \min_{k \in [K]} \mathcal{T}_k$. Therefore, for all $k < K'$, we shall have that $\|\Sigma_k - \Sigma_k^*\|_2 \leq \frac{1}{4}\mathcal{T}_k$. Therefore, for all $k \in [K']$, we can invoke Lemma 3.5 to get that

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2 \leq \frac{2}{\mathcal{T}_k} \|\Sigma_k - \Sigma_k^*\|_F \quad (31)$$

By the assumption of Theorem 3.1, we have

$$\|\boldsymbol{\delta}_k\|_2 \leq \frac{\mathcal{T}_{K,\min}}{20K} \prod_{j=k+1}^{K-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)^{-1} \leq \frac{1}{6}$$

since $\mathcal{T}_{K,\min} \leq \lambda_1^* = 1$ and $K \geq 1$. Therefore, we also have that $\|\boldsymbol{\delta}_k\|_2 \leq \frac{1}{6}$ for all $k \in [K' - 1]$. This allows us to invoke Lemma 3.4 to get that

$$\|\boldsymbol{\Sigma}_{k+1} - \boldsymbol{\Sigma}_{k+1}^*\|_F \leq 3 \|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_F + 2\lambda_k^* \|\mathbf{u}_k - \mathbf{u}_k^*\|_2 + 5\lambda_k^* \|\boldsymbol{\delta}_k\|_2 \quad (32)$$

Plugging (31) into (32) gives

$$\|\boldsymbol{\Sigma}_{k+1} - \boldsymbol{\Sigma}_{k+1}^*\|_F \leq \left(3 + \frac{2\lambda_k^*}{\mathcal{T}_k}\right) \|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_F + 5\lambda_k^* \|\boldsymbol{\delta}_k\|_2$$

Let the sequence $\{Q_k\}_{k=1}^{K'}$ be defined as

$$Q_{k+1} = a_k Q_k + b_k; \quad Q_0 = 0; \quad a_k = 3 + \frac{2\lambda_k^*}{\mathcal{T}_k}; \quad b_k = 5\lambda_k^* \|\boldsymbol{\delta}_k\|_2$$

Then we must have that $\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_F \leq Q_k$ for all $k \in [K']$. Invoking Lemma C.4 thus gives

$$\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_F \leq 5 \sum_{k'=1}^{k-1} \lambda_{k'}^* \|\boldsymbol{\delta}_{k'}\|_2 \prod_{j=k'+1}^{k-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right) \quad (33)$$

Combining with (31) and the fact that $\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \|\boldsymbol{\delta}_k\|_2 + \|\mathbf{u}_k - \mathbf{u}_k^*\|_2$ gives

$$\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \|\boldsymbol{\delta}_k\|_2 + \frac{2}{\mathcal{T}_k} \|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_F \leq \|\boldsymbol{\delta}_k\|_2 + \frac{10}{\mathcal{T}_k} \sum_{k'=1}^{k-1} \lambda_{k'}^* \|\boldsymbol{\delta}_{k'}\|_2 \prod_{j=k'+1}^{k-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right) \quad (34)$$

Notice that

$$\frac{10}{\mathcal{T}_k} \leq \frac{5}{\lambda_k^*} \cdot \left(3 + \frac{2\lambda_k^*}{\mathcal{T}_k}\right)$$

Therefore, (34) becomes

$$\|\mathbf{v}_k - \mathbf{u}_k^*\|_2 \leq \|\boldsymbol{\delta}_k\|_2 + 5 \sum_{k'=1}^{k-1} \frac{\lambda_{k'}^*}{\lambda_k^*} \|\boldsymbol{\delta}_{k'}\|_2 \prod_{j=k'+1}^k \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right) \leq 5 \sum_{k'=1}^k \frac{\lambda_{k'}^*}{\lambda_k^*} \|\boldsymbol{\delta}_{k'}\|_2 \prod_{j=k'+1}^k \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)$$

for all $k \in [K']$. It thus remains to show that $\hat{K} > K$. In this way, we will have $K' = K$, and the theorem is then proved. For the sake of contradiction, that $\hat{K} \leq K$. Then we shall have that there exists $k \in [K]$ such that $\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_2 > \frac{1}{4} \mathcal{T}_{K,\min}$. To reach a contradiction, it thus remains to show that for all $k \in [K]$

$$\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_2 \leq \frac{1}{8} \mathcal{T}_{K,\min}$$

By (33), it suffices to guarantee that

$$5 \sum_{k'=1}^{k-1} \lambda_{k'}^* \|\boldsymbol{\delta}_{k'}\|_2 \prod_{j=k'+1}^{k-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right) \leq \frac{1}{4} \mathcal{T}_{K,\min} \quad (35)$$

Indeed, by the assumption of Theorem 3.1, we require

$$\|\boldsymbol{\delta}_k\|_2 \leq \frac{\mathcal{T}_{K,\min}}{20K} \prod_{j=k+1}^{K-1} \left(3 + \frac{2\lambda_j^*}{\mathcal{T}_j}\right)^{-1}$$

Plugging this requirement into (33) would guarantee that $\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}_k^*\|_2 \leq \frac{1}{4} \mathcal{T}_{K,\min}$ and thus achieves the contradiction. This shows that $\hat{K} > K$ and finishes the proof.

B. Missing Proofs from Section 4

B.1. Proof of Lemma 4.3

Since each $\mathbf{x}_{0,k}$ is generated independently, it suffices to show that for a general \mathbf{x}_0 generated for unit sphere and some unit vector \mathbf{u} , we have $|\mathbf{x}_0^\top \mathbf{u}| \geq$ with high probability, and then apply a union bound for all $k \in [K]$. To start, for a fixed \mathbf{u} , we let \mathbf{U} be an orthogonal matrix with \mathbf{u} being its first row. Since \mathbf{x}_0 is generated uniformly at random from a unit sphere, we must have that $\mathbf{b} = \mathbf{U}\mathbf{x}$ is also generated uniformly at random from a unit sphere. Let b_j denote the j th entry of \mathbf{b} , then $b_1 = \mathbf{x}_0^\top \mathbf{u}$. We first aim at showing

$$\mathbb{P} \left(\sum_{j=1}^d \left(\frac{b_j}{b_1} \right) \leq 2d^3 \right) \geq 1 - \frac{2}{3d}$$

Let $\mathcal{B}_d(1)$ denote the unit ball in \mathbb{R}^d . Define the indicator function $\mathbb{I}(\mathbf{b})$ as

$$\mathbb{I}(\mathbf{b}) = \begin{cases} 1 & \text{if } \sum_{j=1}^d \left(\frac{b_j}{b_1} \right) \leq 2d^3 \\ 0 & \text{otherwise} \end{cases}$$

Then we have

$$\mathbb{P} \left(\sum_{j=1}^d \left(\frac{b_j}{b_1} \right) \leq 2d^3 \right) = \int_{\|\mathbf{x}\|_2=1} \mathbb{I}(\mathbf{b}) d\mu(\mathbf{b})$$

where μ is the probability measure over the unit sphere. Since $\mathbb{I}(\mathbf{b})$ satisfies $\mathbb{I}(\alpha\mathbf{b}) = \mathbb{I}(\mathbf{b})$ for any $\alpha \neq 0$ and $\mathbb{I}(\mathbf{b})$ does not depend on the signs of b_j 's, then we can use Remark 7.2 from (Kuczynski & Woźniakowski, 1992) to get that

$$\begin{aligned} \int_{\|\mathbf{x}\|_2=1} \mathbb{I}(\mathbf{b}) d\mu(\mathbf{b}) &= \frac{1}{V_d} \int_{\mathcal{B}_d(1)} \mathbb{I}(\mathbf{b}) d\mathbf{b} \\ &= 2 \int_0^1 \left(\int_{\sum_{j=2}^d b_j^2 \leq \min\{1-b_1^2, 2b_1 d^3\}} db_2 \dots db_d \right) db_1 \\ &= \frac{2V_{d-1}}{V_d} \int_0^1 \min\{1-b_1^2, 2b_1 d^3\}^{\frac{d-1}{2}} db_1 \end{aligned}$$

where the last equality follows from the observation that the inner integral is exactly the volume of the ball with radius $\min\{1-b_1^2, 2b_1 d^3\}$. We can lower-bound the last integral as

$$\begin{aligned} \int_0^1 \min\{1-b_1^2, 2b_1 d^3\}^{\frac{d-1}{2}} db_1 &\geq \int_{\sqrt{\frac{1}{1+2d^3}}}^1 (1-b_1^2)^{\frac{d-1}{2}} db_1 \\ &= \int_0^1 (1-b_1^2)^{\frac{d-1}{2}} db_1 - \int_0^{\sqrt{\frac{1}{1+2d^3}}} (1-b_1^2)^{\frac{d-1}{2}} db_1 \\ &\geq \frac{V_d}{2V_{d-1}} - \sqrt{\frac{1}{1+2d^3}} \end{aligned}$$

