



Minimax rate-optimal estimation of high-dimensional covariance matrices with incomplete data



T. Tony Cai^a, Anru Zhang^{b,*}

^a Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA, United States

^b University of Wisconsin-Madison, Madison, WI, United states

ARTICLE INFO

Article history:

Received 3 August 2015

Available online 19 May 2016

AMS subject classifications:

primary 62H12

62J10

secondary 62C20

Keywords:

Adaptive thresholding

Bandable covariance matrix

Generalized sample covariance matrix

Missing data

Optimal rate of convergence

Sparse covariance matrix

Thresholding

ABSTRACT

Missing data occur frequently in a wide range of applications. In this paper, we consider estimation of high-dimensional covariance matrices in the presence of missing observations under a general missing completely at random model in the sense that the missingness is not dependent on the values of the data. Based on incomplete data, estimators for bandable and sparse covariance matrices are proposed and their theoretical and numerical properties are investigated.

Minimax rates of convergence are established under the spectral norm loss and the proposed estimators are shown to be rate-optimal under mild regularity conditions. Simulation studies demonstrate that the estimators perform well numerically. The methods are also illustrated through an application to data from four ovarian cancer studies. The key technical tools developed in this paper are of independent interest and potentially useful for a range of related problems in high-dimensional statistical inference with missing data.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

The problem of missing data arises frequently in a wide range of fields, including biomedical studies, social science, engineering, economics, and computer science. Statistical inference in the presence of missing observations has been well studied in classical statistics. See, e.g., Ibrahim and Molenberghs [18] for a review of missing data methods in longitudinal studies and Schafer [26] for literature on handling multivariate data with missing observations. See Little and Rubin [20] and the references therein for a comprehensive treatment of missing data problems.

Missing data also occurs in contemporary high-dimensional inference problems, whose dimension p can be comparable to or even much larger than the sample size n . For example, in large-scale genome-wide association studies (GWAS), it is common for many subjects to have missing values on some genetic markers due to various reasons, including insufficient resolution, image corruption, and experimental error during the laboratory process. Also, different studies may have different volumes of genomic data available by design. For instance, the four genomic ovarian cancer studies discussed in Section 4 have throughput measurements of mRNA gene expression levels, but only one of these also has microRNA measurements (Cancer Genome Atlas Research Network [11], Bonome et al. [4], Tothill et al. [27] and Dressman et al. [15]). Discarding samples with any missingness is highly inefficient and could induce bias due to non-random missingness. It is of significant interest to integrate multiple high-throughput studies of the same disease, not only to boost statistical power

* Corresponding author.

E-mail addresses: tc@wharton.upenn.edu (T.T. Cai), anruzhang@stat.wisc.edu (A. Zhang).

but also to improve the biological interpretability. However, considerable challenges arise when integrating such studies due to missing data.

Although there have been significant recent efforts to develop methodologies and theories for high dimensional data analysis, there is a paucity of methods with theoretical guarantees for statistical inference with missing data in the high-dimensional setting. Under the assumption that the components are missing uniformly and completely at random (MUCR), Loh and Wainwright [21] proposed a non-convex optimization approach to high-dimensional linear regression, Lounici [23] introduced a method for estimating a low-rank covariance matrix and Lounici [22] considered sparse principal component analysis. In these papers, theoretical properties of the procedures were analyzed. These methods and theoretical results critically depend on the MUCR assumption.

Covariance structures play a fundamental role in high-dimensional statistics. It is of direct interest in a wide range of applications including genomic data analysis, particularly for hypothesis generation. Knowledge of the covariance structure is critical to many statistical methods, including discriminant analysis, principal component analysis, clustering analysis, and regression analysis. In the high-dimensional setting with complete data, inference on the covariance structure has been actively studied in recent years. See Cai, Ren and Zhou [7] for a survey of recent results on minimax and adaptive estimation of high-dimensional covariance and precision matrices under various structural assumptions. Estimation of high-dimensional covariance matrices in the presence of missing data also has wide applications in biomedical studies, particularly in integrative genomic analysis which holds great potential in providing a global view of genome function (see Hawkins et al. [17]).

In this paper, we consider estimation of high-dimensional covariance matrices in the presence of missing observations under a general missing completely at random (MCR) model in the sense that the missingness is not dependent on the values of the data. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be n independent copies of a p dimensional random vector \mathbf{X} with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Instead of observing the complete sample $\{\mathbf{X}_1, \dots, \mathbf{X}_n\}$, one observes the sample with missing values, where the observed coordinates of \mathbf{X}_k are indicated by a vector $\mathbf{S}_k \in \{0, 1\}^p$, $k = 1, \dots, n$. That is,

$$X_{jk} \text{ is observed if } S_{jk} = 1 \text{ and } X_{jk} \text{ is missing if } S_{jk} = 0. \quad (1)$$

Here X_{jk} and S_{jk} are respectively the j th coordinate of the vectors \mathbf{X}_k and \mathbf{S}_k . We denote the incomplete sample with missing values by $\mathbf{X}^* = \{\mathbf{X}_1^*, \dots, \mathbf{X}_n^*\}$. The major goal of the present paper is to estimate $\boldsymbol{\Sigma}$, the covariance matrix of \mathbf{X} , with theoretical guarantees based on the incomplete data \mathbf{X}^* in the high-dimensional setting where p can be much larger than n .

This paper focuses on estimation of high-dimensional bandable covariance matrices and sparse covariance matrices in the presence of missing data. These two classes of covariance matrices arise frequently in many applications, including genomics, econometrics, signal processing, temporal and spatial data analyses, and chemometrics. Estimation of these high-dimensional structured covariance matrices have been well studied in the setting of complete data in a number of recent papers, e.g., Bickel and Levina [3,2], Karoui [16], Rothman et al. [24], Cai and Zhou [10], Cai and Liu [5], Cai et al. [6,9] and Cai and Yuan [8]. Given an incomplete sample \mathbf{X}^* with missing values, we introduced a “generalized” sample covariance matrix, which can be viewed as an analog of the usual sample covariance matrix in the case of complete data. For estimation of bandable covariance matrices, where the entries of the matrix decay as they move away from the diagonal, a blockwise tridiagonal estimator is introduced and is shown to be rate-optimal. We then consider estimation of sparse covariance matrices. An adaptive thresholding estimator based on the generalized sample covariance matrix is proposed. The estimator is shown to achieve the optimal rate of convergence over a large class of approximately sparse covariance matrices under mild conditions.

The technical analysis for the case of missing data is much more challenging than that for the complete data, although some of the basic ideas are similar. To facilitate the theoretical analysis of the proposed estimators, we establish two key technical results, first, a large deviation result for a sub-matrix of the generalized sample covariance matrix and second, a large deviation bound for the self-normalized entries of the generalized sample covariance matrix. These technical tools are not only important for the present paper, but also useful for other related problems in high-dimensional statistical inference with missing data.

A simulation study is carried out to examine the numerical performance of the proposed estimation procedures. The results show that the proposed estimators perform well numerically. Even in the MUCR setting, our proposed procedures for estimating bandable, sparse covariance matrices, which do not rely on the information of the missingness mechanism, outperform the ones specifically designed for MUCR. The advantages are more significant under the setting of missing completely at random but not uniformly. We also illustrate our procedure with an application to data from four ovarian cancer studies that have different volumes of genomic data by design. The proposed estimators enable us to estimate the covariance matrix by integrating the data from all four studies and lead to a more accurate estimator. Such high-dimensional covariance matrix estimation with missing data is also useful for other types of data integration. See further discussions in Section 4.4.

The rest of the paper is organized as follows. Section 2 considers estimation of bandable covariance matrices with incomplete data. The minimax rate of convergence is established for the spectral norm loss under regularity conditions. Section 3 focuses on estimation of high-dimensional sparse covariance matrices and introduces an adaptive thresholding estimator in the presence of missing observations. Asymptotic properties of the estimator under the spectral norm loss is also studied. Numerical performance of the proposed methods is investigated in Section 4 through both simulation studies

and an analysis of an ovarian cancer dataset. Section 5 discusses a few related problems. Finally the proofs of the main results are given in Section 6 and the Supplement (see Appendix A).

2. Estimation of bandable covariance matrices

In this section, we consider estimation of bandable covariance matrices with incomplete data. Bandable covariance matrices, whose entries decay as they move away from the diagonal, arise frequently in temporal and spatial data analysis. See, e.g., Bickel and Levina [3] and Cai et al. [7] and the references therein. The procedure relies on a “generalized” sample covariance matrix. We begin with basic notation and definitions that will be used throughout the rest of the paper.

2.1. Notation and definitions

Matrices and vectors are denoted by boldface letters. For a vector $\beta \in \mathbb{R}^p$, we denote the Euclidean q -norm by $\|\beta\|_q$, i.e., $\|\beta\|_q = \sqrt[q]{\sum_{i=1}^p |\beta_i|^q}$. Let $\mathbf{A} = \mathbf{UDV}^\top = \sum_i \lambda_i(\mathbf{A}) \mathbf{u}_i \mathbf{v}_i^\top$ be the singular value decomposition of a matrix $\mathbf{A} \in \mathbb{R}^{p_1 \times p_2}$, where $\mathbf{D} = \text{diag}\{\lambda_1(\mathbf{A}), \dots\}$ with $\lambda_1(\mathbf{A}) \geq \dots \geq 0$ being the singular values. For $1 \leq q \leq \infty$, the Schatten- q norm $\|\mathbf{A}\|_q$ is defined by $\|\mathbf{A}\|_q = \{\sum_i \lambda_i^q(\mathbf{A})\}^{1/q}$. In particular, $\|\mathbf{A}\|_2 = \sqrt{\sum_i \lambda_i^2(\mathbf{A})}$ is the Frobenius norm of \mathbf{A} and will be denoted as $\|\mathbf{A}\|_F$; $\|\mathbf{A}\|_\infty = \lambda_1(\mathbf{A})$ is the spectral norm of \mathbf{A} and will be simply denoted as $\|\mathbf{A}\|$. For $1 \leq q \leq \infty$ and $A \in \mathbb{R}^{p_1 \times p_2}$, we denote the operator ℓ_q norm of \mathbf{A} by $\|\mathbf{A}\|_{\ell_q}$ which is defined as $\|\mathbf{A}\|_{\ell_q} = \max_{\mathbf{x} \in \mathbb{R}^{p_2}} \|\mathbf{Ax}\|_q / \|\mathbf{x}\|_q$. The following are well known facts about the various norms of a matrix $A = (a_{ij})$,

$$\|\mathbf{A}\|_{\ell_1} = \max_j \sum_{i=1}^{p_1} |a_{ij}|, \quad \|\mathbf{A}\|_{\ell_2} = \|\mathbf{A}\| = \lambda_1(\mathbf{A}), \quad \|\mathbf{A}\|_{\ell_\infty} = \max_i \sum_{j=1}^{p_2} |a_{ij}|, \tag{2}$$

and, if \mathbf{A} is symmetric, $\|\mathbf{A}\|_{\ell_1} = \|\mathbf{A}\|_{\ell_\infty} \geq \|\mathbf{A}\|_{\ell_2}$. When R_1, R_2 are two subsets of $\{1, \dots, p_1\}, \{1, \dots, p_2\}$ respectively, we note $\mathbf{A}_{R_1 \times R_2} = (a_{ij})_{i \in R_1, j \in R_2}$ as the sub-matrix of \mathbf{A} with indices R_1 and R_2 . In addition, we simply write $\mathbf{A}_{R_1 \times R_1}$ as \mathbf{A}_{R_1} .

We denote by $\mathbf{X}_1, \dots, \mathbf{X}_n$ a complete random sample (without missing observations) from a p -dimensional distribution with mean μ and covariance matrix Σ . The sample mean and sample covariance matrix are defined as

$$\bar{\mathbf{X}} = \frac{1}{n} \sum_{k=1}^n \mathbf{X}_k, \quad \hat{\Sigma} = \frac{1}{n} \sum_{k=1}^n (\mathbf{X}_k - \bar{\mathbf{X}})(\mathbf{X}_k - \bar{\mathbf{X}})^\top. \tag{3}$$

Now we introduce the notation related to the incomplete data with missing observations. Generally, we use the superscript “*” to denote objects related to missing values. Let $\mathbf{S}_1, \dots, \mathbf{S}_n$ be the indicator vectors for the observed values (see (1)) and let $\mathbf{X}^* = \{\mathbf{X}_1^*, \dots, \mathbf{X}_n^*\}$ be the observed incomplete data where the observed entries are indexed by the vectors $\mathbf{S}_1, \dots, \mathbf{S}_n \in \{0, 1\}^p$. In addition, we define

$$n_{ij}^* = \sum_{k=1}^n S_{ik} S_{jk}, \quad 1 \leq i, j \leq p. \tag{4}$$

Here n_{ij}^* is the number of vectors \mathbf{X}_k^* in which the i th and j th entries are both observed. For convenience, we also denote

$$n_i^* = n_{ii}^*, \quad n_{\min}^* = \min_{i,j} n_{ij}^*. \tag{5}$$

Given a sample $\mathbf{X}^* = \{\mathbf{X}_1^*, \dots, \mathbf{X}_n^*\}$ with missing values, the sample mean and sample covariance matrix can no longer be calculated in the usual way. Instead, we propose the “generalized sample mean” $\bar{\mathbf{X}}^*$ defined by

$$\bar{\mathbf{X}}^* = (\bar{X}_i^*)_{1 \leq i \leq p} \quad \text{with } \bar{X}_i^* = \frac{1}{n_i^*} \sum_{k=1}^n X_{ik} S_{ik}, \quad 1 \leq i \leq p, \tag{6}$$

where X_{ik} is the i th entry of \mathbf{X}_k , and the “generalized sample covariance matrix” $\hat{\Sigma}^*$ defined by

$$\hat{\Sigma}^* = (\hat{\sigma}_{ij}^*)_{1 \leq i, j \leq p} \quad \text{with } \hat{\sigma}_{ij}^* = \frac{1}{n_{ij}^*} \sum_{k=1}^n (X_{ik} - \bar{X}_i^*)(X_{jk} - \bar{X}_j^*) S_{ik} S_{jk}. \tag{7}$$

As will be seen later, the generalized sample mean $\bar{\mathbf{X}}^*$ and the generalized sample covariance matrix $\hat{\Sigma}^*$ play similar roles as those of the conventional sample mean and sample covariance matrix in inference problems, but the technical analysis can be much more involved. Some distinctions between the generalized sample covariance matrix $\hat{\Sigma}^*$ and the usual sample covariance matrix $\hat{\Sigma}$ are that $\hat{\Sigma}^*$ is in general not non-negative definite, and each entry $\hat{\sigma}_{ij}^*$ is the average of a varying number (n_{ij}^*) of samples, which create additional difficulties in the technical analysis.

