

Sparse GCA and Thresholded Gradient Descent

Sheng Gao

Zongming Ma

Department of Statistics and Data Science

University of Pennsylvania

Philadelphia, PA 19104, USA

SHENGGAO@WHARTON.UPENN.EDU

ZONGMING@WHARTON.UPENN.EDU

Editor: Ji Zhu

Abstract

Generalized correlation analysis (GCA) is concerned with uncovering linear relationships across multiple data sets. It generalizes canonical correlation analysis that is designed for two data sets. We study sparse GCA when there are potentially multiple leading generalized correlation tuples in data that are of interest and the loading matrix has a small number of nonzero rows. It includes sparse CCA and sparse PCA of correlation matrices as special cases. We first formulate sparse GCA as a generalized eigenvalue problem at both population and sample levels via a careful choice of normalization constraints. Based on a Lagrangian form of the sample optimization problem, we propose a thresholded gradient descent algorithm for estimating GCA loading vectors and matrices in high dimensions. We derive tight estimation error bounds for estimators generated by the algorithm with proper initialization. We also demonstrate the prowess of the algorithm on a number of synthetic data sets.

Keywords: sparse GCA, sparse CCA, thresholded gradient descent, non-convex optimization, generalized eigenvalue problem

1. Introduction

With the advent of big data acquisition technology, it has become increasingly important to integrate information across multiple data sets collected on a common set of subjects. Canonical correlation analysis (CCA), first proposed by Hotelling (1992), is a widely used statistical tool to integrate information from two data sets: It seeks linear combinations of variables within each data set such that their correlation is maximized.

However, recent advances in fields such as multi-omics and multimodal brain imaging have presented us with new challenges, since scientists are often able to collect more than two data sets on the same set of subjects nowadays. To tackle these challenges, we turn to a useful generalization of CCA called *generalized correlation analysis* (GCA) (Kettenring, 1971) which aims to explore linear relationships across multiple data sources. Kettenring (1971) proposed five different variants for generalized correlation analysis of multiple data sets, where different methods correspond to maximization of different objective functions of covariances and correlations, subject to certain normalization constraints. Tenenhaus and Tenenhaus (2011) extended these approaches to regularized versions by adding l_2 penalties to the objective functions and proposed partial least squares algorithms for solving them.

In addition to unsupervised settings, it has been shown that GCA can be incorporated into supervised learning settings to improve efficiency and accuracy (Shen et al., 2014).

Suppose that the k data sets are i.i.d. realizations of k random vectors $X_{\{1\}} \in \mathbb{R}^{p_1}$, $X_{\{2\}} \in \mathbb{R}^{p_2}$, ..., $X_{\{k\}} \in \mathbb{R}^{p_k}$. Here and after, we use subscript $\{j\}$ to denote the j th set of features and/or the j th data set. At the population level, we propose to seek vectors $a_{\{1\}}, a_{\{2\}}, \dots, a_{\{k\}}$ (called *the first k -tuple of generalized loading vectors*) that solve

$$\begin{aligned} & \underset{l_{\{1\}}, \dots, l_{\{k\}}}{\text{maximize}} && \sum_{i=1}^k \sum_{j=1}^k \text{cov}(l_{\{i\}}^\top X_{\{i\}}, l_{\{j\}}^\top X_{\{j\}}) \\ & \text{subject to} && \sum_{i=1}^k \text{var}(l_{\{i\}}^\top X_{\{i\}}) = 1. \end{aligned} \tag{1}$$

An important difference between the foregoing formulation and those in Kettenring (1971) is the normalization constraint. We propose to normalize the sum of the variances of the linear combinations while Kettenring (1971) requires the individual variance of each linear combination to be one. When $k = 2$, the two normalizations essentially lead to the same solution. One can show that the optimal solution to (1) necessarily has $\text{var}(a_{\{i\}}^\top X_{\{i\}}) = \frac{1}{2}$ for $i = 1$ and 2 , and hence the optimal $a_{\{i\}}$'s are proportional to those from normalizing individual variances. However, when $k \geq 3$, the solutions are usually different and we argue that formulation (1) is more appropriate for two reasons. First, the normalization in (1) conforms with the ‘‘null principle’’ (Donnell et al., 1994) in that if all linear combinations of the X_i 's are uncorrelated then the objective function reduces to the lefthand side of the normalization constraint. In addition, it is more resilient to spurious solutions as illustrated by the following example.

Example 1 *Suppose $k = 3$ and $p_1 = p_2 = p_3 = m$. Let $X_{\{1\}} = Yv_{\{1\}} + Z_{\{1\}}$, $X_{\{2\}} = Yv_{\{2\}} + Z_{\{2\}}$ and $X_{\{3\}} = Z_{\{3\}}$ where $v_{\{1\}}$ and $v_{\{2\}}$ are two deterministic unit vectors in \mathbb{R}^m , $Y \sim N(0, 1)$, $Z_{\{i\}} \stackrel{iid}{\sim} N_m(0, I_m)$, and they are mutually independent. Simple linear algebra shows that the optimal solution to (1) is $a_{\{1\}} = \frac{1}{2}v_{\{1\}}$, $a_{\{2\}} = \frac{1}{2}v_{\{2\}}$ and $a_{\{3\}} = 0$. In contrast, with the individual normalization constraint $\text{var}(l_{\{i\}}^\top X_{\{i\}}) = 1$, the optimal solution would change to $a_{\{1\}} = \frac{1}{\sqrt{2}}v_{\{1\}}$, $a_{\{2\}} = \frac{1}{\sqrt{2}}v_{\{2\}}$ and $a_{\{3\}}$ is any vector in \mathbb{B}^{m-1} where \mathbb{B}^{m-1} is the unit ball in \mathbb{R}^m . Comparing the solutions to two different normalizations, formulation (1) is advantageous in that it provides a meaningful and unique optimal $a_{\{3\}}$.*

The solution to (1) gives the first k -tuple of leading GCA loading vectors. More generally, we want to extract successive k -tuples of leading loading vectors subject to certain additional constraints. In view of (1), suppose we have found $(a_{\{1\}}^{(m)}, \dots, a_{\{k\}}^{(m)})$, for $m = 1, \dots, r - 1$, as the first $(r - 1)$ k -tuples of leading loading vectors. We define the r th k -tuple $(a_{\{1\}}^{(r)}, \dots, a_{\{k\}}^{(r)})$ as the solution to (1) with the following additional constraints:

$$\sum_{i=1}^k \text{cov}((l_{\{i\}}^{(m)})^\top X_{\{i\}}, l_{\{i\}}^\top X_{\{i\}}) = 0, \quad m = 1, \dots, r - 1.$$

To be more concise, let the columns of $L_{\{i\}} \in \mathbb{R}^{p_i \times r}$ be the successive leading loading vectors for $X_{\{i\}}$ for $i = 1, \dots, k$. The foregoing proposal for finding the first r leading k -tuples then reduces to the following optimization problem

$$\begin{aligned} & \underset{L_{\{1\}}, \dots, L_{\{k\}}}{\text{maximize}} && \sum_{i,j} \text{Tr}(L_{\{i\}}^\top \Sigma_{\{ij\}} L_{\{j\}}) \\ & \text{subject to} && \sum_{i=1}^k L_{\{i\}}^\top \Sigma_{\{ii\}} L_{\{i\}} = I_r. \end{aligned} \tag{2}$$

Here $\Sigma_{\{ii\}}$ denotes the covariance matrix of $X_{\{i\}}$ and $\Sigma_{\{ij\}} = \text{cov}(X_{\{i\}}, X_{\{j\}})$ for $i \neq j$. To write (2) more concisely, we define $L = [L_{\{1\}}^\top, \dots, L_{\{k\}}^\top]^\top \in \mathbb{R}^{p \times r}$ where $p = \sum_{i=1}^k p_i$, $\Sigma_0 = \text{diag}(\Sigma_{\{11\}}, \dots, \Sigma_{\{kk\}}) \in \mathbb{R}^{p \times p}$, and $\Sigma \in \mathbb{R}^{p \times p}$ where its (i, j) th block is $\Sigma_{\{ij\}}$ and i th diagonal block is $\Sigma_{\{ii\}}$. Then we can rewrite (2) as

$$\begin{aligned} & \underset{L}{\text{maximize}} && \text{Tr}(L^\top \Sigma L) \\ & \text{subject to} && L^\top \Sigma_0 L = I_r. \end{aligned} \tag{3}$$

In what follows, $A = [A_{\{1\}}^\top, \dots, A_{\{k\}}^\top]^\top$ denotes the solution to (3) where each $A_{\{i\}} \in \mathbb{R}^{p_i \times r}$ and the m th column of $A_{\{i\}}$ is $a_{\{i\}}^{(m)}$ for $i = 1, \dots, k$ and $m = 1, \dots, r$.

Informed readers might have already realized that (3) is closely connected to a generalized eigenvalue problem $|\Sigma - \lambda \Sigma_0| = 0$, where $|M|$ stands for the determinant of a square matrix M . In addition, when Σ_0 is diagonal (i.e., when $k = p$), the problem is equivalent to principal component analysis (PCA) of the correlation matrix: For any non-singular Σ_0 , the columns of $\Sigma_0^{-1/2} A$ are the leading eigenvectors of the correlation matrix $\Sigma_0^{-1/2} \Sigma \Sigma_0^{-1/2}$. When $k = 2$, the solution to (3) is equivalent to CCA up to scaling. Thus, program (3) provides a unified formulation for extracting leading linear covariation within one or across multiple data sets.

In practice, one does not have direct knowledge of the joint sample covariance matrix Σ or even its block diagonal part Σ_0 . The natural choice is to replace them with their sample counterparts. In potential modern applications of GCA, the data set dimensions (i.e., the p_i 's) can be much larger than the sample size n . Hence, one suffers from the curse of dimensionality if no further structural assumption is made (Donoho, 2000; Johnstone, 2001; Bickel and Levina, 2008a,b). Due to its interpretability and practicality, a structural assumption that has been widely adopted in both theory and practice is *sparsity*: Most energy of the solution to (1)–(3) concentrates on a small number of entries (Hastie et al., 2015; Wainwright, 2019). Let A denote the solution to (3). In this manuscript, we adopt the assumption that at most s rows of A contain nonzero entries. In other words, the target of estimation is also the solution to the following *sparse generalized correlation analysis* (Sparse GCA, or SGCA) problem:

$$\begin{aligned} & \underset{L}{\text{maximize}} && \text{Tr}(L^\top \Sigma L) \\ & \text{subject to} && L^\top \Sigma_0 L = I_r \quad \|L\|_{2,0} \leq s. \end{aligned} \tag{4}$$

Here and after, for any matrix L , $\|L\|_{2,0}$ counts the number of nonzero rows in L . In view of the discussion following (3), when $k = p$ and $k = 2$, (4) reduces to sparse PCA of correlation matrix and sparse CCA, respectively.

1.1 Main contributions

The main contributions of the present manuscript are the following.

First, we clarify the target of estimation in generalized correlation analysis (i.e., the solution A to (3)) by considering a natural latent variable model in which an r -dimensional latent variable drives the covariation of k random vectors. Under mild conditions, we show that the linear subspace spanned by the leading r generalized eigenvectors in our GCA formulation (i.e., the columns of A) coincides with the subspace spanned by a concrete functional of parameters in the latent variable model. In addition, we characterize the behavior of generalized eigenvalues under the latent variable model.

Next, for sample sparse GCA, we propose a thresholded gradient descent algorithm for solving a Lagrangian version of the sample counterpart of (4). The algorithm is intuitive and easy to implement. In view of the discussion following (3), such an algorithm provides a unified approach to a number of different sparse unsupervised learning problems, including sparse PCA of correlation matrices and sparse CCA.

Furthermore, we provide a theoretical analysis of the thresholded gradient descent algorithm. We show that with high probability, when initialized properly, estimation errors of the intermediate results after each iteration converge at a geometric rate until they arrive in a tight neighborhood of the population solution to (4). Statistically, we establish tight estimation error bounds for the output of our algorithm as an estimator. Numerically, this implies geometric convergence of our algorithm to such an estimator. Finally, these theoretical findings are corroborated by numerical studies on simulated data sets.

1.2 Related works

In view of the discussion following (3), the present paper is closely connected to the sparse PCA and sparse CCA literature. To date, there is a large literature devoted to various aspects of the sparse PCA problem, including algorithms (Zou et al., 2006; Johnstone and Lu, 2009; Amini and Wainwright, 2009; Witten et al., 2009; Ma, 2013; Yuan and Zhang, 2013; Vu et al., 2013; Wang et al., 2014), information-theoretic limits (Birnbaum et al., 2013; Cai et al., 2013; Vu and Lei, 2012) and computational theoretic limits (Berthet and Rigollet, 2013; Gao et al., 2017). However, this literature has mostly focused on sparse PCA of covariance matrices. Thus, the special case of $k = p$ in our setting complements the existing literature by providing both theory and method for sparse PCA of correlation matrices.

In the case of $k = 2$, the theory and method in this paper specialize to the sparse CCA setting (Chen et al., 2013; Gao et al., 2015a, 2017). In this case, we provide a new iterative algorithm for sparse CCA that achieves optimal estimation rates derived in Gao et al. (2015a). Therefore, for this special case, the present manuscript provides a competitive alternative to existing sparse CCA methods.

When $r = 1$, (4) reduces to the population version of the sparse generalized eigenvalue problem considered in Tan et al. (2018). Tan et al. (2018) proposed a truncated Rayleigh

flow method for estimating the population solution in this special case, and established its rates of convergence. It is not clear how their method can be generalized to estimate successive generalized eigenvectors or eigenspaces. In contrast, our estimator is based on a different algorithm (thresholded gradient descent of a Lagrangian objective function). It not only achieves fast converging estimation rates for estimating the first generalized eigenvector but also works for estimating leading generalized eigenspaces of fixed dimensions that are greater than one.

1.3 Organization of the paper

The rest of this paper is organized as follows. In section 2 we examine the generalized eigenvalue problem underpinning sparse GCA under a latent variable model. Section 3 proposes a thresholded gradient descent algorithm and its initialization via generalized Fantope projection. Section 4 establishes the convergence rate of our algorithm under reasonable initialization. Numerical results are presented in Section 5. Technical proofs are deferred to appendices.

1.4 Notation

For any set S , let $|S|$ denote its size and S^c denote its complement. For any event E , $\mathbf{1}_E$ is its indicator function. For a vector u , $\|u\| = (\sum_i u_i^2)^{1/2}$, $\|u\|_0 = \sum_i \mathbf{1}_{u_i \neq 0}$, $\|u\|_1 = \sum_i |u_i|$, $\|u\|_\infty = \max_i |u_i|$. For any matrix $A = (a_{ij})$, the i th row of A is denoted by A_{i*} and the j th column by A_{*j} . We let $\text{Col}(A)$ denote the span of columns of A . For a positive integer m , $[m]$ denotes the index set $1, 2, \dots, m$. For two subsets I and J of indices, we write A_{IJ} for the $|I| \times |J|$ submatrices formed by a_{ij} with $(i, j) \in I \times J$. When I or J is the whole set, we abbreviate it with an $*$, and so if $A \in \mathbb{R}^{p \times k}$, then $A_{I*} = A_{I[k]}$ and $A_{*J} = A_{[p]J}$. For any square matrix $A = (a_{ij})$, denote its trace by $\text{Tr}(A) = \sum_i a_{ii}$. Moreover, we denote the collection of all $p \times r$ matrices with orthonormal columns by $\mathcal{O}(p, r)$, and abbreviate $\mathcal{O}(r, r)$ by $\mathcal{O}(r)$. The set of $p \times p$ symmetric matrices is denoted by \mathbb{S}^p . Furthermore, $\sigma_i(A)$ stands for the i th largest singular value of A and $\sigma_{\max}(A) = \sigma_1(A)$, $\sigma_{\min}(A) = \sigma_{\min\{p,k\}}(A)$. The Frobenius norm and the operator norm of A are $\|A\|_F = (\sum_{i,j} a_{ij}^2)^{1/2}$ and $\|A\|_{\text{op}} = \sigma_1(A)$, respectively. The infinity norm of A is defined as $\|A\|_\infty = \max_{i,j} |a_{ij}|$. The l_1 norm and the nuclear norm of a matrix A are $\|A\|_1 = \sum_{i,j} |a_{ij}|$ and $\|A\|_* = \sum_i \sigma_i(A)$, respectively. Similarly, $s_i(A)$ stands for the i th largest eigenvalue of $A \in \mathbb{R}^{p \times p}$, while $s_{\max}(A) = s_1(A)$ and $s_{\min}(A) = s_p(A)$. The support of A is defined as $\text{supp}(A) = \{i \in [n] : \|A_{i*}\| > 0\}$. For two symmetric matrices A and B , we write $A \preceq B$ if $B - A$ is positive semidefinite. For any positive semi-definite matrix A , $A^{1/2}$ denotes its principal square root that is positive semi-definite and satisfies $A^{1/2}A^{1/2} = A$. The trace inner product of two matrices $A, B \in \mathbb{R}^{p \times k}$ is $\langle A, B \rangle = \text{Tr}(A^\top B)$. For any real numbers a and b , let $a \wedge b = \min(a, b)$ and $a \vee b = \max(a, b)$. For any two sequences of positive numbers $\{a_n\}$ and $\{b_n\}$, we write $a_n = O(b_n)$ if $\limsup_{n \rightarrow \infty} a_n/b_n$ is finite. Given a random element X , $\mathcal{L}(X)$ denotes its probability distribution. The symbol C and its variants C_1, C' , etc. are generic positive constants and may vary from occurrence to occurrence, unless otherwise specified. The symbols \mathbb{P} and \mathbb{E} stand for generic probability and expectation when the distribution is clear from the context.

2. A Latent Variable Model

In this part, we aim to identify the solutions to (3) as functionals of parameters in the joint distribution of $(X_{\{1\}}^\top, \dots, X_{\{k\}}^\top)^\top$. To this end, we introduce an intuitive latent variable model.

The fundamental assumption underlying generalized correlation analysis is that there exists a shared low-dimensional latent variable which orchestrates the (linear) covariation of observed features across all data sets. Let z be an r -dimensional latent variable. The following latent variable model describes an idealized data generating process:

$$\begin{aligned} X_{\{i\}} &= U_{\{i\}}z + e_{\{i\}}, \quad i = 1, \dots, k, \\ z &\sim N_r(0, I_r), \quad e_{\{i\}} \stackrel{\text{ind}}{\sim} N_{p_i}(0, \Psi_{\{ii\}}). \end{aligned} \quad (5)$$

Here the deterministic matrix $U_{\{i\}} \in \mathbb{R}^{p_i \times r}$ for $i = 1, \dots, k$, $\Psi_{\{ii\}}$ is positive definite, and the latent variable z and the idiosyncratic noises $\{e_{\{i\}}\}_{i=1}^k$ are mutually independent. Under the foregoing latent variable model, the joint covariance matrix of $(X_{\{1\}}^\top, \dots, X_{\{k\}}^\top)^\top \in \mathbb{R}^p$ is given by Σ with the i th diagonal block

$$\Sigma_{\{ii\}} = U_{\{i\}}U_{\{i\}}^\top + \Psi_{\{ii\}}, \quad \text{for } i = 1, \dots, k, \quad (6)$$

and the (i, j) th block

$$\Sigma_{\{ij\}} = U_{\{i\}}U_{\{j\}}^\top, \quad \text{for } 1 \leq i \neq j \leq k. \quad (7)$$

We denote by Ψ the block diagonal matrix with blocks $\Psi_{\{ii\}}$, $i = 1, \dots, k$, on the diagonal. We also let $U = [U_{\{1\}}^\top, U_{\{2\}}^\top, \dots, U_{\{k\}}^\top]^\top \in \mathbb{R}^{p \times r}$. In the rest of this section, we assume that the observed data sets are generated by model (5). In addition, we let

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 0 \quad (8)$$

denote the population generalized eigenvalues of Σ with respect to $\Sigma_0 = \text{diag}(\Sigma_{\{11\}}, \dots, \Sigma_{\{kk\}})$.

The following key lemma identifies the connection between parameters in model (5) and the solution to (3).

Lemma 1 *Suppose that Σ and Σ_0 are specified by model (5)–(7). Let $A = [A_{\{1\}}^\top, \dots, A_{\{k\}}^\top]^\top$ be the solution to (3). If the r th generalized eigenvalue of Σ w.r.t. Σ_0 is larger than 1, i.e., $\lambda_r > 1$, then $\text{Col}(A_{\{i\}}) \subset \text{Col}(\Sigma_{\{ii\}}^{-1}U_{\{i\}})$ for all $i \in [k]$. For any i , if further $\text{rank}(A_{\{i\}}) = r$, then $\text{Col}(A_{\{i\}}) = \text{Col}(\Sigma_{\{ii\}}^{-1}U_{\{i\}})$.*

The next lemma describes the behavior of population generalized eigenvalues under latent variable model (5). It also identifies a sufficient condition for $\lambda_r > 1$. To this end, we start with an assumption motivated by Fan et al. (2019).

Assumption 2 *In latent variable model (5)–(7), we assume the following:*

- The matrix U has full column rank, that is, $\text{rank}(U) = r$;
- Σ_0 and U satisfy $\sigma_r(\Sigma_0^{-1/2}U) \geq 1$.

Lemma 3 *In model (5)–(7), under Assumption 2, we have*

$$r = \max\{j : \lambda_j > 1, j \in [p]\}. \quad (9)$$

In other words, there are exactly r generalized eigenvalues greater than 1. Moreover, define $Y = [U_{\{1\}}^\dagger, U_{\{2\}}^\dagger, \dots, U_{\{k\}}^\dagger]^\top \in \mathbb{R}^{p \times r}$ where $U_{\{i\}}^\dagger$ is the Moore-Penrose inverse of $U_{\{i\}}$. Then the multiplicity of 1 as a generalized eigenvalue is

$$\#\{j : \lambda_j = 1\} = p - \sum_{i=1}^k \text{rank}(U_{\{i\}}) + r - \text{rank}(U - UY^\top U). \quad (10)$$

The foregoing lemma guarantees an eigengap between the r th and the $(r + 1)$ th generalized eigenvalues, and so the leading rank r generalized eigenspace is well-defined at the population level.

By the foregoing lemma, under model (5) we can decompose Σ as

$$\Sigma = \Sigma_0 K \Lambda K^\top \Sigma_0 = \Sigma_0 A \Lambda_r A^\top \Sigma_0 + \Sigma_0 B \bar{\Lambda} B^\top \Sigma_0. \quad (11)$$

Here Λ , Λ_r , and $\bar{\Lambda}$ are all diagonal where $\Lambda_r = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_r)$ collects the first r generalized eigenvalues, $\bar{\Lambda}$ collects the remaining $p - r$ generalized eigenvalues, and $\Lambda = \text{diag}(\Lambda_r, \bar{\Lambda})$. In addition, B collects the eigenvectors associated with the bottom $p - r$ generalized eigenvalues. On the other hand, the decomposition (11) hold as long as the r th and $(r + 1)$ th generalized eigenvalues are distinct and hence is more general than model (5).

Remark 4 *In the case when $k = 2$, The covariance matrices between X and Y can be reparameterized as $\Sigma_{xy} = \Sigma_x V \Theta_r W^\top \Sigma_y$ with $V^\top \Sigma_x V = W^\top \Sigma_y W = I_r$ Gao et al. (2015a, 2017). Here $\Theta_r = \text{diag}(\theta_1, \theta_2, \dots, \theta_r)$ collects the leading r canonical coefficients for CCA and $\Theta_r + I_r = \Lambda_r$. Hence we have $\lambda_r = \theta_r + 1$. The solution $A = [A_{\{1\}}^\top, A_{\{2\}}^\top]^\top$ to (3) satisfies that $A_{\{1\}} = \frac{1}{\sqrt{2}}V, A_{\{2\}} = \frac{1}{\sqrt{2}}W$. Such a relationship does not hold in general for $k \geq 3$.*

Remark 5 *When $k = p$, Σ_0 becomes a diagonal matrix where the diagonal entries are variances of the variables. The problem (3) is then equivalent to finding the leading eigenspace of correlation matrix of the data, defined as $R = \Sigma_0^{-1/2} \Sigma \Sigma_0^{-1/2}$. Let $E_R \in \mathbb{R}^{p \times r}$ be the matrix containing the eigenvectors that span the leading r dimensional eigenspace. If A is the solution to (3), then $\Sigma_0^{1/2} A$ coincides with E_R (up to an $r \times r$ rotation matrix when there is any generalized eigenvalue with multiplicity larger than one). Thus, we essentially estimate the leading eigenspace of the correlation matrix.*

3. Gradient Descent with Hard Thresholding

In this section, we present a thresholded gradient descent algorithm for simultaneously finding multiple leading sparse generalized eigenvectors.

3.1 Motivation

The sample counterpart of (4) can be recast as the following minimization problem:

$$\begin{aligned} & \underset{L}{\text{minimize}} && -\langle \widehat{\Sigma}, LL^\top \rangle \\ & \text{subject to} && L^\top \widehat{\Sigma}_0 L = I_r, \quad \|L\|_{2,0} \leq s. \end{aligned}$$

Here $\widehat{\Sigma}$ is the joint sample covariance matrix of all k sets of features. We propose to solve its Lagrangian version:

$$\underset{L \in \mathbb{R}^{p \times r}}{\text{minimize}} \quad f(L) \quad \text{subject to} \quad \|L\|_{2,0} \leq s, \quad (12)$$

where the objective function $f(L)$ is

$$f(L) = -\langle \widehat{\Sigma}, LL^\top \rangle + \frac{\lambda}{2} \|L^\top \widehat{\Sigma}_0 L - I_r\|_{\text{F}}^2. \quad (13)$$

Effectively, minimizing the first term in (13) maximizes the original objective function whereas minimizing the second term controls the deviation from the normalization constraint in (3). Here λ is a tuning parameter. Intuitively, the larger it is, the more penalty we put on deviating from the normalization constraint and the closer $L^\top \widehat{\Sigma}_0 L$ is to I_r .

To deal with the constraint in (12), we shall perform the following hard thresholding.

Definition 6 *Given a matrix U and a natural number k , we define the output of hard thresholding function $HT(U, k)$ to be the matrix obtained by keeping the k rows with the largest l_2 norms and replacing all other rows with zeros, that is*

$$HT(U, k)_{i*} = \begin{cases} U_{i*}, & \text{if } i \in C_k \\ 0, & \text{otherwise} \end{cases}$$

where C_k is the index set of k rows of U with largest l_2 norms. When there is a tie, we always pick the smaller/smallest index.

3.2 Algorithm

Let $X^{(1)}, X^{(2)}, \dots, X^{(n)}$ be i.i.d. observations generated from the latent variable model with covariance matrix Σ and its block diagonal part Σ_0 . Let $\widehat{\Sigma}$ and $\widehat{\Sigma}_0$ be the sample covariance matrix and its block diagonal part computed on $\{X^{(1)}, X^{(2)}, \dots, X^{(n)}\}$.

Algorithm 1 describes the proposed thresholded gradient descent procedure. Given a proper initial estimator which we shall specify later, each iteration first performs a step of gradient descent on f defined in (12), then it keeps the s' rows with the largest l_2 norms and thresholds the remaining rows to zero. By iteratively performing these two steps, the algorithm can be viewed as heuristics for solving the non-convex optimization problem (4). Here s' is a user-specified tuning parameter and is not necessarily equal to the true sparsity level s . Step 2 in Algorithm 1 is a re-normalization step. This is due to the fact that the stationary point of the objective function $f(L)$ is not the same as the solution to the original optimization problem (4). Thus, we first transform it for iterations and eventually transform back in Step 7 to define the final estimator.

Algorithm 1: Thresholded gradient descent for sparse GCA

Input: Covariance matrix estimator $\widehat{\Sigma}$ and its block diagonal part $\widehat{\Sigma}_0$;
 Initialization \widehat{A}_0 .
Tuning Parameters: Step size η ; Penalty λ ; Sparsity level s' ; Number of iterations T ;
 1: $\widetilde{A}_0 \leftarrow \widehat{A}_0(\widehat{A}_0^\top \widehat{\Sigma}_0 \widehat{A}_0)^{-1/2}$
 2: $\widetilde{V}_1 \leftarrow \widetilde{A}_0(I + \frac{1}{\lambda} \widetilde{A}_0^\top \widehat{\Sigma} \widetilde{A}_0)^{1/2}$
 3: **for** $t = 1, 2, 3, \dots, T$ **do**
 4: $V_{t+1} \leftarrow \widetilde{V}_t - \eta \nabla f(\widetilde{V}_t) = \widetilde{V}_t - 2\eta(-\widehat{\Sigma} \widetilde{V}_t + \lambda \widehat{\Sigma}_0 \widetilde{V}_t(\widetilde{V}_t^\top \widehat{\Sigma}_0 \widetilde{V}_t - I_r))$
 5: $\widetilde{V}_{t+1} \leftarrow HT(V_{t+1}, s')$
 6: **end for**
 7: **Output:** $\widehat{A}_T = \widetilde{V}_T(\widetilde{V}_T^\top \widehat{\Sigma}_0 \widetilde{V}_T)^{-1/2}$

In Algorithm 1, each iteration is computationally efficient: line 4 is essentially matrix multiplication and addition, and line 5 requires calculating and sorting row l_2 norms. Since $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n (X^{(i)} - \bar{X})(X^{(i)} - \bar{X})^\top$ where $\bar{X} = \frac{1}{n} \sum_{i=1}^n X^{(i)}$, an efficient way to calculate $\widehat{\Sigma} \widetilde{V}_t$ would be to calculate $Y_t^{(i)} = (X^{(i)} - \bar{X})^\top \widetilde{V}_t$ first, followed by multiplying $(X^{(i)} - \bar{X}) Y_t^{(i)}$, resulting in $O(npr)$ flops. Line 5 requires $O(rp + p \log p)$ flops for selecting top s' rows. Hence a single iteration of thresholded gradient descent will require $O(npr + p \log p)$ flops in total.