This gives

$$\mathbb{P} \left(\sum_{j=1}^d \left(\frac{b_j}{b_1} \right) \leq 2d^3 \right) \geq 1 - \frac{2V_{d-1}}{V_d} \sqrt{\frac{1}{1+2d^3}}$$

By (Kuczynski & Woźniakowski, 1992), Eq (13), we have $\frac{V_{d-1}}{V_d} \leq 0.412\sqrt{d}$ when $d \geq 8$. Thus, we have

$$\mathbb{P} \left(\sum_{j=1}^d \left(\frac{b_j}{b_1} \right) \leq 2d^3 \right) \geq 1 - 0.824\sqrt{\frac{d}{1+2d^3}} \geq 1 - \frac{2}{3d}$$

Notice that $\sum_{j=1}^d b_j^2 = 1$. Then we have that

$$\sum_{j=2}^d \left(\frac{b_j}{b_1}\right)^2 = \frac{1}{b_1^2} - 1 \Rightarrow |b_1| = \left(\sum_{j=2}^d \left(\frac{b_j}{b_1}\right)^2 + 1\right)^{-\frac{1}{2}}$$

Thus, we have

$$\mathbb{P}\left(|b_1| \geq \frac{1}{\sqrt{1+2d^3}}\right) \geq 1 - \frac{2}{3d}$$

Apply a union bound over all k gives that with probability at least $1 - \frac{2K}{3d}$, we have

$$|\mathbf{x}_{0,k}^\top \mathbf{u}_k| \geq \frac{1}{\sqrt{1+2d^3}}$$

B.2. Proof of Lemma 4.4

For convenience, we define the sequence $\{\hat{\mathbf{x}}_k\}_{k=0}^\infty$ as

$$\hat{\mathbf{x}}_t = \mathbf{M}\hat{\mathbf{x}}_{t-1}; \quad \hat{\mathbf{x}}_0 = \mathbf{x}_0$$

Then \mathbf{x}_t from (14) can be written as

$$\mathbf{x}_t = \frac{\hat{\mathbf{x}}_t}{\|\hat{\mathbf{x}}_t\|_2} = \frac{\hat{\mathbf{x}}_t}{\|\mathbf{M}^t \mathbf{x}_0\|_2}$$

Focusing on $\hat{\mathbf{x}}_t$, since $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$, we have

$$\hat{\mathbf{x}}_t = \mathbf{M}\hat{\mathbf{x}}_{t-1} = \mathbf{M}^* \hat{\mathbf{x}}_{t-1} + \mathbf{H}\hat{\mathbf{x}}_{t-1}$$

Since σ_j^* , \mathbf{a}_j^* are the j th eigenvalue and eigenvector of \mathbf{M}^* , we have that

$$\hat{\mathbf{x}}_t^\top \mathbf{a}_j^* = \hat{\mathbf{x}}_{t-1}^\top \mathbf{M}^* \mathbf{a}_j^* + \hat{\mathbf{x}}_{t-1}^\top \mathbf{H} \mathbf{a}_j^* = \sigma_j^* \hat{\mathbf{x}}_{t-1}^\top \mathbf{a}_j^* + \hat{\mathbf{x}}_{t-1}^\top \mathbf{H} \mathbf{a}_j^*$$

Unrolling the iterates gives

$$\hat{\mathbf{x}}_t^\top \mathbf{a}_j^* = \sigma_j^{*t} \mathbf{x}_0^\top \mathbf{a}_j^* + \sum_{t'=0}^{t-1} \sigma_j^{*t-t'} \hat{\mathbf{x}}_{t'}^\top \mathbf{H} \mathbf{a}_j^* = \sigma_j^{*t} \left(\mathbf{x}_0^\top \mathbf{a}_j^* + \sum_{t'=0}^{t-1} \frac{1}{\sigma_j^{*t'}} \hat{\mathbf{x}}_{t'}^\top \mathbf{H} \mathbf{a}_j^* \right)$$

Recalling that $\hat{\mathbf{x}}_t = \mathbf{M}^t \mathbf{x}_0$, we can see that

$$\hat{\mathbf{x}}_t^\top \mathbf{a}_j^* = \sigma_j^{*t} \left(\mathbf{x}_0^\top \mathbf{a}_j^* + \sum_{t'=0}^{t-1} \left(\frac{\mathbf{M}^{t'}}{\sigma_j^{*t'}} \hat{\mathbf{x}}_0 \right)^\top \mathbf{H} \mathbf{a}_j^* \right) = \sigma_j^{*t} \mathbf{x}_0^\top \left(\mathbf{I} + \sum_{t'=0}^{t-1} \frac{\mathbf{M}^{t'}}{\sigma_j^{*t'}} \cdot \mathbf{H} \right) \mathbf{a}_j^*$$

Plugging in $\mathbf{x}_t = \frac{\hat{\mathbf{x}}_t}{\|\mathbf{M}^t \mathbf{x}_0\|_2}$ gives the first conclusion. Next, without loss of generality, we assume that $\|\mathbf{x}_0\|_2 = 1$, since $|\mathbf{x}_t^\top \mathbf{a}_j^*|$ does not depend on $\|\mathbf{x}_0\|_2$. Since $\|\mathbf{x}_0\|_2 = \|\mathbf{a}_j^*\|_2 = 1$, it holds that

$$\begin{aligned} |\hat{\mathbf{x}}_t^\top \mathbf{a}_j^*| &\leq \sigma_j^{*t} \left(1 + \left\| \sum_{t'=0}^{t-1} \frac{\mathbf{M}^{t'}}{\sigma_j^{*t'}} \cdot \mathbf{H} \mathbf{a}_j^* \right\|_2 \right) \\ &\leq \sigma_j^{*t} \left(1 + \sum_{t'=0}^{t-1} \frac{\sigma_1^{t'}}{\sigma_j^{*t'}} \|\mathbf{H} \mathbf{a}_j^*\|_2 \right) \\ &\leq \sigma_j^{*t} \left(1 + \frac{\left(\frac{\sigma_1}{\sigma_j^*}\right)^t - 1}{\frac{\sigma_1}{\sigma_j^*} - 1} \|\mathbf{H} \mathbf{a}_j^*\|_2 \right) \\ &\leq \sigma_j^{*t} + \frac{\sigma_1^t - \sigma_j^{*t}}{\sigma_j^* - 1} \|\mathbf{H} \mathbf{a}_j^*\|_2 \\ &\leq \sigma_j^{*t} + \frac{\sigma_j^* \sigma_1^t}{\sigma_1 - \sigma_j^*} \|\mathbf{H} \mathbf{a}_j^*\|_2 \end{aligned}$$

Since there exists some $|\mathbf{x}_0^\top \mathbf{a}_1| \geq c_0^{-1} > 0$, we can bound $\|\mathbf{M}^t \mathbf{x}_0\|_2$ as

$$\|\mathbf{M}^t \mathbf{x}_0\|_2 = \left(\sum_{j=1}^r \sigma_j^{2t} (\mathbf{x}_0^\top \mathbf{a}_j)^2 \right)^{\frac{1}{2}} \geq \sigma_1^t |\mathbf{x}_0^\top \mathbf{a}_1| \geq \sigma_1^t c_0^{-1}$$

Therefore, we have

$$|\mathbf{x}_t^\top \mathbf{a}_j^*| \leq \|\mathbf{M}^t \mathbf{x}_0\|_2^{-1} |\hat{\mathbf{x}}_t^\top \mathbf{a}_j^*| \leq c_0 \left(\left(\frac{\sigma_j^*}{\sigma_1} \right)^t + \frac{\sigma_j^*}{\sigma_1 - \sigma_j^*} \|\mathbf{H} \mathbf{a}_j^*\|_2 \right)$$

This shows the second conclusion.