Regarding the mechanism of missingness, the assumption we use for the theoretical analysis is *missing completely at random*. This is a more general setting than the one considered previously by Loh and Wainwright [21] and Lounici [22].

Assumption 2.1 (Missing Completely at Random (MCR)). $\mathbf{S} = \{\mathbf{S}_1, \dots, \mathbf{S}_n\}$ is not dependent on the values of \mathbf{X} . Here \mathbf{S} can be either deterministic or random, but independent of \mathbf{X} .

We adopt Assumption 1 in Chen et al. [13] and assume that the random vector \mathbf{X} is sub-Gaussian satisfying the following assumption.

Assumption 2.2 (Sub-Gaussian Assumption). $\mathbf{X} = \{\mathbf{X}_1, \dots, \mathbf{X}_n\}$. Here the columns \mathbf{X}_k are i.i.d. and can be expressed as

$$\mathbf{X}_k = \mathbf{\Gamma}\mathbf{Z}_k + \boldsymbol{\mu}, \quad k = 1, \dots, n, \tag{8}$$

where $\boldsymbol{\mu}$ is a fixed p -dimensional mean vector, $\mathbf{\Gamma} \in \mathbb{R}^{p \times q}$ is a fixed matrix with $q \geq p$ so that $\mathbf{\Gamma}\mathbf{\Gamma}^\top = \boldsymbol{\Sigma}$, $\mathbf{Z}_k = (Z_{1k}, \dots, Z_{mk})^\top$ is an m -dimensional random vector with the components mean 0, variance 1, and i.i.d. sub-Gaussian, with the exception of i.i.d. Rademacher. More specifically, each Z_{ik} satisfies that $\mathbb{E}Z_{ik} = 0$, $\text{var}(Z_{ik}) = 1$, $0 < \text{var}(Z_{ik}^2) < \infty$, and there exists $\tau > 0$ such that $\mathbb{E}e^{tZ_{ik}} \leq \exp(\tau t^2/2)$ for all $t > 0$.

Note that the exclusion of the Rademacher distribution in Assumption 2.2 is only required for estimation of sparse covariance matrices. See Remark 3.3 for further discussions.

2.2. Rate-optimal blockwise tridiagonal estimator

We follow Bickel [3] and Cai et al. [9] and consider estimating the covariance matrix $\boldsymbol{\Sigma}$ over the parameter space $\mathcal{U}_\alpha = \mathcal{U}_\alpha(M_0, M)$ where

$$\mathcal{U}_\alpha(M_0, M) = \left\{ \boldsymbol{\Sigma} : \max_j \sum_i \{|\sigma_{ij}| : |i - j| > k\} \leq Mk^{-\alpha} \text{ for all } k, \|\boldsymbol{\Sigma}\| \leq M_0 \right\}. \tag{9}$$

Suppose we have n i.i.d. samples with missing values $\mathbf{X}_1^*, \dots, \mathbf{X}_n^*$ with covariance matrix $\boldsymbol{\Sigma} \in \mathcal{U}_\alpha(M_0, M)$. We propose a blockwise tridiagonal estimator $\hat{\boldsymbol{\Sigma}}^{\text{bt}}$ to estimate $\boldsymbol{\Sigma}$. We begin by dividing the generalized sample covariance matrix $\hat{\boldsymbol{\Sigma}}^*$ given by (7) into blocks of size $k \times k$ for some k . More specifically, pick an integer k and let $N = \lceil p/k \rceil$. Set $I_j = \{(j-1)k+1, \dots, jk\}$ for $1 \leq j \leq N-1$, and $I_N = \{(N-1)k+1, \dots, p\}$. For $1 \leq j, j' \leq N$ and $\mathbf{A} = (a_{i_1, i_2})_{p \times p}$, define

$$\mathbf{A}_{I_j \times I_{j'}} = (a_{i_1, i_2})_{i_1 \in I_j, i_2 \in I_{j'}}$$

and define the blockwise tridiagonal estimator $\hat{\boldsymbol{\Sigma}}^{\text{bt}}$ by

$$\hat{\boldsymbol{\Sigma}}_{I_j \times I_{j'}}^{\text{bt}} = \begin{cases} \hat{\boldsymbol{\Sigma}}_{I_j \times I_{j'}}^* & \text{if } |j - j'| \leq 1; \\ \mathbf{0} & \text{otherwise.} \end{cases} \tag{10}$$

That is, $\hat{\boldsymbol{\Sigma}}_{I_j \times I_{j'}}^{\text{bt}}$ is estimated by its sample counterpart if and only if j and j' differ by at most 1. The weight matrix of the blockwise tridiagonal estimator $\hat{\boldsymbol{\Sigma}}^{\text{bt}}$ is illustrated in Fig. 1.

Theorem 2.1. Suppose Assumptions 2.1 and 2.2 hold. Then, conditioning on \mathbf{S} , the blockwise tridiagonal $\hat{\boldsymbol{\Sigma}}^{\text{bt}}$ with $k = (n_{\min}^*)^{1/(2\alpha+1)}$ satisfies

$$\sup_{\boldsymbol{\Sigma} \in \mathcal{U}_\alpha(M, M_0)} \mathbb{E} \|\hat{\boldsymbol{\Sigma}}^{\text{bt}} - \boldsymbol{\Sigma}\|^2 \leq C(n_{\min}^*)^{-2\alpha/(2\alpha+1)} + C \frac{\ln p}{n_{\min}^*}, \tag{11}$$

where C is a constant depending only on M, M_0 , and τ from Assumption 2.2.

The optimal choice of block size k depends on the unknown ‘‘smoothness parameter’’ α . In practice, k can be chosen by cross-validation. See Section 4.1 for further discussions. Moreover, the convergence rate in (11) is optimal as we also have the following lower bound result.

Proposition 2.1. For any $n_0 \geq 1$ such that $p \leq \exp(\gamma n_0)$ for some constant $\gamma > 0$, conditioning on \mathbf{S} we have

$$\inf_{\hat{\boldsymbol{\Sigma}}} \sup_{\substack{\boldsymbol{\Sigma} \in \mathcal{U}_\alpha(M, M_0) \\ S: n_{\min}^* \geq n_0}} \mathbb{E} \left(\|\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|^2 \right) \geq C(n_0)^{-2\alpha/(2\alpha+1)} + C \frac{\ln p}{n_0}.$$

Remark 2.1 (Tapering and Banding Estimators). It should be noted that the same rate of convergence can also be attained by tapering and banding estimators with suitable choices of tapering and banding parameters. Specifically, let $\hat{\boldsymbol{\Sigma}}^{\text{tp}}$ and $\hat{\boldsymbol{\Sigma}}^{\text{bd}}$ be respectively the tapering and banded estimators proposed in Cai et al. [9] and Bickel and Levina [3] with

$$\hat{\boldsymbol{\Sigma}}^{\text{tp}} = \hat{\boldsymbol{\Sigma}}_k^{\text{tp}} = (w_{ij}^{\text{tp}} \hat{\sigma}_{ij}^*)_{1 \leq i, j \leq p} \quad \text{and} \quad \hat{\boldsymbol{\Sigma}}^{\text{bd}} = \hat{\boldsymbol{\Sigma}}_k^{\text{bd}} = (w_{ij}^{\text{bd}} \hat{\sigma}_{ij}^*)_{1 \leq i, j \leq p}, \tag{12}$$

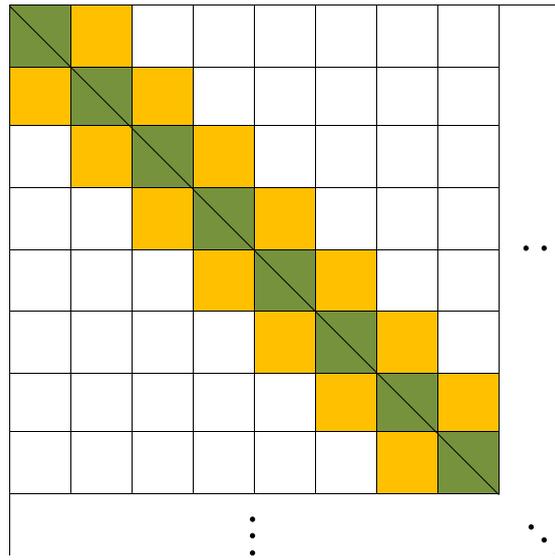


Fig. 1. Weight matrix for the blockwise tridiagonal estimator.

where w_{ij}^{tp} and w_{ij}^{bd} are the weights defined as

$$w_{ij}^{tp} = \begin{cases} 1, & \text{when } |i - j| \leq k/2, \\ 2 - \frac{|i - j|}{k_h}, & \text{when } k/2 < |i - j| < k \\ 0, & \text{otherwise} \end{cases} \quad \text{and} \quad w_{ij}^{bd} = \begin{cases} 1, & \text{when } |i - j| \leq k, \\ 0, & \text{otherwise.} \end{cases} \tag{13}$$

Then the estimators $\hat{\Sigma}^{tp}$ and $\hat{\Sigma}^{bd}$ with $k = (n_{\min}^*)^{1/(2\alpha+1)}$ attains the rate given in (11).

The proof of Theorem 2.1 shares some basic ideas with that for the complete data case (see, e.g. Theorem 2 in Cai et al. [9]). However, it relies on a new key technical tool which is a large deviation result for a sub-matrix of the generalized sample covariance matrix under the spectral norm. This random matrix result for the case of missing data, stated in the following lemma, can be potentially useful for other, related high-dimensional missing data problems. The proof of Lemma 2.1, given in Section 6, is more involved than the complete data case, as in the generalized sample covariance matrix each entry, $\hat{\sigma}_{ij}^*$, is the average of a varying number of samples.

Lemma 2.1. Suppose Assumptions 2.1 and 2.2 hold. Let $\hat{\Sigma}^*$ be the generalized sample covariance matrix defined in (7) and let A and B be two subsets of $\{1, \dots, p\}$. Then, conditioning on \mathbf{S} , the submatrix $\hat{\Sigma}_{A \times B}^*$ satisfies

$$\Pr \left(\|\hat{\Sigma}_{A \times B}^* - \Sigma_{A \times B}\| \leq x \right) \geq 1 - C \cdot (49)^{|A \cup B|} \exp \left\{ -cn_{\min}^* \min \left(\frac{x^2}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right) \right\} \tag{14}$$

for all $x > 0$. Here $C > 0$ and $c > 0$ are two absolute constants.

3. Estimation of sparse covariance matrices

In this section, we consider estimation of high-dimensional sparse covariance matrices in the presence of missing data. We introduce an adaptive thresholding estimator based on incomplete data and investigate its asymptotic properties.

3.1. Adaptive thresholding procedure

Sparse covariance matrices arise naturally in a range of applications including genomics. Estimation of sparse covariance matrices has been considered in several recent papers in the setting of complete data (see, e.g., Bickel and Levina [2], El Karoui [16], Rothman et al. [24], Cai and Zhou [10] and Cai and Liu [5]). Estimation of a sparse covariance matrix is intrinsically a heteroscedastic problem in the sense that the variances of the entries of the sample covariance matrix can vary over a wide range. To treat the heteroscedasticity of the sample covariances, Cai and Liu [5] introduced an adaptive thresholding procedure which adapts to the variability of the individual entries of the sample covariance matrix and outperforms the universal thresholding method. The estimator is shown to be simultaneously rate optimal over collections of sparse covariance matrices.

In the present setting of missing data, the usual sample covariance matrix is not available. Instead we apply the idea of adaptive thresholding to the generalized sample covariance matrix $\hat{\Sigma}^*$. The procedure can be described as follows. Note that $\hat{\Sigma}^*$ defined in (7) is a nearly unbiased estimate of Σ , we may write it element-wise as

$$\hat{\sigma}_{ij}^* \approx \sigma_{ij} + \sqrt{\frac{\theta_{ij}}{n_{ij}^*}} z_{ij}, \quad 1 \leq i, j \leq p,$$

where z_i is approximately normal with mean 0 and variance 1, and θ_{ij} describes the uncertainty of estimator σ_{ij}^* to σ_{ij} such that

$$\theta_{ij} = \text{var} \left\{ (X_i - \mu_i)(X_j - \mu_j) - \sigma_{ij} \right\}.$$

We can estimate θ_{ij} by

$$\hat{\theta}_{ij}^* = \frac{1}{n_{ij}^*} \sum_{k=1}^n \left\{ (X_{ik} - \bar{X}_i^*)(X_{jk} - \bar{X}_j^*) - \hat{\sigma}_{ij}^* \right\}^2 S_{ik} S_{jk}. \quad (15)$$

Lemma 3.1 given at the end of this section shows that $\hat{\theta}_{ij}^*$ is a good estimate of θ_{ij} .

Since the covariance matrix Σ is assumed to be sparse, it is natural to estimate Σ by individually thresholding $\hat{\theta}_{ij}^*$ according to its own variability as measured by $\hat{\theta}_{ij}^*$. Define the thresholding level λ_{ij} by

$$\lambda_{ij} = \delta \sqrt{\frac{\hat{\theta}_{ij}^* \ln p}{n_{ij}^*}}, \quad 1 \leq i, j \leq p,$$

where δ is a thresholding constant which can be taken as 2.

Let T_λ be a thresholding function satisfying the following conditions,

- (1) $|T_\lambda(z)| \leq c_T |y|$ for all z, y such that $|z - y| \leq \lambda$;
- (2) $T_\lambda(z) = 0$ for $|z| \leq \lambda$;
- (3) $|T_\lambda(z) - z| \leq \lambda$, for all $z \in \mathbb{R}$.

These conditions are met by many well-used thresholding functions, including the soft thresholding rule $T_\lambda(z) = \text{sgn}(z)(z - \lambda)_+$, where $\text{sgn}(z)$ is the sign function such that $\text{sgn}(z) = 1$ if $z > 0$, $\text{sgn}(z) = 0$ if $z = 0$, and $\text{sgn}(z) = -1$ if $z < 0$, and the adaptive lasso rule $T_\lambda(z) = z(1 - |\lambda/z|^\eta)_+$ with $\eta \geq 1$ (see Rothman et al. [24]). The hard thresholding function does not satisfy Condition (1), but our analysis also applies to hard thresholding under similar conditions.

