

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

Learning Gaussian graphical models with latent confounders

Permalink

<https://escholarship.org/uc/item/6dj1z2xf>

Authors

Wang, Ke
Franks, Alexander
Oh, Sang-Yun

Publication Date

2023-11-01

DOI

10.1016/j.jmva.2023.105213

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Learning Gaussian graphical models with latent confounders

Ke Wang* Alexander Franks Sang-Yun Oh

Department of Statistics and Applied Probability, University of California Santa Barbara

Abstract

Gaussian Graphical models (GGM) are widely used to estimate network structure in domains ranging from biology to finance. In practice, data is often corrupted by latent confounders which biases inference of the underlying true graphical structure. In this paper, we compare and contrast two strategies for inference in graphical models with latent confounders: Gaussian graphical models with latent variables (LVGGM) and PCA-based removal of confounding (PCA+GGM). While these two approaches have similar goals, they are motivated by different assumptions about confounding. In this paper, we explore the connection between these two approaches and propose a new method, which combines the strengths of these two approaches. We prove the consistency and convergence rate for the PCA-based method and use these results to provide guidance about when to use each method. We demonstrate the effectiveness of our methodology using both simulations and two real-world applications.

1 Introduction

In many domains, it is useful to characterize relationships between features using network models. For example, networks have been used to identify transcriptional patterns and regulatory relationships in genetic networks and applied as a way to characterize functional brain connectivity and cognitive disorders Fox and Raichle (2007); Van Dijk et al. (2012); Barch et al. (2013); Price et al. (2014). One of the most common methods for inferring a network from observations is the Gaussian graphical model (GGM). A GGM is defined with respect to a graph, in which the nodes correspond to joint Gaussian random variables and the edges correspond to the conditional dependencies among pairs of variables. A key property of the GGM is that the presence or absence of edges can be obtained from the precision matrix for multivariate Gaussian random variables Lauritzen (1996). Similar to LASSO regression Tibshirani (1996), we can infer the sparse graph structure via sparse precision matrix estimation with l_1 -regularized maximum likelihood estimation. This family of approaches is called graphical lasso (Glasso) Friedman et al. (2008); Yuan and Lin (2007).

In practice, however, network inference may be complicated due to the presence of latent confounders. For example, when characterizing relationships between the stock prices of publicly trade companies, the existence of overall market and sector factors induces extra correlation between stocks Choi et al. (2011), which can obscure the underlying network structure between companies.

We focus on estimating $\Omega = \Sigma^{-1}$, the precision matrix encoding the graph structure of interest Yuan and Lin (2007); Friedman et al. (2008); Cai et al. (2011); Parsana et al. (2019). When latent confounders are present, the covariance matrix for the observed data, Σ_{obs} can be expressed as

$$\Sigma_{obs} = \Sigma + L_{\Sigma}, \quad (1)$$

where the positive semidefinite matrix L_{Σ} reflects the effect of latent confounders. One approach, which we will call PCA+GGM, is motivated by confounders that affect the marginal correlation between observed variables Parsana et al. (2019) and uses principal component analysis (PCA) as a preprocessing step to remove the effect of these confounders Jolliffe (2005); Bartholomew et al. (2011).

*kewang01@ucsb.edu

PCA removes the leading eigenvectors from Σ_{obs} which are assumed to be L_{Σ} , then a second stage of standard GGM inference follows. PCA+GGM has shown to be useful in estimating gene co-expression networks, where correlated measurement noise and batch effects induce large extraneous marginal correlations between observed variables Geng et al. (2018); Leek and Storey (2007); Stegle et al. (2011); Gagnon-Bartsch et al. (2013); Freytag et al. (2015); Jacob et al. (2016).

Alternatively, equation (1) can be reparametrized as the observed precision matrix by applying the Sherman-Morrison identity Horn and Johnson (2012) as,

$$\Omega_{obs} = \Sigma_{obs}^{-1} = \Omega - L_{\Omega}, \quad (2)$$

where L_{Ω} again reflects the effect of unobserved confounding, i.e., unobserved nodes in a graph Chandrasekaran et al. (2011). One such approach, known as latent variable Gaussian Graphical Models (LVGGM), uses parameterization (2) and involves joint inference for Ω and L_{Ω} . The motivation behind LVGGM is to address the effect of unobserved variables in the complete data graph, which affect the partial correlations of the variables in the observed precision matrix Ω . This perspective can be particularly useful when the unobserved variables would have been included in the graph, had they been observed.

In previous work, either parameterization (1), e.g., PCA+GGM, or (2), e.g., LVGGM, has been used, depending on the source of confounding and motivations as described earlier. Typically, LVGGM is appropriate when confounding is induced by unobserved nodes in a complete data graph of interest, whereas PCA+GGM is more appropriate when confounding corresponds to nuisance variables, e.g., from batch effects.

In practice, the selection between these two methods will depend on user’s belief about the type of confounding present in the observed data. In this paper, our goal is to explore a way to address the effect of confounders in order to obtain the graph structure encoded in Ω without making such selection.

To achieve this goal, we generalize two seemingly different methods, PCA+GGM and LVGGM, into a common framework for addressing the effect of L_{Σ} in order to obtain the graph structure encoded in $\Omega = \Sigma^{-1}$. Based on the generalization, We propose a new method, PCA+LVGGM, to address two different sources of confounding. The combined approach is more general, since PCA+LVGGM contains both LVGGM and PCA+GGM as special cases. To our knowledge, the two methods of addressing confounding have not been discussed together in the literature.