3.3 Initialization via generalized Fantope projection

Success of Algorithm 1 depends crucially on the quality of the initial estimator \widehat{A}_0 . Thus, we need to find an initial estimator that is relatively close to the true generalized eigenspace in some distance so that later iterations could further improve on estimation accuracy. We now introduce such an initial estimator based on Fantope projection (Vu et al., 2013).

Note that

$$\sum_{i,j=1}^k \text{Tr}(A_{\{i\}}^\top \Sigma_{\{ij\}} A_{\{j\}}) = \langle \Sigma, AA^\top \rangle.$$

The idea is to “lift” AA^\top into \mathbb{S}^p , the space of $p \times p$ symmetric matrices, and hence to treat it as a single quantity. Since we assume that A is row sparse, $F = AA^\top$ have at most s^2 nonzero entries which is much smaller than its number of elements. To ensure sparsity of its solution, we impose an entrywise l_1 penalty to define the following objective function for initialization:

$$\underset{F \in \mathbb{S}^p}{\text{minimize}} \quad -\langle \widehat{\Sigma}, F \rangle + \rho \|F\|_1.$$

On the other hand, the normalization constraint on A implies that

$$(\widehat{\Sigma}_0)^{1/2} F (\widehat{\Sigma}_0)^{1/2} \in \mathcal{P}(p, r),$$

where $\mathcal{P}(p, r)$ is defined to be the set of rank r projection matrices

$$\mathcal{P}(p, r) = \{PP^\top, P \in \mathcal{O}(p, r)\}.$$

However, this direct generalization leads to a nonconvex feasible set since $\mathcal{P}(p, r)$ is nonconvex. To obtain a bona fide convex program, in the light of Vu et al. (2013), we use the Fantope set introduced by Dattorro (2003):

$$\mathcal{F}_r = \{X : 0 \preceq X \preceq I \quad \text{and} \quad \text{Tr}(X) = r\}.$$

The motivation is the observation from Overton and Womersley (1991) that $\mathcal{F}_r = \text{conv}(\mathcal{P}(p, r))$, where $\text{conv}(A)$ denotes the convex hull of A . To summarize, our initial estimator is the solution to the following program:

$$\begin{aligned} & \underset{F \in \mathbb{S}^p}{\text{minimize}} && -\langle \widehat{\Sigma}, F \rangle + \rho \|F\|_1, \\ & \text{subject to} && (\widehat{\Sigma}_0)^{1/2} F (\widehat{\Sigma}_0)^{1/2} \in \mathcal{F}_r. \end{aligned} \tag{14}$$

Upon obtaining \widehat{F} as the solution to (14), we collect the leading r eigenvectors of \widehat{F} as $\widetilde{U}_r \in \mathbb{R}^{p \times r}$ and the corresponding leading r eigenvalues as entries of the diagonal matrix $\widetilde{D}_r \in \mathbb{R}^{r \times r}$. Then let $A_0 = \widetilde{U}_r \widetilde{D}_r^{1/2}$ and

$$\widehat{A}_0 = HT(A_0, s'). \tag{15}$$

This finishes initialization for Algorithm 1. The initialization procedure has $O(p^3)$ computation complexity due to the ADMM step involved in solving the generalized Fantope projection. See Gao et al. (2017) for details of the ADMM algorithm.

As we shall show in next section, \widehat{F} suffers a relatively large estimation error rate for estimating AA^\top and hence \widehat{A}_0 for A . However, for Algorithm 1 to work, such an estimator serves well as an initial estimator under mild conditions. In addition to PCA and the current setting, an analogous initialization via convex relaxation idea has appeared in Gao et al. (2017) for performing sparse CCA which is asymmetric.

4. Theoretical Results

We provide theoretical justifications for our algorithms in this section. We first state our main result on how each iteration of Algorithm 1 improves estimation accuracy, followed by a corollary on error bounds achieved by the final estimator. Analysis of generalized Fantope initialization follows our investigation of the main algorithm. In addition, we include a lower bound for the finite k setting at the end of this section.

Parameter space Under covariance structure (11), we define $\mathcal{F}(\{s_i\}_1^k, \{p_i\}_1^k, r, \{\lambda_j\}_1^p; \nu)$ as the collection of all covariance matrices Σ satisfying (11) and the following conditions:

- (i) Sparsity: $A_{\{i\}} \in \mathbb{R}^{p_i \times r}$ with $\|A_{\{i\}}\|_{2,0} \leq s_i$;
- (ii) Bounded spectrum: $\frac{1}{\nu} \leq s_{\min}(\Sigma) < s_{\max}(\Sigma) \leq \nu$, $\frac{1}{\nu} \leq s_{\min}(\Sigma_0) < s_{\max}(\Sigma_0) \leq \nu$; (16)
- (iii) Eigengap: $\lambda_r - \lambda_{r+1} > 0$.

Here λ_i 's are generalized eigenvalues as defined in (8). By Section 2, under model (5) the eigengap condition is satisfied if Assumption 2 holds. The parameter space is then defined

as

$$\mathcal{P}_n(\{s_i\}_1^k, \{p_i\}_1^k, r, \{\lambda_j\}_1^p; \nu) = \{\mathcal{L}(X^{(1)}, X^{(2)}, \dots, X^{(n)}) : X^{(i)} \stackrel{iid}{\sim} N(0, \Sigma), \Sigma \in \mathcal{F}(\{s_i\}_1^k, \{p_i\}_1^k, r, \{\lambda_j\}_1^p; \nu)\}. \quad (17)$$

In what follows, S denotes the true row support of A . While the parameter space requires normality, our theoretical analysis generalizes directly to sub-Gaussian distributions. We omit the generalization in this work as it is mostly formality.

Matrix distance To measure estimation accuracy, we define the distance between two matrices U and $V \in \mathbb{R}^{p \times r}$ as

$$\text{dist}(U, V) = \min_{P \in \mathcal{O}(r)} \|UP - V\|_F. \quad (18)$$

Here $\mathcal{O}(r)$ is the collection of $r \times r$ orthogonal matrices. The matrix distance has been used previously in Ge et al. (2017), Tu et al. (2015) and Golub and Van Loan (2012), among others.

4.1 Main results

The following theorem characterizes numerical convergence of Algorithm 1 when starting at a reasonable initializer.

Theorem 7 *In Algorithm 1, set*

$$\eta \leq \frac{c}{12\lambda_1\nu(1+c)^2}, \quad \lambda = \frac{\lambda_1}{c} \quad (19)$$

for some constant c , and

$$s' \geq \frac{256s\nu^2}{(\lambda_r - \lambda_{r+1})^2\eta^2} \vee s. \quad (20)$$

For initial estimator \hat{A}_0 , suppose that it has row sparsity s' , and that for

$$V = A \left(I + \frac{1}{\lambda} \Lambda_r \right)^{\frac{1}{2}}, \quad (21)$$

\hat{A}_0 is so constructed that after the first two lines of Algorithm 1, we have

$$\text{dist}(V, \bar{V}_1) \leq \frac{1}{8\sqrt{\nu}} \min \left\{ \frac{c(\lambda_r - \lambda_{r+1})}{\sqrt{2}\lambda_1\nu^2(42 + 25c)}, \frac{\sqrt{1+c}}{2} \right\}. \quad (22)$$

If (16) holds and

$$\frac{\sqrt{1+\lambda_1^2}\sqrt{1+\lambda_{r+1}^2}}{(\lambda_r - \lambda_{r+1})^2} \sqrt{\frac{rs \log p}{n}} < c_0 \quad (23)$$

for some sufficiently small constant $c_0 > 0$, then for some constants $C, C' > 0$, uniformly over $\mathcal{P}_n = \mathcal{P}_n(\{s_i\}_1^k, \{p_i\}_1^k, r, \{\lambda_j\}_1^p; \nu)$, with probability at least $1 - \exp(-C'(s' \log(ep/s')))$, for all $t \geq 1$,

$$\text{dist}(V, \bar{V}_{t+1}) \leq \underbrace{C \left(\frac{s'}{s} \right)^{3/2} \frac{\sqrt{1+\lambda_1^2}\sqrt{1+\lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs \log p}{n}}}_{\text{Statistical Error}} + \underbrace{\xi^t \text{dist}(V, \bar{V}_1)}_{\text{Optimization Error}} \quad (24)$$

where $\xi = 1 - \frac{\eta^2(\lambda_r - \lambda_{r+1})^2}{64\nu^2}$.

The foregoing theorem leads to the following corollary on high probability error bounds for estimating A .

Corollary 8 *Suppose the conditions of Theorem 7 hold. For each $t \geq 1$, let $\widehat{A}_t = \overline{V}_t(\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t)^{-1/2}$. Then for some constants $C, C_1, C' > 0$, uniformly over \mathcal{P}_n , with probability at least $1 - \exp(-C'(s' \log(ep/s')))$, for all $t \geq T_0$, where*

$$T_0 = \frac{\log \left(C \left(\frac{s'}{s} \right)^{3/2} \frac{\sqrt{1+\lambda_1^2} \sqrt{1+\lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs \log p}{n}} \right)}{\log \left(1 - \frac{\eta^2(\lambda_r - \lambda_{r+1})^2}{64\nu^2} \right)}, \quad (25)$$

we have

$$\text{dist}(A, \widehat{A}_t) \leq C_1 \left(\frac{s'}{s} \right)^{3/2} \frac{\sqrt{1+\lambda_1^2} \sqrt{1+\lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs \log p}{n}}.$$

In particular, when $T \geq T_0$, the last display holds for the output \widehat{A}_T of Algorithm 1.

We give some brief remarks on Theorem 7 and Corollary 8. The error bound for each step, given by the right side of (24), is composed of two parts: a statistical error term independent of the iteration counter t , and an optimization error that decreases geometrically as t increases. When $t \geq T_0$, the optimization error is dominated by the statistical error and we can achieve the estimation rate in Corollary 8. Thus, T_0 can be interpreted as the minimum number of iterations needed for achieving the estimation error rate in Corollary 8. When the eigengap $\lambda_r - \lambda_{r+1}$ is lower bounded by a positive constant, it is straightforward to verify that $T_0 = O(\log(p+n))$. In other words, thresholded gradient descent drives down the statistical error of the estimator to the desired rate after $O(\log(p+n))$ iterates. It is worth noting that the contraction rate ξ does not depend on ambient dimension p , which can be much larger than r and s .

Condition (22) on initial estimator is related to the notion of “basin of attraction”. It has previously appeared in the literature of machine learning with nonconvex optimization in other contexts, e.g., Chen and Candes (2015); Chi et al. (2019); Wang et al. (2014). The global landscape of (12) could be hard to handle due to non-convexity and sparsity constraints. However, within the basin of attraction, the objective function $f(\cdot)$ is locally smooth and strongly convex, which allows projected gradient descent to find a statistically sound solution at geometric convergence rate.

Outline of proof To prove Theorem 7, we track the estimator trajectory $\{\overline{V}_t : t \geq 0\}$ over iteration. To this end, we first characterize a global high probability event under which the entire trajectory would lie inside the basin of attraction within which the objective function is smooth and strongly convex. On the event, we analyze in sequel the effects of the gradient step and the hard thresholding step in each iteration. In particular, we show that each gradient step drives down the distance between the target of estimation and the current estimator while the hard thresholding step projects the current estimator onto the feasible set. For details, see Propositions 18 and 19, respectively. Combining the two propositions, we obtain a recursive inequality that characterizes the estimator trajectory

as in (24). This completes the major steps in the proof. Finally, Corollary 8 is a direct consequence of normalization on the same event that we performed the foregoing analysis of iteration. For proof details, see Appendix B.

4.2 Analysis of initialization

The following theorem characterizes the estimation accuracy of the proposed initial estimator via generalized Fantope projection.

Theorem 9 *Suppose*

$$\frac{s^2 \log p}{n(\lambda_r - \lambda_{r+1})^2} \leq \epsilon \quad (26)$$

for some sufficiently small $\epsilon > 0$. Let \widehat{F} be the solution to (14) where $\rho = \gamma \sqrt{\frac{\log p}{n}}$ for $\gamma \in [\gamma_1, \gamma_2]$ for some positive constants $\gamma_1 < \gamma_2$. There exist constants $C, C' > 0$ such that uniformly over \mathcal{P}_n , with probability at least $1 - \exp(-C'(s + \log(ep/s)))$,

$$\|\widehat{F} - AA^\top\|_{\mathbb{F}}^2 \leq C \frac{(\sum_{i=1}^k s_i)^2 \log(\sum_{i=1}^k p_i)}{n(\lambda_r - \lambda_{r+1})^2} = O\left(\frac{s^2 \log p}{n(\lambda_r - \lambda_{r+1})^2}\right). \quad (27)$$

The following corollary further bounds the estimation accuracy of \widehat{A}_0 through the lens of \overline{V}_1 defined in line 2 of Algorithm 1.

Corollary 10 *Suppose that the conditions of Theorem 9 hold. Under the choice of s' and condition (23) in Theorem 7, there exist constants $C, C' > 0$ such that uniformly over \mathcal{P}_n , with probability at least $1 - \exp(-C'(s + \log(ep/s)))$,*

$$\text{dist}(V, \overline{V}_1) \leq \frac{Cs}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{\log p}{n}},$$

where V is defined in (21).

Combining Theorems 7 and 9 and Corollaries 8 and 10, we obtain the following corollary on the whole procedure: Algorithm 1 with initialization via generalized Fantope projection.

Corollary 11 *Suppose that*

$$n \geq C_0 \max \left\{ \frac{(1 + \lambda_1^2)(1 + \lambda_{r+1}^2)}{(\lambda_r - \lambda_{r+1})^4} rs \log p, \frac{s^2 \log p}{(\lambda_r - \lambda_{r+1})^2} \right\} \quad (28)$$

for some sufficiently large positive constant C_0 , that $T \geq T_0$ with T_0 in (25), and that all the other conditions in Theorems 7 and 9 are satisfied. There exist constants $C, C' > 0$ such that uniformly over \mathcal{P}_n , with probability at least $1 - \exp(-C'(s + \log(ep/s)))$, the final output satisfies

$$\text{dist}(A, \widehat{A}_T) \leq C_1 \left(\frac{s'}{s}\right)^{3/2} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs \log p}{n}}.$$

4.3 A lower bound for finite k

Corollary 8 provides upper bounds for the proposed procedure in Algorithm 1. For finite k , we have the following information-theoretic lower bound result.

Theorem 12 *Assume that $1 \leq r \leq \frac{\min_i s_i}{2}$, and that*

$$n(\lambda_r - \lambda_{r+1})^2 \geq C_0 \left(r + \max_i \log \frac{ep_i}{s_i} \right) \quad (29)$$

for some sufficiently large positive constant C_0 . Then there exist positive constants c and c_0 such that the minimax risk for estimating A satisfies

$$\inf_{\hat{A}} \sup_{\Sigma \in \mathcal{F}} \mathbb{E}_{\Sigma} \text{dist}^2(A, \hat{A}) \geq c_0 \wedge \frac{c}{n(\lambda_r - \lambda_{r+1})^2} \left(rs + \sum_{i=1}^k s_i \log \frac{ep_i}{s_i} \right). \quad (30)$$

The proof of Theorem 12 is given in Appendix C. Comparing Theorem 12 and Corollary 8, we see that when r is finite and $\lambda_1 \leq C$ for some large positive constant C , as long as there exists $i \in \{1, \dots, k\}$ such that s_i and p_i are of the same order as s and p simultaneously, the lower and upper bounds match. Otherwise, they differ by at most a multiplicative factor of $\log p$.

5. Numerical Results

This section reports numerical results on synthetic data sets. Except for the settings in Section 5.4 and Section 5.6, r , the latent dimension in model (5) is assumed to be known. Choices of tuning parameters are specified in each setting. In practice, they can also be selected using cross-validation on grids. The rest of this section is organized as follows. We first study numerical errors for estimating generalized eigenspaces of different dimensions. For the special case of $r = 1$, we also compare our method with the Rifle method in Tan et al. (2018). Next, we assess the performance of Algorithm 1 in the context of sparse CCA by comparing it to the CoLaR method in Gao et al. (2017). Furthermore, we consider the potential model mis-specification scenario where the input latent dimension of the algorithm is different from the true value. We then apply Algorithm 1 to perform sparse PCA of correlation matrices. Finally, we investigate the performance of Algorithm 1 under a general covariance structure.

5.1 Sparse GCA with different latent dimensions

We first consider sparse GCA of three high dimensional data sets. In particular, we set $n = 500$, $k = 3$, $p_2 = p_3 = 200$, $p_1 = 500$, and $s_1 = s_2 = s_3 = 5$. To generate covariance matrices Σ and Σ_0 , we use the latent variable model specified in Section 2. Specifically, Σ is a block matrix with

$$\begin{aligned} \Sigma_{\{ij\}} &= T_{\{i\}} U_{\{i\}} U_{\{j\}}^{\top} T_{\{j\}}, \quad \Sigma_{\{ii\}} = T_{\{i\}}, \quad \text{for } i \neq j \in \{1, 2, 3\}, \\ U_{\{1\}}^{\top} T_{\{1\}} U_{\{1\}} &= U_{\{2\}}^{\top} T_{\{2\}} U_{\{2\}} = U_{\{3\}}^{\top} T_{\{3\}} U_{\{3\}} = I. \end{aligned}$$

Here each $\Sigma_{\{ii\}} = T_{\{i\}}$ is a Toeplitz matrix, defined by setting $(T_{\{k\}})_{ij} = \sigma_{k_{ij}}$ where $\sigma_{k_{ij}} = a_k^{|i-j|}$ for all $i, j \in [p_k]$ with $a_1 = 0.5, a_2 = 0.7, a_3 = 0.9$. To generate $U_{\{i\}} \in \mathbb{R}^{p_i \times r}$, we first randomly select a support of size 5. For each row in the support, we generate its entries as i.i.d. standard normal random variables. Then all $U_{\{i\}}$'s are normalized with respect to $T_{\{i\}}$. With the foregoing construction, it is straightforward to verify that $\lambda_r - \lambda_{r+1} = 2$ and that $\text{Col}(A_{\{i\}}) = \text{Col}(U_{\{i\}})$. Finally, $\Sigma_0 = \text{diag}(\Sigma_{\{11\}}, \Sigma_{\{22\}}, \Sigma_{\{33\}})$ contains the block diagonal elements of Σ .

We vary r in $\{1, 2, 3, 4, 5\}$ and report squared matrix distances defined as $\text{dist}^2(A, \widehat{A}) = \min_{O \in \mathcal{O}(r)} \|\widehat{A}O - A\|_F^2$ for both initial and final estimators based on 50 repetitions in each setting. For tuning parameters in Algorithm 1, we set $s' = 20$, $\eta = 0.001$, $\lambda = 0.01$, and $T = 15000$. The tuning parameter for generalized Fantope initialization is set to be $\rho = \frac{1}{2} \sqrt{\frac{\log p}{n}}$. The truncation parameter for initialization is also set to be $s' = 20$.

Table 1 reports the results of the aforementioned simulation study. For all latent dimensions, we observe a significant decrease in estimation error after Algorithm 1 is applied. This corroborates the theory in Section 4. To better understand how each iteration of Algorithm 1 improves estimation, we plot $\text{dist}(V, \overline{V}_t)$ and $\text{dist}(A, \widehat{A}_t)$ in logarithmic scale against the iteration counter t for $r = 1, 2, 3, 4, 5$ and $n = 500$. For any t , we set $\widehat{A}_t = \overline{V}_t (\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t)^{-1/2}$ which agrees with the definition of \widehat{A}_T in the last line of Algorithm 1. From Figure 1, we observe an approximate linear decay trend at the beginning in all cases, which corresponds to exponential decay in the original scale. Moreover, after sufficiently many iterations, all error curves plateau, which suggests that the performance of the resulting estimators have stabilized. Both phenomena agree well with the theoretical findings in Theorem 7.

Error	$r = 1$	$r = 2$	$r = 3$	$r = 4$	$r = 5$
Initial	0.1319(0.0737)	0.2308(0.0591)	0.2746(0.0652)	0.2354(0.0483)	0.1969(0.0237)
Final	0.0015(0.0030)	0.0072(0.0209)	0.0098(0.0301)	0.0121(0.0071)	0.0171(0.0061)

Table 1: Median errors of initial (generalized Fantope) and final (Algorithm 1) estimators in squared matrix distance out of 50 repetitions. Median absolute deviations of errors are reported in parentheses.

Tan et al. (2018) proposed a truncated Rayleigh flow method, called Rifle, to solve the sparse generalized eigenvalue problem when $r = 1$. Here we compare our method and Rifle under the same experiment setting as above, with $n = 500$ and 1000. For TGD we set the tuning parameters to be $s' = 20, \eta = 0.01, \lambda = 0.01$ and $T = 15000$. For Rifle we also use $s' = 20$ for truncation, the step size η is set to be the default value 0.01 and we run it for sufficiently many iterations until estimator errors no longer improve. We vary the sample size n in $\{500, 1000\}$ and each experiment setting is repeated 50 times. We report median estimation errors of both initial and final estimators of both methods measured in squared matrix distances in Table 2. Both Rifle and our method produce highly accurate final estimators, which significantly decrease errors of respective initializers. Final estimators by our algorithm are slightly more accurate in both settings.




Figure 1: Plots of $\text{dist}(V, \bar{V}_t)$ (left) and $\text{dist}(A, \hat{A}_t)$ (right) (both in logarithmic scale) vs. iteration counter t .

	Rifle		TGD	
	Initial Error	Final Error	Initial Error	Final Error
$n = 500$	0.1225 (0.0562)	0.0006 (0.0005)	0.1458 (0.0713)	0.0003 (0.0011)
$n = 1000$	0.0805 (0.0449)	0.0003 (0.0012)	0.0920 (0.0504)	0.0002 (0.0003)

Table 2: Median errors of initial and final estimators of both Rifle and Algorithm 1 in squared matrix distance out of 50 repetitions. Median absolute deviations of errors are reported in parentheses.

5.2 Choices of tuning parameters

In this section, we study the selection of tuning parameters. For the sparsity level s' in the hard thresholding step of Algorithm 1, we show that a five-fold cross validation approach works reasonably well. For the Lagrangian multiplier parameter λ , we show that the estimation procedure is robust with respect to it and we recommend fixing it at some small positive constant.

For simulations below, we fix $n = 500$ and $p = 900$. In addition, we fix other tuning parameters throughout the experiments as follows: $\rho = \frac{1}{2} \sqrt{\frac{\log p}{n}}$ for Fantope initialization and $\eta = 0.001$, $T = 15000$ for Algorithm 1.

Choice of sparsity level s' . The procedure for the selection of s' is as follows. We first randomly split the data X into five folds of equal sizes. For $l = 1, \dots, 5$, we use one fold as the test set $X_{(l)}^{\text{test}}$ and the other four folds combined as the training set $X_{(l)}^{\text{train}}$. For each value of s' in a pre-specified grid G' , we apply Algorithm 1 on $X_{(l)}^{\text{train}}$ to obtain an estimator $\hat{A}_{(l)}^{\text{train}}$. Then we compute the test GCA score defined as

$$\text{CV}_{(l)}(s') = \text{Tr} \left(\hat{A}_{(l)}^{\text{train}\top} \hat{\Sigma}_{(l)}^{\text{test}} \hat{A}_{(l)}^{\text{train}} \right).$$

Here $\widehat{\Sigma}_{(l)}^{\text{test}}$ is the sample covariance matrix of the test data $X_{(l)}^{\text{test}}$. Finally, we compute the cross-validation score for s' , defined as $\text{CV}(s') = \frac{1}{5} \sum_{l=1}^5 \text{CV}_{(l)}(s')$. Upon obtaining the cross-validation scores for all $s' \in G'$, we select the sparsity level that maximizes $\text{CV}(s')$.

In the experiments here, we set the grid G' as $\{5i : i = 1, \dots, 20\}$. We fix $r = 3, \lambda = 0.01$ and consider three cases: $s_1 = s_2 = s_3 = 5$ (Case I), $s_1 = s_2 = s_3 = 15$ (Case II) and $s_1 = s_2 = s_3 = 20$ (Case III). The true sparsity levels in three cases are then $s = 15, 45,$ and 60 , respectively. The results from the 5-fold cross validation procedure for the three cases are plotted in Figure 2. Here, we plot $\text{CV}(s')$ against $s' \in G'$ and use error-bars to indicate the range of \pm one standard deviation of cross-validation scores.

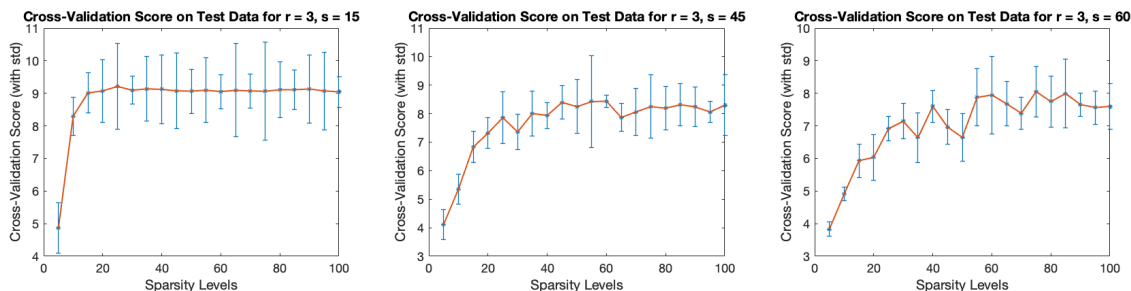


Figure 2: Plots of cross-validation score $\text{CV}(s')$ against $s' \in G'$ for case I ($s = 15$, left), case II ($s = 45$, middle), and case III ($s = 60$, right). Curves indicate average over five folds and error bars indicate one standard deviations.

Figure 2 shows that in all three cases, the cross-validation scores increase with s' until reaching a plateau level and fluctuate around some constant when s' is set to larger values. The sparsity levels selected by this procedure are slightly larger than the true sparsity levels in all three settings. Specifically, the selected sparsity levels for the three cases are 25, 60, and 75, respectively. This suggests that our previous choice of $s' = 20$ for $s = 15$ in Section 5.1 is reasonable.

Choice of penalty λ . We now study the impact of different choices of λ on our estimator. To this end, we vary λ on a log scale and consider all values in the set $\{0.0001, 0.001, 0.01, 0.1, 1\}$ and compare the final estimation errors across these choices. We fix $s' = 20$ and test using the same simulation setting as in Section 5.1 with $r = 3$ and 4. We set all other tuning parameters in the same way and only change λ . Table 3 reports squared distances from estimators to true parameter values based on 50 repetitions in each (r, λ) value combination. From Table 3, we notice that for both $r = 3$ and $r = 4$, the final estimation error is robust with respect to the choice of λ .

5.3 Sparse CCA

To further compare the performance of our algorithm with a benchmark, we compare its performance in sparse CCA problems with the CoLaR method proposed by Gao et al. (2017). Instead of using gradient descent and hard thresholding, Gao et al. (2017) refined the initial estimate by a linear regression with group Lasso penalty. In addition, their

λ	0.0001	0.001	0.01	0.1	1
$r = 3$	0.0121(0.0364)	0.0136(0.0252)	0.0098(0.0301)	0.0093(0.0140)	0.0099(0.0108)
$r = 4$	0.0145(0.0068)	0.0152(0.0102)	0.0121(0.0071)	0.0101(0.0055)	0.0119(0.0093)

Table 3: Median errors of the final (Algorithm 1) estimators in squared matrix distance out of 50 repetitions for $r = 3, 4$ and different choices of λ . Median absolute deviations of errors are reported in parentheses.

estimator was shown to achieve the optimal estimation rate for canonical loading matrix under a prediction loss that is different from (18).

For fair comparison, we use the same simulation settings as in Gao et al. (2017). Using the notation in Remark 4, we set

$$p_1 = p_2, \quad \Sigma_x = \Sigma_y = M, \quad r = 2, \quad \theta_1 = 0.9, \quad \theta_2 = 0.8.$$

The covariance matrix between X and Y is generated from the canonical pair model $\Sigma_{xy} = \Sigma_x V \Theta_2 W^\top \Sigma_y$. Recall that from the relationship between Θ_2 and Λ_2 we obtain that $\lambda_1 = \theta_1 + 1 = 1.9$ and $\lambda_2 = \theta_2 + 1 = 1.8$. Hence the leading two generalized eigenvalues are 1.9 and 1.8. Moreover, the row supports for both V and W are set to be $\{1, 6, 11, 16, 21\}$. The values at the nonzero coordinates are obtained by normalizing (with respect to $M = \Sigma_x = \Sigma_y$) random numbers drawn from the uniform distribution on the finite set $\{-2, -1, 0, 1, 2\}$. We consider three different choices of M as follows: (i) Identity: $M = I$; (ii) Toeplitz: $M = (\sigma_{ij})$ where $\sigma_{ij} = 0.3^{|i-j|}$ for all i, j ; (iii) SparseInv: $M = \sigma_{ij}^0 / \sqrt{\sigma_{ii}^0 \sigma_{jj}^0}$ where $\Sigma^0 = (\sigma_{ij}^0) = \Omega^{-1}$ for $\Omega = \omega_{ij}$ with

$$\omega_{ij} = \mathbf{1}_{i=j} + 0.5 \times \mathbf{1}_{|i-j|=1} + 0.4 \times \mathbf{1}_{|i-j|=2} \quad i, j \in [p_1].$$

We consider the same four configurations of (n, p_1, p_2) as in Gao et al. (2017). For CoLaR and its initialization, we set all tuning parameters to those values used in Gao et al. (2017).