The covariance matrix Σ is estimated by $\hat{\Sigma}^{\text{at}} = (\hat{\sigma}_{ij}^{\text{at}})_{1 \leq i, j \leq p}$ where $\hat{\sigma}_{ij}^{\text{at}}$ is the thresholding estimator defined by

$$\hat{\sigma}_{ij}^{\text{at}} = T_{\lambda_{ij}}(\hat{\sigma}_{ij}^*). \quad (16)$$

Note that here each entry $\hat{\sigma}_{ij}^*$ is thresholded according to its own variability.

3.2. Asymptotic properties

We now investigate the properties of the thresholding estimator $\hat{\Sigma}^{\text{at}}$ over the following parameter space for sparse covariance matrices,

$$\mathcal{H}(c_{n,p}) = \left\{ \Sigma = (\sigma_{ij}) : \max_{1 \leq i \leq p} \sum_{j=1}^p \min \left\{ (\sigma_{ii} \sigma_{jj})^{1/2}, \frac{|\sigma_{ij}|}{\sqrt{(\ln p)/n}} \right\} \leq c_{n,p} \right\}. \quad (17)$$

The parameter space $\mathcal{H}(c_{n,p})$ contains a large collection of sparse covariance matrices and does not impose any constraint on the variances σ_{ii} , $i = 1, \dots, p$. The collection $\mathcal{H}(c_{n,p})$ contains some other commonly used classes of sparse covariance matrices in the literature, including an ℓ_q ball assumption $\max_i \sum_{j=1}^p |\sigma_{ij}|^q \leq s_{n,p}$ in Bickel and Levina [2], and a weak ℓ_q ball assumption $\max_{1 \leq j \leq p} \left\{ |\sigma_{j|k}|^q \right\} \leq s_{n,p}/k$ for each integer k in Cai and Zhou [10] where $|\sigma_{j|k}|$ is the k th largest entry in magnitude of the j th row $(\sigma_{ij})_{1 \leq i \leq p}$. See Cai et al. [7] for more discussions.

We have the following result on the performance of $\hat{\Sigma}^{\text{at}}$ over the parameter space $\mathcal{H}(c_{n,p})$.

Theorem 3.1. Suppose that $\delta \geq 2$, $\ln p = o((n_{\min}^*)^{1/3})$ and Assumptions 2.1 and 2.2 hold. Then, conditioning on \mathbf{S} , there exists a constant C not depending on p , n_{\min}^* or n such that for any $\Sigma \in \mathcal{H}(c_{n,p})$,

$$\Pr \left(\left\| \hat{\Sigma}^{\text{at}} - \Sigma \right\| \leq C c_{n,p} \sqrt{\frac{\ln p}{n_{\min}^*}} \right) \geq 1 - O \left\{ (\ln p)^{-1/2} p^{-\delta+2} \right\}. \tag{18}$$

Moreover, if we further assume that $p \geq (n_{\min}^*)^\xi$ and $\delta \geq 4 + 1/\xi$, we in addition have

$$\mathbb{E} \left(\left\| \hat{\Sigma}^{\text{at}} - \Sigma \right\|^2 \right) \leq C c_{n,p}^2 \frac{\ln p}{n_{\min}^*}. \tag{19}$$

Moreover, the lower bound result below shows that the rate in (19) is optimal.

Proposition 3.1. For any $n_0 \geq 1$ and $c_{n,p} > 0$ such that $c_{n,p} \leq M n_0^{1/2} (\ln p)^{-3/2}$ for some constant $M > 0$, conditioning on \mathbf{S} we have

$$\inf_{\hat{\Sigma}} \sup_{\substack{\Sigma \in \mathcal{H}(c_{n,p}) \\ S: n_{\min}^* \geq n_0}} \mathbb{E} \left(\left\| \hat{\Sigma} - \Sigma \right\|^2 \right) \geq C c_{n,p}^2 \frac{\ln p}{n_0}.$$

Remark 3.1 (ℓ_q Norm Loss). We focus in this paper on estimation under the spectral norm loss. The results given in Theorem 3.1 can be easily generalized to the general matrix ℓ_q norm for $1 \leq q \leq \infty$. The results given in Eqs. (18) and (19) remain valid when the spectral norm is replaced by the matrix ℓ_q norm for $1 \leq q \leq \infty$.

Remark 3.2 (Positive Definiteness). Under mild conditions on Σ , the estimator $\hat{\Sigma}^{\text{at}}$ is positive definite with high probability. However, $\hat{\Sigma}^{\text{at}}$ is not guaranteed to be positive definite for a given dataset. Whenever $\hat{\Sigma}^{\text{at}}$ is not positive semi-definite, a simple extra step can make the final estimator $\hat{\Sigma}_+^{\text{at}}$ positive definite and also rate-optimal.

Write the eigen-decomposition of $\hat{\Sigma}^{\text{at}}$ as $\hat{\Sigma}^{\text{at}} = \sum_{i=1}^p \hat{\lambda}_i \hat{\mathbf{v}}_i \hat{\mathbf{v}}_i^\top$, where $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_p$ are the eigenvalues and $\hat{\mathbf{v}}_i$ are the corresponding eigenvectors. Define the final estimator

$$\hat{\Sigma}_+^{\text{at}} = \hat{\Sigma}^{\text{at}} + \left(|\hat{\lambda}_p| + \frac{\ln p}{n_{\min}^*} \right) I \{ \hat{\lambda}_p < 0 \} \cdot \mathbf{I}_{p \times p},$$

where $\mathbf{I}_{p \times p}$ is the $p \times p$ identity matrix. Then $\hat{\Sigma}_+^{\text{at}}$ is a positive definite matrix with the same structure as that of $\hat{\Sigma}^{\text{at}}$. It is easy to show that $\hat{\Sigma}_+^{\text{at}}$ and $\hat{\Sigma}^{\text{at}}$ attains the same rate of convergence over $\mathcal{H}(c_{n,p})$. See Cai, Ren and Zhou [7] for further discussions.

Remark 3.3 (Exclusion of the Rademacher Distribution). To guarantee that $\hat{\theta}_{ij}^*$ is a good estimate of θ_{ij} , one important condition needed in the theoretical analysis is that $\theta_{ij} / \sqrt{\sigma_{ii}\sigma_{jj}}$ is bounded from below by a positive constant. However when the components of \mathbf{Z}_k in (8) are i.i.d. Rademacher, it is possible that $\theta_{ij} / \sqrt{\sigma_{ii}\sigma_{jj}} = 0$. For example, if Z_1 and Z_2 are i.i.d. Rademacher and $X_i = Z_1 + Z_2$ and $X_j = Z_1 - Z_2$, then $\text{var}(X_i X_j) = \text{var}(Z_1^2 - Z_2^2) = 0$, and this implies $\theta_{ij} / \sqrt{\sigma_{ii}\sigma_{jj}} = 0$.

A key technical tool in the analysis of the adaptive thresholding estimator is a large deviation result for the self-normalized entries of the generalized sample covariance matrix. The following lemma, proved in Section 6, plays a critical role in the proof of Theorem 3.1 and can be useful for other high-dimensional inference problems with missing data.

Lemma 3.1. Suppose $\ln p = o((n_{\min}^*)^{1/3})$ and Assumptions 2.1 and 2.2 hold. For any constants $\delta \geq 2$, $\varepsilon > 0$, $M > 0$, conditioning on \mathbf{S} , we have

$$\Pr \left(\frac{|\hat{\sigma}_{ij}^* - \sigma_{ij}|}{(\hat{\theta}_{ij}^*)^{1/2}} \geq \delta \sqrt{\frac{\ln p}{n_{ij}^*}}, \forall 1 \leq i, j \leq p \right) = O \left\{ (\ln p)^{-1/2} p^{-\delta+2} \right\}, \tag{20}$$

$$\Pr \left(\max_{ij} \frac{|\hat{\theta}_{ij}^* - \theta_{ij}|}{\sigma_{ii}\sigma_{jj}} \geq \varepsilon \right) = O(p^{-M}). \tag{21}$$

In addition to optimal estimation of a sparse covariance matrix Σ under the spectral norm loss, it is also of significant interest to recover the support of Σ , i.e., the locations of the nonzero entries of Σ . The problem has been studied in the case of complete data in, e.g., Cai and Liu [5] and Rothman et al. [24]. With incomplete data, the support can be similarly recovered through adaptive thresholding. Specifically, define the support of $\Sigma = (\sigma_{ij})_{1 \leq i, j \leq p}$ by $\text{supp}(\Sigma) = \{(i, j) : \sigma_{ij} \neq 0\}$. Under the

condition that the non-zero entries of Σ are sufficiently bounded away from zero, the adaptive thresholding estimator $\hat{\Sigma}^{\text{at}}$ recovers the support $\text{supp}(\Sigma)$ consistently. It is noteworthy that in the support recovery analysis, the sparsity assumption is not directly needed.

Theorem 3.2 (Support Recovery). Suppose $\ln p = o((n_{\min}^*)^{1/3})$ and Assumptions 2.1 and 2.2 hold. Let γ be any positive constant. Suppose Σ satisfies

$$|\sigma_{ij}| > (4 + \gamma) \sqrt{\frac{\theta_{ij} \ln p}{n_{ij}^*}}, \quad \text{for all } (i, j) \in \text{supp}(\Sigma). \quad (22)$$

Let $\hat{\Sigma}^{\text{at}}$ be the adaptive thresholding estimator with $\delta = 2$, then, conditioning on \mathbf{S} , we have

$$\Pr \left\{ \text{supp}(\hat{\Sigma}^{\text{at}}) = \text{supp}(\Sigma) \right\} \rightarrow 1 \quad \text{as } n, p \rightarrow \infty. \quad (23)$$

4. Numerical results

We investigate in this section the numerical performance of the proposed estimators through simulations. The proposed adaptive thresholding procedure is also illustrated with an estimation of the covariance matrix based on data from four ovarian cancer studies.

The estimators $\hat{\Sigma}^{\text{bt}}$ and $\hat{\Sigma}^{\text{at}}$ introduced in the previous sections all require specification of the tuning parameters (k or δ). Cross-validation is a simple and practical data-driven method for the selection of these tuning parameters. Numerical results indicate that the proposed estimators with the tuning parameter selected by cross-validation perform well empirically. We begin by introducing the following K -fold cross-validation method for the empirical selection of the tuning parameters.

4.1. Cross-validation

For a pre-specified positive integer N , we construct a grid T of non-negative numbers. For bandable covariance matrix estimation, we set $T = \{1, \lceil p^{1/N} \rceil, \dots, \lceil p^{N/N} \rceil\}$, and for sparse covariance matrix estimation, we let $T = \{0, 1/N, \dots, 4N/N\}$.

Given n samples $\mathbf{X}^* \in \mathbb{R}^{p \times n}$ with missing values, for a given positive integer K , we randomly divide them into two groups of size $n_1 \approx n(K-1)/K$, $n_2 \approx n/K$ for H times. For $h = 1, \dots, H$, we denote by J_1^h and $J_2^h \subseteq \{1, \dots, n\}$ the index sets of the two groups for the h th split. The proposed estimator, $\hat{\Sigma}^{\text{bt}}$ for bandable covariance matrices, or $\hat{\Sigma}^{\text{at}}$ for sparse covariance matrices, is then applied to the first group of data $\mathbf{X}_{J_1^h}^*$ with each value of the tuning parameter $t \in T$ and denote the result by $\hat{\Sigma}_h^{\text{bt}}(t)$ or $\hat{\Sigma}_h^{\text{at}}(t)$ respectively. Denote the generalized sample covariance matrix of the second group of data $\mathbf{X}_{J_2^h}^*$ by $\hat{\Sigma}_h^*$ and set

$$\hat{R}(t) = \frac{1}{H} \sum_{h=1}^H \|\hat{\Sigma}_h(t) - \hat{\Sigma}_h^*\|_F^2, \quad (24)$$

where $\hat{\Sigma}_h(t)$ is either $\hat{\Sigma}^{\text{bt}}(t)$ for bandable covariance matrices, or $\hat{\Sigma}^{\text{at}}(t)$ for sparse covariance matrices.

The final tuning parameter is chosen to be

$$t_* = \underset{T}{\text{argmin}} \hat{R}(t)$$

and the final estimator $\hat{\Sigma}^{\text{bt}}$ (or $\hat{\Sigma}^{\text{at}}$) is calculated using this choice of the tuning parameter t_* . In the following numerical studies, we will use 5-fold cross-validation (i.e., $K = 5$) to select the tuning parameters.

Remark 4.1. The Frobenius norm used in (24) can be replaced by other losses such as the spectral norm. Our simulation results indicate that using the Frobenius norm in (24) works well, even when the true loss is the spectral norm loss.

4.2. Simulation studies

In the simulation studies, we consider the following two settings for the missingness. The first is MUCR where each entry X_{ik} is observed with probability $0 < \rho \leq 1$, and the second is missing not uniformly but completely at random (MCR) where the complete data matrix \mathbf{X} is divided into four equal-size blocks,

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{(11)} & \mathbf{X}_{(12)} \\ \mathbf{X}_{(21)} & \mathbf{X}_{(22)} \end{bmatrix}, \quad \mathbf{X}_{(11)}, \mathbf{X}_{(12)}, \mathbf{X}_{(21)}, \mathbf{X}_{(22)} \in \mathbb{R}^{\frac{p}{2} \times \frac{n}{2}},$$

and each entry of $\mathbf{X}_{(11)}$ and $\mathbf{X}_{(22)}$ is observed with probability $\rho^{(1)}$ and each entry of $\mathbf{X}_{(12)}$ and $\mathbf{X}_{(21)}$ is observed with probability $\rho^{(2)}$, for some $0 < \rho^{(1)}, \rho^{(2)} \leq 1$.