In summary, in this paper,

- we carefully compare PCA+GGM and LVGGM, and illustrate the connection and difference between these two methods. We first theoretically characterize the performance of PCA+GGM. Different from Parsana et al. (2019) who derives asymptotic results, we provide a non-asymptotic convergence result for the performance of PCA+GGM. We observe that the performance of PCA+GGM are largely determined by the spectral structure of Σ and L_{Σ} .
- we propose PCA+LVGGM, which combines elements of PCA+GGM and LVGGM. In simulation, PCA+LVGGM can outperform PCA+GGM or LVGGM when the data is corrupted by multiple confounders. We perform extensive numerical experiments to validate the theory, compare the performance of the three methods, and demonstrate the utility of our approach in two applications.

The remainder of this paper is organized as follows: In Section 2, we introduce the problem definition for GGM, LVGGM and PCA+GGM followed by a brief literature review. Next, we introduce our hybrid method, PCA+LVGGM, and present a novel theoretical results for PCA+GGM in Section 3. We use these result to analyze the similarities and differences between LVGGM and PCA+GGM. In Section 4, we compare the utility of the various approaches in the simulation setting. Finally, in Section 5 we apply the methods on two real world data sets.

We introduce some general notation used throughout the rest of the paper. For a vector $\mathbf{v} = [v_1, \dots, v_p]^T$, define $\|\mathbf{v}\|_2 = \sqrt{\sum_{i=1}^p v_i^2}$, $\|\mathbf{v}\|_1 = \sum_{i=1}^p |v_i|$ and $\|\mathbf{v}\|_{\infty} = \max_i |v_i|$. For a matrix \mathbf{M} , let M_{ij} be its

(i, j) -th entry. Define the Frobenius norm $\|\mathbf{M}\|_F = \sqrt{\sum_i \sum_j M_{ij}^2}$, the element-wise ℓ_1 -norm $\|\mathbf{M}\|_1 = \sum_i \sum_j |M_{ij}|$ and $\|\mathbf{M}\|_\infty = \max_{(i,j)} |M_{ij}|$. We also define the spectral norm $\|\mathbf{M}\|_2 = \sup_{\|\mathbf{v}\|_2 \leq 1} \|\mathbf{M}\mathbf{v}\|_2$ and $\|\mathbf{M}\|_{L_1} = \max_j \sum_i |M_{ij}|$. The nuclear norm $\|\mathbf{M}\|_*$ is defined as the sum of the singular values of \mathbf{M} . When $\mathbf{M} \in \mathbb{R}^{p \times p}$ is symmetric, its eigendecomposition is $\mathbf{M} = \sum_{i=1}^p \lambda_i \mathbf{v}_i \mathbf{v}_i^\top$, where λ_i is the i -th eigenvalue of \mathbf{M} , and \mathbf{v}_i is the i -th eigenvector. We assume that $\lambda_1 \geq \dots \geq \lambda_p$. We call $\lambda_i \mathbf{v}_i \mathbf{v}_i^\top$ the i -th eigencomponent of \mathbf{M} .

2 Problem setup and review

2.1 Gaussian graphical models

Consider a p -dimensional random vector $\mathbf{X} = [X_1, \dots, X_p]^\top$ with covariance matrix Σ and precision matrix Ω . Let $G = (V, E)$ be the graph associated with \mathbf{X} , where V is the set of nodes (or vertices) corresponding to the elements of \mathbf{X} , and E is the set of edges connecting nodes. The graph shows the conditional independence relations between elements of \mathbf{X} . For any pair of connected nodes, the corresponding pairs of variables in \mathbf{X} are conditionally independent given the rest variables, i.e., $X_i \perp\!\!\!\perp X_j | X_{\setminus i,j}$, for all $(i, j) \notin E$. If \mathbf{X} is multivariate Gaussian, then X_i and X_j are conditionally independent given other variables if and only if $\Omega_{ij} = 0$, and thus the graph structure can be recovered from the precision matrix of \mathbf{X} .

Without loss of generality, we assume the variable \mathbf{X} has mean zero in this paper. Assuming that the graph is sparse, given a random sample $\{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(n)}\}$ following the distribution of \mathbf{X} , the Glasso estimate $\hat{\Omega}_{Glasso}$ Yuan and Lin (2007); Friedman et al. (2008) is obtained by solving the following log-likelihood based ℓ_1 -regularized function:

$$\underset{\Omega \succ 0}{\text{minimize}} \quad \text{Tr}(\Omega \Sigma_n) - \text{ln det}(\Omega) + \lambda \|\Omega\|_1, \quad (3)$$

where Tr denotes the trace of a matrix and $\Sigma_n = (1/n) \sum_{k=1}^n \mathbf{X}^{(k)} \mathbf{X}^{(k)\top}$ is the sample covariance matrix. Many alternative objective functions for sparse precision matrix estimation have been proposed Cai et al. (2011); Meinshausen et al. (2006); Peng et al. (2009); Khare et al. (2015). The behavior and convergence rates of these approaches are well studied Bickel et al. (2008b,a); Rothman et al. (2008); Lam and Fan (2009); Ravikumar et al. (2011); Cai et al. (2016).