For tuning parameters in Algorithm 1 and its initialization, we set $\rho = \frac{1}{2} \sqrt{\frac{\log p}{n}}$, $s' = 20$, $\lambda = 0.01$, $\eta = 0.001$, and $T = 10000$. Upon obtaining the final estimate $\hat{A} = [\hat{A}_{\{1\}}^\top, \hat{A}_{\{2\}}^\top]^\top$ from Algorithm 1, we calculate estimators for V and W as $\hat{V} = \hat{A}_{\{1\}} (\hat{A}_{\{1\}}^\top \hat{\Sigma}_x \hat{A}_{\{1\}})^{-1/2}$ and $\hat{W} = \hat{A}_{\{2\}} (\hat{A}_{\{2\}}^\top \hat{\Sigma}_y \hat{A}_{\{2\}})^{-1/2}$.

We summarize simulation results of the foregoing settings in Tables 4–6, each corresponding to a different choice of M . Following Gao et al. (2017), we measure estimation error by prediction loss defined as $L(V, \hat{V}) = \inf_{O \in \mathcal{O}(r)} \|\Sigma_x^{1/2} (\hat{V} O - V)\|_{\text{F}}^2$ and $L(W, \hat{W}) = \inf_{O \in \mathcal{O}(r)} \|\Sigma_y^{1/2} (\hat{W} O - W)\|_{\text{F}}^2$. Each reported number is the median error out of 50 independent repetitions.

In Tables 4–6, the first four columns collect results of CoLaR. Columns “V-GFP” and “W-GFP” collect errors of initial estimators by generalized Fantope projection, while columns “V-TGD” and “W-TGD” report errors of final estimators by Algorithm 1. In all settings, both CoLaR and Algorithm 1 yield good final estimators, which significantly improve over their respective initializers. In addition, Algorithm 1 outperforms CoLaR in most settings.

(n, p_1, p_2)	V-init	W-init	V-CoLaR	W-CoLaR	V-GFP	W-GFP	V-TGD	W-TGD
(300,300,200)	0.2885	0.1706	0.0511	0.0601	0.1391	0.2636	0.0118	0.0118
(600,600,200)	0.3236	0.2004	0.0638	0.0764	0.3420	0.1655	0.0237	0.0485
(300,300,500)	0.1202	0.0664	0.0135	0.0166	0.2646	0.3078	0.0100	0.0076
(600,600,500)	0.1408	0.0811	0.0176	0.0209	0.2459	0.1876	0.0213	0.0157

Table 4: Comparison of CoLaR and Algorithm 1: median errors out of 50 repetitions with Toeplitz covariance structure.

(n, p_1, p_2)	V-init	W-init	V-CoLaR	W-CoLaR	V-GFP	W-GFP	V-TGD	W-TGD
(300,300,200)	0.2653	0.1712	0.0498	0.0646	0.3413	0.5688	0.0454	0.0211
(600,600,200)	0.3167	0.2087	0.0671	0.0776	0.5016	0.7134	0.0231	0.0559
(300,300,500)	0.1207	0.0665	0.0135	0.0159	0.3580	0.1713	0.0124	0.0117
(600,600,500)	0.1448	0.0817	0.0166	0.0203	0.3906	0.4095	0.0082	0.0129

Table 5: Comparison of CoLaR and Algorithm 1: median errors out of 50 repetitions with Identity covariance structure.

(n, p_1, p_2)	V-init	W-init	V-CoLaR	W-CoLaR	V-GFP	W-GFP	V-TGD	W-TGD
(300,300,200)	0.5552	0.5718	0.1568	0.1194	0.2021	0.4924	0.0719	0.0421
(600,600,200)	0.5596	0.6133	0.2123	0.1572	0.4653	0.2222	0.1555	0.0686
(300,300,500)	0.2695	0.1917	0.0242	0.0219	0.1592	0.3058	0.0042	0.0056
(600,600,500)	0.3068	0.2368	0.0338	0.0271	0.1853	0.1404	0.0454	0.0491

Table 6: Comparison of CoLaR and Algorithm 1: median errors out of 50 repetitions with SparseInv covariance structure.

In summary, both CoLaR and Algorithm 1 consistently estimate the leading r canonical loading vectors, while Algorithm 1 has a slight advantage. Moreover, Algorithm 1 is more desirable due to its generality beyond the sparse CCA setting.

5.4 Model mis-specification

To investigate the robustness of our method, we consider one possible mis-specification of the model in the case of two data sets: There are 3 pairs of nontrivial canonical correlations present but we only set $r = 2$. As we are in the CCA setting, we consider three types of covariance matrices specified in the previous subsection. For all three cases, the first two canonical loading vectors and generalized eigenvalues are generated in the same way as the previous subsection. In addition, we also set the support of the third vector to be $\{1, 6, 11, 16, 21\}$. We fix $\theta_3 = 0.3$ (i.e., $\lambda_3 = 1.3$) and focus on the configuration $(n, p_1, p_2) = (300, 300, 500)$. All tuning parameters are set in the same way as in last subsection.

Simulation results under these settings can be found in Table 7, and we continue to use prediction loss as in the previous subsection. Similar to CoLaR, Algorithm 1 continues to

Covariance structure	V-CoLaR	W-CoLaR	V-TGD	W-TGD
Toeplitz	0.0197	0.0197	0.0291	0.0157
Identity	0.0190	0.0195	0.0113	0.0155
SparseInv	0.0348	0.0263	0.0112	0.0488

Table 7: Comparison of CoLaR and Algorithm 1: median errors of estimating first two generalized eigenvectors out of 100 repetitions in mis-specified models.

produce accurate estimates when the input latent dimension r_{in} is smaller than the true value r and there is an eigengap between $\theta_{r_{\text{in}}}$ and $\theta_{r_{\text{in}}+1}$, and thus between $\lambda_{r_{\text{in}}}$ and $\lambda_{r_{\text{in}}+1}$.

5.5 Sparse PCA of correlation matrices

We now apply Algorithm 1 to performing sparse PCA of correlation matrices. See Remark 5 for a detailed account on how sparse PCA of correlation matrices can be cast as a special case of sparse GCA. By Remark 5, let $\widehat{\Sigma}_0$ be the diagonal matrix with p sample variances on the diagonal, we could use $\widehat{\Sigma}_0^{1/2}\widehat{A}$ as our estimator of the r leading eigenvectors of correlation matrix $R = \Sigma_0^{-1/2}\Sigma\Sigma_0^{-1/2}$, where \widehat{A} is the estimator produced by Algorithm 1. We report the squared matrix distance for estimating the leading eigenspace $L(E_R, \widehat{A}) = \min_{O \in \mathcal{O}(r)} \|\widehat{\Sigma}_0^{1/2}\widehat{A}O - E_R\|_{\mathbb{F}}^2$ for initialization and final error, where E_R is the leading eigenspace of R .

Given ambient dimension p , sparsity s , latent dimension r and generalized eigenvalues $\lambda_1 \geq \dots \geq \lambda_r > 1$, we generate a $p \times p$ correlation matrix R with leading eigenvalues $\lambda_1 \geq \dots \geq \lambda_r > \lambda_{r+1} = 1$ and s -sparse leading eigenvectors according to the procedure described in Section D. To further construct the covariance matrix Σ , we generate Σ_0 as a diagonal matrix with i.i.d. uniform random numbers in $[0.1, 1]$ as diagonal elements. Finally, we set $\Sigma = \Sigma_0^{1/2}R\Sigma_0^{1/2}$.

In our simulation study, we set $p = 500, s = 20, r = 3$, and let sample size n be either 500 or 2000. The tuning parameter for initialization is set as $\rho = \frac{1}{2}\sqrt{\frac{\log p}{n}}$. Tuning parameters of Algorithm 1 are set to be $s' = 40, T = 20000, \lambda = 0.01$, and $\eta = 0.001$. The leading generalized eigenvalues are set as $\{5, 5, 5\}$ (case I) or $\{7, 5, 3\}$ (case II). For both cases, our construction of Σ ensures that $\lambda_4 = 1$.

n	Case	Initial Error	Final Error
$n = 500$	I	0.4603 (0.1404)	0.0851 (0.0306)
	II	0.6379 (0.0567)	0.1001 (0.0098)
$n = 2000$	I	0.0882 (0.0177)	0.0184 (0.0025)
	II	0.1060 (0.0697)	0.0258 (0.0025)

Table 8: Sparse PCA of correlation matrices by Algorithm 1: median errors of initial and final estimators out of 50 repetitions. MADs are reported in parentheses.

We report results from 50 repetitions for each case and each sample size in Table 8. From Table 8, it can be seen that in all settings, Algorithm 1 yields accurate estimators of leading eigenspaces of correlation matrices. Due to a larger eigengap in the first case, it has slightly better estimation accuracy which corroborates our theory.

5.6 Sparse GCA with general covariance structure

In this section, we investigate the performance of our algorithm under a general covariance structure and demonstrate empirically that our algorithm works beyond latent variable models as long as the eigengap condition holds.

We consider sparse GCA of three high-dimensional data sets with $p_1 = 500$, $p_2 = p_3 = 200$ and let the rank of all off-diagonal blocks of Σ be $p_{\min} = \min_i p_i = 200$. To generate covariance matrices Σ and Σ_0 , we first define a diagonal matrix Θ with $\Theta_{ii} = 2/i$ for $i = 1, 2, \dots, p_{\min}$. We then define Σ as a block matrix with

$$\Sigma_{\{ij\}} = T_{\{i\}} U_{\{i\}} \Theta U_{\{j\}}^\top T_{\{j\}}, \quad \Sigma_{\{ii\}} = 2T_{\{i\}}, \quad \text{for } i \neq j \in \{1, 2, 3\},$$

$$U_{\{1\}}^\top T_{\{1\}} U_{\{1\}} = U_{\{2\}}^\top T_{\{2\}} U_{\{2\}} = U_{\{3\}}^\top T_{\{3\}} U_{\{3\}} = I.$$

Here, each $\Sigma_{\{ii\}} = 2T_{\{i\}}$ and each $T_{\{i\}}$ is a Toeplitz matrix, defined by setting $(T_{\{k\}})_{ij} = \sigma_{kij}$ where $\sigma_{kij} = a_k^{|i-j|}$ for all $i, j \in [p_k]$ with $a_1 = 0.5, a_2 = 0.7, a_3 = 0.9$. The first five columns of each $U_{\{i\}}$ have row sparsity level $s_i = 5$. As a result, the support sizes of the first three GCA loading vectors are at most $s = 15$. For each $U_{\{i\}} \in \mathbb{R}^{p_i \times p_{\min}}$, we define $U_{\{i\}} = [U_{\{i\}}(1), U_{\{i\}}(2)]$ where $U_{\{i\}}(1) \in \mathbb{R}^{p_i \times s_i}$ and $U_{\{i\}}(2) \in \mathbb{R}^{p_i \times (p_{\min} - s_i)}$. The submatrices $U_{\{i\}}(1), U_{\{i\}}(2)$ are generated as follows. To generate each $U_{\{i\}}(1)$, we first randomly select a support of size s_i . For each row in the support of $U_{\{i\}}(1)$, we generate its entries as i.i.d. standard normal random variables and fill the remaining entries with zeros. Then, we normalize all $U_{\{i\}}(1)$'s with respect to $T_{\{i\}}$. To generate $U_{\{i\}}(2)$, we first perform SVD on $(U_{\{i\}}(1))^\top T_{\{i\}}^{1/2} = P_{\{i\}} D_{\{i\}} Q_{\{i\}}^\top$ to obtain $Q_{\{i\}} \in \mathcal{O}(p_i)$. We then take $\tilde{Q}_{\{i\}} = Q_{\{i\}, *J}$ for $J = \{s_i + 1, \dots, p_{\min}\}$ which is the submatrix of $Q_{\{i\}}$ consisting of the $(s_i + 1)$ th to the p_{\min} th columns of $Q_{\{i\}}$. By construction, we have $(U_{\{i\}}(1))^\top T_{\{i\}}^{1/2} \tilde{Q}_{\{i\}} = 0$ and $\tilde{Q}_{\{i\}} \in \mathbb{R}^{p_i \times (p_{\min} - s_i)}$. Finally, we define $U_{\{i\}}(2) = T_{\{i\}}^{-1/2} \tilde{Q}_{\{i\}}$. By construction, we have $(U_{\{i\}}(2))^\top T_{\{i\}} U_{\{i\}}(2) = I$ and $(U_{\{i\}}(2))^\top T_{\{i\}} U_{\{i\}}(1) = 0$. Thus, the leading five columns of each $U_{\{i\}}$ are sparse and we have $U_{\{i\}}^\top T_{\{i\}} U_{\{i\}} = I$. With the foregoing construction, it is straightforward to verify that the i th generalized eigenvalue is given by $\lambda_i = 1 + 2i^{-1}$ for $i = 1, 2, \dots, p_{\min}$ and the non-trivial GCA loadings are given by $[U_{\{1\}}^\top, U_{\{2\}}^\top, U_{\{3\}}^\top]^\top$. Therefore, there are eigengaps between any consecutive generalized eigenvalues above one and only the first five GCA loading vectors are sparse.

Under the foregoing setup, we aim at estimating the leading GCA loading vector (Case I), the matrix with leading two GCA loading vectors (Case II), and the matrix with leading three GCA loading vectors (Case III). For tuning parameters, we set $s' = 20$ for sparsity level, $\eta = 0.001$, $\lambda = 0.01$, and $T = 15000$ in Algorithm 1, and $\rho = \frac{1}{2} \sqrt{\frac{\log p}{n}}$ in (14). We consider two different sample sizes: $n = 500$ and $n = 2000$.

Table 9 reports the squared matrix distances for the initial and the final estimators from true value defined as $\text{dist}^2(A, \hat{A}) = \min_{O \in \mathcal{O}(r)} \|\hat{A}O - A\|_{\mathbb{F}}^2$ for all three cases and two different sample sizes. Numbers in each cell are the median and the MAD of results over 50 repetitions. Population parameters are also generated independently in each repetition. To gauge the scale of errors, we also report squared Frobenius norm $\|A\|_{\mathbb{F}}^2$ of the estimand in the last column of Table 9.

n	Case	Initial Error	Final Error	$\ A\ _{\mathbb{F}}^2$
$n = 500$	I	0.0500 (0.0358)	0.0005 (0.0007)	0.5669
	II	0.1586 (0.0420)	0.0605 (0.0313)	1.2433
	III	0.3097 (0.1484)	0.1831 (0.1107)	1.8378
$n = 2000$	I	0.0174 (0.0064)	0.0001 (0.0002)	0.5669
	II	0.0390 (0.0268)	0.0136 (0.0063)	1.2433
	III	0.0741 (0.0719)	0.0369 (0.0178)	1.8378

Table 9: Sparse GCA with general covariance structure: median errors of initial and final estimators out of 50 repetitions with MADs reported in parentheses. Medians of squared Frobenius norms of estimands are listed in the last column.

From Table 9, it can be seen that in all settings, Algorithm 1 consistently yields improved estimates of leading GCA loading vectors whenever there is an eigengap.

Appendix A. Proof of Lemmas in Section 2

A.1 Proof of Lemma 1

Proof For $\Lambda_r = \text{diag}(\lambda_1, \dots, \lambda_r)$. Since A solves (3), we have

$$\Sigma_0^{-1} \Sigma A = A \Lambda_r. \quad (31)$$

By model (5)–(7), we have $\Sigma = \Psi + UU^\top$ and $\Sigma_0 = \text{diag}(\Sigma_{\{11\}}, \dots, \Sigma_{\{kk\}})$, and so

$$\Sigma_0^{-1} \Sigma = \begin{bmatrix} I_{p_1} & \Sigma_{\{11\}}^{-1} U_{\{1\}} U_{\{2\}}^\top & \cdots & \Sigma_{\{11\}}^{-1} U_{\{1\}} U_{\{k\}}^\top \\ \Sigma_{\{22\}}^{-1} U_{\{2\}} U_{\{1\}}^\top & I_{p_2} & \cdots & \Sigma_{\{22\}}^{-1} U_{\{2\}} U_{\{k\}}^\top \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{\{kk\}}^{-1} U_{\{k\}} U_{\{1\}}^\top & \Sigma_{\{kk\}}^{-1} U_{\{k\}} U_{\{2\}}^\top & \cdots & I_{p_k} \end{bmatrix}.$$

Thus, (31) leads to

$$\Sigma_0^{-1} \Sigma A = \begin{bmatrix} A_{\{1\}} + \Sigma_{\{11\}}^{-1} U_{\{1\}} \sum_{j \neq 1} U_{\{j\}}^\top A_{\{j\}} \\ \vdots \\ A_{\{k\}} + \Sigma_{\{kk\}}^{-1} U_{\{k\}} \sum_{j \neq k} U_{\{j\}}^\top A_{\{j\}} \end{bmatrix} = \begin{bmatrix} A_{\{1\}} \Lambda_r \\ \vdots \\ A_{\{k\}} \Lambda_r \end{bmatrix} = A \Lambda_r.$$

Collecting terms, we obtain

$$\begin{bmatrix} \Sigma_{\{11\}}^{-1} U_{\{1\}} \sum_{j \neq 1} U_{\{j\}}^\top A_{\{j\}} \\ \vdots \\ \Sigma_{\{kk\}}^{-1} U_{\{k\}} \sum_{j \neq k} U_{\{j\}}^\top A_{\{j\}} \end{bmatrix} = \begin{bmatrix} A_{\{1\}} (\Lambda_r - I_r) \\ \vdots \\ A_{\{k\}} (\Lambda_r - I_r) \end{bmatrix} = A (\Lambda_r - I_r).$$

Since both A and $\Lambda_r - I_r$ are of rank r under the assumption of the lemma, we complete the proof. \blacksquare

A.2 Proof of Lemma 3

Before the proof of the lemma, we recall Weyl's inequality (Weyl, 1912) stated as follows.

Lemma 13 (*Weyl's inequality*) *Let A, B be two $p \times p$ Hermitian matrices and $s_i(A)$ be the i th eigenvalue of A . Then for $1 \leq l, j \leq p$, we have*

$$\begin{aligned} s_l(A + B) &\leq s_j(A) + s_k(B), \quad \text{for } l \geq j + k - 1, \\ s_j(A) + s_l(B) &\leq s_{j+l-p}(A + B), \quad \text{for } j + l \geq p. \end{aligned}$$

Proof The proof is adapted from Fan et al. (2019). It is composed of three parts. First we show that if $i \geq r + 1$, $\lambda_i \leq 1$. Then we will prove that $\lambda_r > 1$. Finally we prove the result on multiplicity of 1. Recall that the generalized eigenvalues are identical to the eigenvalues of $R = \Sigma_0^{-1/2} \Sigma \Sigma_0^{-1/2}$. From now on we study the eigenvalues of R directly.

Step (1): Note that we can write R as

$$R = \Sigma_0^{-1/2} \Sigma \Sigma_0^{-1/2} = \Sigma_0^{-1/2} (UU^\top + \Psi) \Sigma_0^{-1/2} = Q_1 Q_1^\top + Q_2 Q_2^\top$$

where

$$Q_1 = \Sigma_0^{-1/2} U \quad Q_2 = \Sigma_0^{-1/2} \Psi^{1/2}.$$

We first show that $\|Q_2 Q_2^\top\|_{\text{op}} \leq 1$. To this end, note that $\|Q_2 Q_2^\top\|_{\text{op}} = \|Q_2^\top Q_2\|_{\text{op}} = \|\Psi^{1/2} \Sigma_0^{-1} \Psi^{1/2}\|_{\text{op}}$. Note that $\Psi^{1/2} \Sigma_0^{-1} \Psi^{1/2}$ is a block diagonal matrix with the i th block given by $\Psi_{\{ii\}}^{1/2} (\Psi_{\{ii\}} + U_{\{i\}} U_{\{i\}}^\top)^{-1} \Psi_{\{ii\}}^{1/2}$. Thus, it suffices to show that

$$\|\Psi_{\{ii\}}^{1/2} (\Psi_{\{ii\}} + U_{\{i\}} U_{\{i\}}^\top)^{-1} \Psi_{\{ii\}}^{1/2}\|_{\text{op}} \leq 1, \quad \text{for } i \in [k].$$

To this end, by Woodbury matrix identity,

$$\begin{aligned} &\Psi_{\{ii\}}^{1/2} (\Psi_{\{ii\}} + U_{\{i\}} U_{\{i\}}^\top)^{-1} \Psi_{\{ii\}}^{1/2} \\ &= \Psi_{\{ii\}}^{1/2} [\Psi_{\{ii\}}^{-1} - \Psi_{\{ii\}}^{-1} U_{\{i\}} (I_r + U_{\{i\}}^\top \Psi_{\{ii\}}^{-1} U_{\{i\}})^{-1} U_{\{i\}}^\top \Psi_{\{ii\}}^{-1}] \Psi_{\{ii\}}^{1/2} \\ &= I_{p_i} - \Psi_{\{ii\}}^{-1/2} U_{\{i\}} (I_r + U_{\{i\}}^\top \Psi_{\{ii\}}^{-1} U_{\{i\}})^{-1} U_{\{i\}}^\top \Psi_{\{ii\}}^{-1/2}. \end{aligned}$$

Therefore, $0 \preceq \Psi_{\{ii\}}^{1/2} (\Psi_{\{ii\}} + U_{\{i\}} U_{\{i\}}^\top)^{-1} \Psi_{\{ii\}}^{1/2} \preceq I_{p_i}$, and so $\|Q_2 Q_2^\top\|_{\text{op}} \leq 1$.

By Lemma 13, we have, for $i \geq j + k - 1$,

$$\lambda_i = s_i(R) \leq s_j(Q_1 Q_1^\top) + s_k(Q_2 Q_2^\top).$$

Now we set $j = r + 1, k = 1$ and this gives

$$\lambda_i = s_i(R) \leq s_{r+1}(Q_1 Q_1^\top) + s_1(Q_2 Q_2^\top) = s_1(Q_2 Q_2^\top) \quad \text{for } i \geq r + 1.$$

The last equality holds since we have $\text{rank}(Q_1 Q_1^\top) = r$ and so $s_{r+1}(Q_1 Q_1^\top) = 0$. As a result,

$$\lambda_i = s_i(R) \leq s_1(Q_2 Q_2^\top) = \|Q_2 Q_2^\top\|_{\text{op}} \leq 1.$$

This finishes the proof of the first step.

Step (2): By Lemma 13 and Assumption 2, we have

$$\begin{aligned}\lambda_r = s_r(R) &= s_r(Q_1 Q_1^\top + Q_2 Q_2^\top) \geq s_r(Q_1 Q_1^\top) + s_p(Q_2 Q_2^\top) \\ &> s_r(Q_1 Q_1^\top) = \sigma_r^2(\Sigma_0^{-1/2} U) \geq 1.\end{aligned}$$

Here, the strict inequality holds since $Q_2 Q_2^\top$ is of full rank. This finishes the proof for the r th eigenvalue.

Step (3): Now we calculate the multiplicity of 1. To this end, we resort to Theorem 2.2 in Tian (2004), which gives

$$\begin{aligned}\text{rank}(U^\top - U^\top Y U^\top) &= \text{rank}(U - U Y^\top U) = \text{rank}(U^\dagger U - U_{\{1\}}^\dagger U_{\{1\}} - \cdots - U_{\{k\}}^\dagger U_{\{k\}}) \\ &= \text{rank}(\tilde{U}) + \text{rank}(U) - \sum_{i=1}^k \text{rank}(U_{\{i\}}).\end{aligned}$$

Here

$$\tilde{U} = \begin{bmatrix} 0 & \Sigma_{\{12\}} & \Sigma_{\{13\}} & \cdots & \Sigma_{\{1k\}} \\ \Sigma_{\{21\}} & 0 & \Sigma_{\{23\}} & \cdots & \Sigma_{\{2k\}} \\ \Sigma_{\{31\}} & \Sigma_{\{32\}} & 0 & \cdots & \Sigma_{\{3k\}} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \Sigma_{\{k1\}} & \Sigma_{\{k2\}} & \Sigma_{\{k3\}} & \cdots & 0 \end{bmatrix}.$$

Note that we have

$$\Sigma_{\{ij\}} = U_{\{i\}} U_{\{j\}}^\top, \quad \text{rank}(U) = r,$$

by the latent variable model. Moreover, it is easy to verify that multiplicity of 1 as an eigenvalue of R is the same as multiplicity of 0 as an eigenvalue of \tilde{U} , which in turn equals $p - \text{rank}(\tilde{U})$. Thus, multiplicity of 1 is

$$\begin{aligned}p - \text{rank}(\tilde{U}) &= p - \text{rank}(U^\top - U^\top Y U^\top) + \text{rank}(U) - \sum_{i=1}^k \text{rank}(U_{\{i\}}) \\ &= p - \sum_{i=1}^k \text{rank}(U_{\{i\}}) + r - \text{rank}(U - U Y^\top U).\end{aligned}$$

This finishes the proof of the lemma. ■

Appendix B. Proof of Main Results

We first provide a lemma about the distance defined in (18) which is adapted from Lemmas 5.3 and 5.4 from Tu et al. (2015).

Lemma 14 *For any $U, V \in \mathbb{R}^{p \times r}$, we have*

$$\text{dist}^2(U, V) \leq \frac{1}{2(\sqrt{2} - 1)\sigma_r^2(V)} \|UU^\top - VV^\top\|_{\text{F}}^2.$$

If further that $\text{dist}(U, V) \leq \frac{1}{4}\|V\|_{\text{op}}$, we have

$$\|UU^\top - VV^\top\|_{\text{F}} \leq \frac{9}{4}\|U\|_{\text{op}}\text{dist}(U, V).$$

Before diving into the proof, we present a list of nice events on the intersection of which all desired results hold deterministically. Recall that S is the true row support of A with cardinality s . Let $\mathcal{I} \subseteq \{1, 2, 3, \dots, p\}$ be an index set. We use $\widehat{A}(\mathcal{I}) \in \mathbb{R}^{p \times r}$ to denote the leading r generalized eigenvectors that solve

$$\max \langle \widehat{\Sigma}, LL^\top \rangle \quad \text{such that} \quad L^\top \widehat{\Sigma}_0 L = I_r, \quad \text{supp}(L) \subseteq \mathcal{I}. \quad (32)$$

We also denote the diagonal matrix formed by the first r restricted sample generalized eigenvalues by $\widehat{\Lambda}_r(\mathcal{I})$.

For some sufficiently large constant $C > 0$, define

$$B_1 = \left\{ \text{dist}(A, \widehat{A}(\mathcal{I})) \leq C\sqrt{r} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{|\mathcal{I}| \log p}{n}}, \right. \\ \left. \text{for all } \mathcal{I} \supset S \text{ such that } |\mathcal{I}| \leq 2s' + s \right\}. \quad (33)$$

Note that the event includes all index sets containing true support S and the sizes of which are at most $2s' + s$.

We define event B_2 as

$$B_2 = \left\{ \|\widehat{\Sigma}_{\mathcal{I}\mathcal{I}} - \Sigma_{\mathcal{I}\mathcal{I}}\|_{\text{op}} \vee \|\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}} - \Sigma_{0,\mathcal{I}\mathcal{I}}\|_{\text{op}} \leq C\sqrt{\frac{(2s' + s) \log p}{n}}, \right. \\ \left. \text{for all } \mathcal{I} \subseteq [p] \text{ with } |\mathcal{I}| \leq 2s' + s \right\}. \quad (34)$$

Further we define B_3 as

$$B_3 = \left\{ \|\widehat{\Sigma} - \Sigma\|_{\infty} + \|\Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0\|_{\infty} \right. \\ \left. + \lambda_{r+1} \|\Sigma_0 - \widehat{\Sigma}_0\|_{\infty} + \lambda_{r+1} \|\Sigma_0 A A^\top \Sigma_0 - \widehat{\Sigma}_0 A A^\top \widehat{\Sigma}_0\|_{\infty} \leq C\sqrt{\frac{\log p}{n}} \right\}. \quad (35)$$

Let

$$\widetilde{A} = A(A^\top \widehat{\Sigma}_0 A)^{-1/2}, \quad \widetilde{\Lambda}_r = (A^\top \widehat{\Sigma}_0 A)^{1/2} \Lambda_r (A^\top \widehat{\Sigma}_0 A)^{1/2}. \quad (36)$$

The event B_4 is defined as

$$B_4 = \left\{ \|\widetilde{\Lambda}_r - \Lambda_r\|_{\text{F}} + \lambda_{r+1} \|A^\top \widehat{\Sigma}_0 A - I\|_{\text{F}} + \|\Sigma_0^{1/2}(\widetilde{A} - A)\|_{\text{F}} \leq C\sqrt{\frac{r(s + \log(ep/s))}{n}} \right\} \quad (37)$$

for some sufficiently large constant $C > 0$.

The following lemmas guarantee that all these events occur with high probabilities, uniformly over \mathcal{P}_n . In addition, since $s' \geq s$, Lemma 12 in Gao et al. (2015a) implies that B_2 happens with probability at least $1 - \exp(-C's' \log(ep/s'))$ for some positive constant C' , uniformly over \mathcal{P}_n .