B.3. Proof of Lemma 4.5

Recall the definition of the set of empirical deflated matrices

$$\Sigma_{k+1} = \Sigma_k - \mathbf{v}_k \mathbf{v}_k^\top \Sigma_k \mathbf{v}_k \mathbf{v}_k^\top$$

and the set of ground-truth deflated matrices

$$\Sigma_{k+1}^* = \Sigma_k^* - \mathbf{u}_k \mathbf{u}_k^{*\top} \Sigma_k^* \mathbf{u}_k \mathbf{u}_k^{*\top}$$

By Lemma C.1, we have

$$\Sigma_k^* = \sum_{j=k}^n \lambda_j^* \mathbf{u}_j \mathbf{u}_j^{*\top}$$

Therefore, by the orthogonality between \mathbf{u}_i^* and \mathbf{u}_j^* when $i \neq j$, we have

$$\Sigma_k^* \mathbf{u}_j^* = \sum_{i=k}^n \lambda_i^* \mathbf{u}_i^* (\mathbf{u}_i^{*\top} \mathbf{u}_j^*) = \begin{cases} \lambda_j^* \mathbf{u}_j^* & \text{if } j \geq k \\ \mathbf{0} & \text{if } j < k \end{cases} \quad (36)$$

Next, we will focus on $\Sigma_k \mathbf{u}_j^*$. Notice that

$$\Sigma_{k+1} \mathbf{u}_j^* = \Sigma_k \mathbf{u}_j^* - \mathbf{v}_k \mathbf{v}_k^\top \Sigma_k \mathbf{v}_k \mathbf{v}_k^\top \mathbf{u}_j^*$$

Therefore

$$\Sigma_k \mathbf{u}_j^* = \Sigma \mathbf{u}_j^* - \sum_{k'=1}^{k-1} \mathbf{v}_{k'}^\top \mathbf{u}_j^* \cdot \mathbf{v}_{k'} \mathbf{v}_{k'}^\top \Sigma_{k'} \mathbf{v}_{k'} = \lambda_j^* \mathbf{u}_j^* - \sum_{k'=1}^{k-1} \mathbf{v}_{k'}^\top \mathbf{u}_j^* \cdot \mathbf{v}_{k'}^\top \Sigma_{k'} \mathbf{v}_{k'} \cdot \mathbf{v}_{k'} \quad (37)$$

Therefore, combining (36) and (37) gives

$$(\Sigma_k^* - \Sigma_k) \mathbf{u}_j^* = \begin{cases} \sum_{k'=1}^{k-1} \mathbf{v}_{k'}^\top \mathbf{u}_j^* \cdot \mathbf{v}_{k'}^\top \Sigma_{k'} \mathbf{v}_{k'} \cdot \mathbf{v}_{k'} & \text{if } j \geq k \\ \sum_{k'=1}^{k-1} \mathbf{v}_{k'}^\top \mathbf{u}_j^* \cdot \mathbf{v}_{k'}^\top \Sigma_{k'} \mathbf{v}_{k'} \cdot \mathbf{v}_{k'} - \lambda_j^* \mathbf{u}_j^* & \text{if } j < k \end{cases}$$

This shows the first conclusion. Next, we will put more emphasis on the case where $j \geq k$. Under this scenario, we can first show that

$$\mathbf{v}_{k'}^\top \Sigma_{k'} \mathbf{v}_{k'} \leq \lambda_{k'} \|\mathbf{v}_{k'}\|_2^2 = \lambda_{k'}$$

Therefore,

$$\|(\Sigma_k^* - \Sigma_k) \mathbf{u}_j^*\|_2 \leq \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \cdot \|\mathbf{v}_{k'}\|_2 = \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*|$$

B.4. Proof of Lemma 4.6

Let us consider $\mathbf{a}_i^\top (\mathbf{M}^* + \mathbf{H}) \mathbf{a}_j^*$. On one hand, since \mathbf{a}_i is an eigenvector of $\mathbf{M}^* + \mathbf{H}$, we have

$$\mathbf{a}_i^\top (\mathbf{M}^* + \mathbf{H}) \mathbf{a}_j^* = \sigma_i \mathbf{a}_i^\top \mathbf{a}_j^*$$

On the other hand, since \mathbf{a}_i^* is the i th eigenvector of \mathbf{M}^* , we have

$$\mathbf{a}_i^\top (\mathbf{M}^* + \mathbf{H}) \mathbf{a}_j^* = \sigma_j^* \mathbf{a}_i^\top \mathbf{a}_j^* + \mathbf{a}_i^\top \mathbf{H} \mathbf{a}_j^*$$

Therefore,

$$\sigma_i \mathbf{a}_i^\top \mathbf{a}_j^* = \sigma_j^* \mathbf{a}_i^\top \mathbf{a}_j^* + \mathbf{a}_i^\top \mathbf{H} \mathbf{a}_j^*$$

which implies that

$$(\sigma_i - \sigma_j^*) \mathbf{a}_i^\top \mathbf{a}_j^* = \mathbf{a}_i^\top \mathbf{H} \mathbf{a}_j^*$$

Thus, when $\sigma_i \neq \sigma_j^*$, we have

$$\mathbf{a}_i^\top \mathbf{a}_j^* = \frac{\mathbf{a}_i^\top \mathbf{H} \mathbf{a}_j^*}{\sigma_i - \sigma_j^*}$$

B.5. Proof of Lemma 4.7

The proof of Lemma 4.7 is based on the Neumann expansion (Eldridge et al., 2017; Chen et al., 2020), which we state below

Lemma B.1 (Theorem 7 of (Eldridge et al., 2017)). *Let $\mathbf{M}^* \in \mathbb{R}^{d \times d}$ be a rank- r matrix, and let $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$ for some $\mathbf{H} \in \mathbb{R}^{d \times d}$. Let $\sigma_i^*, \mathbf{a}_i^*$ be the i -th eigenvalue and eigenvector of \mathbf{M}^* , and σ_i, \mathbf{a}_i be the i -th eigenvalue and eigenvector of \mathbf{M} . Suppose that $\|\mathbf{H}\|_2 \leq |\sigma_i|$. Then,*

$$\mathbf{a}_i = \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_i} \mathbf{a}_i^\top \mathbf{a}_j^* \left(\sum_{s=0}^{\infty} \frac{1}{\sigma_i^s} \mathbf{H}^s \mathbf{a}_j^* \right).$$

To start, using the Neumann expansion (Lemma B.1), we prove an auxiliary lemma, as stated below.

Lemma B.2. *Under the same setting of Lemma B.1, we further assume that $\sigma_i^* \geq 0$ for all $i \in [r]$, and that $\|\mathbf{H}\|_2 \leq \frac{1}{8} \sigma_1^*$. Then we have*

$$\left\| \mathbf{a}_1 - \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 \leq \frac{4}{3} \cdot \frac{\|\mathbf{H}\|_2}{\sigma_1}$$

Proof. Applying Lemma B.1 with $i = 1$, we have

$$\mathbf{a}_1 = \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \left(\mathbf{a}_j^* + \sum_{s=1}^{\infty} \frac{1}{\sigma_1^s} \mathbf{H}^s \mathbf{a}_j^* \right) = \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* + \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \left(\sum_{s=1}^{\infty} \frac{1}{\sigma_1^s} \mathbf{H}^s \mathbf{a}_j^* \right)$$

Therefore, our quantity of interest can be written as

$$\begin{aligned}
 \left\| \mathbf{a}_1 - \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 &= \left\| \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \left(\sum_{s=1}^{\infty} \frac{1}{\sigma_1^s} \mathbf{H}^s \mathbf{a}_j^* \right) \right\|_2 \\
 &= \left\| \sum_{s=1}^{\infty} \frac{\mathbf{H}^s}{\sigma_1^{s+1}} \sum_{j=1}^r \sigma_j^* \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 \\
 &= \left\| \sum_{s=1}^{\infty} \frac{\mathbf{H}^s}{\sigma_1^{s+1}} \begin{bmatrix} \mathbf{a}_1^* & \dots & \mathbf{a}_r^* \end{bmatrix} \begin{bmatrix} \sigma_1^* \mathbf{a}_1^\top \mathbf{a}_1^* \\ \vdots \\ \sigma_r^* \mathbf{a}_1^\top \mathbf{a}_r^* \end{bmatrix} \right\|_2 \\
 &\leq \left\| \sum_{s=1}^{\infty} \frac{\mathbf{H}^s}{\sigma_1^{s+1}} \right\|_2 \cdot \left\| \begin{bmatrix} \mathbf{a}_1^* & \dots & \mathbf{a}_r^* \end{bmatrix} \right\|_2 \cdot \left\| \begin{bmatrix} \sigma_1^* \mathbf{a}_1^\top \mathbf{a}_1^* \\ \vdots \\ \sigma_r^* \mathbf{a}_1^\top \mathbf{a}_r^* \end{bmatrix} \right\|_2
 \end{aligned} \tag{38}$$

Now, we have a product of three terms to analyze. Recall that $\sigma_1^* > 0$. To start, by Weyl's inequality, we have $\sigma_1 \geq \sigma_1^* - \|\mathbf{H}\|_2$. Since $\|\mathbf{H}\|_2 \leq \frac{1}{8}\sigma_1^*$, we must have that $\|\mathbf{H}\|_2 \leq \sigma_1$. Applying the triangle inequality and the standard result of the geometric series gives

$$\left\| \sum_{s=1}^{\infty} \frac{\mathbf{H}^s}{\sigma_1^{s+1}} \right\|_2 \leq \sum_{s=1}^{\infty} \frac{\|\mathbf{H}\|_2^s}{\sigma_1^{s+1}} = \frac{\|\mathbf{H}\|_2}{\sigma_1(\sigma_1 - \|\mathbf{H}\|_2)} \tag{39}$$