In presence of latent confounders, Glasso and other GGM methods would likely recover a more dense precision matrix owing to spurious partial correlations introduced between observed variables. In other words, even when the underlying graph is sparse conditioned on the latent variables, the observed graph is dense marginally.

2.2 Latent variable Gaussian graphical models

One method for controlling the effects of confounders is the Latent Variable Gaussian Graphical Model (LVGGM) approach first proposed by Chandrasekaran et al. (2012). They assume that the number of latent factors is small compared to the number of observed variables, and that the conditional dependencies among the observed variables conditional on the latent factors is sparse. Consider a $(p+r)$ dimensional mean-zero normal random variable $\mathbf{X} = [\mathbf{X}_O, \mathbf{X}_H]^\top$, where $\mathbf{X}_O \in \mathbb{R}^p$ is observed and $\mathbf{X}_H \in \mathbb{R}^r$ is latent. Let \mathbf{X} have precision matrix $\Omega \in \mathbb{R}^{(p+r) \times (p+r)}$, and the submatrices $\Omega_O \in \mathbb{R}^{p \times p}$, $\Omega_H \in \mathbb{R}^{r \times r}$ and $\Omega_{O,H} \in \mathbb{R}^{p \times r}$ specify the dependencies between observed variables, between latent variables and between the observed and latent variables respectively. By Schur complement, the inverse of the observed covariance matrix satisfies:

$$\Omega_{obs} = \Sigma_{obs}^{-1} = \Omega_O - \Omega_{O,H} \Omega_H^{-1} \Omega_{O,H}^\top = \Omega - \mathbf{L}\Omega. \quad (4)$$

where $\Omega = \Omega_O$ encodes the conditional independence relations of interest and is sparse by assumption. $\mathbf{L}\Omega = \Omega_{O,H} \Omega_H^{-1} \Omega_{O,H}^\top$ reflects the low-rank effect of latent variables \mathbf{X}_H . Based on this sparse plus

low-rank decomposition, Chandrasekaran et al. (2012) proposed the following problem:

$$\begin{aligned} & \underset{\mathbf{\Omega}, \mathbf{L}_{\Omega}}{\text{minimize}} && -\ell(\mathbf{\Omega} - \mathbf{L}_{\Omega}; \mathbf{\Sigma}_n) + \lambda \|\mathbf{\Omega}\|_1 + \gamma \|\mathbf{L}_{\Omega}\|_* \\ & \text{subject to} && \mathbf{L}_{\Omega} \succeq \mathbf{0}, \quad \mathbf{\Omega} - \mathbf{L}_{\Omega} \succ \mathbf{0}, \end{aligned} \quad (5)$$

where $\mathbf{\Sigma}_n$ is the observed sample covariance matrix and $\ell(\mathbf{\Omega}, \mathbf{\Sigma}) = \ln(\det(\mathbf{\Omega})) - \text{Tr}(\mathbf{\Omega}\mathbf{\Sigma})$ is the Gaussian log-likelihood function. The ℓ_1 -norm encourages sparsity on $\mathbf{\Omega}$ and the nuclear norm encourages low-rank structure on \mathbf{L}_{Ω} .

The sparse plus low-rank decomposition is ill-posed if \mathbf{L}_{Ω} is not dense. If \mathbf{L}_{Ω} is sparse, then it is indistinguishable from $\mathbf{\Omega}$, that is, the sparse plus low-rank decomposition works well only when the sparse component is not low-rank and the low-rank component is not sparse Chandrasekaran et al. (2011). In practice, \mathbf{L}_{Ω} is dense if the latent variables have widespread effects.. Identifiability of $\mathbf{\Omega}$ coincides with the incoherence condition in the matrix completion problem Candès and Recht (2009) which requires that $|\mathbf{v}_k^T \mathbf{e}_i|$ is small for all $k \in \{1, \dots, r\}$ and $i \in \{1, \dots, p\}$ where \mathbf{v}_k is the k -th eigenvector of \mathbf{L}_{Ω} and \mathbf{e}_i is the i -th standard basis vector. More analysis on LVGGM can be found in Agarwal et al. (2012); Meng et al. (2014).

Finally, Ren and Zhou (2012) shows that the standard GGM approaches can still recover $\mathbf{\Omega}$ in the presence of latent confounding as long as the spectral norm of the low-rank component is sufficiently small compared to that of $\mathbf{\Sigma}$. This is also verified in our simulations.

2.3 PCA+GGM

Unlike LVGGM, which involves a decomposition of the observed data precision matrix, PCA+GGM involves a decomposition of the observed data covariance matrix:

$$\mathbf{\Sigma}_{obs} = \mathbf{\Omega}_{obs}^{-1} = (\mathbf{\Omega} - \mathbf{L}_{\Omega})^{-1} = \mathbf{\Omega}^{-1} + \mathbf{L}_{\Sigma}. \quad (6)$$