Lemma 15 *Suppose condition (23) holds, there exist constants $C, C' > 0$ such that uniformly over \mathcal{P}_n , B_1 happens with probability at least $1 - \exp(-C' s' \log(ep/s'))$.*

Lemma 16 *Suppose $r\sqrt{\log p/n} \leq c'$ for some sufficiently small constant $c' \in (0, 1)$. Then there exist positive constants C, C' such that uniformly over \mathcal{P}_n , B_3 happens with probability at least $1 - p^{-C'}$.*

Lemma 17 *Suppose $\sqrt{\frac{s+\log(ep/s)}{n}} \leq c'$ for some sufficiently small constant $c' \in (0, 1)$. Then there exist constants $C, C' > 0$ such that uniformly over \mathcal{P}_n , B_4 happens with probability at least $1 - \exp(-C'(s + \log(ep/s)))$.*

In the rest of this section, we will show that Theorem 7 and Corollary 8 hold on event $B_1 \cap B_2$. Theorem 9 holds on event $B_3 \cap B_4$, and Corollary 10 holds on event $B_2 \cap B_3 \cap B_4$. As a result, the entire algorithm with generalized Fantope initialization yields an estimator satisfying the upper bound in Corollary 11 on $B_1 \cap B_2 \cap B_3 \cap B_4$, which holds with probability at least $1 - \exp(-C'(s + \log(ep/s)))$ for some constant $C' > 0$ uniformly over \mathcal{P}_n , by the foregoing lemmas and the union bound. This argument proves Corollary 11 about the whole procedure.

Before proving the main Theorem, we first specifically study effect of the gradient descent step and the hard thresholding step.

The following proposition characterizes the progress in the gradient step.

Proposition 18 *Let $S_t = \text{supp}(\bar{V}_t) \cup \text{supp}(\bar{V}_{t+1}) \cup \text{supp}(A)$ be a super set of the row support of \bar{V}_t . Define*

$$\hat{V}(S_t) = \hat{A}(S_t) \left(I + \frac{\hat{\Lambda}_r(S_t)}{\lambda} \right)^{1/2}.$$

Set

$$\eta \leq \frac{c}{12\lambda_1\nu(c+2)^2} \quad \text{and} \quad \lambda = \frac{\lambda_1}{c}, \quad (38)$$

for some constant $c < 1$. On event B_2 , suppose \bar{V}_t satisfies

$$\text{dist}(\bar{V}_t, \hat{V}(S_t)) \leq \frac{1}{4\sqrt{\nu}} \min \left\{ \frac{c(\lambda_r - \lambda_{r+1})}{\sqrt{2}\lambda_1\nu^2(42 + 25c)}, \frac{\sqrt{1+c}}{2} \right\}, \quad (39)$$

then after the gradient step, we have

$$\text{dist}^2(V_{t+1}^o, \hat{V}(S_t)) \leq \left(1 - \frac{\eta(\lambda_r - \lambda_{r+1})}{4\nu} \right) \text{dist}^2(\bar{V}_t, \hat{V}(S_t)),$$

here we use $V_{t+1}^o \in \mathbb{R}^{p \times r}$ to denote a matrix that has the same entries as those in V_{t+1} on $S_t \times [r]$ and zeros elsewhere.

The following proposition characterizes the effect of hard thresholding.

Proposition 19 *Define V as in (21). Then if we perform hard thresholding by selecting the top s' elements of V_{t+1} , we have*

$$\text{dist}^2(V, \bar{V}_{t+1}) \leq \left(1 + 2\sqrt{\frac{s}{s'}} \left(1 + \sqrt{\frac{s}{s'}} \right) \right) \text{dist}^2(V, V_{t+1}). \quad (40)$$

B.1 Proof of Theorem 7

We first prove the theorem assuming that Propositions 18 and 19 hold. Proofs of the two propositions will be given later in Sections B.3 and B.4. It is worth noting that the results of both propositions hold deterministically on event B_2 .

The first step is to define the effective support in each step. We first define $F_t = \text{supp}(\bar{V}_t)$, recall that S is the true row support of A , then we define the effective restricted set to be

$$S_t = F_t \cup F_{t+1} \cup S.$$

The gradient descent step restricted to S_t can be viewed as

$$V_{t+1, S_t^*} = \bar{V}_{t, S_t^*} - 2\eta(-\hat{\Sigma}_{S_t S_t} \bar{V}_{t, S_t^*} + \lambda \hat{\Sigma}_{0, S_t S_t} \bar{V}_{t, S_t^*} (\bar{V}_{t, S_t^*}^\top \hat{\Sigma}_{0, S_t S_t} \bar{V}_{t, S_t^*} - I_r)).$$

Note that applying hard thresholding on V_{t+1}^o (as defined in Proposition 18) is equivalent to applying hard thresholding on the original V_{t+1} . This allows us to replace the intermediate update by V_{t+1}^o and still obtain the same output sequence \bar{V}_{t+1} . Thus, we will prove instead for the update using $\hat{\Sigma}_{0, S_t S_t}, \hat{\Sigma}_{S_t S_t}$.

Proof Throughout the proof, we assume the event $B_1 \cap B_2$ happens, which occurs with probability at least $1 - \exp(-C' s' \log(ep/s'))$ for some constant $C' > 0$, uniformly over \mathcal{P}_n . As we have mentioned, the conclusions of Propositions 18 and 19 hold on this event, and the remaining arguments in this proof proceed in a deterministic fashion.

We argue by induction on the iteration counter t . Specifically, for $t = 1, 2, \dots$, we will prove that \bar{V}_t satisfies condition (39) and that

$$\text{dist}(\bar{V}_t, V) \leq \xi^{t-1} \text{dist}(\bar{V}_1, V) + \frac{C_1}{1-\xi} \frac{\sqrt{1+\lambda_1^2} \sqrt{1+\lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}} \quad (41)$$

for some positive constant C_1 .

Base Case: By (23), (22), and definition of B_1 in (33), condition (39) is satisfied when $t = 1$. Moreover, when $t = 1$ equation (41) holds trivially.

Induction Step: Suppose that \bar{V}_t satisfies the radius condition (39) and that the induction hypothesis (41) is satisfied at step t . We are to show that (39) and (41) hold for \bar{V}_{t+1} .

In the gradient step, Proposition 18 shows that under radius condition (39) on $\text{dist}(\bar{V}_t, \hat{V}(S_t))$, if we choose the step-size to be $\eta \leq \frac{1}{\beta}$,

$$\text{dist}(V_{t+1}^o, \hat{V}(S_t)) \leq \sqrt{1 - \alpha\eta} \text{dist}(\bar{V}_t, \hat{V}(S_t)) \leq \left(1 - \frac{\alpha\eta}{2}\right) \text{dist}(\bar{V}_t, \hat{V}(S_t)),$$

where

$$\alpha = \frac{(\lambda_r - \lambda_{r+1})}{4\nu}, \quad \beta = 12\lambda\nu \left(2 + \frac{\lambda_1}{\lambda}\right)^2.$$

Recall that

$$V = A \left(I + \frac{1}{\lambda} \Lambda_r\right)^{\frac{1}{2}}, \quad \hat{V}(S_t) = \hat{A}(S_t) \left(I + \frac{1}{\lambda} \hat{\Lambda}_r(S_t)\right)^{\frac{1}{2}}.$$

By (33), on event $B_1 \cap B_2$, since $\lambda = \lambda_1/c$, we have

$$\text{dist}(V, \widehat{V}(S_t)) \leq C_0 \sqrt{r} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{|S_t| \log p}{n}},$$

for some positive constant C_0 . Triangle inequality then leads to

$$\text{dist}(V_{t+1}^o, V) \leq \left(1 - \frac{\alpha\eta}{2}\right) \text{dist}(\bar{V}_t, V) + 2C_0 \sqrt{r} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{|S_t| \log p}{n}}.$$

Turn to the thresholding step. By Proposition 19, we have

$$\begin{aligned} \text{dist}(\bar{V}_{t+1}, V) &\leq \sqrt{1 + 2\sqrt{\frac{s}{s'}} \left(1 + \sqrt{\frac{s}{s'}}\right)} \text{dist}(V_{t+1}^o, V) \\ &\leq \left(1 + \sqrt{\frac{s}{s'}} \left(1 + \sqrt{\frac{s}{s'}}\right)\right) \text{dist}(V_{t+1}^o, V). \end{aligned}$$

Combining the last two displays, we obtain that

$$\begin{aligned} \text{dist}(\bar{V}_{t+1}, V) &\leq \left(1 + \sqrt{\frac{s}{s'}} \left(1 + \sqrt{\frac{s}{s'}}\right)\right) \left[\left(1 - \frac{\alpha\eta}{2}\right) \text{dist}(\bar{V}_t, V) \right. \\ &\quad \left. + 2C_0 \sqrt{r} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{|S_t| \log p}{n}} \right]. \end{aligned}$$

Define

$$\xi = \left(1 + \sqrt{\frac{s}{s'}} \left(1 + \sqrt{\frac{s}{s'}}\right)\right) \left(1 - \frac{\alpha\eta}{2}\right) < 1.$$

Note that we can always ensure $\xi < 1$ by enlarging s' appropriately. Specifically, the choice in the theorem, $s' \geq \frac{16}{\alpha^2 \eta^2} s$, ensures that

$$\xi \leq \left(1 + 2\sqrt{\frac{s}{s'}}\right) \left(1 - \frac{\alpha\eta}{2}\right) \leq \left(1 + \frac{\alpha\eta}{2}\right) \left(1 - \frac{\alpha\eta}{2}\right) = 1 - \frac{\alpha^2 \eta^2}{4}. \quad (42)$$

Since $|S_t| \leq 2s' + s$ and $s' \geq s$, we further obtain that

$$\text{dist}(\bar{V}_{t+1}, V) \leq \xi \text{dist}(\bar{V}_t, V) + C_1 \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}},$$

for some constant $C_1 > 0$. Note that in the above equation C_1 does not depend on t . We now show that it is identical with the constant C_1 in (41). To this end, note that

$$\begin{aligned}
 \text{dist}(\bar{V}_{t+1}, V) &\leq \xi \text{dist}(\bar{V}_t, V) + C_1 \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}} \\
 &\leq \xi \left(\xi^{t-1} \text{dist}(\bar{V}_1, V) + \frac{C_1}{1 - \xi} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}} \right) \\
 &\quad + C_1 \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}} \\
 &\leq \xi^t \text{dist}(\bar{V}_1, V) + \frac{C_1}{1 - \xi} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}}.
 \end{aligned}$$

Here the second inequality is due to the induction hypothesis at the t th iteration. Moreover, it can be ensured that \bar{V}_{t+1} also satisfies (39) since $\xi < 1$ and the extra term is bounded by a sufficiently small constant due to (23). As a result, we have shown that both (39) and (41) are satisfied for all $t \geq 1$.

In summary, on $B_1 \cap B_2$, for any $t \geq 1$,

$$\begin{aligned}
 \text{dist}(\bar{V}_{t+1}, V) &\leq \xi^t \text{dist}(\bar{V}_1, V) + \frac{C_1}{1 - \xi} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}} \\
 &\leq \xi^t \text{dist}(\bar{V}_1, V) + \frac{4C_1}{\alpha^2 \eta^2} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs' \log p}{n}} \\
 &\leq \xi^t \text{dist}(\bar{V}_1, V) + C \left(\frac{s'}{s} \right)^{3/2} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs \log p}{n}}. \quad (43)
 \end{aligned}$$

Here, the second inequality holds since $\frac{1}{1 - \xi} \leq \frac{4}{\alpha^2 \eta^2}$ and the last inequality is due to $s' \geq \frac{16}{\alpha^2 \eta^2} s$. This completes the proof. \blacksquare

B.2 Proof of Corollary 8

Proof We prove the desired result on $B_1 \cap B_2$, which happens with probability at least $1 - \exp(-C'(s' \log(ep/s')))$ for some constant $C' > 0$, uniformly over \mathcal{P}_n . For notational convenience, denote the statistical error rate in (24) by ϵ_n , that is

$$\epsilon_n = \left(\frac{s'}{s} \right)^{3/2} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{rs \log p}{n}}.$$

Under condition (25), the statistical error dominates optimization error in (24). Hence by Theorem 7 $\text{dist}(\bar{V}_t, V)$ is bounded by a constant multiple of ϵ_n , that is,

$$\text{dist} \left(\bar{V}_t, A \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} \right) \leq C_0 \epsilon_n$$

for some constant $C_0 > 0$. Let P be the orthogonal matrix that minimizes the distance, then we can write

$$\bar{V}_t = A \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} P + Q, \quad (44)$$

where $\|Q\|_F \leq C_0 \epsilon_n$.

Since \bar{V}_t is s' row sparse and A is s row sparse, Q is $s + s'$ row sparse. The remaining proof is composed of two steps: (1) bounding $\|(\bar{V}_t^\top \hat{\Sigma}_0 \bar{V}_t)^{-1/2} - P^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P\|_F$ and (2) bounding $\|\bar{V}_t (\bar{V}_t^\top \hat{\Sigma}_0 \bar{V}_t)^{-1/2} - AP\|_F$. The bound in step (2) then gives the desired bound in the statement of the corollary. In the rest of the proof, let $\Delta = (\bar{V}_t^\top \hat{\Sigma}_0 \bar{V}_t)^{-1/2} - P^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P$.

Step (1): By definition we have

$$\begin{aligned} \bar{V}_t^\top \hat{\Sigma}_0 \bar{V}_t &= \left(A \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} P + Q \right)^\top \hat{\Sigma}_0 \left(A \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} P + Q \right) \\ &= P^\top \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} A^\top \hat{\Sigma}_0 A \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} P + Q^\top \hat{\Sigma}_0 A \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} P \\ &\quad + P^\top \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} A^\top \hat{\Sigma}_0 Q + Q^\top \hat{\Sigma}_0 Q. \end{aligned}$$

Since $A^\top \Sigma_0 A = I_r$, we have

$$P^\top \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} A^\top \Sigma_0 A \left(I + \frac{\Lambda_r}{\lambda} \right)^{1/2} P = P^\top \left(I + \frac{\Lambda_r}{\lambda} \right) P.$$

Then we can bound the Frobenius norm for $\bar{V}_t^\top \hat{\Sigma}_0 \bar{V}_t - P^\top (I + \frac{\Lambda_r}{\lambda}) P$ as

$$\begin{aligned} \|\bar{V}_t^\top \hat{\Sigma}_0 \bar{V}_t - P^\top (I + \frac{\Lambda_r}{\lambda}) P\|_F &\leq \underbrace{\|P^\top (I + \frac{\Lambda_r}{\lambda})^{1/2} A^\top (\hat{\Sigma}_0 - \Sigma_0) A (I + \frac{\Lambda_r}{\lambda})^{1/2} P\|_F}_{\text{Term I}} \\ &\quad + 2 \underbrace{\|Q^\top \hat{\Sigma}_0 A (I + \frac{\Lambda_r}{\lambda})^{1/2}\|_F}_{\text{Term II}} + \underbrace{\|Q^\top \hat{\Sigma}_0 Q\|_F}_{\text{Term III}}. \end{aligned}$$

We bound the three terms on the right side of the last display separately. Since A is s sparse, we have that on event B_2 ,

$$\begin{aligned} \|P^\top (I + \frac{\Lambda_r}{\lambda})^{1/2} A^\top (\hat{\Sigma}_0 - \Sigma_0) A (I + \frac{\Lambda_r}{\lambda})^{1/2} P\|_{\text{op}} &\leq \|I + \frac{\Lambda_r}{\lambda}\|_{\text{op}} \|A_{S^*}^\top (\hat{\Sigma}_{0,SS} - \Sigma_{0,SS}) A_{S^*}\|_{\text{op}} \\ &\leq C_1 \sqrt{\frac{s \log p}{n}}. \end{aligned}$$

Since A is of rank r , we can then bound the Frobenius norm of Term I as

$$\|P^\top (I + \frac{\Lambda_r}{\lambda})^{1/2} A^\top (\hat{\Sigma}_0 - \Sigma_0) A (I + \frac{\Lambda_r}{\lambda})^{1/2} P\|_F \leq C_1 \sqrt{\frac{rs \log p}{n}} \leq C_1 \epsilon_n.$$

For Term II, since Q is $s + s'$ sparse, we notice that on event B_2 and under condition (23), we can bound it as

$$\|Q^\top \widehat{\Sigma}_0 A (I + \frac{\Lambda_r}{\lambda})^{1/2}\|_F \leq C_2 \|Q\|_F$$

for some constant $C_2 > 0$. Term III is also dominated by the same upper bound. On B_2 , the operator norms of $\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t$ and $(\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t)^{-1}$ are both upper bounded by a positive constant. In addition, by (19), the operator norms of $I + \Lambda_r/\lambda$ and $(I + \Lambda_r/\lambda)^{-1}$ are also upper bounded by a positive constant. So we conclude that

$$\begin{aligned} \|\Delta\|_F &\leq \|(\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t)^{-1/2}\|_{\text{op}} \|(\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t)^{1/2} - P^\top (I + \frac{\Lambda_r}{\lambda})^{1/2} P\|_F \|P^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P\|_{\text{op}} \\ &\leq C_3 \|\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t - P^\top (I + \frac{\Lambda_r}{\lambda}) P\|_F \leq C_4 \epsilon_n. \end{aligned}$$

where the second to last inequality is due to Lemma 2 in Supplement of Gao et al. (2017).

Step (2): Now we bound $\|\widehat{A}_t - AP\|_F$ as follows. Recall we have defined that $\Delta = (\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t)^{-1/2} - P^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P$.

$$\begin{aligned} \|\widehat{A}_t - AP\|_F &= \|\overline{V}_t (\overline{V}_t^\top \widehat{\Sigma}_0 \overline{V}_t)^{-1/2} - AP\|_F \\ &= \|(A(I + \frac{\Lambda_r}{\lambda})^{1/2} P + Q)(P^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P + \Delta) - AP\|_F \\ &= \|AP + QP^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P + A(I + \frac{\Lambda_r}{\lambda})^{1/2} P\Delta + Q\Delta - AP\|_F \\ &\leq \|QP^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P\|_F + \|A(I + \frac{\Lambda_r}{\lambda})^{1/2} P\Delta\|_F + \|Q\Delta\|_F \\ &\leq \|Q\|_F \|P^\top (I + \frac{\Lambda_r}{\lambda})^{-1/2} P\|_{\text{op}} + \|A(I + \frac{\Lambda_r}{\lambda})^{1/2} P\|_{\text{op}} \|\Delta\|_F + \|Q\|_F \|\Delta\|_{\text{op}} \\ &\leq C_5 \epsilon_n \end{aligned}$$

for some constant $C_5 > 0$, due to the bounds on $\|Q\|_F$ and $\|\Delta\|_F$ that we have established in step (1).

Combining the results above we deduce that

$$\text{dist}(\widehat{A}_t, A) \leq \|\widehat{A}_t - AP\|_F \leq C_5 \epsilon_n$$

with probability at least $1 - \exp(-C'(s' \log(ep/s')))$ for some positive constants C_5 and C' , uniformly over \mathcal{P}_n . This finishes our proof. \blacksquare

B.3 Proof of Proposition 18

Proof Throughout the whole proof, we work on event B_2 defined in (34) which happens with probability at least $1 - \exp(-C' s \log(ep/s'))$ for some constant $C' > 0$, uniformly over \mathcal{P}_n . As mentioned before, the gradient step is equivalent to replacing all covariance matrices by $\widehat{\Sigma}_{S_t, S_t}, \widehat{\Sigma}_{0, S_t, S_t}$ respectively due to sparsity of $\overline{V}_t, \overline{V}_{t+1}$ (since the output sequence \overline{V}_t remains unaltered after this substitution). Effectively, at t th step of gradient descent

the relevant support is S_t since outside this set the output matrix has row equal to 0. We denote principal submatrices $\widehat{\Sigma}_{S_t S_t}, \widehat{\Sigma}_{0, S_t S_t}$ by $\widehat{\Sigma}_t \in \mathbb{R}^{|S_t| \times |S_t|}, \widehat{\Sigma}_{0,t} \in \mathbb{R}^{|S_t| \times |S_t|}$ for simplicity in the proof and define the restricted Lagrangian function f_t to be

$$f_t(L) = -\langle \widehat{\Sigma}_t, LL^\top \rangle + \frac{\lambda}{2} \|L^\top \widehat{\Sigma}_{0,t} L - I_r\|_F^2,$$

where $L \in \mathbb{R}^{|S_t| \times r}$. As a result,

$$\frac{1}{2} \nabla f_t(L) = -\widehat{\Sigma}_t L + \lambda \widehat{\Sigma}_{0,t} L (L^\top \widehat{\Sigma}_{0,t} L - I_r). \quad (45)$$

Throughout the proof of Proposition 18, we will work on the restricted function and its gradient. We denote $L_{t+1} = L_t - \eta \nabla f_t(L_t)$. Then we notice that $L_t = \bar{V}_{t, S_t^*}, L_{t+1} = V_{t+1, S_t^*}$ by our submatrix notation respectively. Our proof will work on distance involving L_t, L_{t+1} which transfers to the desired bound as stated in this proposition.

Before the proof we revisit a lemma characterizing the effect of gradient descent. The following lemma is adapted from Lemma 4 in Chi et al. (2019) and it is an extension of the gradient descent condition from vectors to rank r matrices. For notational convenience, we define L_t^* to be a global minimizer of function $f_t(L)$ (which will be calculated later) and

$$H_X = \operatorname{argmin}_{H \in \mathcal{O}(r)} \|XH - L_t^*\|_F.$$

We define a function $f(L)$ to be β smooth at L if for all Z , we have

$$\operatorname{vec}(Z)^\top \nabla^2 f(L) \operatorname{vec}(Z) \leq \beta \|Z\|_F^2.$$

Lemma 20 *Suppose that f_t is β smooth within a ball $B(L_t^*) = \{L : \|L - L_t^*\|_F \leq R\}$ and that $\nabla f_t(L)P = \nabla f_t(LP)$ for any orthonormal matrix P . Assume that for any $L \in B(L_t^*)$ and any Z , we have*

$$\operatorname{vec}(ZH_Z - L_t^*)^\top \nabla^2 f_t(L) \operatorname{vec}(ZH_Z - L_t^*) \geq \alpha \|ZH_Z - L_t^*\|_F^2.$$

In addition, if $\eta \leq \frac{1}{\beta}$, then using gradient descent with $\operatorname{dist}(L_t, L_t^) \leq R$, we have*

$$\operatorname{dist}^2(L_{t+1}, L_t^*) \leq (1 - \alpha\eta) \operatorname{dist}^2(L_t, L_t^*).$$

Moreover, with $\operatorname{dist}(L_0, L_t^) \leq R$, we have*

$$\operatorname{dist}^2(L_t, L_t^*) \leq (1 - \alpha\eta)^t \operatorname{dist}^2(L_0, L_t^*).$$

In Lemma 4 of Chi et al. (2019) the condition on gradient descent is $L_t \in B(L_t^*)$. Here we generalize it to $\operatorname{dist}(L_t, L_t^*) \leq R$ and the proof follows without change of the original proof. By (45), it is straightforward to verify the condition $\nabla f_t(L)P = \nabla f_t(LP)$. The remaining proof is composed of three steps: (1) deriving the expression for $\operatorname{vec}(Z)^\top \nabla^2 f_t(L) \operatorname{vec}(Z)$, (2) verifying the smoothness condition, and (3) verifying the condition on strong convexity. We check the radius condition at the end of the proof.

Step (1) Recall (45). As a result,

$$\frac{1}{2}\text{vec}\nabla f_t(L) = -(I_r \otimes \widehat{\Sigma}_t)\text{vec}(L) + \lambda(I_r \otimes \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t})\text{vec}(L) - \lambda(I_r \otimes \widehat{\Sigma}_{0,t})\text{vec}(L).$$

The main calculation is to deal with

$$\text{vec}(\widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t}L) = (I \otimes \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t})\text{vec}(L).$$

We now directly compute this expression as follows: since $L \in \mathbb{R}^{|S_t| \times r}$, we have $\text{vec}(L) = [l_1^\top, l_2^\top, \dots, l_r^\top]^\top$ where l_i is the i th column of the matrix. Following this notation, we can write

$$\begin{aligned} \text{vec}(\widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t}L) &= (I \otimes \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t})\text{vec}(L) \\ &= \begin{bmatrix} \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t} & 0 & \dots & 0 \\ 0 & \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t} & \dots & 0 \\ 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t} \end{bmatrix} \begin{bmatrix} l_1 \\ l_2 \\ l_3 \\ \dots \\ l_r \end{bmatrix}. \end{aligned}$$

As a result,

$$\text{vec}(\widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t}L) = \begin{bmatrix} \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t}l_1 \\ \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t}l_2 \\ \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t}l_3 \\ \dots \\ \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t}l_r \end{bmatrix} = \begin{bmatrix} \widehat{\Sigma}_{0,t} \sum_{i=1}^r l_i l_i^\top \widehat{\Sigma}_{0,t}l_1 \\ \widehat{\Sigma}_{0,t} \sum_{i=1}^r l_i l_i^\top \widehat{\Sigma}_{0,t}l_2 \\ \widehat{\Sigma}_{0,t} \sum_{i=1}^r l_i l_i^\top \widehat{\Sigma}_{0,t}l_3 \\ \dots \\ \widehat{\Sigma}_{0,t} \sum_{i=1}^r l_i l_i^\top \widehat{\Sigma}_{0,t}l_r \end{bmatrix}.$$

Now we can calculate the derivative $\frac{\partial \text{vec}\nabla f_t(L)}{\partial \text{vec}(L)}$, note that we can do this block by block: the jk block entry of the Hessian is just $\frac{\partial \text{vec}\nabla f_t(L)_j}{\partial l_k}$. Since

$$\frac{1}{2}\text{vec}\nabla f_t(L) = -(I_r \otimes \widehat{\Sigma}_t)\text{vec}(L) + \lambda(I_r \otimes \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t})\text{vec}(L) - \lambda(I_r \otimes \widehat{\Sigma}_{0,t})\text{vec}(L),$$

the only term we have to deal with is the middle one. By previous calculations, we have

$$\frac{\partial \widehat{\Sigma}_{0,t} \sum_{i=1}^r l_i l_i^\top \widehat{\Sigma}_{0,t}l_j}{\partial l_k} = \frac{\partial \widehat{\Sigma}_{0,t}l_k l_k^\top \widehat{\Sigma}_{0,t}l_j}{\partial l_k} = l_j^\top \widehat{\Sigma}_{0,t}l_k \widehat{\Sigma}_{0,t} + \widehat{\Sigma}_{0,t}l_j l_k^\top \widehat{\Sigma}_{0,t},$$

when $j \neq k$ and

$$\begin{aligned} \frac{\partial \widehat{\Sigma}_{0,t} \sum_{i=1}^r l_i l_i^\top \widehat{\Sigma}_{0,t}l_j}{\partial l_j} &= \sum_{i \neq j} \widehat{\Sigma}_{0,t}l_i l_i^\top \widehat{\Sigma}_{0,t} + 2\widehat{\Sigma}_{0,t}l_j l_j^\top \widehat{\Sigma}_{0,t} + l_j^\top \widehat{\Sigma}_{0,t}l_j \widehat{\Sigma}_{0,t} \\ &= \widehat{\Sigma}_{0,t}LL^\top\widehat{\Sigma}_{0,t} + \widehat{\Sigma}_{0,t}l_j l_j^\top \widehat{\Sigma}_{0,t} + l_j^\top \widehat{\Sigma}_{0,t}l_j \widehat{\Sigma}_{0,t}. \end{aligned}$$

As a result, we can combine the above results to obtain the Hessian

$$\begin{aligned} \frac{1}{2}\nabla^2 f_t(L) &= \frac{1}{2} \frac{\partial \text{vec} \nabla f_t(L)}{\partial \text{vec}(L)} = - \begin{bmatrix} \widehat{\Sigma}_t & \dots & 0 \\ 0 & \widehat{\Sigma}_t & 0 \\ 0 & \dots & \widehat{\Sigma}_t \end{bmatrix} - \lambda \begin{bmatrix} \widehat{\Sigma}_{0,t} & \dots & 0 \\ 0 & \widehat{\Sigma}_{0,t} & 0 \\ 0 & \dots & \widehat{\Sigma}_{0,t} \end{bmatrix} \\ &+ \begin{bmatrix} \widehat{\Sigma}_{0,t} L L^\top \widehat{\Sigma}_{0,t} + \widehat{\Sigma}_{0,t} l_1 l_1^\top \widehat{\Sigma}_{0,t} + l_1^\top \widehat{\Sigma}_{0,t} l_1 \widehat{\Sigma}_{0,t} & \dots & l_1^\top \widehat{\Sigma}_{0,t} l_r \widehat{\Sigma}_{0,t} + \widehat{\Sigma}_{0,t} l_1 l_r^\top \widehat{\Sigma}_{0,t} \\ \dots & \dots & \dots \\ \widehat{\Sigma}_{0,t} l_r l_r^\top \widehat{\Sigma}_{0,t} & \dots & \widehat{\Sigma}_{0,t} L L^\top \widehat{\Sigma}_{0,t} + \widehat{\Sigma}_{0,t} l_r l_r^\top \widehat{\Sigma}_{0,t} \\ + l_r^\top \widehat{\Sigma}_{0,t} l_1 \widehat{\Sigma}_{0,t} & \dots & + l_r^\top \widehat{\Sigma}_{0,t} l_r \widehat{\Sigma}_{0,t} \end{bmatrix}. \end{aligned}$$