As mentioned in the introduction, high-dimensional inference for missing data has been studied in the case of MUCR and we would like to compare our estimators with the corresponding estimators based on a different sample covariance matrix designed for the MUCR case. Under the assumption that $\mathbf{E}\mathbf{X} = \mathbf{0}$ and each entry of \mathbf{X} is observed independently with probability ρ , Wainwright [21] and Lounici [23] introduced the following substitute of the usual sample covariance matrix

$$\hat{\Sigma}^\bullet = (\hat{\sigma}_{ij}^\bullet)_{1 \leq i, j \leq p} \quad \text{with} \quad \hat{\sigma}_{ij}^\bullet = \begin{cases} \frac{1}{n(1-\rho)^2} \sum_{k=1}^n X_{ik}^* X_{jk}^*, & i \neq j \\ \frac{1}{n(1-\rho)} \sum_{k=1}^n X_{ik}^* X_{jk}^*, & i = j \end{cases} \quad (25)$$

where the missing entries of \mathbf{X}^* are replaced by 0's. It is easy to show that $\hat{\Sigma}^\bullet$ is a consistent estimator of Σ under MUCR and can be used similarly as the sample covariance matrix in the complete data setting.

For more general settings where $\mathbf{E}\mathbf{X} \neq \mathbf{0}$ and the coordinates X_1, X_2, \dots, X_p are observed with different probabilities ρ_1, \dots, ρ_p , $\hat{\Sigma}^\bullet$ can be generalized as

$$\hat{\Sigma}^\bullet = (\hat{\sigma}_{ij}^\bullet)_{1 \leq i, j \leq p} \quad \text{with} \quad \hat{\sigma}_{ij}^\bullet = \begin{cases} \frac{1}{n(1-\hat{\rho}_i)(1-\hat{\rho}_j)} \sum_{k=1}^n X_{ik,c}^* X_{jk,c}^*, & i \neq j \\ \frac{1}{n(1-\hat{\rho}_i)} \sum_{k=1}^n X_{ik,c}^* X_{jk,c}^*, & i = j \end{cases} \quad (26)$$

where for $i = 1, \dots, p$ and $k = 1, \dots, n$, $\hat{\rho}_i = \frac{1}{n} \sum_{k=1}^n S_{ik}$ and $X_{ik,c}^* = X_{ik}^* - \bar{X}_i^*$.

Based on $\hat{\Sigma}^\bullet$, we can analogously define the corresponding blockwise tridiagonal estimator $\hat{\Sigma}^{\text{bt}\bullet}$ for bandable covariance matrices, and adaptive thresholding estimator $\hat{\Sigma}^{\text{at}\bullet}$ for sparse covariance matrices.

We first consider estimation of bandable covariance matrices and compare the proposed blockwise tridiagonal estimator $\hat{\Sigma}^{\text{bt}}$ with the corresponding estimator $\hat{\Sigma}^{\text{bt}\bullet}$. For both methods, the tuning parameter k is selected by 5-fold cross-validation with N varying from 20 to 50. The following bandable covariance matrices are considered:

1. (Linear decaying bandable model) $\Sigma = (\sigma_{ij})_{1 \leq i, j \leq p}$ with $\sigma_{ij} = \max\{0, 1 - |i - j|/5\}$.
2. (Squared decaying bandable model) $\Sigma = (\sigma_{ij})_{1 \leq i, j \leq p}$ with $\sigma_{ij} = (|i - j| + 1)^{-2}$.

For missingness, both MUCR and MCR are considered and (25) and (26) are used to calculate $\hat{\Sigma}^\bullet$ respectively. The proposed procedure $\hat{\Sigma}^{\text{bt}}$ is compared with the estimator $\hat{\Sigma}^{\text{bt}\bullet}$, which is based on $\hat{\Sigma}^\bullet$. The results for the spectral norm, ℓ_1 norm and Frobenius norm losses are reported in Table 1. It is easy to see from Table 1 that the proposed estimator $\hat{\Sigma}^{\text{bt}}$ generally outperforms $\hat{\Sigma}^{\text{bt}\bullet}$, especially in the fast decaying setting.

Now we consider estimation of sparse covariance matrices with missing values under the following two models.

1. (Permutation Bandable Model) $\Sigma = (\sigma_{ij})_{1 \leq i, j \leq p}$, where $\sigma_{ij} = \max(0, 1 - 0.2 \cdot |s(i) - s(j)|)$ and $s(i), i = 1, \dots, p$ is a random permutation of $\{1, \dots, p\}$.
2. (Randomly Sparse Model) $\Sigma = \mathbf{I}_p + (\mathbf{D} + \mathbf{D}^\top) / (\|\mathbf{D} + \mathbf{D}^\top\| + 0.01)$, where \mathbf{D} is randomly generated as

$$\mathbf{D} = (d_{ij})_{1 \leq i, j \leq p}, \quad d_{ij} = \begin{cases} 1 & \text{w.p. } 0.1 \\ 0 & \text{w.p. } 0.8 \quad \text{for } i \neq j; \quad d_{ii} = 0. \\ -1 & \text{w.p. } 0.1 \end{cases}$$

Similar to the sparse covariance matrix estimation, for missingness, we consider both MUCR and MCR. The results for the spectral norm, matrix ℓ_1 norm and Frobenius norm losses are summarized in Table 2. It can be seen from Table 2 that, even under the MUCR setting, the proposed estimator $\hat{\Sigma}^{\text{at}}$ based on the generalized sample covariance matrix is uniformly better than the one based on $\hat{\Sigma}^\bullet$. In the more general MCR setting, the difference in the performance between the two estimators is even more significant.

4.3. Comparison with complete samples

For covariance matrix estimation with missing data, an interesting question is: what is the “effective sample size”? That is, for samples with missing values, we would like to know the equivalent size of complete samples such that the accuracy for covariance matrix estimation is approximately the same. We now compare the performance of the proposed estimator based on the incomplete data with the corresponding estimator based on the complete data for various sample sizes. We fix the dimension $p = 100$. For the incomplete data, we consider $n = 1000$ and MUCR with $\rho = 0.5$. The covariance matrix Σ is chosen as

- Linear Decaying Bandable Model (in Bandable Covariance Matrix Estimation);
- Permutation Bandable Model (in Sparse Covariance Matrix Estimation).

Table 1

Comparison between $\hat{\Sigma}^{bt}$ and $\hat{\Sigma}^{bt*}$ in different settings of bandable covariance matrix estimation.

(p, n)	Spectral norm		ℓ_1 norm		Frobenius norm	
	$\hat{\Sigma}^{bt}$	$\hat{\Sigma}^{bt*}$	$\hat{\Sigma}^{bt}$	$\hat{\Sigma}^{bt*}$	$\hat{\Sigma}^{bt}$	$\hat{\Sigma}^{bt*}$
Linear Decay Bandable Model, MUCR $\rho = 0.5$						
(50, 50)	2.78(0.17)	2.88(0.18)	4.37(0.57)	4.57(0.76)	7.73(0.85)	7.85(0.80)
(50, 200)	1.44(0.06)	1.56(0.07)	2.52(0.17)	2.71(0.19)	3.91(0.18)	4.16(0.16)
(200, 100)	2.25(0.13)	2.44(0.16)	3.83(0.32)	4.22(0.46)	10.27(0.29)	10.89(0.29)
(200, 200)	1.67(0.07)	1.82(0.08)	2.81(0.19)	3.08(0.22)	7.19(0.19)	7.68(0.14)
(500, 200)	2.00(0.07)	2.18(0.10)	3.45(0.16)	3.74(0.27)	12.10(0.36)	12.87(0.42)
Squared Decay Bandable Model, MUCR $\rho = 0.5$						
(50, 50)	1.34(0.08)	1.40(0.11)	2.28(0.16)	2.37(0.21)	3.78(0.19)	3.91(0.18)
(50, 200)	0.82(0.01)	0.84(0.01)	1.47(0.03)	1.49(0.02)	2.24(0.02)	2.30(0.02)
(200, 100)	1.13(0.01)	1.17(0.02)	2.12(0.05)	2.18(0.07)	5.74(0.04)	5.91(0.05)
(200, 200)	0.92(0.00)	0.94(0.00)	1.66(0.02)	1.72(0.03)	4.49(0.02)	4.61(0.01)
(500, 200)	0.97(0.00)	0.98(0.00)	1.80(0.02)	1.86(0.02)	7.15(0.01)	7.35(0.01)
Linear Decay Bandable Model, MCR $\rho^{(1)} = 0.8, \rho^{(2)} = 0.2$						
(50, 50)	2.76(0.26)	3.46(1.43)	4.24(0.73)	5.87(2.91)	7.03(1.25)	8.47(1.29)
(50, 200)	1.51(0.11)	2.64(0.40)	2.52(0.30)	4.29(0.99)	3.62(0.30)	5.77(0.45)
(200, 100)	2.32(0.22)	3.93(0.67)	3.73(0.47)	6.21(1.11)	9.04(0.48)	13.47(0.84)
(200, 200)	1.67(0.10)	3.23(0.27)	2.71(0.26)	4.91(0.49)	6.32(0.11)	11.32(0.49)
(500, 200)	1.98(0.09)	3.78(0.20)	3.19(0.20)	5.70(0.42)	10.39(0.12)	18.48(0.49)
Squared Decay Bandable Model, MCR $\rho^{(1)} = 0.8, \rho^{(2)} = 0.2$						
(50, 50)	1.26(0.08)	1.49(0.13)	2.21(0.23)	2.60(0.28)	3.48(0.14)	4.18(0.23)
(50, 200)	0.82(0.01)	0.88(0.04)	1.47(0.05)	1.77(0.11)	2.18(0.04)	2.68(0.11)
(200, 100)	1.06(0.01)	1.30(0.04)	1.96(0.04)	2.44(0.07)	5.32(0.02)	6.51(0.06)
(200, 200)	0.90(0.00)	0.96(0.03)	1.60(0.02)	1.99(0.06)	4.27(0.02)	5.26(0.15)
(500, 200)	0.93(0.00)	1.03(0.01)	1.69(0.01)	2.11(0.03)	6.73(0.01)	8.25(0.04)

Table 2

Comparison between $\hat{\Sigma}^{at}$ and $\hat{\Sigma}^{at*}$ in different settings of sparse covariance matrix estimation.

(p, n)	Spectral norm		ℓ_1 norm		Frobenius norm	
	$\hat{\Sigma}^{at}$	$\hat{\Sigma}^{at*}$	$\hat{\Sigma}^{at}$	$\hat{\Sigma}^{at*}$	$\hat{\Sigma}^{at}$	$\hat{\Sigma}^{at*}$
Permutation Bandable Model, MUCR $\rho = 0.5$						
(50, 50)	4.26(0.24)	4.45(0.41)	5.58(0.58)	6.19(7.54)	11.34(0.79)	11.73(1.08)
(50, 200)	1.70(0.05)	1.74(0.06)	3.31(0.32)	3.42(0.38)	4.93(0.09)	5.07(0.16)
(200, 100)	3.48(0.07)	3.66(0.58)	5.80(0.39)	6.23(14.89)	18.34(0.81)	19.37(5.50)
(200, 200)	2.12(0.04)	2.20(0.03)	4.17(0.29)	4.44(0.32)	11.46(0.14)	11.94(0.13)
(500, 200)	2.28(0.03)	3.51(0.17)	4.17(0.15)	6.55(0.72)	16.85(0.10)	21.96(0.49)
Randomly Sparse Model, MUCR $\rho = 0.5$						
(50, 50)	1.76(0.07)	1.96(0.62)	3.69(0.24)	4.20(5.89)	5.75(0.51)	6.27(2.95)
(50, 200)	1.05(0.00)	1.06(0.00)	2.73(0.04)	2.74(0.05)	3.75(0.03)	3.77(0.04)
(200, 100)	1.40(0.01)	1.45(0.01)	4.88(0.08)	4.94(0.09)	8.34(0.07)	8.50(0.07)
(200, 200)	1.07(0.00)	1.09(0.01)	4.44(0.03)	4.46(0.03)	7.42(0.02)	7.43(0.02)
(500, 200)	1.14(0.01)	1.31(0.01)	6.39(0.04)	6.65(0.08)	11.73(0.01)	12.23(0.05)
Permutation Bandable Model, MCR $\rho^{(1)} = 0.8, \rho^{(2)} = 0.2$						
(50, 50)	4.23(0.38)	4.71(1.17)	6.67(2.30)	7.46(8.92)	11.22(1.34)	11.71(2.01)
(50, 200)	1.64(0.05)	2.79(0.39)	2.94(0.21)	4.52(0.95)	4.41(0.13)	6.29(0.46)
(200, 100)	3.17(0.06)	4.16(0.57)	5.73(0.66)	8.11(1.87)	15.93(0.53)	18.03(0.77)
(200, 200)	2.00(0.03)	3.22(0.18)	3.65(0.16)	5.70(0.60)	9.83(0.11)	13.29(0.55)
(500, 200)	2.22(0.03)	3.45(0.17)	4.09(0.17)	6.44(0.96)	16.80(0.14)	21.93(0.45)
Randomly Sparse Model, MCR $\rho^{(1)} = 0.8, \rho^{(2)} = 0.2$						
(50, 50)	2.15(0.46)	2.19(0.49)	4.21(0.94)	4.47(4.65)	6.36(0.96)	7.25(1.57)
(50, 200)	1.09(0.02)	1.16(0.04)	2.82(0.19)	2.99(0.32)	3.83(0.10)	4.00(0.20)
(200, 100)	1.46(0.02)	1.82(0.03)	4.96(0.12)	5.61(0.21)	8.45(0.07)	10.10(0.14)
(200, 200)	1.08(0.00)	1.20(0.01)	4.46(0.04)	4.57(0.05)	7.43(0.02)	7.66(0.04)
(500, 200)	1.12(0.01)	1.33(0.01)	6.35(0.04)	6.60(0.07)	11.71(0.02)	12.20(0.06)

Correspondingly, we consider the similar settings for the complete data with the same Σ and p , but different sample size n_c , where n_c can be one of the following three values,

- $\overline{n_{pair}^*} = \sum_{i,j=1}^n n_{ij}^*/p^2$: the average number of pairs of (x_i, x_j) 's that can be observed within the same sample;
- $\overline{n_s^*} = \sum_{i=1}^n n_i^*/p$: the average number of single x_i 's that can be observed;
- n : the same number of samples with the missing values.

Table 3
Comparison between incomplete samples and complete samples.