Motivated by confounding from measurement error and batch effects, Parsana et al. (2019) proposed the principal components correction (PC-correction) for removing \mathbf{L}_{Σ} . Consider observed data \mathbf{X}_{obs} , such that

$$\mathbf{X}_{obs} = \mathbf{X} + \mathbf{AZ}, \quad (7)$$

where $\mathbf{X} \sim N(\mathbf{0}, \mathbf{\Sigma})$ and $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{I}_r)$. Matrix $\mathbf{A} \in \mathbb{R}^{p \times r}$ is non-random so that $\mathbf{L}_{\Sigma} = \mathbf{AA}^T$. In general, additional structural assumptions are needed to distinguish \mathbf{L}_{Σ} from $\mathbf{\Sigma}$. As we will discuss in Section 3, one of our contributions is to show that if the spectral norm of \mathbf{L}_{Σ} is large relative to that of $\mathbf{\Sigma}$, then under mild conditions, \mathbf{L}_{Σ} is close to the sum of the first few eigencomponents of $\mathbf{\Sigma}_{obs}$. Therefore, one can remove the first r eigencomponents from $\mathbf{\Sigma}_{obs}$ Parsana et al. (2019). This PCA+GGM method is described in Procedure 1. Note that the number of principal components needs to be determined a priori, which we discuss in subsequent sections.

2.4 Combining PCA+GGM and LVGGM

As previously mentioned, while LVGGM and PCA+GGM solve the same problem, they are motivated by different sources of confounding. In applications, the observed data may be corrupted by multiple sources of confounding, and thus elements from both methods are needed. For example, in the biological application discussed in Section 5.1, both batch effects and unmeasured biological variables likely confound estimates of graph structure between observed variables. This motivates us to propose the PCA+LVGGM strategy described below.

As (4) illustrated, the observed precision matrix $\mathbf{\Omega}'$ may have been corrupted by a latent factor \mathbf{L}_{Ω} :

$$\mathbf{\Omega}' = \mathbf{\Omega} - \mathbf{L}_{\Omega}. \quad (8)$$

Now, rewriting (8) in terms of $\mathbf{\Sigma} = \mathbf{\Omega}^{-1}$ and $\mathbf{\Sigma}' = \mathbf{\Omega}'^{-1}$, applying the Sherman-Morrison identity on $\mathbf{\Omega}'$ gives,

$$\mathbf{\Sigma}' = \mathbf{\Sigma} + \mathbf{L}'_{\Omega}, \quad (9)$$

Procedure 1 PCA+GGM

Input: Sample covariance matrix, $\hat{\Sigma}_{obs} = (1/n) \sum_{k=1}^n \mathbf{X}_{obs}^{(k)} \mathbf{X}_{obs}^{(k)\top}$; rank of $\hat{\mathbf{L}}_{\Sigma}$, r

Output: Precision matrix estimate, $\hat{\Omega}$

1: Estimate $\hat{\mathbf{L}}_{\Sigma}$ from eigencomponents:

$$\hat{\Sigma}_{obs} = \sum_{i=1}^p \hat{\lambda}_i \hat{\theta}_i \hat{\theta}_i^{\top}, \quad \hat{\mathbf{L}}_{\Sigma} = \sum_{i=1}^r \hat{\lambda}_i \hat{\theta}_i \hat{\theta}_i^{\top}$$

2: Remove $\hat{\mathbf{L}}_{\Sigma}$:

$$\hat{\Sigma} = \hat{\Sigma}_{obs} - \hat{\mathbf{L}}_{\Sigma}.$$

3: Using $\hat{\Sigma}$, compute $\hat{\Omega}$ as solution to (3)

Procedure 2 PCA+LVGGM

Input: Sample covariance matrix, $\hat{\Sigma}_{obs}$; rank of $\hat{\mathbf{L}}_{\Sigma}$, r_P ; rank of $\hat{\mathbf{L}}_{\Omega}$, r_L

Output: Precision matrix estimate, $\hat{\Omega}$

1: Estimate $\hat{\mathbf{L}}_{\Sigma}$ from eigencomponents:

$$\hat{\Sigma}_{obs} = \sum_{i=1}^p \hat{\lambda}_i \hat{\theta}_i \hat{\theta}_i^{\top}, \quad \hat{\mathbf{L}}_{\Sigma} = \sum_{i=1}^{r_P} \hat{\lambda}_i \hat{\theta}_i \hat{\theta}_i^{\top}$$

2: Remove $\hat{\mathbf{L}}_{\Sigma}$:

$$\hat{\Sigma} = \hat{\Sigma}_{obs} - \hat{\mathbf{L}}_{\Sigma}.$$

3: Using $\hat{\Sigma}$, compute $\hat{\Omega}$ as solution to (5) with γ such that $\text{rank}(\hat{\mathbf{L}}_{\Omega}) = r_L$

where \mathbf{L}'_{Ω} is still a low-rank matrix. If Σ' is further corrupted by an additive latent factor represented by \mathbf{L}_{Σ} , the following equation described the observed matrix Σ_{obs} :

$$\Sigma_{obs} = \Sigma' + \mathbf{L}_{\Sigma} = \Sigma + \mathbf{L}'_{\Omega} + \mathbf{L}_{\Sigma} \quad (10)$$

In the above example, following our theoretical analysis in Section 3, if the spectral norm of \mathbf{L}_{Σ} is much larger than that of Σ and \mathbf{L}'_{Ω} , then removing \mathbf{L}_{Σ} using the PC-correction is likely to be effective. If the spectral norm of \mathbf{L}'_{Ω} is not much larger than that of Σ , then PC-correction is not a good choice to remove \mathbf{L}'_{Ω} . If \mathbf{L}'_{Ω} is dense, then Ω and \mathbf{L}_{Ω} can be well estimated by LVGGM. In (10), the overall confounding $\mathbf{L}'_{\Omega} + \mathbf{L}_{\Sigma}$ is the sum of two low-rank components with different norms, we can consider using both methods: first remove \mathbf{L}_{Σ} via eigendecomposition, then apply LVGGM to estimate Ω and \mathbf{L}_{Ω} . We call this procedure PCA+LVGGM and it is shown in Procedure 2. We discuss methods for setting the ranks for \mathbf{L}_{Σ} (defined as r_P) and \mathbf{L}'_{Ω} (defined as r_L) in Section 3.5.