Now we proceed to calculate $\text{vec}(Z)^\top \nabla^2 f_t(L) \text{vec}(Z)$. We have

$$\text{vec}(Z)^\top \nabla^2 f_t(L) \text{vec}(Z) = \sum_{i,j} z_i^\top \nabla^2 f_t(L)_{ij} z_j$$

where z_i is the i th column of Z . Substituting the expression on the Hessian, we have

$$\begin{aligned} \frac{1}{2} \text{vec}(Z)^\top \nabla^2 f_t(L) \text{vec}(Z) &= - \sum_{i=1}^r z_i^\top \widehat{\Sigma}_t z_i - \lambda \left(\sum_{i=1}^r z_i^\top \widehat{\Sigma}_{0,t} z_i \right) + \lambda \left(\sum_{i \neq j} l_i^\top \widehat{\Sigma}_{0,t} l_j z_i^\top \widehat{\Sigma}_{0,t} z_j \right. \\ &+ \sum_{i \neq j} z_i^\top \widehat{\Sigma}_{0,t} l_j l_i^\top \widehat{\Sigma}_{0,t} z_j + \sum_{i=1}^r (z_i^\top \widehat{\Sigma}_{0,t} l_i)^2 \\ &\left. + \sum_{i=1}^r z_i^\top \widehat{\Sigma}_{0,t} L L^\top \widehat{\Sigma}_{0,t} z_i + \sum_{i=1}^r (l_i^\top \widehat{\Sigma}_{0,t} l_i) (z_i^\top \widehat{\Sigma}_{0,t} l_i) \right). \end{aligned}$$

Now we claim the following simplification:

$$\begin{aligned} \frac{1}{2} \text{vec}(Z)^\top \nabla^2 f_t(L) \text{vec}(Z) &= g_t(L, Z) \\ &= - \langle \widehat{\Sigma}_t, Z Z^\top \rangle - \lambda \langle \widehat{\Sigma}_{0,t}, Z Z^\top \rangle + \lambda \langle Z Z^\top, \widehat{\Sigma}_{0,t} L L^\top \widehat{\Sigma}_{0,t} \rangle \\ &\quad + \lambda \langle L^\top \widehat{\Sigma}_{0,t} L, Z^\top \widehat{\Sigma}_{0,t} Z \rangle + \lambda \langle Z^\top \widehat{\Sigma}_{0,t} L, L^\top \widehat{\Sigma}_{0,t} Z \rangle. \end{aligned}$$

Now we begin to prove this claim. We expand the above expression as follows:

$$\begin{aligned}
 g_t(L, Z) &= -\langle \widehat{\Sigma}_t, ZZ^\top \rangle - \lambda(\langle \widehat{\Sigma}_{0,t}, ZZ^\top \rangle - \langle ZZ^\top, \widehat{\Sigma}_{0,t}LL^\top \widehat{\Sigma}_{0,t} \rangle \\
 &\quad - \langle L^\top \widehat{\Sigma}_{0,t}L, Z^\top \widehat{\Sigma}_{0,t}Z \rangle - \langle Z^\top \widehat{\Sigma}_{0,t}L, L^\top \widehat{\Sigma}_{0,t}Z \rangle) \\
 &= -\sum_{i=1}^r z_i^\top \widehat{\Sigma}_t z_i - \lambda \sum_{i=1}^r z_i^\top \widehat{\Sigma}_{0,t} z_i + \lambda \sum_{i,j} z_i^\top \widehat{\Sigma}_{0,t} l_j l_j^\top \widehat{\Sigma}_{0,t} z_i + \lambda \sum_{i,j} z_i^\top \widehat{\Sigma}_{0,t} l_j l_i^\top \widehat{\Sigma}_{0,t} z_j \\
 &\quad + \lambda \sum_{i \neq j} l_i^\top \widehat{\Sigma}_{0,t} l_j z_i^\top \widehat{\Sigma}_{0,t} z_j + \lambda \sum_{i=1}^r (l_i^\top \widehat{\Sigma}_{0,t} l_i)(z_i^\top \widehat{\Sigma}_{0,t} z_i) \\
 &= -\sum_{i=1}^r z_i^\top \widehat{\Sigma}_t z_i - \lambda \sum_{i=1}^r z_i^\top \widehat{\Sigma}_{0,t} z_i + \lambda (\sum_{i \neq j} l_i^\top \widehat{\Sigma}_{0,t} l_j z_i^\top \widehat{\Sigma}_{0,t} z_j + \sum_{i \neq j} z_i^\top \widehat{\Sigma}_{0,t} l_j l_i^\top \widehat{\Sigma}_{0,t} z_j \\
 &\quad + \sum_{i=1}^r (l_i^\top \widehat{\Sigma}_{0,t} l_i)(z_i^\top \widehat{\Sigma}_{0,t} z_i) + \sum_{i=1}^r z_i^\top \widehat{\Sigma}_{0,t} \sum_{j=1}^r l_j l_j^\top \widehat{\Sigma}_{0,t} z_i + \sum_{i=1}^r (z_i^\top \widehat{\Sigma}_{0,t} l_i)^2) \\
 &= -\sum_{i=1}^r z_i^\top \widehat{\Sigma}_t z_i - \lambda (\sum_{i=1}^r z_i^\top \widehat{\Sigma}_{0,t} z_i) + \lambda (\sum_{i \neq j} l_i^\top \widehat{\Sigma}_{0,t} l_j z_i^\top \widehat{\Sigma}_{0,t} z_j + \sum_{i \neq j} z_i^\top \widehat{\Sigma}_{0,t} l_j l_i^\top \widehat{\Sigma}_{0,t} z_j \\
 &\quad + \sum_{i=1}^r (z_i^\top \widehat{\Sigma}_{0,t} l_i)^2 + \sum_{i=1}^r z_i^\top \widehat{\Sigma}_{0,t} LL^\top \widehat{\Sigma}_{0,t} z_i + \sum_{i=1}^r (l_i^\top \widehat{\Sigma}_{0,t} l_i)(z_i^\top \widehat{\Sigma}_{0,t} l_i)) \\
 &= \frac{1}{2} \text{vec}(Z)^\top \nabla^2 f_t(L) \text{vec}(Z).
 \end{aligned}$$

This completes the first step.

Step (2) In view of the lemma, our next task would be to bound the smoothness parameter in a neighborhood of L_t^* . The neighborhood will be defined by the distance $\|\widehat{\Sigma}_{0,t}^{1/2} L_t^* - \widehat{\Sigma}_{0,t}^{1/2} L\|_F \leq \delta$, (we define in this unusual way due to the normalization constraint and specific δ given by the condition will be explained later), which by triangle inequality gives

$$\|\widehat{\Sigma}_{0,t}^{1/2} L_t^*\|_{\text{op}} - \delta \leq \|\widehat{\Sigma}_{0,t}^{1/2} L\|_{\text{op}} \leq \|\widehat{\Sigma}_{0,t}^{1/2} L_t^*\|_{\text{op}} + \delta.$$

We first find this global minimizer L_t^* up to a rotation matrix. Setting the gradient equal to 0, we have any critical point must satisfy the following equation

$$\widehat{\Sigma}_t L = \lambda \widehat{\Sigma}_{0,t} L (L^\top \widehat{\Sigma}_{0,t} L - I_r).$$

From the above equation, we deduce that the column space of global minimizer should coincide with some generalized eigenvectors (not necessarily leading ones). Hence without loss of generality we assume the global minimizer of function is achieved at when $L_t^* = \widehat{L}_t D_t$ for D_t being an invertible matrix and \widehat{L}_t being the generalized eigenvectors with eigenvalues $\widehat{\lambda}_{\mathcal{I}}$ for $\mathcal{I} \subseteq [p]$, $|\mathcal{I}| = r$ for sample covariance matrices. Then we have

$$\widehat{\Sigma}_t \widehat{L}_t D_t = \lambda \widehat{\Sigma}_{0,t} \widehat{L}_t D_t (D_t^2 - I_r)$$

and this gives

$$\widehat{\Sigma}_{0,t} \widehat{L}_t \widehat{\Lambda}_{\mathcal{I}} D_t = \lambda \widehat{\Sigma}_{0,t} \widehat{L}_t D_t (D_t^2 - I_r).$$

Here we abuse the notation a bit to denote diagonal matrices with entries $\widehat{\lambda}_{\mathcal{I}}$ to be $\widehat{\Lambda}_{\mathcal{I}}$. We deduce that $D_t = (I_r + \frac{1}{\lambda}\widehat{\Lambda}_{\mathcal{I}})^{\frac{1}{2}}$. This is true for any critical point and now we will show that the global minimizer is achieved at \mathcal{I} being $\{1, 2, \dots, r\}$. To see this, note that

$$\begin{aligned} f_t(L^*) &= f_t(\widehat{L}_t D_t) = -\langle \widehat{\Sigma}_t, \widehat{L}_t D_t D_t^\top \widehat{L}_t^\top \rangle + \frac{\lambda}{2} \|D_t^\top \widehat{L}_t^\top \widehat{\Sigma}_0 \widehat{L}_t D_t - I_r\|_{\mathbb{F}}^2 \\ &= -\text{Tr}(D_t^\top \widehat{\Lambda}_{\mathcal{I}} D_t) + \frac{\lambda}{2} \|D_t^\top D_t - I_r\|_{\mathbb{F}}^2 = -\text{Tr}(\widehat{\Lambda}_{\mathcal{I}}) - \frac{1}{\lambda} \text{Tr}(\widehat{\Lambda}_{\mathcal{I}}^2) + \frac{1}{2\lambda} \text{Tr}(\widehat{\Lambda}_{\mathcal{I}}^2). \end{aligned}$$

We notice that the above quantity is minimized only when $\widehat{\Lambda}_{\mathcal{I}} = \widehat{\Lambda}_r(S_t)$, that is, when we are selecting the leading r generalized eigenvectors of $\widehat{\Sigma}_t$ with respect to $\widehat{\Sigma}_{0,t}$. As a result,

$$D_t = \left(I_r + \frac{1}{\lambda} \widehat{\Lambda}_r(S_t) \right)^{\frac{1}{2}}, \quad L_t^* = \widehat{A}(S_t)_{S_t^*} \left(I_r + \frac{1}{\lambda} \widehat{\Lambda}_r(S_t) \right)^{\frac{1}{2}} = \widehat{V}(S_t)_{S_t^*},$$

where $\widehat{\Lambda}_r(S_t)$ is the diagonal matrix with entries being first r generalized eigenvalues for sample covariance matrices as specified before. According to our definition, $\widehat{A}(S_t)_{S_t^*}$ and $\widehat{V}(S_t)_{S_t^*}$ are of size $|S_t| \times r$. In the rest of the proof, we denote $\widehat{V}(S_t)_{S_t^*}$ by \widehat{V} for simplicity. Similarly, we slightly abuse notation and abbreviate $\widehat{\Lambda}_r(S_t), \widehat{A}(S_t)_{S_t^*}, \widehat{B}(S_t)_{S_t^*}$ as $\widehat{\Lambda}_r, \widehat{A}, \widehat{B}$ in the rest of this proof. Then we have $\|\widehat{\Sigma}_{0,t}^{1/2} L_t^*\|_{\text{op}} = \|\widehat{\Sigma}_{0,t}^{1/2} \widehat{V}\|_{\text{op}} = \sqrt{1 + \frac{\widehat{\lambda}_1}{\lambda}} \leq 1 + \frac{\lambda_1}{\lambda}$, on event B_2 and assumption in the theorem. Here with slight abuse of notation, we use $\widehat{\lambda}_i$ to denote the i th restricted sample generalized eigenvalue.

Now we can bound the smoothness parameter from above by controlling each term in Hessian matrix. Since we are working on event B_2 , we have that $\frac{1}{2}\nu \leq \|\widehat{\Sigma}_t\|_{\text{op}} \leq 2\nu$ and $\frac{1}{2}\nu \leq \|\widehat{\Sigma}_{0,t}\|_{\text{op}} \leq 2\nu$. Similar bound holds for minimum restricted sample generalized eigenvalue. When $\|\widehat{\Sigma}_{0,t}^{1/2} L_t^* - \widehat{\Sigma}_{0,t}^{1/2} L\|_{\mathbb{F}} \leq \delta$, we have

$$\begin{aligned} g_t(L, Z) &= -\langle \widehat{\Sigma}_t, ZZ^\top \rangle - \lambda \langle \widehat{\Sigma}_{0,t}, ZZ^\top \rangle + \lambda \langle ZZ^\top, \widehat{\Sigma}_{0,t} L L^\top \widehat{\Sigma}_{0,t} \rangle + \lambda \langle L^\top \widehat{\Sigma}_{0,t} L, Z^\top \widehat{\Sigma}_{0,t} Z \rangle \\ &\quad + \lambda \langle Z^\top \widehat{\Sigma}_{0,t} L, L^\top \widehat{\Sigma}_{0,t} Z \rangle \\ &\leq 0 + 0 + \lambda \|Z\|_{\mathbb{F}}^2 \|L^\top \widehat{\Sigma}_{0,t}\|_{\text{op}}^2 + \lambda \|L^\top \widehat{\Sigma}_{0,t} L Z^\top\|_{\mathbb{F}} \|\widehat{\Sigma}_{0,t} Z\|_{\mathbb{F}} + \lambda \|Z^\top \widehat{\Sigma}_{0,t} L\|_{\mathbb{F}} \|L^\top \widehat{\Sigma}_{0,t} Z\|_{\mathbb{F}} \\ &\leq 4\lambda\nu(\delta + \|\widehat{\Sigma}_{0,t}^{1/2} \widehat{V}\|_{\text{op}})^2 \|Z\|_{\mathbb{F}}^2 + 2\lambda(\delta + \|\widehat{\Sigma}_{0,t}^{1/2} \widehat{V}\|_{\text{op}})^2 \nu \|Z\|_{\mathbb{F}}^2 \\ &\leq (6\lambda\nu(\delta + \|\widehat{\Sigma}_{0,t}^{1/2} \widehat{V}\|_{\text{op}})^2) \|Z\|_{\mathbb{F}}^2 = \frac{1}{2}\beta \|Z\|_{\mathbb{F}}^2. \end{aligned}$$

Then we can upper bound the Hessian eigenvalue by

$$\beta = 12\lambda\nu \left(\delta + 1 + \frac{\lambda_1}{\lambda} \right)^2.$$

Step (3) To derive the strong convexity parameter α , we start from the function evaluated at the global minimizer $L = \widehat{V}$. Define $\widetilde{Z} = ZH_Z - \widehat{V}$, we now lower bound $g_t(\widehat{V}, \widetilde{Z})$ by a

constant multiple of $\frac{1}{2}\|\tilde{Z}\|_{\mathbb{F}}^2$. To do this, we substitute the expression into g_t and obtain

$$\begin{aligned}
 g_t(\widehat{V}, \tilde{Z}) &= -\langle \widehat{\Sigma}_t + \lambda \widehat{\Sigma}_{0,t}, \tilde{Z} \tilde{Z}^\top \rangle + \lambda \langle \widehat{V}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \tilde{Z}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle + \lambda \langle \tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V} \rangle \\
 &\quad + \lambda \langle \tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle \\
 &= -\langle \widehat{\Sigma}_t + \lambda \widehat{\Sigma}_{0,t}, \tilde{Z} \tilde{Z}^\top \rangle + \lambda \langle I + \frac{1}{\lambda} \widehat{\Lambda}_r, \tilde{Z}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle + \lambda \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t} \widehat{V} \widehat{V}^\top \widehat{\Sigma}_{0,t} \rangle \\
 &\quad + \lambda \langle \tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle \\
 &= \underbrace{-\langle \widehat{\Sigma}_t, \tilde{Z} \tilde{Z}^\top \rangle + \langle \widehat{\Lambda}_r, \tilde{Z}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle + \lambda \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t} \widehat{V} \widehat{V}^\top \widehat{\Sigma}_{0,t} \rangle}_{\text{Term I}} + \underbrace{\lambda \langle \tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle}_{\text{Term II}}.
 \end{aligned}$$

We deal with Term I and Term II separately. To simplify Term I, we recall that

$$\widehat{\Sigma}_{0,t}^{-1} \widehat{\Sigma}_t \widehat{\Sigma}_{0,t}^{-1} = \widehat{A} \widehat{\Lambda}_r \widehat{A}^\top + \widehat{B} \widehat{\Lambda} \widehat{B}^\top,$$

where we have defined the remaining generalized eigenvectors by \widehat{B} and the rest of eigenvalues by diagonal entries of $\widehat{\Lambda}$, both on the restricted set S_t . Here we also omit the dependence on t . By the definition of restricted sample generalized eigenvectors, we have

$$\begin{aligned}
 \text{Term I} &= \langle \widehat{\Lambda}_r, \tilde{Z}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle + \langle \tilde{Z} \tilde{Z}^\top, \lambda \widehat{\Sigma}_{0,t} \widehat{V} \widehat{V}^\top \widehat{\Sigma}_{0,t} - \widehat{\Sigma}_t \rangle \\
 &= \langle \tilde{Z} \widehat{\Lambda}_r \tilde{Z}^\top, \widehat{\Sigma}_{0,t} \rangle + \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t} (\lambda \widehat{A} (I + \widehat{\Lambda}_r / \lambda) \widehat{A}^\top - \widehat{\Sigma}_{0,t}^{-1} \widehat{\Sigma} \widehat{\Sigma}_{0,t}^{-1}) \widehat{\Sigma}_{0,t} \rangle \\
 &\geq \widehat{\lambda}_r \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t} \rangle + \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t} (\lambda \widehat{A} (I + \widehat{\Lambda}_r / \lambda) \widehat{A}^\top - (\widehat{A} \widehat{\Lambda}_r \widehat{A}^\top + \widehat{B} \widehat{\Lambda} \widehat{B}^\top)) \widehat{\Sigma}_{0,t} \rangle \\
 &= \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t}^{1/2} (\widehat{\lambda}_r I + \widehat{\Sigma}_{0,t}^{1/2} \widehat{A} (\lambda I) \widehat{A}^\top \widehat{\Sigma}_{0,t}^{1/2} - \widehat{\Sigma}_{0,t}^{1/2} \widehat{B} \widehat{\Lambda} \widehat{B} \widehat{\Sigma}_{0,t}^{1/2}) \widehat{\Sigma}_{0,t}^{1/2} \rangle \\
 &= \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t}^{1/2} (\widehat{\lambda}_r (\widehat{\Sigma}_{0,t}^{1/2} \widehat{A} \widehat{A}^\top \widehat{\Sigma}_{0,t}^{1/2} + \widehat{\Sigma}_{0,t}^{1/2} \widehat{B} \widehat{B}^\top \widehat{\Sigma}_{0,t}^{1/2}) + \lambda \widehat{\Sigma}_{0,t}^{1/2} \widehat{A} \widehat{A}^\top \widehat{\Sigma}_{0,t}^{1/2} \\
 &\quad - \widehat{\Sigma}_{0,t}^{1/2} \widehat{B} \widehat{\Lambda} \widehat{B}^\top \widehat{\Sigma}_{0,t}^{1/2}) \widehat{\Sigma}_{0,t}^{1/2} \rangle \\
 &= \langle \tilde{Z} \tilde{Z}^\top, \widehat{\Sigma}_{0,t}^{1/2} (\widehat{\Sigma}_{0,t}^{1/2} \widehat{A} (\lambda + \widehat{\lambda}_r) I \widehat{A}^\top \widehat{\Sigma}_{0,t}^{1/2} + \widehat{\Sigma}_{0,t}^{1/2} \widehat{B} (\widehat{\lambda}_r - \widehat{\Lambda}) \widehat{B}^\top \widehat{\Sigma}_{0,t}^{1/2}) \widehat{\Sigma}_{0,t}^{1/2} \rangle \\
 &\geq \frac{1}{2\nu} \langle \tilde{Z} \tilde{Z}^\top, (\widehat{\lambda}_r - \widehat{\lambda}_{r+1}) I \rangle = \frac{1}{2\nu} (\widehat{\lambda}_r - \widehat{\lambda}_{r+1}) \|\tilde{Z}\|_{\mathbb{F}}^2.
 \end{aligned}$$

The first inequality is due to the fact we only select the first r generalized eigenvalues, the second inequality follows from the fact that the second term can be bounded using the eigengap since $\widehat{\Lambda}$ is a diagonal matrix containing the last $|S_t| - r$ generalized eigenvalues. As a result, Term I can be bounded in terms of eigengap at sample level. Note that the above inequality holds for any \tilde{Z} .

Now we bound Term II. Here we use the fact that H_Z is the solution to $\min_{P \in \mathcal{O}(r)} \|ZP - \widehat{V}\|_{\mathbb{F}}$. By definition of \tilde{Z} , we hope to bound $\lambda \langle (ZH_Z - \widehat{V})^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} (ZH_Z - \widehat{V}) \rangle$. We will prove now that this can be lower bounded by 0, as is the case in Example 1 (46) in Chi et al. (2019). Note that we have

$$\begin{aligned}
 &\langle \tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle \\
 &= \text{Tr}(\widehat{V}^\top \widehat{\Sigma}_{0,t} (ZH_Z - \widehat{V}) \widehat{V}^\top \widehat{\Sigma}_{0,t} (ZH_Z - \widehat{V})) \\
 &= \text{Tr}(\widehat{\Sigma}_{0,t} (ZH_Z - \widehat{V}) \widehat{V}^\top \widehat{\Sigma}_{0,t} (ZH_Z - \widehat{V}) \widehat{V}^\top) \\
 &= \text{Tr}(\widehat{\Sigma}_{0,t} ZH_Z \widehat{V}^\top \widehat{\Sigma}_{0,t} ZH_Z \widehat{V}^\top) - 2 \text{Tr}(\widehat{\Sigma}_{0,t} \widehat{V} \widehat{V}^\top \widehat{\Sigma}_{0,t} ZH_Z \widehat{V}^\top) + \text{Tr}(\widehat{\Sigma}_{0,t} \widehat{V} \widehat{V}^\top \widehat{\Sigma}_{0,t} \widehat{V} \widehat{V}^\top).
 \end{aligned}$$

By Lemma 2 in Ten Berge (1977), we know that $ZH_Z\widehat{V}^\top \succeq 0$ and it is also symmetric. If we denote $ZH_Z\widehat{V}^\top = L_0L_0^\top$ and write $Z_1 = \widehat{\Sigma}_{0,t}^{1/2}L_0$, $Z_2 = \widehat{\Sigma}_{0,t}^{1/2}\widehat{V}$, we have

$$\langle \widetilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} \widetilde{Z} \rangle = \|Z_1 Z_1^\top\|_{\mathbb{F}}^2 + \|Z_2 Z_2^\top\|_{\mathbb{F}}^2 - 2\langle Z_1 Z_1^\top, Z_2 Z_2^\top \rangle \geq 0$$

by Cauchy Schwarz Inequality. This proves that Term II is non-negative. Finally we conclude that we have

$$g_t(\widehat{V}, \widetilde{Z}) \geq \frac{1}{2\nu}(\widehat{\lambda}_r - \widehat{\lambda}_{r+1})\|\widetilde{Z}\|_{\mathbb{F}}^2.$$

Now we argue for general L . First we show that for any L in the neighborhood of \widehat{V}

$$\begin{aligned} |g_t(L, \widetilde{Z}) - g_t(\widehat{V}, \widetilde{Z})| &= |\lambda \langle \widetilde{Z} \widetilde{Z}^\top, \widehat{\Sigma}_{0,t}(LL^\top - \widehat{V}\widehat{V}^\top)\widehat{\Sigma}_{0,t} \rangle + \lambda \langle L^\top \widehat{\Sigma}_{0,t} L - \widehat{V}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widetilde{Z}^\top \widehat{\Sigma}_{0,t} \widetilde{Z} \rangle \\ &\quad + \lambda (\langle \widetilde{Z}^\top \widehat{\Sigma}_{0,t} L, L^\top \widehat{\Sigma}_{0,t} \widetilde{Z} \rangle - \langle \widetilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} \widetilde{Z} \rangle) | \\ &\leq c_1 \|L - \widehat{V}\|_{\mathbb{F}} \|\widetilde{Z}\|_{\mathbb{F}}^2, \end{aligned}$$

for some positive constant c_1 depending on λ and ν . Specifically, this bound can be obtained by bounding each term. The first term can be bounded by

$$\begin{aligned} \lambda |\langle \widetilde{Z} \widetilde{Z}^\top, \widehat{\Sigma}_{0,t}(LL^\top - \widehat{V}\widehat{V}^\top)\widehat{\Sigma}_{0,t} \rangle| &\leq \lambda \|\widetilde{Z}\|_{\mathbb{F}}^2 \|\widehat{\Sigma}_{0,t}(LL^\top - \widehat{V}\widehat{V}^\top)\widehat{\Sigma}_{0,t}\|_{\text{op}} \\ &\leq 2\lambda\nu \|\widehat{\Sigma}_{0,t}^{1/2} LL^\top \widehat{\Sigma}_{0,t}^{1/2} - \widehat{\Sigma}_{0,t}^{1/2} \widehat{V} \widehat{V}^\top \widehat{\Sigma}_{0,t}^{1/2}\|_{\text{op}} \|\widetilde{Z}\|_{\mathbb{F}}^2 \\ &\leq \frac{9}{2} \lambda\nu \|\widehat{\Sigma}_{0,t}^{1/2} L\|_{\text{op}} \|\widehat{\Sigma}_{0,t}^{1/2} L - \widehat{\Sigma}_{0,t}^{1/2} \widehat{V}\|_{\mathbb{F}} \|\widetilde{Z}\|_{\mathbb{F}}^2 \\ &\leq \frac{9}{2} \sqrt{2} \lambda\nu^{3/2} \left(1 + \frac{\lambda_1}{\lambda} + \delta\right) \|L - \widehat{V}\|_{\mathbb{F}} \|\widetilde{Z}\|_{\mathbb{F}}^2, \end{aligned}$$

under the assumption that $\text{dist}(\widehat{\Sigma}_{0,t}^{1/2} L, \widehat{\Sigma}_{0,t}^{1/2} \widehat{V}) \leq \frac{1}{4} \|\widehat{\Sigma}_{0,t}^{1/2} \widehat{V}\|_{\text{op}} = \frac{1}{4} \sqrt{1 + \frac{\widehat{\lambda}_1}{\lambda}}$. The first inequality is the property of matrix norm, the second inequality is due to the bound on sample covariance matrix. The third inequality follows from Lemma 14 and the last inequality follows from the neighborhood assumption. The second term can be bounded in a similar fashion as

$$\begin{aligned} &\lambda |\langle L^\top \widehat{\Sigma}_{0,t} L - \widehat{V}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widetilde{Z}^\top \widehat{\Sigma}_{0,t} \widetilde{Z} \rangle| \\ &\leq \lambda \|(L^\top \widehat{\Sigma}_{0,t} L - \widehat{V}^\top \widehat{\Sigma}_{0,t} \widehat{V}) \widetilde{Z}^\top\|_{\mathbb{F}} \|\widehat{\Sigma}_{0,t} \widetilde{Z}\|_{\mathbb{F}} \\ &\leq 2\lambda\nu \|L^\top \widehat{\Sigma}_{0,t} L - \widehat{V}^\top \widehat{\Sigma}_{0,t} \widehat{V}\|_{\text{op}} \|\widetilde{Z}\|_{\mathbb{F}}^2 \\ &= 2\lambda\nu \|\widehat{V}^\top \widehat{\Sigma}_{0,t} (L - \widehat{V}) + (L - \widehat{V})^\top \widehat{\Sigma}_{0,t} \widehat{V} + (L - \widehat{V})^\top \widehat{\Sigma}_{0,t} (L - \widehat{V})\|_{\text{op}} \|\widetilde{Z}\|_{\mathbb{F}}^2 \\ &\leq 2\lambda\nu (2\|\widehat{\Sigma}_{0,t}^{1/2} \widehat{V}\|_{\text{op}} \|\widehat{\Sigma}_{0,t}^{1/2} (L - \widehat{V})\|_{\text{op}} + \delta \|\widehat{\Sigma}_{0,t}^{1/2} (L - \widehat{V})\|_{\text{op}}) \|\widetilde{Z}\|_{\mathbb{F}}^2 \\ &= 2 \left(\delta + 2\sqrt{1 + \frac{\widehat{\lambda}_1}{\lambda}} \right) \lambda\nu \|\widehat{\Sigma}_{0,t}^{1/2} (L - \widehat{V})\|_{\text{op}} \|\widetilde{Z}\|_{\mathbb{F}}^2 \\ &\leq 2\sqrt{2} \left(\delta + 2 \left(1 + \frac{\lambda_1}{\lambda}\right) \right) \lambda\nu^{3/2} \|L - \widehat{V}\|_{\mathbb{F}} \|\widetilde{Z}\|_{\mathbb{F}}^2. \end{aligned}$$

The first line is standard matrix inequality and the rest follows from the neighborhood assumption as well as the assumption on the operator norm of sample covariance matrices on event B_2 . Finally we bound the last term as

$$\begin{aligned}
 & \lambda |(\langle \tilde{Z}^\top \widehat{\Sigma}_{0,t} L, L^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle - \langle \tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V}, \widehat{V}^\top \widehat{\Sigma}_{0,t} \tilde{Z} \rangle)| \\
 &= \lambda |\text{Tr}((\tilde{Z}^\top \widehat{\Sigma}_{0,t} L)^2 - (\tilde{Z}^\top \widehat{\Sigma}_{0,t} \widehat{V})^2)| \\
 &= \lambda |\text{Tr}(\tilde{Z}^\top \widehat{\Sigma}_{0,t} (L + \widehat{V}) \tilde{Z}^\top \widehat{\Sigma}_{0,t} (L - \widehat{V}))| \\
 &\leq \lambda \|\tilde{Z}^\top \widehat{\Sigma}_{0,t} (L + \widehat{V})\|_{\text{F}} \|\tilde{Z}^\top \widehat{\Sigma}_{0,t} (L - \widehat{V})\|_{\text{F}} \\
 &\leq 2\sqrt{2}\lambda\nu^{3/2} \|\tilde{Z}\|_{\text{F}}^2 \|L - \widehat{V}\|_{\text{F}} \|\widehat{\Sigma}_{0,t}^{1/2} (L + \widehat{V})\|_{\text{op}} \\
 &\leq 2\sqrt{2}\lambda\nu^{3/2} \left(2 \left(1 + \frac{\lambda_1}{\lambda}\right) + \delta\right) \|L - \widehat{V}\|_{\text{F}} \|\tilde{Z}\|_{\text{F}}^2.
 \end{aligned}$$

In this way, we see that we can choose the constant to be

$$c_1 = 4\sqrt{2} \left(\delta + 2 \left(1 + \frac{\lambda_1}{\lambda}\right) \right) \lambda\nu^{3/2} + \frac{9}{2} \sqrt{2}\lambda\nu^{3/2} \left(1 + \frac{\lambda_1}{\lambda} + \delta\right) = \sqrt{2}\lambda\nu^{3/2} \left(\frac{25}{2} \left(1 + \frac{\lambda_1}{\lambda}\right) + \frac{17}{2}\delta \right).$$

Thus for any L within a δ neighborhood of $\widehat{V} = L_t^*$ as defined above, we have, by triangle inequality,

$$\begin{aligned}
 g_t(L, \tilde{Z}) &\geq g_t(\widehat{V}, \tilde{Z}) - |g_t(L, \tilde{Z}) - g_t(\widehat{V}, \tilde{Z})| \\
 &\geq \frac{1}{2\nu} (\widehat{\lambda}_r - \widehat{\lambda}_{r+1}) \|\tilde{Z}\|_{\text{F}}^2 - c_1 \|L - \widehat{V}\|_{\text{F}} \|\tilde{Z}\|_{\text{F}}^2 \\
 &\geq \frac{1}{4\nu} (\widehat{\lambda}_r - \widehat{\lambda}_{r+1}) \|\tilde{Z}\|_{\text{F}}^2 \geq \frac{1}{8\nu} (\lambda_r - \lambda_{r+1}) \|\tilde{Z}\|_{\text{F}}^2,
 \end{aligned}$$

as long as $\|L - \widehat{V}\|_{\text{F}} \leq \frac{\widehat{\lambda}_r - \widehat{\lambda}_{r+1}}{4\nu c_1}$, which is guaranteed by the assumption in the theorem and the event B_2 .