Setting	Sample size	Spectral norm	ℓ_1 norm	Frobenius norm
Bandable Covariance Matrix Estimation				
Missing data	$n = 1000$	0.72(0.01)	1.25(0.03)	2.40(0.01)
Complete data	$n_c = \overline{n_{\text{pair}}^*}$	0.97(0.03)	1.49(0.05)	2.48(0.04)
Complete data	$n_c = \overline{n_s^*}$	0.65(0.01)	1.01(0.03)	1.69(0.03)
Complete data	$n_c = n$	0.48(0.01)	0.73(0.01)	1.22(0.01)
Sparse Covariance Matrix Estimation				
Missing data	$n = 1000$	0.75(0.01)	1.37(0.04)	2.90(0.02)
Complete data	$n_c = \overline{n_{\text{pair}}^*}$	0.83(0.02)	1.31(0.05)	2.94(0.04)
Complete data	$n_c = \overline{n_s^*}$	0.65(0.01)	1.01(0.03)	1.86(0.04)
Complete data	$n_c = n$	0.45(0.01)	0.64(0.01)	1.12(0.01)

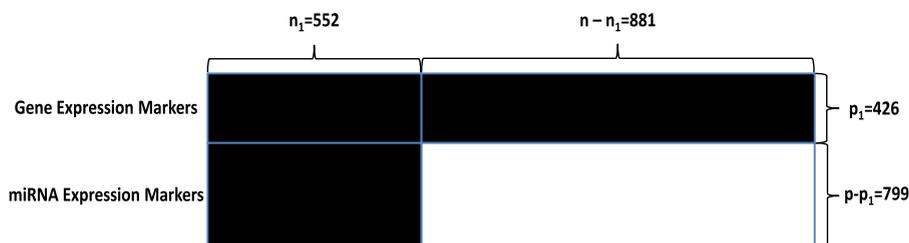


Fig. 2. Illustration of the ovarian cancer dataset. Black block = completely observed; White block = completely missing.

The results for all the settings are summarized in Table 3. It can be seen that the equivalent sample size depends on the loss function and in general it is between $\overline{n_{\text{pair}}^*}$ and $\overline{n_s^*}$. Overall, the average risk under the missing data setting is most comparable to that under the complete data setting for the sample size of $n_c = \overline{n_{\text{pair}}^*}$, the average number of observed pairs.

4.4. Analysis of ovarian cancer data

In this section, we illustrate the proposed adaptive thresholding procedure with an application to data from four ovarian cancer genomic studies, Cancer Genome Atlas Research Network [11] (TCGA), Bonome et al. [4] (BONO), Dressman et al. [15] (DRES) and Tothill et al. [27] (TOTH). The method introduced in Section 3 enables us to estimate the covariance matrix by integrating data from all four studies and thus yields a more accurate estimator. The data structure is illustrated in Fig. 2. The gene expression markers (the first 426 rows) are observed in all four studies without any missingness (the top black block in Fig. 2). The miRNA expression markers are observed in 552 samples from the TCGA study (the bottom left block in Fig. 2) and completely missing in the 881 samples from the TOTH, DRES, BONO and part of TCGA studies (the white block in Fig. 2).

Our goal is to estimate the covariance matrix Σ of the 1225 variables with the particular interest in the cross-covariances between the gene and miRNA expression markers. It is clear that the missingness here is not uniformly at random. On the other hand, it is reasonable to assume the missingness does not depend on the value of the data and thus missing completely at random (Assumption 2.1) can be assumed. We apply the adaptive thresholding procedure with $\delta = 2$ to estimate the covariance matrix and recover its support based on all the observations. The support of the estimate is shown in a heatmap in Fig. 3. The left panel is for the whole covariance matrix and the right panel zooms into the cross-covariances between the gene and miRNA expression markers.

It can be seen from Fig. 3 that the two diagonal blocks, with 12.24% and 8.39% nonzero off-diagonal entries respectively, are relatively dense, indicating that the relationships among the gene expression markers and those among the miRNA expression markers, as measured by their covariances, are relatively close. In contrast, the cross-covariances between gene and miRNA expression markers are very sparse with only 0.38% of significant gene-miRNA pairs. The gene and miRNA expression markers affect each other through different mechanisms, the cross-covariances between the gene and miRNA markers are of significant interest (see Ko et al. [19]). It is worthwhile to take a closer look at the cross-covariance matrix displayed on the right panel in Fig. 3. For each given gene, we count the number of miRNAs whose covariances with this gene are significant, and then rank all the genes by the counts. Similarly, we rank all the miRNAs. The top 5 genes and the top 5 miRNA expression markers are shown in Table 4.

Many of these gene and miRNA expression markers have been studied before in the literature. For example, the miRNA expression markers hsa-miR-142-5p and hsa-miR-142-3p have been demonstrated in Andreopoulos and Anastassiou [1] as standing out among the miRNA markers as having higher correlations with more genes, as well as methylation sites. Carraro et al. [12] finds that inhibition of miR-142-3p leads to ectopic expression of the gene marker ACTA2. This indicates strong interaction between miR-142-3p and ACTA2.

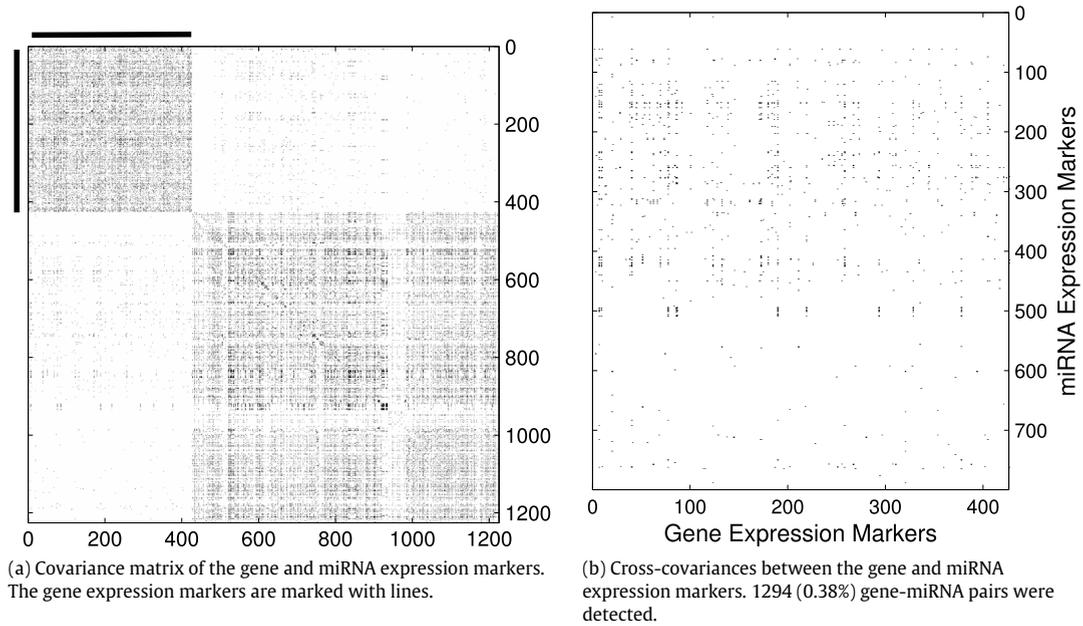


Fig. 3. Heatmaps of the covariance matrix estimate with all the observed data.

Table 4
Genes and miRNA's with most selected pairs.

Gene expression marker	Counts	miRNA expression marker	Counts
ACTA2	61	hsa-miR-142-5p	31
INHBA	57	hsa-miR-142-3p	29
COL10A1	53	hsa-miR-22	26
BGN	46	hsa-miR-21*	24
NID1	41	hsa-miR-146a	21

Table 5
Genes and miRNA's with most selected pairs after masking.

Gene expression marker	Counts	miRNA expression marker	Counts
ACTA2	60	hsa-miR-142-3p	31
INHBA	56	hsa-miR-142-5p	30
COL10A1	50	hsa-miR-146a	21
BGN	43	hsa-miR-150	21
NID1	40	hsa-miR-21*	21

To further demonstrate the robustness of our proposed procedure against missingness, we consider a setting with additional missing observations. We first randomly select half of the 552 complete samples (where both gene and miRNA expression markers are observed) and half of the 881 incomplete samples (where only gene expression markers are observed), and then independently mask each entry of the selected samples with probability 0.05. The proposed adaptive thresholding procedure is then applied to the data with these additional missing values. The estimated covariance matrix is shown in heatmaps in Fig. 4. These additional missing observations do not significantly affect the estimation accuracy. Fig. 4 is visually very similar to Fig. 3. To quantify the similarity between the two estimates, we calculate the Matthews correlation coefficient (MCC) between them. The value of MCC is equal to 0.9441, which indicates that the estimate based on the data with the additional missingness is very close to the estimate based on the original samples. We also pay close attention to the cross-covariance matrix displayed on the right panel in Fig. 4 and rank the gene and miRNA expression markers in the same way as before. The top 5 genes and the top 5 miRNA expression markers, listed in Table 5, are nearly identical to those given in Table 4, which are based on the original samples. These results indicate that the proposed method is robust against additional missingness.

5. Discussions

We considered in the present paper estimation of bandable and sparse covariance matrices in the presence of missing observations. The pivotal quantity is the generalized sample covariance matrix defined in (7). The technical analysis is

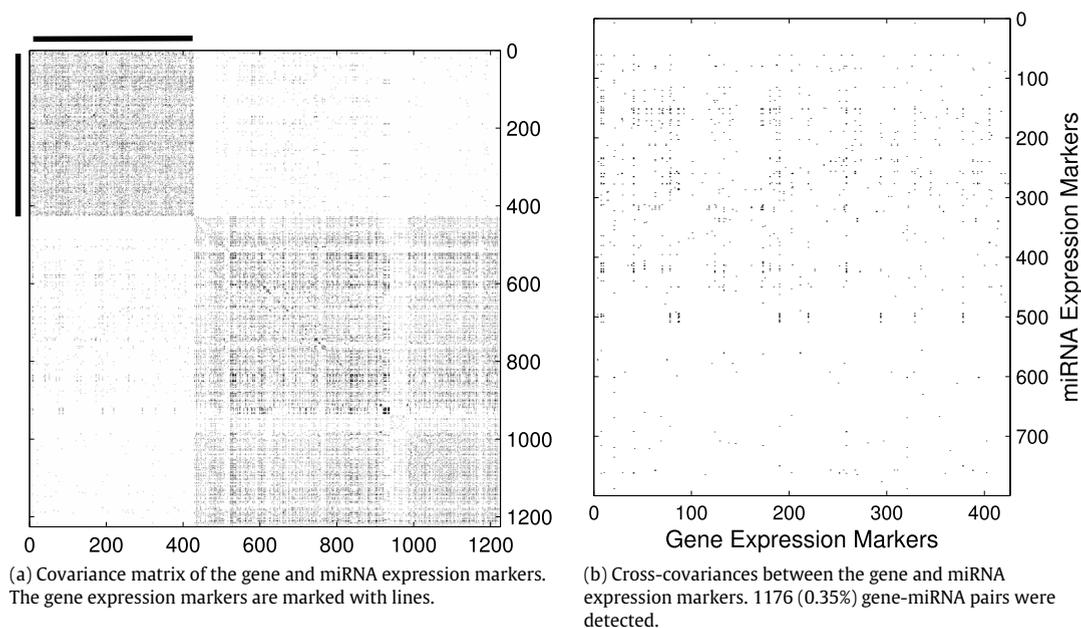


Fig. 4. Heatmaps of the covariance matrix estimate with additional missing values.

more challenging due to the missing data. We have mainly focused on the spectral norm loss in the theoretical analysis. Performance under other losses such as the Frobenius norm can also be analyzed.

To illustrate the proposed methods, we integrated four ovarian cancer studies. These methods for high-dimensional covariance matrix estimation with missing data are also useful for other types of data integration. For example, linking multiple data sources such as electronic data records, medicare data, registry data and patient reported outcomes could greatly increase the power of exploratory studies such as phenome-wide association studies (Denny et al. [14]). However, missing data inevitably arises and may hinder the potential of integrative analysis. In addition to random missingness due to unavailable information on a small fraction of patients, many variables such as the genetic measurements may only exist in one or two data sources and are hence structurally missing for other data sources. Our proposed methods could potentially provide accurate recovery of the covariance matrix in the presence of missingness.

In this paper, we allowed the proportion of missing values to be non-negligible as long as the minimum number of occurrences of any pair of variables n_{\min}^* is of order n . An interesting question is what happens when the number of observed values is large but n_{\min}^* is small (or even zero). We believe that the covariance matrix Σ can still be well estimated under certain global structural assumptions. This is out of the scope of the present paper and is an interesting problem for future research.

The key ideas and techniques developed in this paper can be used for a range of other related problems in high-dimensional statistical inference with missing data. For example, the same techniques can also be applied to estimation of other structured covariance matrices such as Toeplitz matrices, which have been studied in the literature in the case of complete data. When there are missing data, we can construct similar estimators using the generalized sample covariance matrix. The large deviation bounds for a sub-matrix and self-normalized entries of the generalized sample covariance matrix developed in Lemmas 3.1 and 2.1 would be helpful for analyzing the properties of the estimators.

The techniques can also be used on two-sample problems such as estimation of differential correlation matrices and hypothesis testing on the covariance structures. The generalized sample covariance matrix can be standardized to form the generalized sample correlation matrix which can then be used to estimate the differential correlation matrix in the two-sample case. It is also of significant interest in some applications to test the covariance structures in both one- and two-sample settings based on incomplete data. In the one-sample case, it is of interest to test the hypothesis $\{H_0 : \Sigma = \mathbf{I}\}$ or $\{H_0 : \mathbf{R} = \mathbf{I}\}$, where \mathbf{R} is the correlation matrix. In the two-sample case, one wishes to test the equality of two covariance matrices $\{H_0 : \Sigma_1 = \Sigma_2\}$. These are interesting problems for further exploration in the future.

6. Proofs

We prove Theorem 2.1 and the key technical result Lemma 6.1 for the bandable covariance matrix estimation in this section.

6.1. Proof of Lemma 2.1

To prove this lemma, we first introduce the following technical tool for the spectral norm of the sub-matrices.