3 Theoretical analysis and model comparisons

In this section, we investigate the theoretical properties of PCA+GGM. Our results reveal precisely how the eigenstructure of the observed covariance matrix affects the the performance of PCA+GGM. The theoretical analysis provides practical insights into when each graph estimation method should (or should not) be applied. Specifically, we derive the convergence rate of PCA+GGM and compare it to that of LVGGM. As shown in theoretical analysis by Parsana et al. (2019), the low-rank confounder can be well estimated by PC-correction if the number of features $p \rightarrow \infty$ with the number of observations n fixed. We provide a non-asymptotic analysis depending on p and n and our result shows that the graph can be recovered exactly when $n \rightarrow \infty$ with fixed p . When additional assumptions are satisfied, e.g. spiky covariance structure and incoherent eigenvectors, the convergence rate can be improved to $O(\sqrt{\ln p/n})$.

3.1 Convergence analysis on PCA+GGM

Without loss of generality, we consider the case of a rank-one confounder. Assume that we have a random sample of p -dimensional random vectors:

$$\mathbf{X}_{obs}^{(i)} = \mathbf{X}^{(i)} + \sigma \mathbf{v} \mathbf{Z}^{(i)}, \quad i \in \{1, \dots, n\}, \quad (11)$$

where $Cov(\mathbf{X}^{(i)}) = \Sigma$ and $\mathbf{Z}^{(i)}$ is a univariate standard normal random variable. $\mathbf{v} \in \mathbb{R}^p$ is a non-random vector with unit norm, and σ is a non-negative scalar constant. Without loss of generality, we

assume that $\mathbf{X}^{(i)} \perp\!\!\!\perp \mathbf{Z}^{(i)}$. To see how \mathbf{v} affects estimation, we assume that \mathbf{v} is the k -th eigenvector of Σ . The discussion on general \mathbf{v} is deferred to Section 3.2. Therefore, the covariance matrix of $\mathbf{X}_{obs}^{(i)}$ is:

$$\Sigma_{obs} = \Sigma + \sigma^2 \mathbf{v} \mathbf{v}^T = \Sigma_{-k} + (\lambda_k(\Sigma) + \sigma^2) \mathbf{v} \mathbf{v}^T, \quad (12)$$

where Σ_{-k} is the matrix Σ without the k -th eigencomponent and $\lambda_k(\Sigma)$ is the k -th eigenvalue of Σ . When $\sigma^2 > \lambda_1(\Sigma)$, $\lambda_k(\Sigma) + \sigma^2$ becomes the first eigenvalue of Σ_{obs} , and \mathbf{v} is the corresponding first eigenvector. We remove the first principal component from the sample covariance matrix $\hat{\Sigma}_{obs}$:

$$\hat{\Sigma} = \hat{\Sigma}_{obs} - \hat{\lambda}_1 \hat{\theta}_1 \hat{\theta}_1^T, \quad (13)$$

where $\hat{\lambda}_1$ is the first eigenvalue of $\hat{\Sigma}_{obs}$ and $\hat{\theta}_1$ is the first eigenvector of $\hat{\Sigma}_{obs}$. Then we use $\hat{\Sigma}$ to estimate Ω . We first show that under mild conditions, $\hat{\Sigma}$ is close to Σ . Following (Bickel et al., 2008b, 3.1), we assume that Σ_{obs} is well-conditioned such that:

$$0 < \epsilon \leq \lambda_p(\Sigma_{obs}) \leq \lambda_1(\Sigma_{obs}) \leq \frac{1}{\epsilon}, \quad (14)$$

where ϵ is independent of p .

Theorem 1. *Let λ_i be the i -th eigenvalue of Σ_{obs} , θ_i be the i -th eigenvector of Σ_{obs} , and $\nu = \lambda_1 - \lambda_2$ be the eigengap of Σ_{obs} . Suppose Σ_{obs} satisfies condition (14) and $\mathbf{X}_{obs}^{(i)}$ is generated as (11). Further assume that $\sigma^2 > \lambda_1(\Sigma)$. Suppose $n \geq p$ and $\|\Sigma\|_2 \sqrt{(\nu+1)/\nu^2} \sqrt{p/n} \leq 1/128$, then:*

$$\|\hat{\Sigma} - \Sigma\|_\infty \leq C_1 \sqrt{\frac{\ln p}{n}} + C_2 \sqrt{\frac{\nu+1}{\nu^2}} \sqrt{\frac{p}{n}} + C_3 \sqrt{\frac{p}{n}} + \lambda_k(\Sigma) \|\theta_1 \theta_1^T\|_\infty,$$

with probability greater than $1 - C_4/p$ for constants C_i 's > 1 .