Recall that $g_t(L, \tilde{Z}) = \frac{1}{2} \text{vec}(\tilde{Z})^\top \nabla^2 f_t(L) \text{vec}(\tilde{Z})$. This motivates us to pick

$$\alpha = \frac{(\lambda_r - \lambda_{r+1})}{4\nu}$$

under appropriate radius conditions.

To finally find the radius of attraction region such that $\text{dist}(L_t, L_t^*) = \text{dist}(\bar{V}_t, \widehat{V}(S_t)) \leq R$, we notice that throughout the proof for smoothness and strongly convexity to hold, we require 3 conditions on the distance

$$\|L - L_t^*\|_{\text{F}} \leq \frac{(\widehat{\lambda}_r - \widehat{\lambda}_{r+1})}{4\nu c_1}, \quad \|\widehat{\Sigma}_{0,t}^{1/2} L - \widehat{\Sigma}_{0,t}^{1/2} L_t^*\|_{\text{F}} \leq \delta, \quad \|\widehat{\Sigma}_{0,t}^{1/2} L - \widehat{\Sigma}_{0,t}^{1/2} L_t^*\|_{\text{F}} \leq \frac{1}{4} \sqrt{1 + \frac{\lambda_1}{\lambda}}.$$

Hence to ensure all of the above three conditions to hold, we can just set

$$R = \min \left\{ \frac{(\lambda_r - \lambda_{r+1})}{8\nu c_1}, \frac{1}{8} \frac{\sqrt{1 + \frac{\lambda_1}{\lambda}}}{\sqrt{\nu}}, \frac{\delta}{\sqrt{2\nu}} \right\}. \quad (46)$$

Here recall that $\frac{1}{\nu}$ is the lower bound of minimum eigenvalue of Σ_0 .

In summary, we conclude that there is a constant α such that for any L in the neighborhood of L_t^* and for any Z ,

$$\text{vec}(ZH_Z - L_t^*)^\top \nabla^2 f_t(L) \text{vec}(ZH_Z - L_t^*) \geq \alpha \|ZH_Z - L_t^*\|_F^2.$$

Combine with the upper bound and resort to Lemma 20, we conclude that if we choose the step-size to be $\eta \leq \frac{1}{\beta}$, under the radius conditions,

$$\text{dist}^2(L_{t+1}, \widehat{V}) \leq (1 - \alpha\eta) \text{dist}^2(L_t, \widehat{V}).$$

Moreover, we have

$$\alpha = \frac{(\lambda_r - \lambda_{r+1})}{4\nu}, \quad \beta = 12\lambda\nu \left(\delta + 1 + \frac{\lambda_1}{\lambda} \right)^2.$$

This finishes our analysis for the gradient descent step. We comment here that Proposition 18 is then proved by simply taking $\delta = 1$ and calculate the radius in (46) and α, β accordingly. We keep δ in the proof for the sake of generality.

Specifically, under the choice that $\lambda = \frac{\lambda_1}{c}$, we have

$$\eta \leq \frac{c}{12\lambda_1\nu(c+2)^2} = \frac{1}{\beta},$$

and we have, by lifting the matrix of size $|S_t| \times r$ to $p \times r$ by filling in 0s,

$$\text{dist}^2(V_{t+1}^o, \widehat{V}(S_t)) \leq \left(1 - \frac{\eta(\lambda_r - \lambda_{r+1})}{4\nu} \right) \text{dist}^2(\bar{V}_t, \widehat{V}(S_t)).$$

■

B.4 Proof of Proposition 19

Proof Recall that we have defined

$$\text{supp}(V) = S, \quad \text{supp}(\bar{V}_{t+1}) = F_{t+1}.$$

Now we define sets $F_1 = S \setminus F_{t+1}$, $F_2 = S \cap F_{t+1}$ and $F_3 = F_{t+1} \setminus S$, the sizes of which are denoted by k_1, k_2 and k_3 respectively. We also define

$$x_1 = \|V_{F_1^*}\|_F, \quad x_2 = \|V_{F_2^*}\|_F,$$

$$y_1 = \|V_{t+1, F_1^*}\|_F, \quad y_2 = \|V_{t+1, F_2^*}\|_F, \quad y_3 = \|V_{t+1, F_3^*}\|_F.$$

Finally, we define

$$x_1^2 + x_2^2 = \|V\|_F^2 = X^2, \quad y_1^2 + y_2^2 + y_3^2 \leq \|V_{t+1}\|_F^2 = Y^2.$$

Let $\Delta = \text{Tr}(|V^\top V_{t+1}|)$ be the quantity we are interested in. Here we let $\text{Tr}(|M|)$ denote the trace norm $\text{Tr}(|M|) = \text{Tr}(\sqrt{M^\top M})$ which is simply the sum of singular values (a.k.a. nuclear

norm). This is a special case of Schatten p norm, defined as $\|T\|_p = \text{Tr}(|T|^p)^{\frac{1}{p}}$ which is essentially vector p norm of a vector composed of singular values.

We first prove the following claim about the effect of truncation step:

$$\text{Tr}(|V^\top V_{t+1}|) - \text{Tr}(|V^\top \bar{V}_{t+1}|) \leq \sqrt{\frac{s}{s'}} \min \left(\sqrt{X^2 Y^2 - \Delta^2}, \frac{1 + \sqrt{s/s'}}{XY} (X^2 Y^2 - \Delta^2) \right). \quad (47)$$

Proof Since in the hard thresholding step we greedily pick the rows with largest l_2 norms, we have

$$\frac{y_1^2}{k_1} \leq \frac{y_3^2}{k_3}.$$

In addition, since $k_1 + k_2 = s \leq s' = k_2 + k_3$, we also have $k_1 \leq k_3$. By applying Holder's inequality on Schatten norm and the definition of support set, we have

$$\begin{aligned} \Delta^2 &\leq (\|V_{F_1^*}\|_F \|V_{t+1, F_1^*}\|_F + \|V_{F_2^*}\|_F \|V_{t+1, F_2^*}\|_F)^2 = (x_1 y_1 + x_2 y_2)^2 \\ &\leq (y_1^2 + y_2^2) X^2 \leq (Y^2 - y_3^2) X^2 \leq X^2 Y^2 - \frac{k_3}{k_1} X^2 y_1^2, \end{aligned}$$

and this gives

$$y_1^2 \leq \frac{k_1}{k_3 X^2} (X^2 Y^2 - \Delta^2) \leq \frac{k_1 + k_2}{k_3 + k_2} (Y^2 - \Delta^2 / X^2) = \frac{s}{s'} (Y^2 - \Delta^2 / X^2),$$

where the second inequality holds since $k_1 \leq k_3$. Now we split the arguments into two cases.

Case I: $\Delta < XY \sqrt{\frac{s}{s+s'}}$. In this case, we obtain that

$$\Delta < \sqrt{\frac{s}{s'}} \sqrt{X^2 Y^2 - \Delta^2}.$$

This implies (47) immediately.

Case II: We can now assume that $\Delta \geq XY \sqrt{\frac{s}{s+s'}}$. Then we have

$$y_1^2 \leq \frac{s}{s'} (Y^2 - \Delta^2 / X^2) \leq \frac{\Delta^2}{X^2}.$$

By definition we have

$$x_1 y_1 + \sqrt{Y^2 - y_1^2} \sqrt{X^2 - x_1^2} \geq x_1 y_1 + x_2 y_2 \geq \Delta.$$

Solving this inequality as a quadratic inequality in x_1 yields

$$x_1 \leq \frac{\Delta y_1 + \sqrt{(X^2 Y^2 - \Delta^2)(Y^2 - y_1^2)}}{Y^2}.$$

By definition of $X^2 = x_1^2 + x_2^2$ we have $x_1 \leq X$. Also we have $\Delta \leq XY$ by Holder inequality on Schatten norm.

Combining the above inequalities together, we have

$$x_1 \leq \min \left(X, \frac{Xy_1 + \sqrt{X^2Y^2 - \Delta^2}}{Y} \right).$$

Note that $y_1 \leq \sqrt{\frac{s}{s'}} \frac{\sqrt{X^2Y^2 - \Delta^2}}{X}$ by previous calculation. Substituting it into the above equation, we have

$$x_1 \leq \min \left(X, \frac{\sqrt{\frac{s}{s'}} \sqrt{X^2Y^2 - \Delta^2} + \sqrt{X^2Y^2 - \Delta^2}}{Y} \right).$$

Finally, we can compute that

$$\begin{aligned} x_1 y_1 &\leq \sqrt{\frac{s}{s'}} \frac{1}{X} \sqrt{X^2Y^2 - \Delta^2} \min \left(X, \frac{1}{Y} \sqrt{X^2Y^2 - \Delta^2} (1 + \sqrt{\frac{s}{s'}}) \right) \\ &= \sqrt{\frac{s}{s'}} \min \left(\sqrt{X^2Y^2 - \Delta^2}, \frac{1 + \sqrt{\frac{s}{s'}}}{XY} (X^2Y^2 - \Delta^2) \right). \end{aligned}$$

This is the error induced by the truncation step, finally we have

$$\begin{aligned} \text{Tr}(|V^\top V_{t+1}|) - \text{Tr}(|V^\top \bar{V}_{t+1}|) &\leq \text{Tr}(|V^\top (\bar{V}_{t+1} - V_{t+1})|) \\ &\leq x_1 y_1 \leq \sqrt{\frac{s}{s'}} \min \left(\sqrt{X^2Y^2 - \Delta^2}, \frac{1 + \sqrt{\frac{s}{s'}}}{XY} (X^2Y^2 - \Delta^2) \right). \end{aligned}$$

This finishes the proof for the claim. \blacksquare

We now switch to work with the distance metric $\text{dist}(U, V) = \min_{P \in \mathcal{O}(r)} \|UP - V\|_F$. We first expand the expression

$$\begin{aligned} \|UP - V\|_F^2 &= \|UP\|_F^2 + \|V\|_F^2 - 2\langle UP, V \rangle \\ &= \|U\|_F^2 + \|V\|_F^2 - 2\text{Tr}(P^\top U^\top V). \end{aligned}$$

So, to minimize the distance metric we defined is to maximize $\text{Tr}(P^\top U^\top V)$. To this end, let ADB^\top denote the singular value decomposition of $U^\top V$, then we have

$$\text{Tr}(P^\top U^\top V) = \text{Tr}(P^\top ADB^\top) = \text{Tr}(B^\top P^\top AD) = \text{Tr}(ZD)$$

for Z being an orthogonal matrix. Since D is a diagonal matrix, we have $\text{Tr}(ZD) = \sum_i Z_{ii} D_{ii}$ and to maximize this value we would like to have $Z_{ii} = 1$ for all i since D_{ii} are non-negative. As a result, we have the optimal P being AB^\top and the optimal value is given by

$$\min \|UP - V\|_F^2 = \|U\|_F^2 + \|V\|_F^2 - 2\text{Tr}(D) = \|U\|_F^2 + \|V\|_F^2 - 2\text{Tr}(|U^\top V|).$$

Now we consider the subspace distance we have defined. Recall that $\Delta = \text{Tr}(|V^\top V_{t+1}|)$. Then the subspace distance can be written as

$$\begin{aligned}
 & \text{dist}^2(\bar{V}_{t+1}, V) \\
 &= \|V\|_{\mathbb{F}}^2 + \|\bar{V}_{t+1}\|_{\mathbb{F}}^2 - 2 \text{Tr} |V^\top \bar{V}_{t+1}| \\
 &\leq \|V\|_{\mathbb{F}}^2 + \|\bar{V}_{t+1}\|_{\mathbb{F}}^2 - 2 \text{Tr} |V^\top V_{t+1}| + 2\sqrt{\frac{s}{s'}} \min \left(\sqrt{X^2 Y^2 - \Delta^2}, \frac{1 + \sqrt{\frac{s}{s'}}}{XY} (X^2 Y^2 - \Delta^2) \right) \\
 &\leq \|V\|_{\mathbb{F}}^2 + \|V_{t+1}\|_{\mathbb{F}}^2 - 2 \text{Tr} |V^\top V_{t+1}| + 2\sqrt{\frac{s}{s'}} \frac{1 + \sqrt{\frac{s}{s'}}}{XY} (X^2 Y^2 - \Delta^2) \\
 &= \text{dist}^2(V_{t+1}, V) + 2\sqrt{\frac{s}{s'}} \frac{1 + \sqrt{\frac{s}{s'}}}{XY} (X^2 Y^2 - \Delta^2),
 \end{aligned}$$

where the first equality follows from the previous expansion of distance, the first inequality follows from Claim (47) and we use the fact that truncation reduces the Frobenius norm in the second inequality. In the last equality, we use again the relationship between trace norm and distance. To deal with the extra error term, we use the following bound

$$\begin{aligned}
 X^2 Y^2 - \Delta^2 &= (XY + \Delta)(XY - \Delta) \leq 2XY(XY - \Delta) \leq XY(X^2 + Y^2 - 2\Delta) \\
 &= XY(\|V\|_{\mathbb{F}}^2 + \|V_{t+1}\|_{\mathbb{F}}^2 - 2 \text{Tr}(|V^\top V_{t+1}|)) = XY \text{dist}^2(V_{t+1}, V).
 \end{aligned}$$

Combining the inequalities, we have

$$\begin{aligned}
 \text{dist}^2(\bar{V}_{t+1}, V) &\leq \text{dist}^2(V_{t+1}, V) + 2\sqrt{\frac{s}{s'}} \frac{1 + \sqrt{\frac{s}{s'}}}{XY} (X^2 Y^2 - \Delta^2) \\
 &\leq \text{dist}^2(V_{t+1}, V) + 2\sqrt{\frac{s}{s'}} \frac{1 + \sqrt{\frac{s}{s'}}}{XY} XY \text{dist}^2(V_{t+1}, V) \\
 &= \left(1 + 2\sqrt{\frac{s}{s'}} \left(1 + \sqrt{\frac{s}{s'}} \right) \right) \text{dist}^2(V_{t+1}, V).
 \end{aligned}$$

This finishes our analysis for the hard thresholding step. ■

B.5 Proof of Theorem 9

Proof We present the proof of initialization using generalized Fantope in this section. Some proof arguments originate from Gao et al. (2017).

Recall that we have defined

$$\tilde{A} = A(A^\top \hat{\Sigma}_0 A)^{-1/2}, \quad \tilde{\Lambda}_r = (A^\top \hat{\Sigma}_0 A)^{1/2} \Lambda_r (A^\top \hat{\Sigma}_0 A)^{1/2}, \quad \text{and} \quad \tilde{F} = \tilde{A} \tilde{A}^\top. \quad (48)$$

Throughout the proof, we work on the event $B_3 \cap B_4$, which by Lemmas 16 and 17 happens with probability at least $1 - \exp(-C'(s + \log(ep/s)))$ for some constant $C' > 0$, uniformly over \mathcal{P}_n .

We shall need the following lemma, which is a simplified version of Lemma 6.3 in Gao et al. (2017). The proof is essentially the same and so we omit it.

Lemma 21 (*Curvature of Fantope*) Let $P \in \mathcal{O}(p, r)$ and $D = \text{diag}(d_1, d_2, \dots, d_r)$ with $d_1 \geq d_2 \geq \dots \geq d_r \geq 0$. If $F \in \mathcal{F}_r$, then

$$\langle PDP^\top, PP^\top - F \rangle \geq \frac{d_r}{2} \|PP^\top - F\|_F^2.$$

Recall that

$$\Sigma = \Sigma_0 K \Lambda K^\top \Sigma_0 = \Sigma_0 A \Lambda_r A^\top \Sigma_0 + \Sigma_0 B \bar{\Lambda} B^\top \Sigma_0.$$

For notational simplicity, for any positive semi-definite matrix B , we define

$$\phi_{\max}^B(k) = \max_{\|u\|_0 \leq k, u \neq 0} \frac{u^\top B u}{u^\top u}, \quad \phi_{\min}^B(k) = \max_{\|u\|_0 \leq k, u \neq 0} \frac{u^\top B u}{u^\top u}.$$

In the rest of this proof, let $\Delta = \hat{F} - \tilde{F}$ with \tilde{F} defined in (48). As in Gao et al. (2017), the main proof consists of two steps. The first step is to derive upper bound of $\|\hat{\Sigma}_0^{1/2} \Delta \hat{\Sigma}_0^{1/2}\|_F$ and the second step is to lower bound $\|\hat{\Sigma}_0^{1/2} \Delta \hat{\Sigma}_0^{1/2}\|_F$ by $\|\Delta\|_F$.

Step (1) First note that by Lemma 17, \tilde{A} is well defined on event B_4 , so $\tilde{F} = \tilde{A} \tilde{A}^\top$ is also well defined on event B_4 . In addition, on event B_4 , we have

$$\|\Sigma_0^{1/2} (\tilde{F} - A A^\top) \Sigma_0^{1/2}\|_F \leq C \sqrt{\frac{r(s + \log p)}{n}}. \quad (49)$$

We first show that \tilde{F} is a feasible solution, that is, $\hat{\Sigma}_0^{1/2} \tilde{A} \tilde{A}^\top \hat{\Sigma}_0^{1/2} \in \mathcal{F}_r$. By (48),

$$\tilde{F} = \tilde{A} \tilde{A}^\top = A (A^\top \hat{\Sigma}_0 A)^{-1} A^\top.$$

Let $M = \hat{\Sigma}_0^{1/2} \tilde{F} \hat{\Sigma}_0^{1/2} = \hat{\Sigma}_0^{1/2} A (A^\top \hat{\Sigma}_0 A)^{-1} A^\top \hat{\Sigma}_0^{1/2}$. Then

$$\text{Tr}(M) = \text{Tr}(\hat{\Sigma}_0^{1/2} A (A^\top \hat{\Sigma}_0 A)^{-1} A^\top \hat{\Sigma}_0^{1/2}) = r.$$

Next we check that $0 \preceq M \preceq I$. It is obvious that $0 \preceq M$. Moreover,

$$\|M\|_{\text{op}} \leq \|\hat{\Sigma}_0^{1/2} \tilde{A}\|_{\text{op}} \|\hat{\Sigma}_0^{1/2} \tilde{A}\|_{\text{op}} \leq 1$$

by the construction of \tilde{A} . Hence, $\tilde{F} \in \mathcal{F}_r$.

Recall that \hat{F} is the solution to (14). The basic inequality implies

$$-\langle \hat{\Sigma}, \hat{F} \rangle + \rho \|\hat{F}\|_1 \leq -\langle \hat{\Sigma}, \tilde{F} \rangle + \rho \|\tilde{F}\|_1.$$

Rearranging terms, we have

$$0 \leq \rho (\|\tilde{F}\|_1 - \|\tilde{F} + \Delta\|_1) + \langle \hat{\Sigma}, \Delta \rangle. \quad (50)$$

We deal with each term on the right side separately. For the first term on the right, we have

$$\begin{aligned} \|\tilde{F}\|_1 - \|\tilde{F} + \Delta\|_1 &= \|\tilde{F}_{SS}\|_1 - \|\tilde{F}_{SS} + \Delta_{SS}\|_1 - \|\Delta_{(SS)^c}\|_1 \\ &\leq \|\Delta_{SS}\|_1 - \|\Delta_{(SS)^c}\|_1. \end{aligned} \quad (51)$$

Here the first equality holds since \widehat{F} is supported on $S \times S$. For the second term, we have

$$\langle \widehat{\Sigma}, \Delta \rangle = \langle \widehat{\Sigma} - \Sigma, \Delta \rangle + \underbrace{\langle \Sigma_0 A \Lambda_r A^\top \Sigma_0, \Delta \rangle}_{\text{Term I}} + \underbrace{\langle \Sigma_0 B \bar{\Lambda} B^\top \Sigma_0, \Delta \rangle}_{\text{Term II}}. \quad (52)$$

We now bound Term I and Term II separately.

Bound for Term I: For Term I, we decompose it in the following way:

$$\langle \Sigma_0 A \Lambda_r A^\top \Sigma_0, \Delta \rangle = \langle \Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle + \langle \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle.$$

Now we bound the second term on the right as follows

$$\begin{aligned} \langle \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle &= \langle \widehat{\Sigma}_0^{1/2} A \Lambda_r A^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2} (\widehat{F} - \widetilde{F}) \widehat{\Sigma}_0^{1/2} \rangle \\ &= \langle \widehat{\Sigma}_0^{1/2} \widetilde{A} \widetilde{\Lambda}_r \widetilde{A}^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2} (\widehat{F} - \widetilde{F}) \widehat{\Sigma}_0^{1/2} \rangle \\ &= \langle \widehat{\Sigma}_0^{1/2} \widetilde{A} \Lambda_r \widetilde{A}^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2} (\widehat{F} - \widetilde{F}) \widehat{\Sigma}_0^{1/2} \rangle + \langle \widehat{\Sigma}_0^{1/2} \widetilde{A} (\widetilde{\Lambda}_r - \Lambda_r) \widetilde{A}^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2} \rangle \\ &\leq \langle \widehat{\Sigma}_0^{1/2} \widetilde{A} \Lambda_r \widetilde{A}^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2} (\widehat{F} - \widetilde{F}) \widehat{\Sigma}_0^{1/2} \rangle + \|\widetilde{\Lambda}_r - \Lambda_r\|_{\text{F}} \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\text{F}}. \end{aligned}$$

As a result, Term I can be bounded in the following way

$$\text{Term I} \leq \langle \Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle + \langle \widehat{\Sigma}_0^{1/2} \widetilde{A} \Lambda_r \widetilde{A}^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2} (\widehat{F} - \widetilde{F}) \widehat{\Sigma}_0^{1/2} \rangle + \delta_1 \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\text{F}},$$

for $\delta_1 = \|\widetilde{\Lambda}_r - \Lambda_r\|_{\text{F}}$.

Bound for Term II: To bound Term II, we first notice that by definition, $B^\top \Sigma_0 A$ is zero matrix since A, B are normalized with respect to Σ_0 . As a result, we have

$$\begin{aligned} \langle \Sigma_0 B \bar{\Lambda} B^\top \Sigma_0, \Delta \rangle &= \langle \Sigma_0 B \bar{\Lambda} B^\top \Sigma_0, \widehat{F} \rangle \\ &\leq \lambda_{r+1} \langle \Sigma_0 B B^\top \Sigma_0, \widehat{F} \rangle \\ &= \lambda_{r+1} \langle \Sigma_0 B B^\top \Sigma_0, \widehat{F} - \widetilde{A} \widetilde{A}^\top \rangle \\ &= \lambda_{r+1} \langle \Sigma_0^{1/2} (I - \Sigma_0^{1/2} A A^\top \Sigma_0^{1/2}) \Sigma_0^{1/2}, \widehat{F} - \widetilde{A} \widetilde{A}^\top \rangle \\ &= \lambda_{r+1} \langle \Sigma_0 - \Sigma_0 A A^\top \Sigma_0, \widehat{F} - \widetilde{A} \widetilde{A}^\top \rangle. \end{aligned}$$

Here, the second last equality holds since $\Sigma_0^{1/2} A A^\top \Sigma_0^{1/2} + \Sigma_0^{1/2} B B^\top \Sigma_0^{1/2} = I$. We further decompose the rightmost side as

$$\begin{aligned} &\lambda_{r+1} \langle \Sigma_0 - \Sigma_0 A A^\top \Sigma_0, \widehat{F} - \widetilde{A} \widetilde{A}^\top \rangle \\ &= \lambda_{r+1} \langle \Sigma_0 - \widehat{\Sigma}_0, \Delta \rangle - \lambda_{r+1} \langle \Sigma_0 A A^\top \Sigma_0 - \widehat{\Sigma}_0 A A^\top \widehat{\Sigma}_0, \Delta \rangle \\ &\quad + \underbrace{\lambda_{r+1} \langle \widehat{\Sigma}_0, \widehat{F} - \widetilde{A} \widetilde{A}^\top \rangle}_{\text{Term A}} - \underbrace{\lambda_{r+1} \langle \widehat{\Sigma}_0^{1/2} A A^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2} (\widehat{F} - \widetilde{A} \widetilde{A}^\top) \widehat{\Sigma}_0^{1/2} \rangle}_{\text{Term B}}. \end{aligned}$$

We now deal with Term A and Term B separately. By definition we have $\widehat{\Sigma}_0^{1/2} \widehat{F} \widehat{\Sigma}_0^{1/2} \in \mathcal{F}_r$, hence

$$\langle \widehat{\Sigma}_0, \widehat{F} - \widetilde{A} \widetilde{A}^\top \rangle = \text{Tr}(\widehat{\Sigma}_0^{1/2} \widehat{F} \widehat{\Sigma}_0^{1/2}) - \text{Tr}(\widetilde{A}^\top \widehat{\Sigma}_0 \widetilde{A}) = r - r = 0.$$

For Term B, we have

$$\begin{aligned}
& -\lambda_{r+1}\langle \widehat{\Sigma}_0^{1/2}AA^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{A}\widetilde{A}^\top)\widehat{\Sigma}_0^{1/2} \rangle \\
& = -\lambda_{r+1}\langle \widehat{\Sigma}_0^{1/2}\widetilde{A}\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{A}\widetilde{A}^\top)\widehat{\Sigma}_0^{1/2} \rangle \\
& \quad - \lambda_{r+1}\langle \widehat{\Sigma}_0^{1/2}\widetilde{A}(A^\top\widehat{\Sigma}_0A - I)\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{A}\widetilde{A}^\top)\widehat{\Sigma}_0^{1/2} \rangle \\
& \leq -\lambda_{r+1}\langle \widehat{\Sigma}_0^{1/2}\widetilde{A}\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{A}\widetilde{A}^\top)\widehat{\Sigma}_0^{1/2} \rangle + \lambda_{r+1}\|A^\top\widehat{\Sigma}_0A - I\|_{\text{F}}\|\widehat{\Sigma}_0^{1/2}\Delta\widehat{\Sigma}_0^{1/2}\|_{\text{F}}.
\end{aligned}$$