Lemma 6.1. Suppose $\Sigma \in \mathbb{R}^{p \times p}$ is any positive semi-definite matrix, $A, B \in \{1, \dots, p\}$, then

$$\|\Sigma_{A \times B}\| \leq (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}. \tag{27}$$

The proof of Lemma 6.1 is provided later and now we move back to the proof of Lemma 2.1. Without loss of generality, we assume that $\mu = \mathbf{E}\mathbf{X} = 0$. We further define

$$\check{\Sigma}^* = (\check{\sigma}_{ij}^*)_{1 \leq i, j \leq p}, \quad \check{\sigma}_{ij}^* = \frac{1}{n_{ij}^*} \sum_{k=1}^n X_{ik} X_{jk} S_{ik} S_{jk}. \tag{28}$$

Also for convenience of presentation, we use C, C_1, c, \dots to denote uniform constants, whose exact values may vary in different scenarios. The lemma is now proved in the following steps:

1. We first consider for fixed unit vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^p$ with $\text{supp}(\mathbf{a}) \subseteq A, \text{supp}(\mathbf{b}) \subseteq B$, the tail bound of $\mathbf{a}^\top (\hat{\Sigma}^* - \Sigma) \mathbf{b}$. We would like to show that there exist uniform constants $C_1, c > 0$ such that for all $x > 0$,

$$\Pr \left\{ \left| \mathbf{a}^\top (\hat{\Sigma}^* - \Sigma) \mathbf{b} \right| \geq x \right\} \leq C_1 \exp \left\{ -cn_{\min}^* \min \left(\frac{x^2}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right) \right\}. \tag{29}$$

Specifically, we will bound $\mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}) \mathbf{b}$ and $\mathbf{a}^\top (\check{\Sigma}^* - \Sigma) \mathbf{b}$ separately in the next two steps.

2. We consider $\mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}) \mathbf{b}$ first. Since

$$\check{\sigma}_{ij}^* - \hat{\sigma}_{ij}^* = \frac{1}{n_{ij}^*} \sum_{k=1}^n (X_{jk} \bar{X}_i^* + X_{ik} \bar{X}_j^*) S_{ik} S_{jk} - \bar{X}_i^* \bar{X}_j^*,$$

$\mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}) \mathbf{b}$ can be written as

$$\begin{aligned} \mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}) \mathbf{b} &= \sum_{i,j=1}^p a_i b_j (\check{\sigma}_{ij}^* - \hat{\sigma}_{ij}^*) \\ &= \sum_{i,j=1}^p a_i b_j \left(\frac{\sum_{k=1}^n X_{ik} S_{ik}}{n_i^*} \cdot \frac{\sum_{l=1}^n X_{jl} S_{jl}}{n_j^*} + \frac{\sum_{k=1}^n X_{ik} S_{ik} S_{jk}}{n_{ij}^*} - \frac{\sum_{l=1}^n X_{jl} S_{jl}}{n_j^*} - \frac{\sum_{k=1}^n X_{ik} S_{ik}}{n_i^*} \cdot \frac{\sum_{l=1}^n X_{jl} S_{jl}}{n_j^*} \right) \\ &= \sum_{i,j=1}^p \sum_{k,l=1}^n X_{ik} X_{jl} a_i b_j \left(\frac{S_{ik} S_{jl} S_{jl}}{n_i^* n_{ij}^*} + \frac{S_{ik} S_{jk} S_{jl}}{n_{ij}^* n_j^*} - \frac{S_{ik} S_{jl}}{n_i^* n_j^*} \right). \end{aligned} \tag{30}$$

We can calculate from (30) that

$$\begin{aligned} \left| \mathbf{E} \mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}) \mathbf{b} \right| &= \left| \sum_{i,j=1}^p \sum_{k=1}^n \sigma_{ij} a_i b_j \left(\frac{S_{ik} S_{ik} S_{jk}}{n_i^* n_{ij}^*} + \frac{S_{ik} S_{jk} S_{jk}}{n_{ij}^* n_j^*} - \frac{S_{ik} S_{jk}}{n_i^* n_j^*} \right) \right| \\ &= \left| \sum_{i,j=1}^p \sigma_{ij} \frac{a_i}{n_i^*} b_j + \sum_{i,j=1}^p a_i \frac{b_j}{n_j^*} \sigma_{ij} - \sum_{k=1}^n \sum_{i,j=1}^p \frac{S_{ik} a_i}{n_i^*} \frac{S_{jk} b_j}{n_j^*} \sigma_{ij} \right| \\ &\leq \left| \left(\frac{a_1}{n_1^*}, \dots, \frac{a_p}{n_p^*} \right) \Sigma \mathbf{b} \right| + \left| \mathbf{a}^\top \Sigma \left(\frac{b_1}{n_1^*}, \dots, \frac{b_p}{n_p^*} \right)^\top \right| \\ &\quad + \sum_{k=1}^n \left| \left(\frac{S_{1k} a_1}{n_1^*}, \dots, \frac{S_{pk} a_p}{n_p^*} \right) \Sigma \left(\frac{S_{1k} b_1}{n_1^*}, \dots, \frac{S_{pk} b_p}{n_p^*} \right)^\top \right| \\ &\leq \|\Sigma_{A \times B}\| \frac{\|\mathbf{a}\|_2 \|\mathbf{b}\|_2}{n_{\min}^*} + \|\Sigma_{A \times B}\| \frac{\|\mathbf{a}\|_2 \|\mathbf{b}\|_2}{n_{\min}^*} \\ &\quad + \sum_{k=1}^n \|\Sigma_{A \times B}\| \cdot \frac{1}{2} \left\{ \left\| \left(\frac{S_{1k} a_1}{n_1^*}, \dots, \frac{S_{pk} a_p}{n_p^*} \right) \right\|_2^2 + \left\| \left(\frac{S_{1k} b_1}{n_1^*}, \dots, \frac{S_{pk} b_p}{n_p^*} \right) \right\|_2^2 \right\}. \end{aligned} \tag{31}$$

For the last term in (31), we have the following bound,

$$\begin{aligned} & \sum_{k=1}^n \|\Sigma_{A \times B}\| \cdot \frac{1}{2} \left\{ \left\| \left(\frac{S_{1k}a_1}{n_1^*}, \dots, \frac{S_{pk}a_p}{n_p^*} \right) \right\|_2^2 + \left\| \left(\frac{S_{1k}b_1}{n_1^*}, \dots, \frac{S_{pk}b_p}{n_p^*} \right) \right\|_2^2 \right\} \\ &= \|\Sigma_{A \times B}\| \sum_{k=1}^n \sum_{i=1}^p \frac{1}{2} \left(\frac{S_{ik}a_i^2}{n_i^{*2}} + \frac{S_{ik}b_i^2}{n_i^{*2}} \right) \\ &= \|\Sigma_{A \times B}\| \sum_{i=1}^p \frac{1}{2} \left(\frac{a_i^2 + b_i^2}{n_i^*} \right) \\ &\leq \|\Sigma_{A \times B}\| \sum_{i=1}^p \frac{a_i^2 + b_i^2}{2n_{\min}^*} \leq \frac{(\|\Sigma_A\| \|\Sigma_B\|)^{1/2}}{n_{\min}^*}. \end{aligned}$$

Thus, by (31) and the inequality above, we have

$$\left| \mathbf{Ea}^\top (\check{\Sigma}^* - \hat{\Sigma}^*) \mathbf{b} \right| \leq \frac{3(\|\Sigma_A\| \|\Sigma_B\|)^{1/2}}{n_{\min}^*}. \tag{32}$$

The last term of (30) can be treated as a quadratic form of the vectorization of \mathbf{X} : $\text{vec}(\mathbf{X}) \in \mathbb{R}^{pn}$. We note the last term as $\text{vec}(\mathbf{X})^\top \mathbf{Q} \text{vec}(\mathbf{X})$, where $\mathbf{Q} \in \mathbb{R}^{pn \times pn}$ and

$$\mathbf{Q}_{(i,k),(j,l)} = a_i b_j \left(\frac{S_{ik} S_{il} S_{jl}}{n_i^* n_{ij}^*} + \frac{S_{ik} S_{jk} S_{jl}}{n_{ij}^* n_j^*} - \frac{S_{ik} S_{jl}}{n_i^* n_j^*} \right), \quad 1 \leq i, j \leq p, 1 \leq k, l \leq n.$$

\mathbf{Q} has the following properties,

$$\begin{aligned} \|\mathbf{Q}\|_F^2 &= \sum_{i,j=1}^p \sum_{k,l=1}^n a_i^2 b_j^2 \left(\frac{S_{ik} S_{il} S_{jl}}{n_i^* n_{ij}^*} + \frac{S_{ik} S_{jk} S_{jl}}{n_{ij}^* n_j^*} - \frac{S_{ik} S_{jl}}{n_i^* n_j^*} \right)^2 \\ &\leq \sum_{i,j=1}^p a_i^2 b_j^2 \sum_{k,l=1}^n \left(2 \frac{S_{ik} S_{il} S_{jl}}{n_i^{*2} n_{ij}^{*2}} + 2 \frac{S_{ik} S_{jk} S_{jl}}{n_{ij}^{*2} n_j^{*2}} + \frac{S_{ik} S_{jl}}{n_i^{*2} n_j^{*2}} \right), \quad \text{since } S_{ik} \in \{0, 1\}; \\ &\leq \sum_{i,j=1}^p a_i^2 b_j^2 \frac{5}{n_{\min}^{*2}} = \frac{5 \|\mathbf{a}\|_2^2 \|\mathbf{b}\|_2^2}{n_{\min}^{*2}} = \frac{5}{n_{\min}^{*2}}; \end{aligned} \tag{33}$$

$$\|\mathbf{Q}\| \leq \|\mathbf{Q}\|_F \leq \frac{\sqrt{5} \|\mathbf{a}\|_2 \|\mathbf{b}\|_2}{n_{\min}^*} \leq \frac{\sqrt{5}}{n_{\min}^*}. \tag{34}$$

For $\text{vec}(\mathbf{X}) \in \mathbb{R}^{pn}$, since its segments $\{\mathbf{X}_k, k = 1, \dots, p\}$ are independent and $\mathbf{X}_k = \mathbf{\Gamma} \mathbf{Z}_k$, we can further write $\text{vec}(\mathbf{X}) = \mathbf{D}_\Gamma \text{vec}(\mathbf{Z})$, where $\mathbf{D}_\Gamma \in \mathbb{R}^{pn \times qn}$ is with n diagonal blocks of $\mathbf{\Gamma}$, $\text{vec}(\mathbf{Z})$ is a (qn) -dimensional i.i.d. sub-Gaussian random vector. Based on Hanson-Wright’s inequality (Theorem 1.1 in Rudelson and Vershynin [25]),

$$\begin{aligned} & \Pr \left\{ \left| \mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}^*) \mathbf{b} - \mathbf{Ea}^\top (\check{\Sigma}^* - \hat{\Sigma}^*) \mathbf{b} \right| \geq x \right\} \\ &= \Pr \left\{ \left| \text{vec}(\mathbf{X})^\top \mathbf{Q} \text{vec}(\mathbf{X}) - \mathbf{E} \text{vec}(\mathbf{X})^\top \mathbf{Q} \text{vec}(\mathbf{X}) \right| \geq x \right\} \\ &= \Pr \left[\left| \text{vec}(\mathbf{Z})^\top \mathbf{D}_\Gamma^\top \mathbf{Q} \mathbf{D}_\Gamma \text{vec}(\mathbf{Z}) - \mathbf{E} \left\{ \text{vec}(\mathbf{Z})^\top \mathbf{D}_\Gamma^\top \mathbf{Q} \mathbf{D}_\Gamma \text{vec}(\mathbf{Z}) \right\} \right| \geq x \right] \\ &\leq 2 \exp \left\{ -c \min \left(\frac{x^2}{\tau^4 \|\mathbf{D}_\Gamma^\top \mathbf{Q} \mathbf{D}_\Gamma\|_F^2}, \frac{x}{\tau^2 \|\mathbf{D}_\Gamma \mathbf{Q} \mathbf{D}_\Gamma\|} \right) \right\}. \end{aligned} \tag{35}$$

Here $c > 0$ is a uniform constant. Since \mathbf{Q} is supported on $\{(i, k), (j, l) : i \in A, j \in B\}$, we have $\mathbf{D}_\Gamma^\top \mathbf{Q} \mathbf{D}_\Gamma = \mathbf{D}_{\Gamma_A}^\top \mathbf{Q}_{A \times B} \mathbf{D}_{\Gamma_B}$. Here $\mathbf{D}_{\Gamma_A} \in \mathbb{R}^{|A|n \times qn}$, $\mathbf{D}_{\Gamma_B} \in \mathbb{R}^{|B|n \times qn}$ are with n diagonal block $\mathbf{\Gamma}_{A \times [q]}$ and $\mathbf{\Gamma}_{B \times [q]}$, respectively, where $[q] = \{1, \dots, q\}$. Since $\mathbf{\Gamma}_{A \times [q]} \mathbf{\Gamma}_{A \times [q]}^\top = \Sigma_A$, $\mathbf{\Gamma}_{B \times [q]} \mathbf{\Gamma}_{B \times [q]}^\top = \Sigma_B$, we know

$$\|\mathbf{D}_{\Gamma_A}\| = \|\mathbf{\Gamma}_{A \times [q]}\| \leq \|\Sigma_A\|^{1/2}, \quad \|\mathbf{\Gamma}_{B \times [q]}\| \leq \|\mathbf{D}_{\Gamma_B}\| \leq \|\Sigma_B\|^{1/2}.$$

Then we further have

$$\begin{aligned} & \Pr \left\{ \left| \mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}^*) \mathbf{b} - \mathbf{Ea}^\top (\check{\Sigma}^* - \hat{\Sigma}^*) \mathbf{b} \right| \geq x \right\} \\ &\leq 2 \exp \left\{ -c \min \left(\frac{x^2}{\tau^4 \|\mathbf{D}_{\Gamma_A}^\top \mathbf{Q}_{A \times B} \mathbf{D}_{\Gamma_B}\|_F^2}, \frac{x}{\tau^2 \|\mathbf{D}_{\Gamma_A}^\top \mathbf{Q}_{A \times B} \mathbf{D}_{\Gamma_B}\|} \right) \right\} \end{aligned}$$

$$\begin{aligned}
&\leq 2 \exp \left\{ -c \min \left(\frac{x^2}{\tau^4 \|\mathbf{D}_{\Gamma_B}\|^2 \|\mathbf{D}_{\Gamma_A}^\top\|^2 \|\mathbf{Q}\|_F^2}, \frac{x}{\tau^2 \|\mathbf{D}_{\Gamma_B}\| \|\mathbf{D}_{\Gamma_A}^\top\| \|\mathbf{Q}\|} \right) \right\} \\
&\leq 2 \exp \left[-c \min \left\{ \frac{x^2}{\tau^4 \|\Sigma_A\| \|\Sigma_B\| \|\mathbf{Q}\|_F^2}, \frac{x}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2} \|\mathbf{Q}\|} \right\} \right] \\
&\leq 2 \exp \left[-c \min \left\{ \frac{x^2 n_{\min}^{*2}}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x n_{\min}^*}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right\} \right]. \tag{36}
\end{aligned}$$

We define $x' = \max \{x - 3(\|\Sigma_A\| \|\Sigma_B\|)^{1/2}/n_{\min}^*, 0\}$, combining the inequality above and (32), we have

$$\begin{aligned}
\Pr \left\{ \mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}^*) \mathbf{b} \geq x \right\} &\leq \Pr \left\{ \mathbf{a}^\top (\check{\Sigma}^* - \hat{\Sigma}^*) \mathbf{b} - \mathbf{E} \mathbf{a}^\top \check{\Sigma}^* \mathbf{b} \geq x' \right\} \\
&\leq 2 \exp \left[-c \min \left\{ \frac{(x')^2 n_{\min}^{*2}}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x' n_{\min}^*}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right\} \right] \\
&\leq 2 \exp \left[-c' \min \left\{ \frac{x^2 n_{\min}^{*2}}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x n_{\min}^*}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right\} \right] + C \max \left(\frac{1}{\tau^4}, \frac{1}{\tau^2} \right) \\
&\leq C \exp \left[-c' \min \left\{ \frac{x^2 n_{\min}^{*2}}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x n_{\min}^*}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right\} \right]. \tag{37}
\end{aligned}$$

In the last inequality above, we used a fact that τ is lower bounded by a uniform constant. This is due to Assumption 2.2 that $\mathbf{E}(Z) = 0$, $\text{var}(Z) = 1$, $\mathbf{E} \exp(tZ) \leq \exp(t^2\tau^2/2)$. Then,

$$\exp(4\tau^2/2) \geq \frac{1}{2} \{ \mathbf{E} \exp(2Z) + \mathbf{E} \exp(-2Z) \} = \sum_{k=0}^{\infty} \frac{2^{2k} \mathbf{E} Z^{2k}}{(2k)!} \geq 2\mathbf{E} Z^2 = 2,$$

which implies $\tau^2 \geq \frac{1}{2} \ln(2)$.