Proof: Recall when $\sigma^2 > \lambda_1(\Sigma)$, $\lambda_k(\Sigma) + \sigma^2$ becomes the first eigenvalue of Σ_{obs} , and \mathbf{v} becomes its first eigenvector. By (12) and (13),

$$\begin{aligned} \hat{\Sigma} - \Sigma &= (\hat{\Sigma}_{obs} - \hat{\lambda}_1 \hat{\theta}_1 \hat{\theta}_1^T) - (\Sigma_{obs} - \lambda_1 \theta_1 \theta_1^T + \lambda_k(\Sigma) \theta_1 \theta_1^T) \\ &= (\hat{\Sigma}_{obs} - \Sigma_{obs}) + (\lambda_1 \theta_1 \theta_1^T - \hat{\lambda}_1 \hat{\theta}_1 \hat{\theta}_1^T) - \lambda_k(\Sigma) \theta_1 \theta_1^T. \end{aligned}$$

At a high level, we bound $\|\hat{\Sigma} - \Sigma\|_\infty$ by bounding the norms of $\Sigma_{obs} - \hat{\Sigma}_{obs}$, $\lambda_1 - \hat{\lambda}_1$ and $\theta_1 - \hat{\theta}_1$. The details of the complete proof is in Appendix A. \square

The bound in Theorem 1 can be further simplified as $C_s \sqrt{p/n} + \lambda_k(\Sigma) \|\theta_1 \theta_1^T\|_\infty$ for some large constant C_s . We express it in the above form because it provides more insight on how each term affects the result. Now we analyze the bound in Theorem 1 in detail.

The error bound in Theorem 1 depends on the largest eigenvalue of Σ_{obs} , the eigengap $\nu = \lambda_1(\Sigma_{obs}) - \lambda_2(\Sigma_{obs})$, the eigenvector of the confounder and n and p . The term $\sqrt{(\nu+1)/\nu^2}$ shows that if the eigengap ν is larger, the estimation error bound will be smaller. Recall that when $\sigma^2 > \lambda_1(\Sigma)$, $\lambda_k(\Sigma) + \sigma^2$ becomes the first eigenvalue of Σ_{obs} . Hence if $\sigma^2 \gg \lambda_1(\Sigma)$, then the eigengap ν is large. The fact that a larger eigengap leads to a better convergence rate is closely related to the concept of ‘‘effective dimension’’ (also known as ‘‘effective rank’’). The effective rank, $r(\mathbf{M})$, of any positive semidefinite matrix $\mathbf{M} \in \mathbb{R}^{p \times p}$, is defined as:

$$r(\mathbf{M}) := \frac{\text{Tr}(\mathbf{M})}{\lambda_1(\mathbf{M})} = \frac{\sum_{i=1}^p \lambda_i(\mathbf{M})}{\lambda_1(\mathbf{M})} \leq C, \quad (15)$$

where $C \geq 1$ can be viewed as the effective dimension of \mathbf{M} Meng et al. (2014); Koltchinskii and Lounici (2017); Wainwright (2019). \mathbf{M} is approximately low-rank if the first few eigenvalues are much larger than the rest, and $r(\mathbf{M})$ will be much smaller than the observed dimension p . In this case, we

can significantly reduce the magnitude of the dependence on $O(\sqrt{p/n})$ by replacing p with effective dimension C , in (15). We provide a sharper bound for matrices with small effective rank in Theorem 2.

Next, we reason about the last term in the error bound, $\lambda_k(\mathbf{\Sigma})\|\boldsymbol{\theta}_1\boldsymbol{\theta}_1^T\|_\infty$. In practice, in many sparse graphs inferred from real world data, the first few eigenvalues of $\mathbf{\Sigma}$ are much larger than the rest, i.e., $\lambda_1(\mathbf{\Sigma}) \gg \lambda_k(\mathbf{\Sigma})$ for large enough $k > 1$. This is also true for many common graph data generating models (see Appendix C). This means that if the eigenvector of the low-rank component is one of the first few eigenvectors of $\mathbf{\Sigma}$, then the error bound will be much larger. This result shows that the first few eigencomponents play a more important role in determining the structure of $\mathbf{\Sigma}$ and its inverse. Thus, the error of the PCA+GGM estimator will be large if those first few eigencomponents are removed by PC-correction.

Note that $\|\boldsymbol{\theta}_1\boldsymbol{\theta}_1^T\|_\infty$ is upper bounded by 1, since $\boldsymbol{\theta}_1$ is the eigenvector of some matrix, and thus has unit Euclidean norm; however, $\|\boldsymbol{\theta}_1\boldsymbol{\theta}_1^T\|_\infty$ can be much smaller than 1 when $\boldsymbol{\theta}_1$ is incoherent with standard basis, e.g. dense. One extreme case is when all the elements of $\boldsymbol{\theta}_1$ are $1/\sqrt{p}$, in which case $\|\boldsymbol{\theta}_1\boldsymbol{\theta}_1^T\|_\infty = 1/p$. This setup corresponds to a scenario in which the confounder has a widespread effect over all the p variables in the signal, which is in accordance with one requirement in LVGGM. LVGGM requires the low-rank component to be dense. For both PCA+GGM and LVGGM, more "widespread" confounding implies smaller estimation error. Based on these observations, we provide a tighter bound under small effective rank and incoherent $\boldsymbol{\theta}_1$.