Therefore, Term II has the following bound

$$\begin{aligned}
\langle \Sigma_0 B \bar{\Lambda} B^\top \Sigma_0, \Delta \rangle & \leq \lambda_{r+1}\langle \Sigma_0 - \widehat{\Sigma}_0, \Delta \rangle - \lambda_{r+1}\langle \Sigma_0 AA^\top \Sigma_0 - \widehat{\Sigma}_0 AA^\top \widehat{\Sigma}_0, \Delta \rangle \\
& \quad - \lambda_{r+1}\langle \widehat{\Sigma}_0^{1/2}\widetilde{A}\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{A}\widetilde{A}^\top)\widehat{\Sigma}_0^{1/2} \rangle + \delta_2\|\widehat{\Sigma}_0^{1/2}\Delta\widehat{\Sigma}_0^{1/2}\|_{\text{F}},
\end{aligned}$$

for $\delta_2 = \lambda_{r+1}\|A^\top\widehat{\Sigma}_0A - I\|_{\text{F}}$. Now we combine the results for Term I and Term II to obtain

$$\begin{aligned}
& \langle \Sigma_0 A \Lambda_r A^\top \Sigma_0, \Delta \rangle + \langle \Sigma_0 B \bar{\Lambda} B^\top \Sigma_0, \Delta \rangle \\
& \leq \langle \Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle \\
& \quad + \langle \widehat{\Sigma}_0^{1/2}\widetilde{A}\Lambda_r\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{F})\widehat{\Sigma}_0^{1/2} \rangle + \delta_1\|\widehat{\Sigma}_0^{1/2}\Delta\widehat{\Sigma}_0^{1/2}\|_{\text{F}} \\
& \quad + \lambda_{r+1}\langle \Sigma_0 - \widehat{\Sigma}_0, \Delta \rangle - \lambda_{r+1}\langle \Sigma_0 AA^\top \Sigma_0 - \widehat{\Sigma}_0 AA^\top \widehat{\Sigma}_0, \Delta \rangle \\
& \quad - \lambda_{r+1}\langle \widehat{\Sigma}_0^{1/2}\widetilde{A}\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{A}\widetilde{A}^\top)\widehat{\Sigma}_0^{1/2} \rangle + \delta_2\|\widehat{\Sigma}_0^{1/2}\Delta\widehat{\Sigma}_0^{1/2}\|_{\text{F}} \\
& = \langle \Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle + \lambda_{r+1}\langle \Sigma_0 - \widehat{\Sigma}_0, \Delta \rangle \\
& \quad - \lambda_{r+1}\langle \Sigma_0 AA^\top \Sigma_0 - \widehat{\Sigma}_0 AA^\top \widehat{\Sigma}_0, \Delta \rangle + \delta\|\widehat{\Sigma}_0^{1/2}\Delta\widehat{\Sigma}_0^{1/2}\|_{\text{F}} \\
& \quad + \langle \widehat{\Sigma}_0^{1/2}\widetilde{A}(\Lambda_r - \lambda_{r+1}I)\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{F})\widehat{\Sigma}_0^{1/2} \rangle,
\end{aligned}$$

where $\delta = \delta_1 + \delta_2$. Plugging back to (50), together (51) with (52), we have

$$\begin{aligned}
0 & \leq \rho(\|\widetilde{F}\|_1 - \|\widehat{F} + \Delta\|_1) + \langle \widehat{\Sigma}, \Delta \rangle \\
& \leq \rho(\|\Delta_{SS}\|_1 - \|\Delta_{(SS)^c}\|_1) + \langle \widehat{\Sigma} - \Sigma, \Delta \rangle \\
& \quad + \langle \Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle + \lambda_{r+1}\langle \Sigma_0 - \widehat{\Sigma}_0, \Delta \rangle \\
& \quad - \lambda_{r+1}\langle \Sigma_0 AA^\top \Sigma_0 - \widehat{\Sigma}_0 AA^\top \widehat{\Sigma}_0, \Delta \rangle \\
& \quad + \delta\|\widehat{\Sigma}_0^{1/2}\Delta\widehat{\Sigma}_0^{1/2}\|_{\text{F}} + \langle \widehat{\Sigma}_0^{1/2}\widetilde{A}(\Lambda_r - \lambda_{r+1}I)\widetilde{A}^\top\widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{F})\widehat{\Sigma}_0^{1/2} \rangle.
\end{aligned} \tag{53}$$

By Holder's inequality, we have

$$\begin{aligned}
& \langle \widehat{\Sigma} - \Sigma, \Delta \rangle + \langle \Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0, \Delta \rangle \\
& \quad + \lambda_{r+1}\langle \Sigma_0 - \widehat{\Sigma}_0, \Delta \rangle - \lambda_{r+1}\langle \Sigma_0 AA^\top \Sigma_0 - \widehat{\Sigma}_0 AA^\top \widehat{\Sigma}_0, \Delta \rangle \\
& \leq \|\widehat{\Sigma} - \Sigma\|_{\infty}\|\Delta\|_1 + \|\Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0\|_{\infty}\|\Delta\|_1 \\
& \quad + \lambda_{r+1}\|\Sigma_0 - \widehat{\Sigma}_0\|_{\infty}\|\Delta\|_1 + \lambda_{r+1}\|\Sigma_0 AA^\top \Sigma_0 - \widehat{\Sigma}_0 AA^\top \widehat{\Sigma}_0\|_{\infty}\|\Delta\|_1.
\end{aligned} \tag{54}$$

On event B_3 , we can pick $\rho = \gamma\sqrt{\frac{\log p}{n}}$ for some large constant $\gamma > 0$ such that

$$\begin{aligned}
& \|\widehat{\Sigma} - \Sigma\|_{\infty} + \|\Sigma_0 A \Lambda_r A^\top \Sigma_0 - \widehat{\Sigma}_0 A \Lambda_r A^\top \widehat{\Sigma}_0\|_{\infty} \\
& \quad + \lambda_{r+1}\|\Sigma_0 - \widehat{\Sigma}_0\|_{\infty} + \lambda_{r+1}\|\Sigma_0 AA^\top \Sigma_0 - \widehat{\Sigma}_0 AA^\top \widehat{\Sigma}_0\|_{\infty} \leq \frac{\rho}{2},
\end{aligned} \tag{55}$$

thus the lefthand side of (54) is bounded by $\frac{\rho}{2}\|\Delta\|_1$. Furthermore, by Lemma 21, we have

$$\begin{aligned} & \langle \widehat{\Sigma}_0^{1/2} \widetilde{A}(\Lambda_r - \lambda_{r+1}I) \widetilde{A}^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widehat{F} - \widetilde{F}) \widehat{\Sigma}_0^{1/2} \rangle \\ &= -\langle \widehat{\Sigma}_0^{1/2} \widetilde{A}(\Lambda_r - \lambda_{r+1}I) \widetilde{A}^\top \widehat{\Sigma}_0^{1/2}, \widehat{\Sigma}_0^{1/2}(\widetilde{A}\widetilde{A}^\top - \widehat{F}) \widehat{\Sigma}_0^{1/2} \rangle \\ &\leq -\frac{\lambda_r - \lambda_{r+1}}{2} \|\widehat{\Sigma}_0^{1/2}(\widetilde{A}\widetilde{A}^\top - \widehat{F}) \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}^2. \end{aligned}$$

Together with (53) and (54), the last display implies that when (55) holds,

$$0 \leq \rho(\|\Delta_{SS}\|_1 - \|\Delta_{(SS)^c}\|_1) + \frac{\rho}{2}\|\Delta\|_1 - \frac{\lambda_r - \lambda_{r+1}}{2} \|\widehat{\Sigma}_0^{1/2}(\widetilde{A}\widetilde{A}^\top - \widehat{F}) \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}^2 + \delta \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}.$$

Rearranging terms and multiplying both side by 2, we obtain

$$\begin{aligned} (\lambda_r - \lambda_{r+1}) \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}^2 &\leq 3\rho \|\Delta_{SS}\|_1 - \rho \|\Delta_{(SS)^c}\|_1 + 2\delta \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}} \\ &\leq 3\rho \|\Delta_{SS}\|_1 + 2\delta \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}. \end{aligned}$$

This can be view as a quadratic equation, which, by Lemma 2 in Cai et al. (2013), yields

$$\|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}^2 \leq \frac{4\delta^2}{(\lambda_r - \lambda_{r+1})^2} + \frac{6\rho}{\lambda_r - \lambda_{r+1}} \|\Delta_{SS}\|_1. \quad (56)$$

Combining the last two displays, we have

$$\begin{aligned} 0 &\leq 3\rho \|\Delta_{SS}\|_1 - \rho \|\Delta_{(SS)^c}\|_1 + \frac{\delta^2}{\lambda_r - \lambda_{r+1}} + (\lambda_r - \lambda_{r+1}) \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}^2 \\ &\leq 9\rho \|\Delta_{SS}\|_1 - \rho \|\Delta_{(SS)^c}\|_1 + \frac{5\delta^2}{\lambda_r - \lambda_{r+1}}. \end{aligned}$$

This can be viewed as a version of generalized cone condition. Finally, using Cauchy-Schwarz inequality, $\|\Delta_{SS}\|_1 \leq s \|\Delta_{SS}\|_{\mathbb{F}}$, and so (56) leads to

$$\|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}^2 \leq \frac{4\delta^2}{(\lambda_r - \lambda_{r+1})^2} + \frac{6\rho s}{\lambda_r - \lambda_{r+1}} \|\Delta_{SS}\|_{\mathbb{F}}. \quad (57)$$

and this is the end of the first step.

Step (2) Recall that we have established the generalized cone condition

$$\|\Delta_{(SS)^c}\|_1 \leq 9\|\Delta_{SS}\|_1 + \frac{5\delta^2}{(\lambda_r - \lambda_{r+1})\rho}. \quad (58)$$

In this step we lower bound $\|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\mathbb{F}}^2$ by a function of $\|\Delta\|_{\mathbb{F}}$ on this cone.

Adapting the ‘‘peeling’’ argument in Bickel et al. (2009), we define the index set $J_1 = \{(i_l, j_l)\}_{l=1}^t$ in $(S \times S)^c$ correspond to the entries with the t largest absolute values in Δ , and also define $\widetilde{J} = (S \times S) \cup J_1$. Next we partition \widetilde{J}^c into disjoint subsets J_2, \dots, J_M of

size t and possibly $|J_M| < t$ such that each J_m is the set of indices corresponding the entries of the t largest absolute values in Δ outside $\tilde{J} \cup \bigcup_{j=2}^{m-1} J_j$. Then by triangle inequality

$$\begin{aligned} \|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\text{F}} &\geq \|\widehat{\Sigma}_0^{1/2} \Delta_{\tilde{J}} \widehat{\Sigma}_0^{1/2}\|_{\text{F}} - \sum_{m=2}^M \|\widehat{\Sigma}_0^{1/2} \Delta_{J_m} \widehat{\Sigma}_0^{1/2}\|_{\text{F}} \\ &\geq \phi_{\min}^{\widehat{\Sigma}_0}(s+t) \|\Delta_{\tilde{J}}\|_{\text{F}} - \phi_{\max}^{\widehat{\Sigma}_0}(t) \sum_{m=2}^M \|\Delta_{J_m}\|_{\text{F}}. \end{aligned}$$

In addition, by our construction of the index sets,

$$\begin{aligned} \sum_{m=2}^M \|\Delta_{J_m}\|_{\text{F}} &\leq \sqrt{t} \sum_{m=2}^M \|\Delta_{J_m}\|_{\infty} \leq \frac{\sqrt{t}}{t} \sum_{m=2}^M \|\Delta_{J_{m-1}}\|_1 \leq t^{-1/2} \|\Delta_{(SS)^c}\|_1 \\ &\leq t^{-1/2} (9 \|\Delta_{SS}\|_1 + \frac{5\delta^2}{\lambda_r - \lambda_{r+1}}) \\ &\leq \frac{9s}{\sqrt{t}} \|\Delta_{\tilde{J}}\|_{\text{F}} + \frac{5\delta^2}{(\lambda_r - \lambda_{r+1})\rho\sqrt{t}}. \end{aligned}$$

The second last inequality follows from the generalized cone condition (58). Hence combining the results above, we have

$$\|\widehat{\Sigma}_0^{1/2} \Delta \widehat{\Sigma}_0^{1/2}\|_{\text{F}} \geq \kappa_1 \|\Delta_{\tilde{J}}\|_{\text{F}} - \kappa_2 \frac{\delta^2}{(\lambda_r - \lambda_{r+1})\rho\sqrt{t}}, \quad (59)$$

where

$$\begin{aligned} \kappa_1 &= \phi_{\min}^{\widehat{\Sigma}_0}(s+t) - \frac{9s}{\sqrt{t}} \phi_{\max}^{\widehat{\Sigma}_0}(t), \\ \kappa_2 &= 5\phi_{\max}^{\widehat{\Sigma}_0}(t). \end{aligned}$$

Taking $t = c_1 s^2$ for c_1 sufficiently large, using the same argument as in Gao et al. (2017), under condition (26) we can lower bound κ_1 by some constant C_1 and upper bound κ_2 by some constant C_2 . Combining (57) and (59), we obtain

$$\|\Delta_{\tilde{J}}\|_{\text{F}}^2 \leq C_1 \frac{s\rho}{\lambda_r - \lambda_{r+1}} \|\Delta_{\tilde{J}}\|_{\text{F}} + C_2 \left(\frac{\delta^2}{(\lambda_r - \lambda_{r+1})^2} + \left[\frac{\delta^2}{\rho(\lambda_r - \lambda_{r+1})\sqrt{t}} \right]^2 \right).$$

Solving this equation gives

$$\|\Delta_{\tilde{J}}\|_{\text{F}}^2 \leq C_3 \left(\frac{s^2 \rho^2}{(\lambda_r - \lambda_{r+1})^2} + \frac{\delta^2}{(\lambda_r - \lambda_{r+1})^2} + \left(\frac{\delta^2}{\rho(\lambda_r - \lambda_{r+1})\sqrt{t}} \right)^2 \right)$$

for some positive constant C_3 that is sufficiently large. Also we have

$$\|\Delta_{\tilde{J}^c}\|_{\text{F}} \leq \sum_{m=2}^M \|\Delta_{J_m}\|_{\text{F}} \leq \frac{9s}{\sqrt{t}} \|\Delta_{\tilde{J}}\|_{\text{F}} + \frac{5\delta^2}{(\lambda_r - \lambda_{r+1})\rho\sqrt{t}}.$$

Combining the last two displays and using $t = c_1 s^2$, we obtain that

$$\|\Delta\|_{\mathbb{F}} \leq C_4 \left(\frac{s\rho + \delta}{\lambda_r - \lambda_{r+1}} + \frac{\delta^2}{s\rho(\lambda_r - \lambda_{r+1})} \right).$$

Recall that on event B_3 we pick $\rho = \gamma \sqrt{\frac{\log p}{n}}$. Then on event B_4 , we have

$$\delta \leq C_0 \rho \sqrt{t}.$$

Hence

$$\|\Delta\|_{\mathbb{F}} \leq C_5 \frac{s\rho}{\lambda_r - \lambda_{r+1}} = C_5 \gamma \frac{s\sqrt{\log p}}{\sqrt{n}(\lambda_r - \lambda_{r+1})}$$

on event $B_3 \cap B_4$.

Step (3) The last step follows from $\|AA^\top - \widehat{F}\|_{\mathbb{F}} \leq \|\Delta\|_{\mathbb{F}} + \|\widehat{F} - AA^\top\|_{\mathbb{F}}$. We combine the last display and (49) to obtain the desired bounds for each term on the right side. This finishes our proof. \blacksquare

B.6 Proof of Corollary 10

Proof Define

$$\tau_n^2 = \frac{s^2 \log p}{n(\lambda_r - \lambda_{r+1})^2}.$$

We will bound $\text{dist}(\bar{V}_1, V)$ by a constant multiple of τ_n on event $B_2 \cap B_3 \cap B_4$.

Notice that since $A_0 A_0^\top$ is the best rank r approximation, and AA^\top also has rank r . Thus, we obtain that $\|A_0 A_0^\top - \widehat{F}\|_{\mathbb{F}} \leq \|AA^\top - \widehat{F}\|_{\mathbb{F}}$. Triangle inequality further leads to

$$\|A_0 A_0^\top - AA^\top\|_{\mathbb{F}} \leq \|A_0 A_0^\top - \widehat{F}\|_{\mathbb{F}} + \|AA^\top - \widehat{F}\|_{\mathbb{F}} \leq 2\|AA^\top - \widehat{F}\|_{\mathbb{F}}.$$

Together with Lemma 14, the last display implies $\text{dist}(A, A_0) \leq C_0 \tau_n$ for some positive constant C_0 .

By definition $\widehat{A}_0 = HT(A_0, s')$, \bar{V}_1 and \widetilde{A}_0 are both s' sparse. Moreover, following the lines of the proof of Proposition 19 and Corollary 8, we have

$$\text{dist}(\widetilde{A}_0, A) \leq C_1 \tau_n$$

for some constant $C_1 > 0$. Let P be the orthogonal matrix such that $\|\widetilde{A}_0 P - A\|_{\mathbb{F}} = \text{dist}(\widetilde{A}_0, A)$. Then we have

$$\begin{aligned} \bar{V}_1 P &= \widetilde{A}_0 (I + \widetilde{A}_0^\top \widehat{\Sigma} \widetilde{A}_0 / \lambda)^{1/2} P = \widetilde{A}_0 P P^\top (I + \widetilde{A}_0^\top \widehat{\Sigma} \widetilde{A}_0 / \lambda)^{1/2} P \\ &= \widetilde{A}_0 P (I + P^\top \widetilde{A}_0^\top \widehat{\Sigma} \widetilde{A}_0 P / \lambda)^{1/2}. \end{aligned}$$

Since

$$\text{dist}(\bar{V}_1, V) \leq \|\bar{V}_1 P - V\|_{\mathbb{F}} = \left\| \widetilde{A}_0 P (I + P^\top \widetilde{A}_0^\top \widehat{\Sigma} \widetilde{A}_0 P / \lambda)^{1/2} - A (I + \Lambda_r / \lambda)^{1/2} \right\|_{\mathbb{F}},$$

we turn to bound the rightmost side. To this end, triangle inequality gives

$$\begin{aligned} & \|\tilde{A}_0 P(I + P^\top \tilde{A}_0^\top \widehat{\Sigma} \tilde{A}_0 P / \lambda)^{1/2} - A(I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}} \\ & \leq \underbrace{\|\tilde{A}_0 P(I + P^\top \tilde{A}_0^\top \widehat{\Sigma} \tilde{A}_0 P / \lambda)^{1/2} - \tilde{A}_0 P(I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}}}_{\text{Term I}} + \underbrace{\|(\tilde{A}_0 P - A)(I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}}}_{\text{Term II}}. \end{aligned}$$

We now bound each term separately. For Term II, we have

$$\|(\tilde{A}_0 P - A)(I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}} \leq \|\tilde{A}_0 P - A\|_{\text{F}} \|(I + \Lambda_r / \lambda)^{1/2}\|_{\text{op}} \leq C_2 \tau_n.$$

The last inequality holds since $\lambda = \lambda_1 / c$. To deal with Term I, we define $\Delta = \tilde{A}_0 P - A$ and notice that

$$\begin{aligned} & \|\tilde{A}_0 P(I + P^\top \tilde{A}_0^\top \widehat{\Sigma} \tilde{A}_0 P / \lambda)^{1/2} - \tilde{A}_0 P(I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}} \\ & \leq \|\tilde{A}_0 P\|_{\text{op}} \|(I + P^\top \tilde{A}_0^\top \widehat{\Sigma} \tilde{A}_0 P / \lambda)^{1/2} - (I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}}. \end{aligned}$$

By matrix root perturbation bound (e.g., Lemma 2 in Gao et al. (2017)), we have

$$\begin{aligned} & \|(I + P^\top \tilde{A}_0^\top \widehat{\Sigma} \tilde{A}_0 P / \lambda)^{1/2} - (I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}} \\ & \leq C_3 \|P^\top \tilde{A}_0^\top \widehat{\Sigma} \tilde{A}_0 P - \Lambda_r\|_{\text{F}} \\ & \leq C_3 \|(A + \Delta)^\top \widehat{\Sigma} (A + \Delta) - A^\top \Sigma A\|_{\text{F}} \\ & = C_3 (\|A^\top (\widehat{\Sigma} - \Sigma) A\|_{\text{F}} + 2\|\Delta^\top \widehat{\Sigma} A\|_{\text{F}} + \|\Delta^\top \widehat{\Sigma} \Delta\|_{\text{F}}) \end{aligned}$$

for some positive constant C_3 . Since \tilde{A}_0 is s' sparse and A is s sparse, $\Delta = \tilde{A}_0 P - A$ is $s + s'$ sparse. By definition of P , $\|\Delta\|_{\text{F}} = \text{dist}(\tilde{A}_0, A) \leq C_1 \tau_n$. By a similar argument to the proof of Corollary 8, on event B_2 , we have

$$\|(I + P^\top \tilde{A}_0^\top \widehat{\Sigma} \tilde{A}_0 P / \lambda)^{1/2} - (I + \Lambda_r / \lambda)^{1/2}\|_{\text{F}} \leq C_4 \sqrt{\frac{rs \log p}{n}} \leq C_4 \tau_n.$$

Consequently, Term I is also dominated by a constant multiple of τ_n . Thus, we conclude that on the intersection of $B_2 \cap B_3 \cap B_4$,

$$\text{dist}(\bar{V}_1, V) \leq C_5 \tau_n$$

for some positive constant C_5 . By union bound, the intersection of the three events holds with probability at least $1 - \exp(-C'(s + \log(ep/s)))$, uniformly over \mathcal{P}_n . This completes the proof. \blacksquare

B.7 Proof of Lemma 15

Proof In this proof, we work on event B_2 defined in (34), which occurs with probability at least $1 - \exp(-C' s' \log(ep/s'))$ for some constant $C' > 0$, uniformly over \mathcal{P}_n . The proof relies on Theorem 3.1 of Sun (1983), a matrix perturbation bound for generalized eigenspaces. Fix any \mathcal{I} such that $S \subset \mathcal{I}$ and $|\mathcal{I}| \leq 2s' + s$. We shall apply the theorem on matrix pair $(\widehat{\Sigma}_{\mathcal{I}\mathcal{I}}, \Sigma_{0,\mathcal{I}\mathcal{I}})$ and its perturbation $(\widehat{\Sigma}_{\mathcal{I}\mathcal{I}}, \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}})$.

To this end, for the fixed set \mathcal{I} , define

$$\delta = \min_{\lambda, \lambda'} \frac{\lambda - \lambda'}{\sqrt{(1 + \lambda^2)(1 + \lambda'^2)}} \quad \text{for } \lambda \in \{\lambda_1, \dots, \lambda_r\} \text{ and } \lambda' \in \{\widehat{\lambda}_{r+1}, \dots, \widehat{\lambda}_{|\mathcal{I}|}\},$$

where λ_i denotes the i th generalized eigenvalue of matrix pair $(\Sigma_{\mathcal{I}\mathcal{I}}, \Sigma_{0,\mathcal{I}\mathcal{I}})$ and $\widehat{\lambda}_i$ denotes the i th sample generalized eigenvalue of matrix pair $(\widehat{\Sigma}_{\mathcal{I}\mathcal{I}}, \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}})$. Define event

$$\widetilde{B}_2 = \left\{ \widehat{\lambda}_{r+1} \leq 2\lambda_{r+1} \wedge \frac{\lambda_r + \lambda_{r+1}}{2} \right\}.$$

On event \widetilde{B}_2 , since generalized eigenvalues are in decreasing order, we have

$$\delta = \frac{\lambda_r - \widehat{\lambda}_{r+1}}{\sqrt{(1 + \lambda_1^2)(1 + \widehat{\lambda}_{r+1}^2)}} \geq \frac{1}{2} \frac{\lambda_r - \lambda_{r+1}}{\sqrt{(1 + \lambda_1^2)(1 + \widehat{\lambda}_{r+1}^2)}}.$$

Moreover, since $\widehat{\lambda}_{r+1}^2 \leq 4\lambda_{r+1}^2$ on \widetilde{B}_2 , we further obtain that

$$\delta \geq \frac{1}{4} \frac{\lambda_r - \lambda_{r+1}}{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}.$$

On the other hand, \widetilde{B}_2 can be equivalently defined as

$$\widetilde{B}_2 = \left\{ \widehat{\lambda}_{r+1} - \lambda_{r+1} \leq \lambda_{r+1} \wedge \frac{\lambda_r - \lambda_{r+1}}{2} \right\}.$$

Since λ_i is also the i th eigenvalue of $\Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2} \Sigma_{\mathcal{I}\mathcal{I}} \Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2}$ and $\widehat{\lambda}_i$ the i th eigenvalue of $\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2} \widehat{\Sigma}_{\mathcal{I}\mathcal{I}} \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2}$, Weyl's inequality then implies that on event B_2 ,

$$\begin{aligned} |\widehat{\lambda}_{r+1} - \lambda_{r+1}| &\leq \|\Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2} \Sigma_{\mathcal{I}\mathcal{I}} \Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2} - \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2} \widehat{\Sigma}_{\mathcal{I}\mathcal{I}} \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} \\ &\leq \|(\Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2} - \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2}) \Sigma_{\mathcal{I}\mathcal{I}} \Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} + \|\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2} (\Sigma_{\mathcal{I}\mathcal{I}} - \widehat{\Sigma}_{\mathcal{I}\mathcal{I}}) \Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} \\ &\quad + \|\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2} \widehat{\Sigma}_{\mathcal{I}\mathcal{I}} (\Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2} - \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2})\|_{\text{op}} \\ &\leq \|\Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2} - \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} \|\Sigma_{\mathcal{I}\mathcal{I}} \Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} + \|\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} \|\Sigma_{\mathcal{I}\mathcal{I}} - \widehat{\Sigma}_{\mathcal{I}\mathcal{I}}\|_{\text{op}} \|\Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} \\ &\quad + \|\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2} \widehat{\Sigma}_{\mathcal{I}\mathcal{I}}\|_{\text{op}} \|\Sigma_{0,\mathcal{I}\mathcal{I}}^{-1/2} - \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}^{-1/2}\|_{\text{op}} \\ &\leq C_0 \sqrt{\frac{|\mathcal{I}| \log p}{n}} \end{aligned}$$

for some constant C_0 depending on ν . Here the last inequality is due to the matrix root bound from Lemma 2 in Gao et al. (2015a). Under condition (23), by the choice of s' in (20), the last display implies $|\widehat{\lambda}_{r+1} - \lambda_{r+1}| \leq (\frac{\lambda_r - \lambda_{r+1}}{2} \wedge \lambda_{r+1})$. Hence \widetilde{B}_2 holds on event B_2 under condition (23).

Furthermore, we have

$$\gamma(\Sigma_{\mathcal{I}\mathcal{I}}, \Sigma_{0,\mathcal{I}\mathcal{I}}) = \min_{\|x\|=1} \sqrt{(x^\top \Sigma_{\mathcal{I}\mathcal{I}} x)^2 + (x^\top \Sigma_{0,\mathcal{I}\mathcal{I}} x)^2} \geq \frac{1}{\nu},$$

and

$$\gamma(\widehat{\Sigma}_{\mathcal{I}\mathcal{I}}, \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}}) = \min_{\|x\|=1} \sqrt{(x^\top \widehat{\Sigma}_{\mathcal{I}\mathcal{I}} x)^2 + (x^\top \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}} x)^2} \geq \frac{1}{2\nu},$$

on event B_2 . Now we apply Theorem 3.1 of Sun (1983) for perturbation on restricted covariance matrices and obtain that

$$\begin{aligned} \text{dist}(A, \widehat{A}(\mathcal{I})) &\leq C_1 \frac{\sqrt{\|\Sigma_{\mathcal{I}\mathcal{I}}^2 + \Sigma_{0,\mathcal{I}\mathcal{I}}^2\|_{\text{op}}}}{\gamma(\Sigma_{\mathcal{I}\mathcal{I}}, \Sigma_{0,\mathcal{I}\mathcal{I}}) \gamma(\widehat{\Sigma}_{\mathcal{I}\mathcal{I}}, \widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}})} \frac{\sqrt{\|(\widehat{\Sigma}_{\mathcal{I}\mathcal{I}} - \Sigma_{\mathcal{I}\mathcal{I}})A_{\mathcal{I}^*}\|_{\text{F}}^2 + \|(\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}} - \Sigma_{0,\mathcal{I}\mathcal{I}})A_{\mathcal{I}^*}\|_{\text{F}}^2}}{\delta} \\ &\leq C_2 \nu^3 \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\|A_{\mathcal{I}^*}\|_{\text{F}}^2 (\|\widehat{\Sigma}_{\mathcal{I}\mathcal{I}} - \Sigma_{\mathcal{I}\mathcal{I}}\|_{\text{op}}^2 + \|\widehat{\Sigma}_{0,\mathcal{I}\mathcal{I}} - \Sigma_{0,\mathcal{I}\mathcal{I}}\|_{\text{op}}^2)} \\ &\leq C_3 \nu^3 \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{r \frac{|\mathcal{I}| \log p}{n}} \\ &\leq C_4 \sqrt{r} \frac{\sqrt{1 + \lambda_1^2} \sqrt{1 + \lambda_{r+1}^2}}{\lambda_r - \lambda_{r+1}} \sqrt{\frac{|\mathcal{I}| \log p}{n}}, \end{aligned}$$

for some positive constant C_4 on event B_2 . This finishes the proof of Lemma 15. \blacksquare

B.8 Proofs of Lemma 16 and Lemma 17

Proof Note that $\lambda_{r+1} < 1$. Hence, it suffices to show that each of the four infinity norm terms on the left side in the definition of event B_3 in (35) is upper bounded by a constant multiple of $\sqrt{\log p/n}$ with probability at least $1 - p^{-C'}$ for some positive constant C' , uniformly over \mathcal{P}_n . This can be achieved by following the lines of the proof of Lemma 6.4 in Gao et al. (2017) and so we omit the details. \blacksquare

Proof We first bound the operator norm of each term on the left side of (37) by a constant multiple of $\sqrt{\frac{1}{n}(s + \log(ep/s))}$, and the result on Frobenius follows directly since each term is of rank r .