Next, since $\mathbf{a}_1^*, \dots, \mathbf{a}_r^*$ are orthonormal vectors, we must have that

$$\left\| \begin{bmatrix} \mathbf{a}_1^* & \dots & \mathbf{a}_r^* \end{bmatrix} \right\|_2 \leq 1 \tag{40}$$

For the same reason, and since $\sigma_i^* \geq 0$ for all $i \in [r]$, we also have that

$$\left\| \begin{bmatrix} \sigma_1^* \mathbf{a}_1^\top \mathbf{a}_1^* \\ \vdots \\ \sigma_r^* \mathbf{a}_1^\top \mathbf{a}_r^* \end{bmatrix} \right\|_2 \leq \sigma_1^* \left\| \begin{bmatrix} \mathbf{a}_1^\top \mathbf{a}_1^* \\ \vdots \\ \mathbf{a}_1^\top \mathbf{a}_r^* \end{bmatrix} \right\|_2 = \sigma_1^* \left\| \mathbf{a}_1^\top \begin{bmatrix} \mathbf{a}_1^* & \dots & \mathbf{a}_r^* \end{bmatrix} \right\|_2 \leq \sigma_1^* \tag{41}$$

To put things together, we plug (39), (40), and (41) into (38) to get that

$$\left\| \mathbf{a}_1 - \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 \leq \frac{\|\mathbf{H}\|_2}{\sigma_1(\sigma_1 - \|\mathbf{H}\|_2)} \cdot 1 \cdot \sigma_1^* = \frac{\sigma_1^* \|\mathbf{H}\|_2}{\sigma_1(\sigma_1 - \|\mathbf{H}\|_2)} \tag{42}$$

Apply again $\|\mathbf{H}\|_2 \leq \frac{1}{8}\sigma_1^*$, we have

$$\sigma_1 - \|\mathbf{H}\|_2 \geq \sigma_1^* - 2\|\mathbf{H}\|_2 \geq \frac{3}{4}\sigma_1^*$$

Therefore, (42) becomes

$$\left\| \mathbf{a}_1 - \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 \leq \frac{4}{3} \cdot \frac{\|\mathbf{H}\|_2}{\sigma_1} \tag{43}$$

□

Now, we will use Lemma B.2 to prove the desired statements. Notice that the Euclidean projection of \mathbf{a}_1 onto the subspace spanned by $\{\mathbf{a}_j^*\}_{j=1}^r$ (namely \mathcal{S}^*) is given by $\sum_{j=1}^r \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^*$. Therefore, we have

$$\begin{aligned} \left(1 - \sum_{j=1}^r (\mathbf{a}_1^\top \mathbf{a}_j^*)^2\right)^{\frac{1}{2}} &= \left\| \mathbf{a}_1 - \sum_{j=1}^r \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 \\ &= \min_{\mathbf{z} \in \mathcal{S}^*} \|\mathbf{a}_1 - \mathbf{z}\|_2 \\ &\leq \left\| \mathbf{a}_1 - \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 \end{aligned} \quad (44)$$

where the inequality follows from the fact that $\sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \in \mathcal{S}^*$. With the assumption of the lemma, we can invoke Lemma B.2 to get that

$$\left\| \mathbf{a}_1 - \sum_{j=1}^r \frac{\sigma_j^*}{\sigma_1} \mathbf{a}_1^\top \mathbf{a}_j^* \cdot \mathbf{a}_j^* \right\|_2 \leq \frac{4}{3} \cdot \frac{\|\mathbf{H}\|_2}{\sigma_1} \leq \frac{32}{21} \cdot \frac{\|\mathbf{H}\|_2}{\sigma_1^*}$$

Combining with (44) gives

$$\sum_{j=1}^r (\mathbf{a}_1^\top \mathbf{a}_j^*)^2 \geq 1 - \left(\frac{32}{21} \cdot \frac{\|\mathbf{H}\|_2}{\sigma_1^*}\right)^2 \geq 1 - 2.4 \cdot \frac{\|\mathbf{H}\|_2^2}{\sigma_1^{*2}}$$

Moving $\sum_{j=2}^r (\mathbf{a}_1^\top \mathbf{a}_j^*)^2$ to the right-hand side of the inequality gives

$$(\mathbf{a}_1^\top \mathbf{a}_1^*)^2 \geq 1 - 2.4 \cdot \frac{\|\mathbf{H}\|_2^2}{\sigma_1^{*2}} - \sum_{j=2}^r (\mathbf{a}_1^\top \mathbf{a}_j^*)^2$$

The proof then conclude by letting $\mathbf{a}_1 = \mathbf{u}_k$, $\sigma_1^* = \lambda_k^*$, $\mathbf{a}_j^* = \mathbf{u}_{k+j-1}^*$ for all $j \in [d - k + 1]$, and $\mathbf{H} = \Sigma_k - \Sigma_k^*$.

B.6. Proof of Lemma 4.8

To start, by the definition of c_0 , we invoke Lemma 4.4 with $\mathbf{M} = \Sigma_k$, $\mathbf{M}^* = \Sigma_k^*$ to obtain that

$$|\mathbf{v}_k^\top \mathbf{u}_j^*| \leq c_0 \left(\left(\frac{\lambda_j^*}{\lambda_k}\right)^t + \frac{\lambda_j^*}{\lambda_k - \lambda_j^*} \|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2 \right) \quad (45)$$

Also, invoking Lemma 4.5, we have that when $j \geq k$,

$$\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2 \leq \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \quad (46)$$

Plugging the upper bound of $\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2$ from (46) into (45), we have that when $j \geq k$

$$|\mathbf{v}_k^\top \mathbf{u}_j^*| \leq c_0 \left(\left(\frac{\lambda_j^*}{\lambda_k}\right)^t + \frac{\lambda_j^*}{\lambda_k - \lambda_j^*} \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \right)$$

Equivalently, we can multiply both sides of the inequality with λ_k to obtain that

$$\lambda_k |\mathbf{v}_k^\top \mathbf{u}_j^*| \leq \lambda_k c_0 \left(\left(\frac{\lambda_j^*}{\lambda_k}\right)^t + \frac{\lambda_j^*}{\lambda_k - \lambda_j^*} \sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \right) \quad (47)$$

Notice that, the form of the inequality in (47) allows us to invoke Lemma C.3 to analyze the growth of $\lambda_k |\mathbf{v}_k^\top \mathbf{u}_j^*|$. In particular, setting $a_k = c_0 \lambda_k \left(\frac{\lambda_j^*}{\lambda_k}\right)^t$ for $k > 1$ and $a_1 = \lambda_1 |\mathbf{v}_1^\top \mathbf{u}_j^*|$, and $b_k = \frac{c_0 \lambda_k \lambda_j^*}{\lambda_k - \lambda_j^*}$, we can have that $\lambda_k |\mathbf{v}_k^\top \mathbf{u}_j^*| \leq Q_k$

for all k for Q_k defined in Lemma C.3. Therefore, Lemma C.3 implies that

$$\sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \leq c_0 \sum_{k'=1}^{k-1} \lambda_{k'} \left(\frac{\lambda_j^*}{\lambda_k} \right)^t \prod_{\ell=k'+1}^{k-1} \left(1 + \frac{c_0 \lambda_\ell \lambda_j^*}{\lambda_\ell - \lambda_j^*} \right) \quad (48)$$

Now, we focus on the term $\left(1 + \frac{c_0 \lambda_\ell \lambda_j^*}{\lambda_\ell - \lambda_j^*} \right)$ in (48). In particular, we should notice that $\left(1 + \frac{c_0 \lambda_\ell \lambda_j^*}{\lambda_\ell - \lambda_j^*} \right)$ is only used in (48) for $\ell < k$. We can observe that

$$\frac{c_0 \lambda_\ell \lambda_j^*}{\lambda_\ell - \lambda_j^*} = c_0 \left(\frac{1}{\lambda_j^*} - \frac{1}{\lambda_\ell} \right)^{-1}$$

This implies that $\left(1 + \frac{c_0 \lambda_\ell \lambda_j^*}{\lambda_\ell - \lambda_j^*} \right)$ increases monotonically as λ_ℓ decrease. Recall that $|\lambda_\ell^* - \lambda_\ell| \leq \|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{8} \mathcal{T}_{k,\min}$. This implies that for all $\ell < k$, we must have that $\lambda_\ell \geq \lambda_{\ell+1}^* \geq \lambda_{k+1}^* \geq \lambda_k^*$. Thus, we must have that

$$1 + \frac{c_0 \lambda_\ell \lambda_j^*}{\lambda_\ell - \lambda_j^*} \leq 1 + \frac{c_0 \lambda_k^* \lambda_j^*}{\lambda_k^* - \lambda_j^*} \leq 1 + \frac{c_0 \lambda_k^* \lambda_{k+1}^*}{\lambda_k^* - \lambda_{k+1}^*} \quad (49)$$