3. It is easy to see that $\mathbf{E} \check{\Sigma}^* = \Sigma$, so $\mathbf{E} \mathbf{a}^\top (\check{\Sigma}^* - \Sigma) \mathbf{b} = 0$. Then

$$\begin{aligned}
\mathbf{a}^\top (\check{\Sigma}^* - \Sigma) \mathbf{b} &= \sum_{i,j=1}^p a_i b_j \left(\frac{1}{n_{ij}^*} \sum_{k=1}^n X_{ik} X_{jk} S_{ik} S_{jk} \right) - \mathbf{E} \sum_{i,j=1}^p a_i b_j \left(\frac{1}{n_{ij}^*} \sum_{k=1}^n X_{ik} X_{jk} S_{ik} S_{jk} \right) \\
&= \sum_{k=1}^n \sum_{i,j=1}^p \left(\frac{a_i b_j S_{ik} S_{jk}}{n_{ij}^*} X_{ik} X_{jk} - \mathbf{E} \frac{a_i b_j S_{ik} S_{jk}}{n_{ij}^*} X_{ik} X_{jk} \right) \\
&\triangleq \sum_{k=1}^n (\mathbf{X}_k^\top \mathbf{C}^k \mathbf{X}_k - \mathbf{E} \mathbf{X}_k^\top \mathbf{C}^k \mathbf{X}_k) \\
&= \sum_{k=1}^n (\mathbf{Z}_k^\top \Gamma^\top \mathbf{C}^k \Gamma \mathbf{Z}_k - \mathbf{E} \mathbf{Z}_k^\top \Gamma^\top \mathbf{C}^k \Gamma \mathbf{Z}_k). \tag{38}
\end{aligned}$$

Here $\mathbf{C}^k \in \mathbb{R}^{p \times p}$ is a matrix such that $C_{ij}^k = a_i b_j S_{ik} S_{jk} / n_{ij}^*$. Note that \mathbf{C}^k is supported on $A \times B$, we can prove the following properties of \mathbf{C}^k .

$$\begin{aligned}
\|\Gamma^\top \mathbf{C}^k \Gamma\|_F &= \sqrt{\text{tr}(\mathbf{C}^k \Gamma \Gamma^\top \mathbf{C}^{k\top} \Gamma \Gamma^\top)} \\
&= \sqrt{\text{tr}(\mathbf{C}_{A \times B}^k \Gamma_{B \times [q]} \Gamma_{B \times [q]}^\top \mathbf{C}_{A \times B}^{k\top} \Gamma_{A \times [q]} \Gamma_{A \times [q]}^\top)} \\
&\leq \|\Gamma_{B \times [q]}\| \|\Gamma_{A \times [q]}\| \sqrt{\text{tr}(\mathbf{C}^k \mathbf{C}^{k\top})} \\
&\leq (\|\Sigma_A\| \|\Sigma_B\|)^{1/2} \sqrt{\text{tr}(\mathbf{C}^k \mathbf{C}^{k\top})}; \tag{39}
\end{aligned}$$

$$\begin{aligned}
\sum_{k=1}^n \|\Gamma^\top \mathbf{C}^k \Gamma\|_F^2 &\leq \|\Sigma_A\| \|\Sigma_B\| \|\mathbf{C}^k\|_F^2 = \|\Sigma_A\| \|\Sigma_B\| \sum_{k=1}^n \sum_{i,j=1}^p \left(\frac{a_i b_j S_{ik} S_{jk}}{n_{ij}^*} \right)^2 \\
&= \|\Sigma_A\| \|\Sigma_B\| \sum_{k=1}^n \sum_{i,j=1}^p \frac{S_{ik} S_{jk} a_i^2 b_j^2}{n_{ij}^{*2}} = \|\Sigma_A\| \|\Sigma_B\| \sum_{i,j=1}^p \frac{a_i^2 b_j^2}{n_{ij}^*} \\
&\leq \|\Sigma_A\| \|\Sigma_B\| \frac{\|\mathbf{a}\|_2^2 \|\mathbf{b}\|_2^2}{n_{\min}^*} = \frac{\|\Sigma_A\| \|\Sigma_B\|}{n_{\min}^*}; \tag{40}
\end{aligned}$$

$$\begin{aligned}
 \|\Gamma^\top \mathbf{C}^k \Gamma\| &\leq \|\Gamma^\top \mathbf{C}^k \Gamma\|_F \leq (\|\Sigma_A\| \|\Sigma_B\|)^{1/2} \sqrt{\text{tr}(\mathbf{C}^k \mathbf{C}^{k\top})} \\
 &\leq (\|\Sigma_A\| \|\Sigma_B\|)^{1/2} \sqrt{\sum_{i,j=1}^p \left(\frac{a_i b_j S_{ik} S_{jk}}{n_{ij}^*}\right)^2} \\
 &\leq (\|\Sigma_A\| \|\Sigma_B\|)^{1/2} \sqrt{\sum_{i,j=1}^p \frac{a_i^2 b_j^2}{n_{\min}^{*2}}} \leq \frac{\sqrt{\|\Sigma_A\| \|\Sigma_B\|}}{n_{\min}^*}.
 \end{aligned} \tag{41}$$

Now, note that the last line of (38) can be also equivalently written as

$$\text{vec}(\mathbf{Z})^\top \mathbf{C}^{\text{con}} \text{vec}(\mathbf{Z})^\top - \text{Evec}(\mathbf{Z})^\top \mathbf{C}^{\text{con}} \text{vec}(\mathbf{Z})^\top,$$

$$\mathbf{C}^{\text{con}} = \begin{bmatrix} \Gamma^\top \mathbf{C}^1 \Gamma & & \\ & \ddots & \\ & & \Gamma^\top \mathbf{C}^n \Gamma \end{bmatrix} \in \mathbb{R}^{(qn) \times (qn)},$$

where $\text{vec}(\mathbf{Z})$ is the vectorization of \mathbf{Z} , which is an qn -dimensional i.i.d. sub-Gaussian vector. Based on the properties of \mathbf{C}^k above, we have

$$\begin{aligned}
 \|\mathbf{C}^{\text{con}}\|_F^2 &= \sum_{k=1}^n \|\Gamma^\top \mathbf{C}^k \Gamma\|_F^2 \stackrel{(40)}{\leq} \frac{\|\Sigma_A\| \|\Sigma_B\|}{n_{\min}^*}, \\
 \|\mathbf{C}^{\text{con}}\| &\leq \max_{1 \leq k \leq n} \|\Gamma^\top \mathbf{C}^k \Gamma\| \stackrel{(41)}{\leq} \frac{(\|\Sigma_A\| \|\Sigma_B\|)^{1/2}}{n_{\min}^*}.
 \end{aligned}$$

Now applying Hanson-Wright's inequality (Theorem 1.1 in Rudelson and Vershynin [25]), we have

$$\begin{aligned}
 \Pr \left\{ \left| \text{vec}(\mathbf{Z}_k)^\top \mathbf{C}^{\text{con}} \text{vec}(\mathbf{Z}_k) - \text{Evec}(\mathbf{Z}_k)^\top \mathbf{C}^{\text{con}} \text{vec}(\mathbf{Z}_k) \right| \geq x \right\} \\
 \leq 2 \exp \left\{ -c \min \left(\frac{x^2}{\tau^4 \|\mathbf{C}^{\text{con}}\|_F^2}, \frac{x}{\tau^2 \|\mathbf{C}^{\text{con}}\|} \right) \right\} \\
 \leq 2 \exp \left\{ -cn_{\min}^* \min \left(\frac{x^2}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right) \right\}.
 \end{aligned} \tag{42}$$

Thus,

$$\Pr \left\{ \left| \mathbf{a}^\top (\check{\Sigma}^* - \Sigma) \mathbf{b} \right| \geq x \right\} \leq 2 \exp \left\{ -cn_{\min}^* \min \left(\frac{x^2}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x}{\tau^2 (\|\Sigma_A\| \|\Sigma_B\|)^{1/2}} \right) \right\}. \tag{43}$$

Here c is a uniform constant. Combining (43) and (37), we have (29).

4. Next, we use the ε -net technique to give the bound on $\|\hat{\Sigma}_{A \times B}^* - \Sigma_{A \times B}\|$. Denote $\mathbf{D}^* = \hat{\Sigma}_{A \times B}^* - \Sigma_{A \times B}$. Suppose $S_{1/3}^A$ is the $(1/3)$ -net for all unit vectors in $\mathbb{R}^{|A|}$; similarly $S_{1/3}^B$ is the $(1/3)$ -net for all unit vectors in $\mathbb{R}^{|B|}$. Based on the proof of Lemma 3 in Cai et al. [9], we can let $\text{Card}(S_{1/3}^A) \leq 7^k$, $\text{Card}(S_{1/3}^B) \leq 7^k$. Since for all $\mathbf{a}, \mathbf{a}_0 \in \mathbb{R}^{|A|}$, $\mathbf{b}, \mathbf{b}_0 \in \mathbb{R}^{|B|}$,

$$\begin{aligned}
 |\mathbf{a}^\top \mathbf{D}^* \mathbf{b}| - |\mathbf{a}_0^\top \mathbf{D}^* \mathbf{b}_0| &\leq |\mathbf{a}^\top \mathbf{D}^* \mathbf{b} - \mathbf{a}_0^\top \mathbf{D}^* \mathbf{b}_0| \leq |(\mathbf{a} - \mathbf{a}_0)^\top \mathbf{D}^* \mathbf{b}| + |\mathbf{a}_0^\top \mathbf{D}^* (\mathbf{b} - \mathbf{b}_0)| \\
 &\leq (\|\mathbf{a} - \mathbf{a}_0\|_2 + \|\mathbf{b} - \mathbf{b}_0\|_2) \|\mathbf{D}^*\|,
 \end{aligned} \tag{44}$$

we have for all $\mathbf{a} \in \mathbb{R}^{|A|}$, $\mathbf{b} \in \mathbb{R}^{|B|}$, $\|\mathbf{a}\|_2 = \|\mathbf{b}\|_2 = 1$, we can find $\mathbf{a}_0 \in S_{1/3}^A$, $\mathbf{b}_0 \in S_{1/3}^B$ such that $\|\mathbf{a}_0 - \mathbf{a}\|_2 \leq 1/3$, $\|\mathbf{b}_0 - \mathbf{b}\|_2 \leq 1/3$, then

$$\begin{aligned}
 |\mathbf{a}^\top \mathbf{D}^* \mathbf{b}| &\leq |\mathbf{a}_0^\top \mathbf{D}^* \mathbf{b}_0| + \frac{2}{3} \|\mathbf{D}^*\| \leq \sup_{\mathbf{a}_0 \in S_{1/3}^A, \mathbf{b}_0 \in S_{1/3}^B} |\mathbf{a}_0^\top \mathbf{D}^* \mathbf{b}_0| + \frac{2}{3} \|\mathbf{D}^*\|, \\
 \|\mathbf{D}^*\| &= \sup_{\substack{\mathbf{a} \in \mathbb{R}^{|A|}, \mathbf{b} \in \mathbb{R}^{|B|}, \\ \|\mathbf{a}\|_2 = \|\mathbf{b}\|_2 = 1}} |\mathbf{a}^\top \mathbf{D}^* \mathbf{b}| \leq \sup_{\mathbf{a}_0 \in S_{1/3}^A, \mathbf{b}_0 \in S_{1/3}^B} |\mathbf{a}_0^\top \mathbf{D}^* \mathbf{b}_0| + \frac{2}{3} \|\mathbf{D}^*\|,
 \end{aligned}$$

which yields

$$\|\hat{\Sigma}_{A \times B}^* - \Sigma_{A \times B}\| = \|\mathbf{D}^*\| \leq 3 \sup_{\mathbf{a}_0 \in S_{1/3}^A, \mathbf{b}_0 \in S_{1/3}^B} |\mathbf{a}_0^\top \mathbf{D}^* \mathbf{b}_0|. \tag{45}$$

Finally, by combining (29) and the inequality above, we know there exist uniform constants $C_1, c > 0$ such that for all $t > 0$,

$$\begin{aligned} \Pr\left(\|\hat{\Sigma}_{A \times B}^* - \Sigma_{A \times B}\| \geq x\right) &\leq \Pr\left(\sup_{\mathbf{a}_0 \in S_{1/3}^A, \mathbf{b}_0 \in S_{1/3}^B} |\mathbf{a}_0^\top \mathbf{D}^* \mathbf{b}_0| \geq \frac{x}{3}\right) \\ &\leq C_1(7)^{|A|+|B|} \exp\left[-cn_{\min}^* \min\left\{\frac{x^2}{\tau^4 \|\Sigma_A\| \|\Sigma_B\|}, \frac{x}{\tau^2(\|\Sigma_A\| \|\Sigma_B\|)^{1/2}}\right\}\right]. \end{aligned} \tag{46}$$