As in the proof of Lemma 6.1 of Gao et al. (2017), we have

$$\|\Sigma_0^{1/2}(\widetilde{A} - A)\|_{\text{op}} \leq \|\Sigma_0^{1/2} A\|_{\text{op}} \|(A^\top \widehat{\Sigma}_0 A)^{1/2} - I\|_{\text{op}} \|(A^\top \widehat{\Sigma}_0 A)^{-1/2}\|_{\text{op}},$$

$$\|\widetilde{\Lambda}_r - \Lambda_r\|_{\text{op}} \leq \|(A^\top \widehat{\Sigma}_0 A)^{1/2} - I\|_{\text{op}} \|\Lambda_r (A^\top \widehat{\Sigma}_0 A)^{1/2}\|_{\text{op}} + \|\Lambda_r\|_{\text{op}} \|(A^\top \widehat{\Sigma}_0 A)^{1/2} - I\|_{\text{op}}.$$

So it remains to bound $\|(A^\top \widehat{\Sigma}_0 A)^{1/2} - I\|_{\text{op}}$. By Lemma 2 in Gao et al. (2015a), it suffices to bound $\|A^\top \widehat{\Sigma}_0 A - I\|_{\text{op}}$. Note that

$$\begin{aligned} \|A^\top \widehat{\Sigma}_0 A - I\|_{\text{op}} &= \|A^\top (\widehat{\Sigma}_0 - \Sigma_0) A\|_{\text{op}} \\ &= \|A_{S^*}^\top (\widehat{\Sigma}_{0,SS} - \Sigma_{0,SS}) A_{S^*}\|_{\text{op}} \\ &= \sup_{\|v\|=1} (A_{S^*} v)^\top (\widehat{\Sigma}_{0,SS} - \Sigma_{0,SS}) (A_{S^*} v) \\ &\leq \|\Sigma_{0,SS}^{1/2} A_{S^*}\|_{\text{op}}^2 \|\Sigma_{0,SS}^{-1/2} \widehat{\Sigma}_{0,SS} \Sigma_{0,SS}^{-1/2} - I\|_{\text{op}} \\ &\leq \|\Sigma_{0,SS}^{-1/2} \widehat{\Sigma}_{0,SS} \Sigma_{0,SS}^{-1/2} - I\|_{\text{op}} \leq C \|\widehat{\Sigma}_{0,SS} - \Sigma_{0,SS}\|_{\text{op}}. \end{aligned}$$

By Lemma 3 of Gao et al. (2015a), we then have

$$\|A^\top \widehat{\Sigma}_0 A - I\|_{\text{op}} \leq C_1 \sqrt{\frac{1}{n} \left(s + \log \frac{ep}{s} \right)},$$

with probability at least $1 - \exp(-C'(s + \log(ep/s)))$ for some constant $C' > 0$, uniformly over \mathcal{P}_n . This completes the proof. \blacksquare

Appendix C. Proof of Lower Bound

We first present a lemma on Kullback-Leibler divergence between data distributions from a special kind of covariance matrix. The Lemma can be viewed as a general case of Lemma 1 in Gao et al. (2015b).

Lemma 22 *For $t = 1, 2$, define $\Sigma^{(t)}$ to be a block matrix whose (i, j) th block is given by $\lambda U_{\{i\}}^{(t)} U_{\{j\}}^{(t)\top}$ for $i \neq j$, where $U_{\{i\}}^{(t)} \in O(p_i, r)$ and whose i th diagonal block is given by I_{p_i} . Further let $P^{(t)}$ denote the distribution of a random i.i.d sample of size n from the $N_p(0, \Sigma^{(t)})$ distribution where $p = \sum_{i=1}^k p_i$. Then*

$$D(P^{(1)} \| P^{(2)}) = \frac{\lambda^2 n}{4(-(k-1)\lambda^2 + (k-2)\lambda + 1)} \sum_{i \neq j} \left\| U_{\{i\}}^{(1)} U_{\{j\}}^{(1)\top} - U_{\{i\}}^{(2)} U_{\{j\}}^{(2)\top} \right\|_{\text{F}}^2.$$

Proof The first step is to find the eigenvalues of $\Sigma^{(t)}$. To this end, define

$$\begin{aligned} U^{(t)} &= \left[U_{\{1\}}^{(t)\top}, U_{\{2\}}^{(t)\top}, \dots, U_{\{k\}}^{(t)\top} \right]^\top, \\ K_i^{(t)} &= \left[-U_{\{1\}}^{(t)\top}, -U_{\{2\}}^{(t)\top}, \dots, -U_{\{i-1\}}^{(t)\top}, (k-1)U_{\{i\}}^{(t)\top}, -U_{\{i+1\}}^{(t)\top}, \dots, -U_{\{k\}}^{(t)\top} \right]^\top. \end{aligned}$$

That is, $K_i^{(t)} \in \mathbb{R}^{p \times r}$ is a block matrix with j th blocks being $-U_{\{j\}}^{(t)\top}$ for $j \neq i$ and the i th block being $(k-1)U_{\{i\}}^{(t)\top}$. It is straightforward to verify that $\Sigma^{(t)}$ yields the following decomposition

$$\Sigma^{(t)} = I + \frac{(k-1)\lambda}{k} U^{(t)} U^{(t)\top} - \frac{\lambda}{k^2} \sum_{i=1}^k K_i^{(t)} K_i^{(t)\top}.$$

We then deduce that the eigen-structure of $\Sigma^{(t)}$ is given by r eigenvalues equal to $1 + \lambda(k-1)$, $(k-1)r$ eigenvalues equal to $1 - \lambda$ and $p - kr$ eigenvalues equal to 1. Furthermore, we notice that

$$\Sigma^{(t)}U^{(t)} = (1 + (k-1)\lambda)U^{(t)},$$

hence the leading r dimensional eigenspace is given by the matrix $U^{(t)}$. This in particular implies that

$$\det \Sigma^{(1)} = \det \Sigma^{(2)}.$$

Now the KL divergence is given by

$$\begin{aligned} D(P^{(1)}||P^{(2)}) &= \frac{n}{2} \left(\text{Tr} \left((\Sigma^{(2)})^{-1} \Sigma^{(1)} \right) - \sum_{i=1}^k p_i - \log \det \left((\Sigma^{(2)})^{-1} \Sigma^{(1)} \right) \right) \\ &= \frac{n}{2} \left(\text{Tr} \left((\Sigma^{(2)})^{-1} \Sigma^{(1)} \right) - p \right) \\ &= \frac{n}{2} \left(\text{Tr} \left((\Sigma^{(2)})^{-1} \left(\Sigma^{(1)} - \Sigma^{(2)} \right) \right) \right). \end{aligned}$$

We now calculate the inverse of the block matrix $\Sigma^{(2)}$. To this end, we first guess the form of the inverse and then determine the coefficients. Specifically, we try to solve for the inverse with the following form

$$(\Sigma^{(2)})^{-1} = \begin{bmatrix} I + aU_{\{1\}}^{(2)}U_{\{1\}}^{(2)\top} & bU_{\{1\}}^{(2)}U_{\{2\}}^{(2)\top} & \cdots & \cdots & bU_{\{1\}}^{(2)}U_{\{k\}}^{(2)\top} \\ bU_{\{2\}}^{(2)}U_{\{1\}}^{(2)\top} & I + aU_{\{2\}}^{(2)}U_{\{2\}}^{(2)\top} & bU_{\{2\}}^{(2)}U_{\{3\}}^{(2)\top} & \cdots & bU_{\{2\}}^{(2)}U_{\{k\}}^{(2)\top} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ bU_{\{k\}}^{(2)}U_{\{1\}}^{(2)\top} & \cdots & \cdots & \cdots & I + aU_{\{k\}}^{(2)}U_{\{k\}}^{(2)\top} \end{bmatrix}$$

with the (i, j) th block being $bU_{\{i\}}^{(2)}U_{\{j\}}^{(2)\top}$ for $i \neq j$ and i th diagonal block being $I + aU_{\{i\}}^{(2)}U_{\{i\}}^{(2)\top}$. To determine the values of (a, b) , we solve for the equation $(\Sigma^{(2)})^{-1}\Sigma^{(2)} = I_p$. This requires two conditions on the matrix product: the off-diagonal block to be 0 and the diagonal block to be I . For the i th diagonal block, we require

$$I + (k-1)b\lambda U_{\{i\}}^{(2)}(U_{\{i\}}^{(2)})^\top + aU_{\{i\}}^{(2)}(U_{\{i\}}^{(2)})^\top = I,$$

and for the (i, j) th off-diagonal block, we require

$$\lambda b(k-2)U_{\{i\}}^{(2)}(U_{\{j\}}^{(2)})^\top + bU_{\{i\}}^{(2)}(U_{\{j\}}^{(2)})^\top + \lambda U_{\{i\}}^{(2)}(U_{\{j\}}^{(2)})^\top + a\lambda U_{\{i\}}^{(2)}(U_{\{j\}}^{(2)})^\top = 0.$$

Solving the above equations, we have

$$\begin{aligned} a &= -(k-1)b\lambda, \\ b &= \frac{\lambda}{(k-1)\lambda^2 - \lambda(k-2) - 1}. \end{aligned}$$

As a result, we can simplify the expression for KL divergence as follows

$$\begin{aligned}
 D(P^{(1)}||P^{(2)}) &= \frac{n}{2} \left(\text{Tr} \left((\Sigma^{(2)})^{-1} \left(\Sigma^{(1)} - \Sigma^{(2)} \right) \right) \right) \\
 &= \frac{\lambda n}{2} \left(\text{Tr} \left(\sum_{i=1}^k \sum_{j \neq i} b U_{\{i\}}^{(2)} (U_{\{j\}}^{(2)})^\top \left(U_{\{j\}}^{(1)} (U_{\{i\}}^{(1)})^\top - U_{\{j\}}^{(2)} (U_{\{i\}}^{(2)})^\top \right) \right) \right) \\
 &= -\frac{b\lambda n}{2} \left(\sum_{i,j:i \neq j} \text{Tr} \left(I - U_{\{i\}}^{(2)} (U_{\{j\}}^{(2)})^\top U_{\{j\}}^{(1)} (U_{\{i\}}^{(1)})^\top \right) \right) \\
 &= -\frac{b\lambda n}{4} \sum_{i,j:i \neq j} \left\| U_{\{i\}}^{(1)} (U_{\{j\}}^{(1)})^\top - U_{\{i\}}^{(2)} (U_{\{j\}}^{(2)})^\top \right\|_{\text{F}}^2 \\
 &= \frac{\lambda^2 n}{4(-(k-1)\lambda^2 + (k-2)\lambda + 1)} \sum_{i,j:i \neq j} \left\| U_{\{i\}}^{(1)} (U_{\{j\}}^{(1)})^\top - U_{\{i\}}^{(2)} (U_{\{j\}}^{(2)})^\top \right\|_{\text{F}}^2.
 \end{aligned}$$

This finishes our proof for the Lemma. Note that when $k = 2$, we recover Lemma 1 in Gao et al. (2015b) as a special case. ■

C.1 Proof of Theorem 12

Proof The main body of the proof is adapted from Gao et al. (2015a). The main tool for our proof is Fano's Lemma. For the sake of completeness, we provide the following version of Fano's Lemma from Yu et al. (1997).

Lemma 23 *Let (Θ, ρ) be a metric space and $\{P_\theta : \theta \in \Theta\}$ a collection of probability measures. For any totally bounded $T \in \Theta$, denoted by $\mathcal{M}(T, \rho, \epsilon)$ the ϵ -packing number of T with respect to ρ , i.e., the maximal number of points in T whose pairwise maximum distance in ρ is at least ϵ . Define the Kullback-Leibler diameter T by*

$$d_{\text{KL}}(T) \triangleq \sup_{\theta, \theta' \in T} D(P_\theta || D_{\theta'}).$$

Then,

$$\inf_{\theta} \sup_{\theta \in \Theta} \mathbb{E}_\theta [\rho^2(\hat{\theta}(X), \theta)] \geq \sup_{T \in \Theta} \sup_{\epsilon > 0} \frac{\epsilon^2}{4} \left(1 - \frac{d_{\text{KL}}(T) + \log 2}{\log \mathcal{M}(T, \rho, \epsilon)} \right).$$

The proof is composed of two steps corresponding to different terms in the lower bound. Throughout the proof, we define $\lambda = \lambda_r - \lambda_{r+1}$ for simplicity.

Step I. We first establish the term involving $r \sum_{i=1}^k s_i$. To this end, let $U_{\{i\}0} = \begin{bmatrix} I_r \\ \mathbf{0} \end{bmatrix} \in O(p_i, r)$ for each $i = 1, 2, \dots, k$. For some $\epsilon_0 \in (0, \sqrt{r \wedge (s_1 - r)}]$ to be specified later, let

$$B(\epsilon_0) = \{U_{\{1\}} \in O(p_1, r) : \text{supp}(U_{\{1\}}) \subset [s_1], \|U_{\{1\}} - U_{\{1\}0}\|_{\text{F}} \leq \epsilon_0\}$$

and

$$T_0 = \left\{ \Sigma = \begin{pmatrix} I_{p_1} & \lambda U_{\{1\}} U_{\{2\}0}^\top & \lambda U_{\{1\}} U_{\{3\}0}^\top \cdots & \lambda U_{\{1\}} U_{\{k\}0}^\top \\ \lambda U_{\{2\}0} U_{\{1\}}^\top & I_{p_2} & \cdots & \lambda U_{\{2\}0} U_{\{k\}0}^\top \\ \cdots & \cdots & \cdots & \cdots \\ \lambda U_{\{k\}0} U_{\{1\}}^\top & \lambda U_{\{k\}0} U_{\{2\}0}^\top & \cdots & I_{p_k} \end{pmatrix} : U_{\{1\}} \in B(\epsilon_0) \right\},$$

with $U = [U_{\{1\}}^\top, U_{\{2\}0}^\top, \dots, U_{\{k\}0}^\top]^\top$. Since our target of estimation is A instead of U , we first establish the relationship between U and A under T_0 . Note that under the construction of Σ , we have $\Sigma_0 = I_p$ so the generalized eigenspace coincides with eigenspace. From Lemma 22, we deduce that the leading r dimensional eigenspace of Σ is given by $\text{span}(U)$. From the normalization constraint such that $A^\top A = I$, we conclude that $\text{dist}(A, \frac{1}{\sqrt{k}}U) = 0$. From here on, we first derive the minimax lower bound for the estimation of U and the lower bound for estimating A under the matrix distance follows immediately by scaling.

The above analysis also implies that $T_0 \subset \mathcal{F}$, where \mathcal{F} is our original parameter space. By Lemma 22,

$$\begin{aligned} d_{\text{KL}}(T_0) &= \sup_{U_{\{1\}}^{(1)}, U_{\{1\}}^{(2)} \in B(\epsilon_0)} D(P^{(1)} \| P^{(2)}) \\ &= \sup_{U_{\{1\}}^{(1)}, U_{\{1\}}^{(2)} \in B(\epsilon_0)} \frac{\lambda^2 n}{4(-(k-1)\lambda^2 + (k-2)\lambda + 1)} \sum_{i \neq j} \left\| U_{\{i\}}^{(1)} U_{\{j\}}^{(1)\top} - U_{\{i\}}^{(2)} U_{\{j\}}^{(2)\top} \right\|_{\text{F}}^2 \\ &= \sup_{U_{\{1\}}^{(1)}, U_{\{1\}}^{(2)} \in B(\epsilon_0)} \frac{\lambda^2 n}{4(-(k-1)\lambda^2 + (k-2)\lambda + 1)} \sum_{j=2}^k 2 \left\| U_{\{1\}}^{(1)} U_{\{j\}0}^{(1)\top} - U_{\{1\}}^{(2)} U_{\{j\}0}^{(2)\top} \right\|_{\text{F}}^2 \\ &= \sup_{U_{\{1\}}^{(1)}, U_{\{1\}}^{(2)} \in B(\epsilon_0)} \frac{\lambda^2 n}{2(-(k-1)\lambda^2 + (k-2)\lambda + 1)} \sum_{j=2}^k \left\| U_{\{1\}}^{(1)} - U_{\{1\}}^{(2)} \right\|_{\text{F}}^2 \\ &= \frac{2\lambda^2 n(k-1)\epsilon_0^2}{-(k-1)\lambda^2 + (k-2)\lambda + 1}. \end{aligned}$$

Here, the second to last equality follows from the definition of $B(\epsilon_0)$.

We now establish a lower bound for the packing number of T_0 . For some $\alpha \in (0, 1)$ which shall be specified later, we define $\{\tilde{U}_{\{1\}}(1), \tilde{U}_{\{1\}}(2), \dots, \tilde{U}_{\{1\}}(N)\} \subset O(p_1, r)$ to be a maximal set such that $\text{supp}(\tilde{U}_{\{1\}}(i)) \subset [s_1]$ and for $\forall i \neq j$,

$$\left\| \tilde{U}_{\{1\}}(i) \tilde{U}_{\{1\}}^\top(i) - U_{\{1\}0} U_{\{1\}0}^\top \right\|_{\text{F}} \leq \epsilon_0, \quad \left\| \tilde{U}_{\{1\}}(i) \tilde{U}_{\{1\}}^\top(i) - \tilde{U}_{\{1\}}(j) \tilde{U}_{\{1\}}^\top(j) \right\|_{\text{F}} \geq \sqrt{2}\alpha\epsilon_0.$$

Then by Lemma 1 in Cai et al. (2013), for some absolute constant $C > 1$,

$$N \geq \left(\frac{1}{C\alpha} \right)^{r(s_1-r)}.$$

For each $\tilde{U}_{\{1\}}(i)$, we define $\bar{U}_{\{1\}}(i)$ to be the matrix such that

$$\left\| U_{\{1\}0} - \bar{U}_{\{1\}}(i) \right\|_{\text{F}}^2 = \text{dist}^2 \left(\tilde{U}_{\{1\}}(i), U_{\{1\}0} \right).$$

Then for any i , we deduce that $\bar{U}_{\{1\}}(i) \in O(p_1, r)$, $\text{supp}(\bar{U}_{\{1\}}(i)) \subset [s_1]$ and $\bar{U}_{\{1\}}(i)\bar{U}_{\{1\}}(i)^\top = \tilde{U}_{\{1\}}(i)\tilde{U}_{\{1\}}(i)^\top$. In addition, Lemma 6.6 in Gao et al. (2017) implies that

$$\|\bar{U}_{\{1\}}(i) - U_{\{1\}0}\|_F = \text{dist}\left(\tilde{U}_{\{1\}}(i), U_{\{1\}0}\right) = \frac{1}{\sqrt{2}}\|\tilde{U}_{\{1\}}(i)\tilde{U}_{\{1\}}^\top(i) - U_{\{1\}0}U_{\{1\}0}^\top\|_F \leq \epsilon_0$$

hence $\bar{U}_{\{1\}}(i) \in B(\epsilon_0)$. On the other hand, note that from Lemma 6.6 in Gao et al. (2017),

$$\begin{aligned} \text{dist}(\bar{U}_{\{1\}}(i), \bar{U}_{\{1\}}(j)) &= \frac{1}{\sqrt{2}}\|\bar{U}_{\{1\}}(i)\bar{U}_{\{1\}}^\top(i) - \bar{U}_{\{1\}}(j)\bar{U}_{\{1\}}^\top(j)\|_F \\ &= \frac{1}{\sqrt{2}}\|\tilde{U}_{\{1\}}(i)\tilde{U}_{\{1\}}^\top(i) - \tilde{U}_{\{1\}}(j)\tilde{U}_{\{1\}}^\top(j)\|_F \geq \alpha\epsilon_0. \end{aligned}$$

Define the metric to be $\rho(\Sigma^{(1)}, \Sigma^{(2)}) = \text{dist}(U_{\{1\}}^{(1)}, U_{\{1\}}^{(2)})$. The above argument implies that for $\epsilon = \alpha\epsilon_0$,

$$\log \mathcal{M}(T_0, \rho, \epsilon) \geq r(s_1 - r) \log \frac{1}{C\alpha}.$$

Setting

$$\epsilon = c_0 \left[\sqrt{r \wedge (s_1 - r)} \wedge \sqrt{\frac{-(k-1)\lambda^2 + (k-2)\lambda + 1}{n(k-1)\lambda^2} r(s_1 - r)} \right]$$

for sufficiently small constants c_0, α , we obtain a lower bound of the order

$$r \wedge (s_1 - r) \wedge \frac{-(k-1)\lambda^2 + (k-2)\lambda + 1}{n(k-1)\lambda^2} r(s_1 - r)$$

by applying Lemma 23. By symmetry, we also have the lower bound with $r(s_1 - r)$ replaced by $r(s_i - r)$ for $i = 2, 3, \dots, k$. Furthermore, recall that we have $r \leq \frac{1}{2} \min_i s_i$ hence $r \leq (s_i - r)$, we obtain the lower bound of the order

$$r \wedge \frac{r \sum_{i=1}^k s_i}{n\lambda^2}$$

for the estimation of U with metric $\text{dist}^2(U, \hat{U})$, when k is finite.

Step II We now turn to establish the lower bound term involving $s_i \log \frac{ep_i}{s_i}$. The step follows from the rank-one argument from Chen et al. (2013). Without loss of generality, we may assume that $s_i \leq \frac{p_i}{2}$ for any i , we then consider the following subset of the parameter space:

$$T_1 = \left\{ \Sigma = \begin{pmatrix} I_{p_1} & \lambda U_{\{1\}}U_{\{2\}0}^\top & \lambda U_{\{1\}}U_{\{3\}0}^\top \cdots & \lambda U_{\{1\}}U_{\{k\}0}^\top \\ \lambda U_{\{2\}0}U_{\{1\}}^\top & I_{p_2} & \cdots & \lambda U_{\{2\}0}U_{\{k\}0}^\top \\ \cdots & \cdots & \cdots & \cdots \\ \lambda U_{\{k\}0}U_{\{1\}}^\top & \lambda U_{\{k\}0}U_{\{2\}0}^\top & \cdots & I_{p_k} \end{pmatrix} : U_{\{1\}} = \begin{bmatrix} I_{r-1} & 0 \\ 0 & u_r \end{bmatrix} \right\},$$

for $u_r \in \mathbb{S}^{p_1-r+1}$ and $|\text{supp}(u_r)| \leq s_1 - r + 1$. Restricting on the set T_1 , the minimax risk for estimating U is the same as the minimax risk for estimating u_r under the squared error

loss $\|u_r - \hat{u}_r\|_{\mathbb{F}}^2$. Let $X_{\{i\}} = [X_{\{i\}1}, X_{\{i\}2}]$ for $X_{\{i\}1} \in \mathbb{R}^{n \times (r-1)}$ and $X_{\{i\}2} \in \mathbb{R}^{n \times (p_i - r + 1)}$. By the same argument in Gao et al. (2015a), it is further equivalent to estimating u_r under the squared loss based on the observations from $X_{\{i\}2}$ since $X_{\{i\}2}$ for $i = 1, 2, \dots, k$ is a sufficient statistic for u_r . Applying the argument in Chen et al. (2013), we obtain the lower bound for $\text{dist}^2(U, \hat{U})$ with the following term

$$\frac{1}{n\lambda^2} \left(s_1 \log \frac{ep_1}{s_1} \right) \wedge 1.$$

By symmetry the same lower bound holds if we replace s_1, p_1 by s_i, p_i for $i = 2, 3, \dots, k$.

Combining all the above steps and noticing the $\frac{1}{\sqrt{k}}$ scaling between the estimation of A and U , we finally deduce that the lower bound for $\text{dist}^2(A, \hat{A})$ is given by (30). This finishes our proof. \blacksquare

Appendix D. Simulation Details for Sparse PCA of Correlation Matrices

Here we describe the procedure of generating a $p \times p$ correlation matrix R with eigenvalues $\lambda_1 \geq \dots \geq \lambda_r > \lambda_{r+1} = 1 \geq \dots \geq \lambda_p > 0$ and s -sparse leading eigenvectors. Necessarily, $\sum_{i=1}^r \lambda_i \leq s < p - 1$ and $s \geq r$. Without loss of generality, we assume that $s = m \times 2^l$ where m and l are positive integers and $2^l \geq r$.

Since $2^l \geq r$, there is a $2^l \times r$ matrix T_0 such that $T_{0,ij} \in \{\pm 1\}$ and that the columns of T_0 are orthogonal. For example, we may take the first r columns of a $2^l \times 2^l$ Hadamard matrix. Fix such a T_0 , we generate an $s \times r$ matrix T_s as

$$T_s = \frac{1}{\sqrt{s}} [T_0^\top \quad \dots \quad T_0^\top]^\top.$$

By our construction, each row of T_s has the same l_2 norm and the columns are orthonormal. Next, we define

$$R_s = T_s \text{diag}(\theta_1, \dots, \theta_r) T_s^\top + \left(1 - \frac{\sum_{i=1}^r \theta_i}{s} \right) I_s.$$

This decomposition ensures that the diagonal of R_s is equal to 1. We then solve for θ_i such that the eigenvalues of R_s are given by $\lambda_1, \dots, \lambda_r$. Then we augment this $s \times s$ positive definite matrix to a $p \times p$ correlation matrix $R = \text{diag}(R_s, I_{p-s})$. It is straightforward to verify that R has leading eigenvalues $\lambda_1 \geq \dots \geq \lambda_r > 1$ and the columns of $T = [T_s^\top \quad O_{(p-s) \times r}^\top]^\top$ are the corresponding eigenvectors.

References

- A.A. Amini and M.J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *The Annals of Statistics*, 37(5B):2877–2921, 2009.
- Q. Berthet and P. Rigollet. Computational lower bounds for sparse pca. *arXiv preprint arXiv:1304.0828*, 2013.

- Peter J Bickel and Elizaveta Levina. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008a.
- Peter J Bickel and Elizaveta Levina. Regularized estimation of large covariance matrices. *The Annals of Statistics*, 36(1):199–227, 2008b.
- Peter J Bickel, Ya’acov Ritov, and Alexandre B Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- A. Birnbaum, I.M. Johnstone, B. Nadler, and D. Paul. Minimax bounds for sparse PCA with noisy high-dimensional data. *The Annals of Statistics*, 41(3):1055–1084, 2013.
- T Tony Cai, Zongming Ma, and Yihong Wu. Sparse pca: Optimal rates and adaptive estimation. *The Annals of Statistics*, 41(6):3074–3110, 2013.
- Mengjie Chen, Chao Gao, Zhao Ren, and Harrison H Zhou. Sparse cca via precision adjusted iterative thresholding. *arXiv preprint arXiv:1311.6186*, 2013.
- Yuxin Chen and Emmanuel Candes. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Advances in Neural Information Processing Systems*, pages 739–747, 2015.
- Yuejie Chi, Yue M Lu, and Yuxin Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.
- Jon Dattorro. Convex optimization and euclidean distance geometry, 2005. *Meboo, Palo Alto*, 2003.
- Deborah J Donnell, Andreas Buja, and Werner Stuetzle. Analysis of additive dependencies and concavities using smallest additive principal components. *The Annals of Statistics*, pages 1635–1668, 1994.
- David L Donoho. High-dimensional data analysis: The curses and blessings of dimensionality. *AMS math challenges lecture*, 1(2000):32, 2000.
- Jianqing Fan, Jianhua Guo, and Shurong Zheng. Estimating number of factors by adjusted eigenvalues thresholding. *arXiv preprint arXiv:1909.10710*, 2019.
- Chao Gao, Zongming Ma, Zhao Ren, and Harrison H Zhou. Minimax estimation in sparse canonical correlation analysis. *The Annals of Statistics*, 43(5):2168–2197, 2015a.
- Chao Gao, Zongming Ma, Zhao Ren, and Harrison H Zhou. Supplement to” minimax estimation in sparse canonical correlation analysis”. *The Annals of Statistics*, 2015b.
- Chao Gao, Zongming Ma, and Harrison H Zhou. Sparse cca: Adaptive estimation and computational barriers. *The Annals of Statistics*, 45(5):2074–2101, 2017.
- Rong Ge, Chi Jin, and Yi Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1233–1242. JMLR. org, 2017.

- Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHU press, 2012.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. CRC press, 2015.
- Harold Hotelling. Relations between two sets of variates. In *Breakthroughs in statistics*, pages 162–190. Springer, 1992.
- Iain M Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of statistics*, pages 295–327, 2001.
- I.M. Johnstone and A.Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.
- Jon R Kettenring. Canonical analysis of several sets of variables. *Biometrika*, 58(3):433–451, 1971.
- Zongming Ma. Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*, 41(2):772–801, 2013.
- Michael L. Overton and Robert S. Womersley. On the sum of the largest eigenvalues of a symmetric matrix. *SIAM J. Matrix Anal. Appl.*, 1991.
- Cencheng Shen, Ming Sun, Minh Tang, and Carey E Priebe. Generalized canonical correlation analysis for classification. *Journal of Multivariate Analysis*, 130:310–322, 2014.
- Ji-guang Sun. The perturbation bounds for eigenspaces of a definite matrix-pair. *Numerische Mathematik*, 41(3):321–343, 1983.
- Kean Ming Tan, Zhaoran Wang, Han Liu, and Tong Zhang. Sparse generalized eigenvalue problem: Optimal statistical rates via truncated rayleigh flow. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(5):1057–1086, 2018.
- Jos MF Ten Berge. Orthogonal procrustes rotation for two or more matrices. *Psychometrika*, 42(2):267–276, 1977.
- Arthur Tenenhaus and Michel Tenenhaus. Regularized generalized canonical correlation analysis. *Psychometrika*, 76(2):257, 2011.
- Yongge Tian. Rank equalities for block matrices and their moore-penrose inverses. *Houston J. Math*, 30(4):483–510, 2004.
- Stephen Tu, Ross Boczar, Max Simchowitz, Mahdi Soltanolkotabi, and Benjamin Recht. Low-rank solutions of linear matrix equations via procrustes flow. *arXiv preprint arXiv:1507.03566*, 2015.
- Vincent Q Vu, Juhee Cho, Jing Lei, and Karl Rohe. Fantope projection and selection: A near-optimal convex relaxation of sparse pca. In *Advances in neural information processing systems*, pages 2670–2678, 2013.

- V.Q. Vu and J. Lei. Minimax sparse principal subspace estimation in high dimensions. *arXiv preprint arXiv:1211.0373*, 2012.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Zhaoran Wang, Huanran Lu, and Han Liu. Tighten after relax: Minimax-optimal sparse pca in polynomial time. In *Advances in neural information processing systems*, pages 3383–3391, 2014.
- Hermann Weyl. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, 1912.
- D.M. Witten, R. Tibshirani, and T. Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10:515–534, 2009.
- Bin Yu, Fano Assouad, and Lucien Le Cam. Festschrift for lucien le cam, 1997.
- Xiao-Tong Yuan and Tong Zhang. Truncated power method for sparse eigenvalue problems. *Journal of Machine Learning Research*, 14(Apr):899–925, 2013.
- H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15:265–286, 2006.