Recall that in the statement of Lemma 4.8 we defined $\mathcal{G}_k = 1 + \frac{c_0 \lambda_k^* \lambda_{k+1}^*}{\lambda_k^* - \lambda_{k+1}^*}$. Then (49) implies that (48) can be simplified. Suppose that $\lambda_k \geq \lambda_j^*$, then we have

$$\sum_{k'=1}^{k-1} \lambda_{k'} |\mathbf{v}_{k'}^\top \mathbf{u}_j^*| \leq c_0 \sum_{k'=1}^{k-1} \lambda_{k'} \mathcal{G}_k^{k-k'-1} \left(\frac{\lambda_j^*}{\lambda_k} \right)^t \leq c_0 \sum_{k'=1}^{k-1} \mathcal{G}_k^{k-k'-1} \left(\frac{\lambda_j^*}{\lambda_k} \right)^t \quad (50)$$

where the last inequality follows from the fact that

$$\lambda_{k'} \leq \lambda_{k'}^* + \|\Sigma_{k'} - \Sigma_{k'}^*\|_2 \leq \lambda_2^* + \frac{1}{8} \mathcal{T}_{k',\min} \leq \lambda_1^* = 1$$

for all $k \geq 2$, and $\lambda_1 = \lambda_1^* = 1$ by definition. Plugging (50) into the bound on $\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2$ in (46) gives

$$\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2 \leq c_0 \sum_{k'=1}^{k-1} \mathcal{G}_k^{k-k'-1} \left(\frac{\lambda_j^*}{\lambda_k} \right)^t \quad (51)$$

Now, we can invoke Lemma 4.6 with $\mathbf{M} = \Sigma_k$, $\mathbf{M}^* = \Sigma_k^*$, and correspondingly $\mathbf{a}_1 = \mathbf{u}_k$, $\mathbf{a}_j = \mathbf{u}_j^*$ to get that

$$|\mathbf{u}_k^\top \mathbf{u}_j^*| = \left| \frac{\mathbf{u}_k^\top (\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*}{\lambda_k - \lambda_j^*} \right| \leq \frac{\|(\Sigma_k - \Sigma_k^*) \mathbf{u}_j^*\|_2}{|\lambda_k - \lambda_j^*|} \quad (52)$$

This allows us to plug (51) into (52) to obtain that

$$|\mathbf{u}_k^\top \mathbf{u}_j^*| \leq \frac{c_0}{|\lambda_k - \lambda_j^*|} \sum_{k'=1}^{k-1} \mathcal{G}_k^{k-k'-1} \left(\frac{\lambda_j^*}{\lambda_k} \right)^t$$

Since $\|\Sigma_{k'} - \Sigma_{k'}^*\|_F \leq \frac{1}{8} \mathcal{T}_{k,\min}$ for all $k' \in [k-1]$, then we have that $\lambda_{k-1} \geq \lambda_{k-1}^* - \frac{1}{8} \mathcal{T}_{k-1} \geq \lambda_k^* + \frac{1}{8} \mathcal{T}_{k-1} \geq \lambda_k$. Apply Lemma C.5, we have that as long as $t \geq \frac{\log \mathcal{G}_k}{\log \lambda_{k'} - \lambda_{k'+1}}$ for all $k' \in [k-1]$,

$$\begin{aligned} |\mathbf{u}_k^\top \mathbf{u}_j^*| &\leq \frac{c_0}{|\lambda_k - \lambda_j^*|} \sum_{k'=1}^{k-1} \mathcal{G}_k^{k-k'-1} \left(\frac{\lambda_j^*}{\lambda_k} \right)^t \\ &= \frac{c_0 \lambda_j^{*t} \mathcal{G}_k^{k-1}}{|\lambda_k - \lambda_j^*|} \sum_{k'=1}^{k-1} \frac{(\lambda_{k'}^{-1})^t}{\mathcal{G}_k^{k'}} \\ &\leq \frac{c_0 \lambda_j^{*t} \mathcal{G}_k^{k-1}}{|\lambda_k - \lambda_j^*|} \cdot \left(1 - \mathcal{G}_k \cdot \max_{k' \in [k-1]} \left(\frac{\lambda_{k'+1}}{\lambda_{k'}} \right)^t \right)^{-1} \frac{\lambda_{k-1}^{-t}}{\mathcal{G}_k^{k-1}} \\ &= \frac{c_0}{|\lambda_k - \lambda_j^*|} \left(1 - \mathcal{G}_k \cdot \max_{k' \in [k-1]} \left(\frac{\lambda_{k'+1}}{\lambda_{k'}} \right)^t \right)^{-1} \left(\frac{\lambda_j^*}{\lambda_{k-1}} \right)^t \end{aligned}$$

In fact, as long as $t \geq \frac{\log 2\mathcal{G}_k}{\log \frac{\lambda_{k'} - \lambda_{k'+1}}{\lambda_{k'}}}$, we must have that for all $k' \in [k-1]$,

$$\mathcal{G}_k \cdot \max_{k' \in [k-1]} \left(\frac{\lambda_{k'+1}}{\lambda_{k'}} \right)^t \leq \frac{1}{2}$$

which implies that for all $j > k$

$$|\mathbf{u}_k^\top \mathbf{u}_j^*| \leq \frac{2c_0}{|\lambda_k - \lambda_j^*|} \left(\frac{\lambda_j^*}{\lambda_{k-1}} \right)^t \leq \frac{16c_0}{7\mathcal{T}_k} \left(\frac{\lambda_j^*}{\lambda_k^*} \right)^t \quad (53)$$

where the last inequality follows from

$$|\lambda_k - \lambda_j^*| \geq \left| \lambda_k^* - \frac{1}{8}\mathcal{T}_k - \lambda_j^* \right| \geq \frac{7}{8}\mathcal{T}_k$$

since $\lambda_k \geq \lambda_k^* - \|\Sigma_k - \Sigma_k^*\|_2 \geq \lambda_k^* - \frac{1}{8}\mathcal{T}_k$, and $\lambda_k^* \geq \lambda_j^* + \mathcal{T}_k$ for $j > k$. Applying Lemma 4.7, we have that

$$(\mathbf{u}_k^\top \mathbf{u}_k^*)^2 \geq 1 - \frac{2.4}{\lambda_k^{*2}} \|\Sigma_k - \Sigma_k^*\|_2^2 - \sum_{j=k+1}^d (\mathbf{u}_k^\top \mathbf{u}_j^*)^2$$

Therefore, $\|\mathbf{u}_k - \mathbf{u}_k^*\|_2^2$ can be bounded as

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2^2 \leq 2 \left(1 - (\mathbf{u}_k^\top \mathbf{u}_k^*)^2 \right) \leq \frac{4.8}{\lambda_k^*} \|\Sigma_k - \Sigma_k^*\|_2^2 + \sum_{j=k+1}^d (\mathbf{u}_k^\top \mathbf{u}_j^*)^2 \quad (54)$$

We finally plug (53) into (54) to get

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2^2 \leq \frac{4.8}{\lambda_k^*} \|\Sigma_k - \Sigma_k^*\|_2^2 + \frac{11c_0^2}{2\mathcal{T}_k^2} \sum_{j=k+1}^d \left(\frac{\lambda_j^*}{\lambda_k^*} \right)^{2t}$$

as desired. Lastly, we should notice that since $\|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{8}\mathcal{T}_{k,\min} \leq \lambda_{k'}^* - \lambda_{k'+1}^*$ for all $k' \in [k]$, it must hold that $\lambda_{k'} \geq \lambda_{k'}^* - \|\Sigma_k - \Sigma_k^*\|_F \geq \frac{1}{8}\lambda_{k'+1}^* + \frac{7}{8}\lambda_{k'}^*$ and $\lambda_{k'+1} \leq \lambda_{k'+1}^* + \|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{8}\lambda_{k'}^* + \frac{7}{8}\lambda_{k'+1}^*$ for all $k' \in [k-1]$. Thus, the condition that $t \geq \frac{\log 2\mathcal{G}_k}{\log \frac{\lambda_{k'} - \lambda_{k'+1}}{\lambda_{k'}}}$ can be satisfied as long as

$$t \geq \log 2\mathcal{G}_k \left(\log \frac{\frac{1}{8}\lambda_{k'+1}^* + \frac{7}{8}\lambda_{k'}^*}{\frac{1}{8}\lambda_{k'}^* + \frac{7}{8}\lambda_{k'+1}^*} \right)^{-1} = \log 2\mathcal{G}_k \left(\log \frac{\lambda_{k'+1}^* + 7\lambda_{k'}^*}{\lambda_{k'}^* + 7\lambda_{k'+1}^*} \right)^{-1}$$

which is provided as an assumption in the lemma. This then finishes the proof.