Theorem 2. *Following the same notations and assumptions for $\mathbf{\Sigma}_{obs}$ in Theorem 1 and again assume that $\sigma^2 > \lambda_1(\mathbf{\Sigma})$. Further assume that the effective rank of $\mathbf{\Sigma}_{obs}$ (defined in (15)) $r(\mathbf{\Sigma}_{obs}) \asymp p/\lambda_1$, the eigengap ν satisfies $\lambda_1 \asymp \nu \gg p\sqrt{\ln p/n}$, and $\boldsymbol{\theta}_1$ is incoherent, i.e. $\|\boldsymbol{\theta}\|_\infty \leq C_1/\sqrt{p}$. Then:*

$$\|\hat{\mathbf{\Sigma}} - \mathbf{\Sigma}\|_\infty \leq C_2\left(\sqrt{\frac{\ln p}{n}} + \frac{1}{p}\right),$$

with probability greater than $1 - C_3/p$ for some C_i 's > 1 .

Proof: The complete proof is in Appendix A. □

After obtaining $\hat{\mathbf{\Sigma}}$, we can use Glasso, CLIME Cai et al. (2011) or any sparse GGM estimation approach to estimate $\mathbf{\Omega}$. We can have a good estimate of $\mathbf{\Omega}$ when $\|\hat{\mathbf{\Sigma}} - \mathbf{\Sigma}\|_\infty$ is small. With the same input $\hat{\mathbf{\Sigma}}$, the theoretical convergence rate of the estimate obtained from CLIME is of the same order as the Glasso estimate. The derivation of the error bound of Glasso requires the irrepresentability condition and restricted eigenvalue conditions (see Ravikumar et al. (2011)). Due to the length of the article, we only show the proof of the edge selection consistency for CLIME, meaning that for the theoretical analysis, we apply CLIME method after obtaining $\hat{\mathbf{\Sigma}}$.

The CLIME estimator $\hat{\mathbf{\Omega}}_1$ is obtained by solving:

$$\underset{\mathbf{\Omega}}{\text{minimize}} \quad \|\mathbf{\Omega}\|_1, \quad \text{subject to} \quad \|\hat{\mathbf{\Sigma}}\mathbf{\Omega} - \mathbf{I}\|_\infty \leq \lambda_n. \quad (16)$$

Since $\hat{\mathbf{\Omega}}_1$ might not be symmetric, we need the symmetrization step to obtain $\hat{\mathbf{\Omega}}$.

Following Cai et al. (2011), we assume that $\mathbf{\Omega}$ is in the following class:

$$U(s_0, M_0) = \{\mathbf{\Omega} = \omega_{ij} : \mathbf{\Omega} \succ \mathbf{0}, \|\mathbf{\Omega}\|_{L_1} < M_0, \max_{1 \leq i \leq p} \sum_{j=1}^p I_{\{\omega_{ij} \neq 0\}} \leq s_0(p)\}, \quad (17)$$

where we allow s_0 and M_0 to grow as p and n increase. With $\hat{\mathbf{\Sigma}}$ obtained from equation (13) as the input of (16), we have the following result.

Theorem 3. Suppose that assumptions in Theorem 1 hold, $\mathbf{\Omega} \in U(s_0, M_0)$, and λ_n is chosen as

$$M_0(C_1\sqrt{\frac{\ln p}{n}} + C_2\sqrt{\frac{p}{n}} + C_3\sqrt{\frac{\nu+1}{\nu^2}}\sqrt{\frac{p}{n}} + \lambda_k(\mathbf{\Sigma})\|\boldsymbol{\theta}_1\boldsymbol{\theta}_1^T\|_\infty),$$

then:

$$\|\mathbf{\Omega} - \hat{\mathbf{\Omega}}\|_\infty \leq 2M_0\lambda_n,$$

with probability greater than $1 - C_4/p$. C_i 's are defined the same as in Theorem 1.

When assumptions in Theorem 2 hold, $\mathbf{\Omega} \in U(s_0, M_0)$, and λ'_n is chosen as $M_0(C_5\sqrt{\ln p/n})$, then:

$$\|\mathbf{\Omega} - \hat{\mathbf{\Omega}}\|_\infty \leq 2M_1\lambda'_n,$$

with probability greater than $1 - C_6/p$ for C_i 's > 1 .

Proof: The main steps follow the proof of Theorem 6 in Cai et al. (2011). The complete proof is in Appendix A. \square

Therefore, if the minimum magnitude of $\mathbf{\Omega}$ is larger than the error bounds above, then we can have exact edge selection with high probability.

3.2 Generalizations

The analysis in previous sections assumes that the low-rank confounder has rank 1, is independent of \mathbf{X} and the eigenvector of the covariance of the \mathbf{L}_Σ is one of the eigenvectors of $\mathbf{\Sigma}$. We now comment on more general settings.