Since $|A| + |B| \leq 2|A \cup B|$, we have finished the proof of Lemma 2.1. \square

Proof of Lemma 6.1. Since Σ is positive semi-definite, we can find the Cholesky decomposition such that $\Sigma = \mathbf{V}\mathbf{V}^\top$. Then $\Sigma_{A \times B} = \mathbf{V}_{A \times [p]} \mathbf{V}_{B \times [p]}^\top$ and

$$\begin{aligned} \|\Sigma_{A \times B}\| &= \max_{\substack{\mathbf{x} \in \mathbb{R}^{|A|}, \mathbf{y} \in \mathbb{R}^{|B|} \\ \|\mathbf{x}\|_2 = \|\mathbf{y}\|_2 = 1}} \mathbf{x}^\top \mathbf{V}_{A \times [p]} \mathbf{V}_{B \times [p]}^\top \mathbf{y} \\ &\leq \max_{\substack{\mathbf{x} \in \mathbb{R}^{|A|}, \mathbf{y} \in \mathbb{R}^{|B|} \\ \|\mathbf{x}\|_2 = \|\mathbf{y}\|_2 = 1}} (\mathbf{x}^\top \mathbf{V}_{A \times [p]} \mathbf{V}_{A \times [p]}^\top \mathbf{x})^{1/2} (\mathbf{y}^\top \mathbf{V}_{B \times [p]} \mathbf{V}_{B \times [p]}^\top \mathbf{y})^{1/2} \\ &= \max_{\substack{\mathbf{x} \in \mathbb{R}^{|A|}, \mathbf{y} \in \mathbb{R}^{|B|} \\ \|\mathbf{x}\|_2 = \|\mathbf{y}\|_2 = 1}} (\mathbf{x}^\top \Sigma_A \mathbf{x})^{1/2} (\mathbf{y}^\top \Sigma_B \mathbf{y})^{1/2} = \|\Sigma_A\| \|\Sigma_B\|. \end{aligned}$$

Here we have used the Cauchy–Schwarz inequality. \square

6.2. Proof of Theorem 2.1

Define $\mathbf{B} = (b_{ij})_{1 \leq i, j \leq p}$ such that $b_{ij} = \sigma_{ij}$ if $i \in I_s, j \in I_{s'}$ and $|s - s'| \leq 1$, and 0 otherwise. Let $\Delta = \Sigma - \mathbf{B}$. Then

$$\|\hat{\Sigma}^{\text{bt}} - \Sigma\| \leq \|\hat{\Sigma}^{\text{bt}} - \mathbf{B}\| + \|\Delta\|.$$

It is easy to see that

$$\|\Delta\| \leq \|\Delta\|_{\ell_1} \leq \max_i \sum_{j: |i-j| \geq k} |\sigma_{ij}| \leq Mk^{-\alpha}.$$

To bound $\|\hat{\Sigma}^{\text{bt}} - \mathbf{B}\|$, note that

$$\|\hat{\Sigma}^{\text{bt}} - \mathbf{B}\| = \sup_{u \in \mathbb{R}^p: \|u\|_2 = 1} |\langle u, (\hat{\Sigma}^{\text{bt}} - \mathbf{B})u \rangle|.$$

For any $u \in \mathbb{R}^p, \|u\|_2 = 1$, we have

$$\begin{aligned} |\langle u, (\hat{\Sigma}^{\text{bt}} - \mathbf{B})u \rangle| &\leq \sum_{s, s': |s-s'| \leq 1} \left| \langle u_{I_s}, (\hat{\Sigma}_{I_s \times I_{s'}}^* - \Sigma_{I_s \times I_{s'}}) u_{I_{s'}} \rangle \right| \\ &\leq \sum_{s, s': |s-s'| \leq 1} \|u_{I_s}\|_2 \|u_{I_{s'}}\|_2 \|\hat{\Sigma}_{I_s \times I_{s'}}^* - \Sigma_{I_s \times I_{s'}}\| \\ &\leq \left(\sum_{s, s': |s-s'| \leq 1} \|u_{I_s}\|_2 \|u_{I_{s'}}\|_2 \right) \left(\max_{|s-s'| \leq 1} \|\hat{\Sigma}_{I_s \times I_{s'}}^* - \Sigma_{I_s \times I_{s'}}\| \right). \end{aligned}$$

The Cauchy–Schwarz inequality yields

$$\sum_{s, s': |s-s'| \leq 1} \|u_{I_s}\|_2 \|u_{I_{s'}}\|_2 \leq \frac{1}{2} \sum_{s, s': |s-s'| \leq 1} (\|u_{I_s}\|_2^2 + \|u_{I_{s'}}\|_2^2) \leq 3 \sum_{s=1}^N \|u_{I_s}\|_2^2 = 3. \tag{47}$$

Therefore,

$$\|\hat{\Sigma}^{\text{bt}} - \Sigma\| \leq \|\hat{\Sigma}^* - \mathbf{B}\| + \|\Delta\| \leq 3 \max_{|s-s'| \leq 1} \|\hat{\Sigma}_{I_s \times I_{s'}}^* - \Sigma_{I_s \times I_{s'}}\| + Mk^{-\alpha},$$

which yields

$$\mathbb{E} \|\hat{\Sigma}^{\text{bt}} - \Sigma\|^2 \leq 18 \mathbb{E} \left(\max_{|s-s'| \leq 1} \|\hat{\Sigma}_{I_s \times I_{s'}}^* - \Sigma_{I_s \times I_{s'}}\| \right)^2 + 2M^2 k^{-2\alpha}.$$

According to Lemma 2.1, there exists constant $C, c > 0$ which only depend on τ such that for all $x > 0$,

$$\Pr \left(\max_{|s-s'|\leq 1} \|\hat{\Sigma}_{I_s \times I_{s'}} - \Sigma_{I_s \times I_{s'}}\| \geq x \right) \leq C \left\lceil \frac{p}{k} \right\rceil (49)^k \exp \left\{ -cn_{\min}^* \min \left(\frac{x^2}{\|\Sigma\|^2}, \frac{x}{\|\Sigma\|} \right) \right\}. \quad (48)$$

Now we set $t = C'(k + \ln p)/n_{\min}^*$ for C' large enough. The spectral norm risk satisfies

$$\begin{aligned} E\|\hat{\Sigma}^{bt} - \Sigma\|^2 &\leq 18E \max_{|s-s'|\leq 1} \|\hat{\Sigma}_{I_s \times I_{s'}}^*\| + 2M^2k^{-2\alpha} \\ &\leq 18 \int_0^\infty \Pr \left(\max_{|s-s'|\leq 1} \|\hat{\Sigma}_{I_s \times I_{s'}} - \Sigma_{I_s \times I_{s'}}\|^2 \geq x \right) dx + 2M^2k^{-2\alpha} \\ &\leq 18t + 18 \int_t^\infty \Pr \left(\max_{|s-s'|\leq 1} \|\hat{\Sigma}_{I_s \times I_{s'}} - \Sigma_{I_s \times I_{s'}}\|^2 \geq x \right) dx + 2M^2k^{-2\alpha} \\ &\leq 18t + C \left\lceil \frac{p}{k} \right\rceil (49)^k \int_t^\infty \exp \left\{ -c'n_{\min}^* \min \left(x, x^{\frac{1}{2}} \right) \right\} dx + 2M^2k^{-2\alpha} \\ &\leq 18t + C \left\lceil \frac{p}{k} \right\rceil (49)^k \frac{1}{n_{\min}^*} \exp(-c'n_{\min}^* t) + 2M^2k^{-2\alpha}, \end{aligned} \quad (49)$$

then (49) yields

$$E\|\hat{\Sigma}^{bt} - \Sigma\|^2 \leq C \left(\frac{k + \ln p}{n_{\min}^*} + k^{-2\alpha} \right), \quad (50)$$

where C only depends on τ, M, M_0 . We can finally finish the proof of Theorem 2.1 by taking $k = (n_{\min}^*)^{1/(2\alpha+1)}$. \square

Acknowledgments

We thank Tianxi Cai for the ovarian cancer dataset and for helpful discussions. We also thank the Editor, the Associate editor, one referee and Zoe Russek for useful comments which have helped to improve the presentation of the paper.

The research of Tony Cai and Anru Zhang was supported in part by NSF Grant DMS-1208982 and DMS-1403708, and NIH Grant R01 CA127334.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <http://dx.doi.org/10.1016/j.jmva.2016.05.002>.

References

- [1] B. Andreopoulos, D. Anastassiou, Integrated analysis reveals hsa-mir-142 as a representative of a lymphocyte-specific gene expression and methylation signature, *Cancer Inform.* 11 (2012) 61–75.
- [2] P.J. Bickel, E. Levina, Covariance regularization by thresholding, *Ann. Statist.* 36 (2008) 2577–2604.
- [3] P.J. Bickel, E. Levina, Regularized estimation of large covariance matrices, *Ann. Statist.* 36 (2008) 199–227.
- [4] T. Bonome, J.-Y. Lee, D.-C. Park, M. Radonovich, C. Pise-Masison, J. Brady, G.J. Gardner, K. Hao, W.H. Wong, J.C. Barrett, et al., Expression profiling of serous low malignant potential, low-grade, and high-grade tumors of the ovary, *Cancer Res.* 65 (2005) 10602–10612.
- [5] T.T. Cai, W. Liu, Adaptive thresholding for sparse covariance matrix estimation, *J. Amer. Statist. Assoc.* 106 (2011) 672–684.
- [6] T.T. Cai, Z. Ma, Y. Wu, Optimal estimation and rank detection for sparse spiked covariance matrices, *Probab. Theory Related Fields* 161 (2015) 781–815.
- [7] T.T. Cai, Z. Ren, H.H. Zhou, Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation, *Electron. J. Stat.* 10 (2016) 1–59.
- [8] T.T. Cai, M. Yuan, Adaptive covariance matrix estimation through block thresholding, *Ann. Statist.* 40 (2012) 2014–2042.
- [9] T.T. Cai, C.-H. Zhang, H. Zhou, Optimal rates of convergence for covariance matrix estimation, *Ann. Statist.* 38 (2010) 2118–2144.
- [10] T.T. Cai, H. Zhou, Optimal rates of convergence for sparse covariance matrix estimation, *Ann. Statist.* 40 (2012) 2389–2420.
- [11] Cancer Genome Atlas Research Network, Integrated genomic analyses of ovarian carcinoma, *Nature* 474 (2011) 609–615.
- [12] G. Carraro, A. Shrestha, J. Rostkovius, A. Contreras, C.-M. Chao, E. El Agha, B. MacKenzie, S. Dilai, D. Guidolin, M.M. Taketo, et al., mir-142-3p balances proliferation and differentiation of mesenchymal cells during lung development, *Development* 141 (6) (2014) 1272–1281.
- [13] S.X. Chen, L.X. Zhang, P.S. Zhong, Tests for high-dimensional covariance matrices, *J. Amer. Statist. Assoc.* 105 (2010) 810–819.
- [14] J.C. Denny, M.D. Ritchie, M.A. Basford, J.M. Pulley, L. Bastarache, K. Brown-Gentry, D. Wang, D.R. Masys, D.M. Roden, D.C. Crawford, Phewas: demonstrating the feasibility of a phenome-wide scan to discover gene-disease associations, *Bioinformatics* 26 (2010) 1205–1210.
- [15] H.K. Dressman, A. Berchuck, G. Chan, J. Zhai, A. Bild, R. Sayer, J. Cragun, J. Clarke, R.S. Whitaker, L. Li, J. Gray, J. Marks, G.S. Ginsburg, A. Potti, M. West, J.R. Nevins, J.M. Lancaster, An integrated genomic-based approach to individualized treatment of patients with advanced-stage ovarian cancer, *J. Clin. Oncol.* 25 (2007) 517–525.
- [16] N. El Karoui, Operator norm consistent estimation of large-dimensional sparse covariance matrices, *Ann. Statist.* 36 (2008) 2717–2756.
- [17] R.D. Hawkins, G.C. Hon, B. Ren, Next-generation genomics: an integrative approach, *Nat. Rev. Genet.* 11 (2010) 476–486.
- [18] J.G. Ibrahim, G. Molenberghs, Missing data methods in longitudinal studies: a review, *TEST* 18 (2009) 1–43.
- [19] S.Y. Ko, N. Barengo, A. Ladanyi, J.S. Lee, F. Marini, E. Lengyel, H. Naora, Hoxa9 promotes ovarian cancer growth by stimulating cancer-associated fibroblasts, *J. Clin. Invest.* 122 (2012) 3603–3617.
- [20] R.J.A. Little, D.B. Rubin, *Statistical Analysis with Missing Data*, second ed., John Wiley & Sons, New York, 2002.

- [21] P.-L. Loh, M. Wainwright, High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity, *Ann. Statist.* 40 (2012) 1637–1664.
- [22] K. Lounici, Sparse Principal Component Analysis with Missing Observations. High Dimensional Probability VI, in: *Prog. Proba.*, vol. 66, Institute of Mathematical Statistics (IMS) Collections, 2013, pp. 327–356.
- [23] K. Lounici, High-dimensional covariance matrix estimation with missing observations, *Bernoulli* 20 (3) (2014) 1029–1058.
- [24] A.J. Rothman, E. Levina, J. Zhu, Generalized thresholding of large covariance matrices, *J. Amer. Statist. Assoc.* 104 (2009) 177–186.
- [25] M. Rudelson, R. Vershynin, Hanson-wright inequality and sub-gaussian concentration, *Electron. Comm. Probab.* 18 (2013) 1–9.
- [26] J.L. Schafer, *Analysis of Incomplete Multivariate Data*, CRC press, 2010.
- [27] R.W. Tothill, A.V. Tinker, J. George, R. Brown, S.B. Fox, S. Lade, D.S. Johnson, M.K. Trivett, D. Etemadmoghadam, B.e.a. Locandro, Novel molecular subtypes of serous and endometrioid ovarian cancer linked to clinical outcome, *Clin. Cancer Res.* 14 (2008) 5198–5208.