B.7. Proof of Theorem 4.1

Let \hat{K} be the smallest integer such that $\|\Sigma_{\hat{K}} - \Sigma_{\hat{K}}^*\|_2 > \frac{1}{8}\mathcal{T}_{K,\min}$, and let $K' = \min\{\hat{K}, K+1\}$. Therefore, for all $k < K'$, we shall have that $\|\Sigma_k - \Sigma_k^*\|_2 \leq \frac{1}{8}\mathcal{T}_{K,\min}$. Moreover, since $t \geq \log 2\mathcal{G}_k \left(\log \frac{\lambda_{k'+1}^* + 7\lambda_{k'}^*}{\lambda_{k'}^* + 7\lambda_{k'+1}^*} \right)^{-1}$, we can invoke Lemma 4.8 to have that

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2^2 \leq \frac{4.8}{\lambda_k^*} \|\Sigma_k - \Sigma_k^*\|_2^2 + \frac{11c_0^2}{2\mathcal{T}_k^2} \sum_{j=k+1}^d \left(\frac{\lambda_j^*}{\lambda_k^*} \right)^{2t} \quad (55)$$

Applying Lemma C.5 again with $g = 1$ and $p_{k'} = \lambda_{d-k'+1}^*$ yields that, as long as $t \geq \frac{1}{\log \frac{\lambda_j^* - \lambda_{j+1}^*}}{\lambda_j^*}$ for all $j > k$, it holds that

$$\sum_{j=k+1}^d \lambda_j^{*2t} \leq \left(1 - \max_{j>k} \left(\frac{\lambda_j^*}{\lambda_{j+1}^*} \right)^{2t} \right) \lambda_{k+1}^{*2t} \leq 2\lambda_{k+1}^{*2t}$$

This transforms (55) into

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2^2 \leq \frac{4.8}{\lambda_k^*} \|\Sigma_k - \Sigma_k^*\|_2^2 + \frac{11c_0^2}{\mathcal{T}_k^2} \left(\frac{\lambda_{k+1}^*}{\lambda_k^*} \right)^{2t} \quad (56)$$

Now, we invoke Lemma 3.4 to obtain that, as long as $\|\delta_k\| \leq \frac{1}{6}$, we shall have that

$$\|\Sigma_{k+1} - \Sigma_{k+1}^*\|_F \leq 3 \|\Sigma_k - \Sigma_k^*\|_F + 2\lambda_k^* \|\mathbf{u}_k - \mathbf{u}_k^*\|_2 + 5\lambda_k^* \|\delta_k\|_2 \quad (57)$$

We first plug (56) into (57) by applying $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ to have that

$$\|\Sigma_{k+1} - \Sigma_{k+1}^*\|_F \leq 8 \|\Sigma_k - \Sigma_k^*\|_F + \lambda_k^* \left(5 \|\delta_k\|_2 + \frac{7c_0}{\mathcal{T}_k} \left(\frac{\lambda_{k+1}^*}{\lambda_k^*} \right)^t \right) \quad (58)$$

Now, we can apply Lemma C.4 to unroll (58) to get that

$$\|\Sigma_k - \Sigma_k^*\|_F \leq \sum_{k'=1}^{k-1} 8^{k-k'} \lambda_{k'}^* \left(5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*} \right)^t \right) \quad (59)$$

Next, we can plug (59) into (56) to get that

$$\|\mathbf{u}_k - \mathbf{u}_k^*\|_2 \leq 3 \sum_{k'=1}^{k-1} 8^{k-k'} \frac{\lambda_{k'}^*}{\lambda_k^*} \left(5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*} \right)^t \right) + \frac{4c_0}{\mathcal{T}_k} \left(\frac{\lambda_{k+1}^*}{\lambda_k^*} \right)^t \quad (60)$$

Finally, recall the definition of $\mathbf{v}_k = \mathbf{u}_k + \delta_k$. Therefore, we have

$$\begin{aligned} \|\mathbf{v}_k - \mathbf{u}_k^*\|_2 &\leq 3 \sum_{k'=1}^{k-1} 8^{k-k'} \frac{\lambda_{k'}^*}{\lambda_k^*} \left(5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*} \right)^t \right) + \frac{4c_0}{\mathcal{T}_k} \left(\frac{\lambda_{k+1}^*}{\lambda_k^*} \right)^t + \|\delta_k\|_2 \\ &\leq 3 \sum_{k'=1}^k 8^{k-k'} \frac{\lambda_{k'}^*}{\lambda_k^*} \left(5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*} \right)^t \right) \end{aligned}$$

for all $k \in [K']$. It thus remains to show that $\hat{K} > K$. In this way, we will have $K' = K$ and the theorem is then proved. Assume, for the sake of contradiction, that $\hat{K} \leq K$. Then we shall have that there exists $k \in [K]$ such that $\|\Sigma_k - \Sigma_k^*\|_2 > \frac{1}{8} \mathcal{T}_{K,\min}$. To reach a contradiction, it thus remains to show that for all $k \in [K]$

$$\|\Sigma_k - \Sigma_k^*\|_2 \leq \frac{1}{8} \mathcal{T}_{K,\min}$$

By (59), it suffices to guarantee that

$$\sum_{k'=1}^{K-1} 8^{K-k'} \lambda_{k'}^* \left(5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*} \right)^t \right) \leq \frac{1}{8} \mathcal{T}_{K,\min} \quad (61)$$

Moreover, recall that \mathbf{v}_k is the output of the k th power iteration procedure. Therefore, we can invoke Lemma C.6 to have that

$$\|\delta_k\|_2 \leq 2 \left(\frac{\lambda_2(\Sigma_k)}{\lambda_k} \right)^t$$

Let $r_k = \max \left\{ \frac{\lambda_2(\Sigma_k)}{\lambda_k}, \frac{\lambda_{k+1}^*}{\lambda_k^*} \right\}$. Therefore,

$$5 \|\delta_{k'}\|_2 + \frac{7c_0}{\mathcal{T}_{k'}} \left(\frac{\lambda_{k'+1}^*}{\lambda_{k'}^*} \right)^t \leq \frac{17c_0}{\mathcal{T}_{k'}} r_{k'}^t$$

Thus, to guarantee (61), it suffices to guarantee that

$$\sum_{k'=1}^{K-1} 8^{K-k'} \frac{\lambda_{k'}^*}{\mathcal{T}_{k'}} r_{k'}^t \leq \frac{\mathcal{T}_{K,\min}}{140c_0}$$

which, since $\lambda_k \leq \lambda_1 = \lambda_1^* = 1$, is satisfied as long as

$$t \geq \frac{3(K-k) + \log \frac{140c_0K}{\mathcal{T}_{K,\min}^2}}{\log r_k^{-1}}$$

Since $\|\Sigma_k - \Sigma_k^*\|_F \leq \frac{1}{8}\mathcal{T}_{k,\min}$, we have that $r_k \leq \frac{7\lambda_{k+1}^* + \lambda_k}{7\lambda_k^* + \lambda_{k+1}^*}$. Therefore, the requirement on t becomes

$$t \geq \left(\log \frac{7\lambda_k^* + \lambda_{k+1}^*}{7\lambda_{k+1}^* + \lambda_k^*} \right)^{-1} \left(3(K-k) + \log \frac{140c_0K}{\mathcal{T}_{K,\min}^2} \right)$$

for all $k \in [K]$, which is provided as an assumption of the Theorem. This finishes the proof.

C. Auxiliary Lemmas

Lemma C.1. *Let $\{\Sigma_k^*\}_{k=1}^d$ be defined in (4). For all $k \in [d]$, we have*

$$\Sigma_k^* = \sum_{j=k}^d \lambda_j^* \mathbf{u}_j^* \mathbf{u}_j^{*\top}; \quad \Sigma_k^* \mathbf{u}_\ell^* = \begin{cases} \lambda_\ell^* \mathbf{u}_\ell^* & \text{if } \ell \geq k \\ \mathbf{0} & \text{if } \ell < k \end{cases}$$

Proof. We shall prove the statement by induction. For the base case, when $k = 1$, we have $\Sigma_k^* = \Sigma$, and, by the form of the eigen-decomposition, we have

$$\Sigma_k^* = \sum_{j=1}^d \lambda_j^* \mathbf{u}_j^* \mathbf{u}_j^{*\top}; \quad \mathbf{u}_k^* \mathbf{u}_k^{*\top} \Sigma_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top} = \lambda_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top}$$

Now, suppose that we have

$$\Sigma_k^* = \sum_{j=k}^d \lambda_j^* \mathbf{u}_j^* \mathbf{u}_j^{*\top}$$

We can see that

$$\mathbf{u}_k^* \mathbf{u}_k^{*\top} \Sigma_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top} = \sum_{j=k}^d \lambda_j^* (\mathbf{u}_k^{*\top} \mathbf{u}_j^*)^2 \mathbf{u}_k^* \mathbf{u}_k^{*\top} = \lambda_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top}$$

where the last equality follows from the orthogonality between \mathbf{u}_k^* and \mathbf{u}_j^* when $j \neq k$. Therefore

$$\Sigma_{k+1}^* = \sum_{j=k}^d \lambda_j^* \mathbf{u}_j^* \mathbf{u}_j^{*\top} - \lambda_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top} = \sum_{j=k+1}^d \lambda_j^* \mathbf{u}_j^* \mathbf{u}_j^{*\top}$$

which shows the inductive step and thus proves the first statement. To see the second statement, we notice that $\mathbf{u}_j^{*\top} \mathbf{u}_\ell^* = 0$ for all $j \neq \ell$ and $\mathbf{u}_j^{*\top} \mathbf{u}_\ell^* = 0$ if $j = \ell$. Therefore

$$\Sigma_k^* \mathbf{u}_\ell^* = \sum_{j=k}^d \mathbf{u}_j^{*\top} \mathbf{u}_\ell^* \cdot \lambda_j^* \mathbf{u}_j^* = \begin{cases} \lambda_\ell^* \mathbf{u}_\ell^* & \text{if } \ell \geq k \\ \mathbf{0} & \text{if } \ell < k \end{cases}$$

□

Lemma C.2. *Consider the power iteration procedure in (14). Let $\{\sigma_j\}_{j=1}^r$ be the eigenvalues of \mathbf{M} in (14). Then we have*

$$\min_{s \in \{\pm 1\}} \|s \cdot \mathbf{x}_t - \mathbf{a}_1\|_2^2 \leq 2 \left(\frac{\sigma_2}{\sigma_1} \right)^t$$

Proof. For convenience, we define the sequence $\{\hat{\mathbf{x}}_k\}_{k=0}^\infty$ as

$$\hat{\mathbf{x}}_t = \mathbf{M}\hat{\mathbf{x}}_{t-1}; \quad \hat{\mathbf{x}}_0 = \mathbf{x}_0$$

Then \mathbf{x}_t from (14) can be written as

$$\mathbf{x}_t = \frac{\hat{\mathbf{x}}_t}{\|\hat{\mathbf{x}}_t\|_2} = \frac{\hat{\mathbf{x}}_t}{\|\mathbf{M}^t \mathbf{x}_0\|_2}$$

Let $\{\mathbf{a}_j\}_{j=1}^d$ be an extended orthogonal basis of $\{\mathbf{a}_j\}_{j=1}^r$. Then we can write

$$\mathbf{x}_0 = \sum_{j=1}^d \mathbf{x}_0^\top \mathbf{a}_j \cdot \mathbf{a}_j$$

Let $\sigma_j = 0$ for $j \geq r + 1$. Then we have that $\mathbf{M}\mathbf{a}_j = \sigma_j \mathbf{a}_j$. Therefore

$$\mathbf{M}^t \mathbf{x}_0 = \sum_{j=1}^d \mathbf{x}_0^\top \mathbf{a}_j \cdot \mathbf{M}^t \mathbf{a}_j = \sum_{j=1}^d \mathbf{x}_0^\top \mathbf{a}_j \cdot \sigma_j^t \mathbf{a}_j = \sum_{j=1}^r \mathbf{x}_0^\top \mathbf{a}_j \cdot \sigma_j^t \mathbf{a}_j$$

Since $\|\mathbf{M}^t \mathbf{x}_0\|_2 \leq \sigma_1^t$, we must have that

$$|\mathbf{x}_t^\top \mathbf{a}_j| = \frac{|\hat{\mathbf{x}}_t^\top \mathbf{a}_j|}{\|\mathbf{M}^t \mathbf{x}_0\|_2} \leq \left(\frac{\sigma_j}{\sigma_1}\right)^t |\mathbf{x}_0^\top \mathbf{a}_j|$$

Since $\{\mathbf{a}_j\}_{j=1}^d$ is an orthogonal basis of \mathbb{R}^d , we must have that $\|\mathbf{x}_t\|_2^2 = \sum_{j=1}^d (\mathbf{x}_t^\top \mathbf{a}_j)^2$. Thus

$$(\mathbf{x}_t^\top \mathbf{a}_1)^2 = 1 - \sum_{j=2}^d (\mathbf{x}_t^\top \mathbf{a}_j)^2 \geq 1 - \sum_{j=2}^d \left(\frac{\sigma_j}{\sigma_1}\right)^{2t} (\mathbf{x}_0^\top \mathbf{a}_j)^2 \geq 1 - \left(\frac{\sigma_2}{\sigma_1}\right)^{2t}$$

where the last inequality follows from $\sigma_j \leq \sigma_2$ for $j \geq 2$ and $\sum_{j=1}^d (\mathbf{x}_0^\top \mathbf{a}_j)^2 = \|\mathbf{x}_0\|_2^2 = 1$. Therefore, we can conclude that

$$\min_{s \in \{\pm 1\}} \|s \cdot \mathbf{x}_t - \mathbf{a}_1\|_2^2 \leq 2 \left(1 - (\mathbf{x}_t^\top \mathbf{a}_1)^2\right) \leq 2 \left(\frac{\sigma_2}{\sigma_1}\right)^{2t}$$

Taking a square root for both sides gives the desired result. □

Lemma C.3. Consider a sequence of quantities $\{Q_k\}_{k=1}^\infty$ satisfying

$$Q_k = a_k + b_k \sum_{k'=1}^{k-1} Q_{k'}$$

with some $a_k, b_k \geq 0$ for all $k \in \mathbb{Z}_+$. Setting $Q_1 = a_1$, then we shall have that

$$Q_k = a_k + b_k \sum_{k'=1}^{k-1} a_{k'} \prod_{\ell=k'+1}^{k-1} (1 + b_\ell)$$

Proof. We shall prove by induction start from $Q_1 = a_1$. Suppose that Q_1, \dots, Q_k satisfies the characterization. To start, we can see that

$$\sum_{k'=1}^{k-1} Q_{k'} = \frac{1}{b_k} (Q_k - a_k) = \sum_{k'=1}^{k-1} a_{k'} \prod_{\ell=k'+1}^{k-1} (1 + b_\ell)$$

Then, by the inductive hypothesis, we have

$$\begin{aligned}
 \sum_{k'=1}^k Q_{k'} &= Q_k + \sum_{k'=1}^{k-1} Q_{k'} \\
 &= a_k + (b_k + 1) \sum_{k'=1}^{k-1} a_{k'} \prod_{\ell=k'+1}^{k-1} (1 + b_\ell) \\
 &= a_k + \sum_{k'=1}^{k-1} a_{k'} \prod_{\ell=k'+1}^k (1 + b_\ell) \\
 &= \sum_{k'=1}^k a_{k'} \prod_{\ell=k'+1}^k (1 + b_\ell)
 \end{aligned}$$

This implies that

$$Q_{k+1} = a_{k+1} + b_{k+1} \sum_{k'=1}^k Q_{k'} = a_{k+1} + b_{k+1} \sum_{k'=1}^k a_{k'} \prod_{\ell=k'+1}^k (1 + b_\ell)$$

which finishes the proof. \square

Lemma C.4. Consider a sequence of quantities $\{Q_k\}_{k=1}^\infty$ satisfying

$$Q_{k+1} = a_k Q_k + b_k$$

with some $a_k, b_k \geq 0$ for all $k \in \mathbb{Z}_+$. Set $b_0 = Q_1$. Then we have that

$$Q_k = \sum_{k'=0}^{k-1} b_{k'} \prod_{j=k'+1}^{k-1} a_j$$

Proof. We shall prove by induction. For the base case, let $k = 1$. In this case, we have that

$$Q_1 = \sum_{k'=0}^0 b_{k'} \prod_{j=k'+1}^0 a_j = b_0 = Q_1$$

For the inductive case, assume that the property holds for k . Then we have that

$$Q_{k+1} = a_k Q_k + b_k = a_k \cdot \sum_{k'=0}^{k-1} b_{k'} \prod_{j=k'+1}^{k-1} a_j + b_k = \sum_{k'=0}^k b_{k'} \prod_{j=k'+1}^k a_j$$

This proves the inductive step and finishes the proof. \square

Lemma C.5. Let $\{p_k\}_{k=1}^\infty$ be a positive increasing sequence, and $g, t > 0$. Let the series S_k be defined as

$$S_k = \sum_{k'=1}^{k-1} \frac{p_{k'}^t}{g^{k'}}$$

As long as $t > \frac{\log g}{\log p_{k'} - \log p_{k'+1}}$ for all $k' \in [k-1]$, we shall have that

$$S_k \leq \left(1 - \max_{k' \in [k-1]} g \left(\frac{p_{k'}}{p_{k'+1}} \right)^t \right)^{-1} \frac{p_{k-1}^t}{g^{k-1}}$$

Proof. Set $\gamma = \max_{k' \in [k-1]} g \left(\frac{p_{k'}}{p_{k'+1}} \right)^t$. Since $t > \frac{\log g}{\log p_{k'+1} - \log p_{k'}}$ for all $k' \in [k-1]$, we must have that $\gamma < 1$. Moreover, by definition of γ , it must hold that for all $k' \in [k-1]$

$$\frac{p_{k'}^t}{g^{k'}} \leq \gamma \cdot \frac{p_{k'+1}^t}{g^{k'+1}}$$

Thus, we have that for all $k' \in [k-1]$

$$\frac{p_{k'}^t}{g^{k'}} \leq \gamma^{k-k'-1} \cdot \frac{p_{k-1}^t}{g^{k-1}}$$

Therefore S_k can be upper-bounded as

$$S_k \leq \sum_{k'=1}^{k-1} \gamma^{k-k'-1} \cdot \frac{p_{k-1}^t}{g^{k-1}} = \frac{p_{k-1}^t}{g^{k-1}} \sum_{k'=0}^{k-2} \gamma^{k'} \leq \frac{p_{k-1}^t}{g^{k-1}(1-\gamma)}$$

□

Lemma C.6. Let \mathbf{x}_t be the result of running power iteration starting from \mathbf{x}_0 for t iterations, as defined in (14). Let σ_j, \mathbf{a}_j be the j -th eigenvalue and eigenvector of \mathbf{M} in (14). Assume that there exists $c_0 > 0$ such that $|\mathbf{x}_0^\top \mathbf{a}_1| \geq c_0^{-1}$. Then for all $j = 2, \dots, r$, we have:

$$|\mathbf{x}_t^\top \mathbf{a}_j| = |(\mathbf{x}_t - \mathbf{a}_1)^\top \mathbf{a}_j| \leq c_0 \left(\frac{\sigma_j}{\sigma_1} \right)^t.$$

Proof. We notice that for $j \geq 2$, it holds that $(\mathbf{x}_t - \mathbf{a}_1)^\top \mathbf{a}_j = \mathbf{x}_t^\top \mathbf{a}_j$. Since \mathbf{M} is symmetric, we have that

$$\mathbf{x}_t^\top \mathbf{a}_j = \alpha \mathbf{x}_0^\top \mathbf{M}^t \mathbf{a}_j = \alpha \sigma_j^t \mathbf{x}_0^\top \mathbf{a}_j = \frac{\sigma_j^t}{\|\mathbf{M}^t \mathbf{x}_0\|_2} \mathbf{x}_0^\top \mathbf{a}_j$$

This brings the focus to $\|\mathbf{M}^t \mathbf{x}_0\|_2$. Let r be the rank of \mathbf{M} . Let $\{\mathbf{a}_j\}_{j=1}^d$ be an extended orthogonal basis of $\{\mathbf{a}_j\}_{j=1}^r$. Then we can write

$$\mathbf{x}_0 = \sum_{j=1}^d \mathbf{x}_0^\top \mathbf{a}_j \cdot \mathbf{a}_j$$

Let $\sigma_j = 0$ for $j \geq r+1$. Then we have that $\mathbf{M} \mathbf{a}_j = \sigma_j \mathbf{a}_j$. Therefore

$$\mathbf{M}^t \mathbf{x}_0 = \sum_{j=1}^d \mathbf{x}_0^\top \mathbf{a}_j \cdot \mathbf{M}^t \mathbf{a}_j = \sum_{j=1}^d \mathbf{x}_0^\top \mathbf{a}_j \cdot \sigma_j^t \mathbf{a}_j = \sum_{j=1}^r \mathbf{x}_0^\top \mathbf{a}_j \cdot \sigma_j^{*t} \mathbf{a}_j$$

where the last equality follows from $\sigma_j^t = 0$ when $t > 0$. Therefore, according to the Pythagorean Theorem,

$$\|\mathbf{M}^t \mathbf{x}_0\|_2 = \left(\sum_{j=1}^r \sigma_j^{2t} (\mathbf{x}_0^\top \mathbf{a}_j)^2 \right)^{\frac{1}{2}} \geq \sigma_1^t |\mathbf{x}_0^\top \mathbf{a}_1|$$

Recall that $|\mathbf{x}_0^\top \mathbf{a}_1| \geq c_0^{-1} > 0$. Thus, we have

$$|(\mathbf{x}_t - \mathbf{a}_1)^\top \mathbf{a}_j| = \frac{\sigma_j^t}{\|\mathbf{M}^t \mathbf{x}_0\|_2} |\mathbf{x}_0^\top \mathbf{a}_j| \leq \left(\frac{\sigma_j}{\sigma_1} \right)^t \cdot \frac{|\mathbf{x}_0^\top \mathbf{a}_j|}{|\mathbf{x}_0^\top \mathbf{a}_1|} \leq |\mathbf{x}_0^\top \mathbf{a}_1|^{-1} \left(\frac{\sigma_j}{\sigma_1} \right)^t \leq c_0 \left(\frac{\sigma_j}{\sigma_1} \right)^t$$

□

D. Experiment Details of Figure 1

We conduct spectral clustering (Belkin & Niyogi, 2003) on a subset of the MNIST dataset (Deng, 2012), namely $\mathcal{S} = \{\mathbf{x}_i\}_{i=1}^n$. In particular, our experiment starts with building a weighted graph $\mathcal{G} = (V, E)$ where each node in the graph \mathcal{G} represents a sample image, and each edge weight is chosen according to

$$e_{ij} = \begin{cases} \exp\left(-\frac{1}{2} \|\mathbf{x}_i - \mathbf{x}_j\|_2^2\right) & \text{if } \mathbf{x}_i \in r\text{NN}(\mathbf{x}_j) \text{ or } \mathbf{x}_j \in r\text{NN}(\mathbf{x}_i) \\ 0 & \text{otherwise} \end{cases}$$

Here $r\text{NN}(\mathbf{x}_i)$ represents the set of r Nearest Neighbors of \mathbf{x}_i . To be more specific, two nodes in the graph are considered adjacent if the two corresponding samples are top- r nearest neighbors, and the weight of the edge between the two nodes is computed using the RBF kernel with unit variance. Next, we compute the degree-normalized graph Laplacian as our similarity matrix

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$$

where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is the diagonal degree matrix, and $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the adjacency matrix. We extract the top- k eigenvectors of the similarity matrix \mathbf{L} with the deflation algorithm in Algorithm 1 with t power iteration steps in each call to the PCA sub-routine, which results in the approximate eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_k$. We form a matrix $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_k] \in \mathbb{R}^{n \times k}$, and assign the i th sample with a new feature vector \mathbf{v}_i as the i th row of \mathbf{U} . Lastly, we perform k' -means clustering with the newly extracted features $\{\mathbf{v}_i\}_{i=1}^n$, forming clusters $C = \{C_j\}_{j=1}^{k'}$, where $C_j \subseteq [n]$. The clustering result is evaluated using the mutual information metric (Hubert & Arabie, 1985).

$$\text{MI}(C, C^*) = \sum_{j, j'=1}^{k'} \mathbb{P}(C_j \cap C_{j'}^*) \log \frac{\mathbb{P}(C_j \cap C_{j'}^*)}{\mathbb{P}(C_j) \mathbb{P}(C_{j'}^*)} \quad (62)$$

where $\mathbb{P}(C_j) = \frac{|C_j|}{n}$ and, similarly, $\mathbb{P}(C_j \cap C_{j'}^*) = \frac{|C_j \cap C_{j'}^*|}{n}$. Here $C^* = \{C_j^*\}_{j=1}^{k'}$ is the ground-truth clustering with $C_j^* = \{i \in [n] : \text{the label of sample } i \text{ is } j\}$.

In our experiment, we use a subset of MNIST with $n = 1000$ randomly drawn samples. When constructing the graph, we use $r = 10$, that is, we connect each node to ten of its nearest neighbors. We cluster the samples into $k' = 10$ clusters. In Figure 1, the x-axis represents the number of power iteration steps t in each call to the sub-routine PCA, and the y-axis represents the mutual information score in (62). Each line in the figure represents different k , which is the number of eigenvectors extracted from \mathbf{L} .