- *Higher rank:* For ease of interpretation, we assume that the confounder can be expressed as $\sum_{i=1}^r \sigma_i \mathbf{v}_i$. If $\min_i \sigma_i^2 > \lambda_1(\mathbf{\Sigma})$, then when running PCA+GGM, the low-rank component can be removed due to its larger norm compared to that of $\mathbf{\Sigma}$. According to Theorem 1, PCA+GGM can still perform well if \mathbf{v}_i 's are not the top eigenvectors of $\mathbf{\Sigma}$.
- *Arbitrary \mathbf{v} :* In Theorem 1, \mathbf{v} is assumed to be an eigenvector of $\mathbf{\Sigma}$, but the result in Theorem 1 provides insights about the bound with an arbitrary \mathbf{v} . Assume that \mathbf{v}_1 is the eigenvector of confounder 1 and \mathbf{v}_2 is the eigenvector of confounder 2. \mathbf{v}_1 is the i -th eigenvector of $\mathbf{\Sigma}$ and \mathbf{v}_2 is the j -th eigenvector of $\mathbf{\Sigma}$. In Theorem 1, we show that the bound for estimating $\mathbf{\Sigma}$ with confounder 1 is larger than the bound with confounder 2 if $i < j$ because \mathbf{v}_1 is associated with a larger eigenvalue. If a confounder's eigenvector \mathbf{v} is a linear combination of \mathbf{v}_1 and \mathbf{v}_2 , the bound should be between the bound with confounder 1 and the bound with confounder 2. Following this idea, we now show how our analysis works for an arbitrary \mathbf{v} in more detail. When the eigenvector of the low-rank component is not one of the eigenvectors of $\mathbf{\Sigma}$, we can express that vector using the eigenvectors of $\mathbf{\Sigma}$ as basis. For example, we assume that in (11), $\mathbf{v} = \sum_{i=1}^p a_i \boldsymbol{\theta}_i$, where $\boldsymbol{\theta}_i$ means the i -th eigenvector of $\mathbf{\Sigma}$. We say \mathbf{v} is closely aligned with $\boldsymbol{\theta}_1$ if $|a_1|$ is significantly large compared with other $|a_i|$'s. Equivalently, $|\mathbf{v}^T \boldsymbol{\theta}_1| \gg |\mathbf{v}^T \boldsymbol{\theta}_i|$ for $i \neq 1$, if \mathbf{v} is closely aligned with $\boldsymbol{\theta}_1$. In this case, the first eigencomponent of $\mathbf{\Sigma}$ will be removed, thus leading to a poor estimate of $\mathbf{\Sigma}$ using PC-correction. If the eigenvector of the low-rank component is not closely aligned with the first few eigenvectors of $\mathbf{\Sigma}$, then we won't lose too much useful information when removing the top principal components and PCA+GGM can still perform well. In general, \mathbf{v} can hardly be closely aligned with the first few eigenvectors of $\mathbf{\Sigma}$. We verify this with simulations in Section 4.2. We randomly generate multiple confounders and notice that their eigenvectors are not closely aligned with the top eigenvectors of $\mathbf{\Sigma}$. Thus, PCA-based methods are effective to remove the confounders.

3.3 Comparison with LVGGM

Now we compare LVGGM to PCA+GGM in more detail. We observe that PCA+GGM can be viewed as a supplement to LVGGM. The assumptions of PCA+GGM can be well satisfied when the assumptions of LVGGM cannot be satisfied. In (12), now let \mathbf{v} be the k -th eigenvector of $\mathbf{\Omega}$ (thus the $(p - k + 1)$ -th eigenvector of $\mathbf{\Sigma}$), the Sherman-Morrison identity gives

$$\mathbf{\Sigma}_{obs}^{-1} = \mathbf{\Omega} - \frac{\lambda_k(\mathbf{\Omega})^2}{\lambda_k(\mathbf{\Omega}) + (1/\sigma^2)} \mathbf{v}\mathbf{v}^T = \mathbf{\Omega} - \mathbf{L}_{\mathbf{\Omega}}. \quad (18)$$

We can see that as σ increases, $\lambda_k(\mathbf{\Omega})^2/(\lambda_k(\mathbf{\Omega}) + (1/\sigma^2))$ increases. In the simulations in Section 4, we observe that LVGGM performs poorly when \mathbf{v} is closely aligned with the first few eigenvectors of $\mathbf{\Omega}$ (thus the last few eigenvectors of $\mathbf{\Sigma}$). One way to interpret why LVGGM does not work well under this setting is because the nuclear norm penalty in LVGGM will shrink large eigenvalues. Specifically, when k is small and σ^2 is large, $\lambda_k(\mathbf{\Omega})^2/(\lambda_k(\mathbf{\Omega}) + (1/\sigma^2))$ is large. Therefore, the nuclear norm regularization in LVGGM introduces larger bias. Additionally, when k is small, \mathbf{v} is one of the top eigenvectors of $\mathbf{\Omega}$. We empirically observe that the top eigenvectors of $\mathbf{\Omega}$ can be coherent with standard basis. This observation coincides with the identifiability issue of LVGGM, i.e., a sparse low-rank component worsens the performance of LVGGM.

This observation is consistent with the conclusion in Agarwal et al. (2012). They impose a spikiness condition, which is a weaker condition than the incoherence condition in Chandrasekaran et al. (2012). The spikiness condition requires that $\|\mathbf{L}_{\mathbf{\Omega}}\|_{\infty}$ is not too large. (18) shows that $\mathbf{L}_{\mathbf{\Omega}}$ tends to have a larger spectral norm when \mathbf{v} is aligned with the first few eigenvectors of $\mathbf{\Omega}$ and σ is large, since in this case, $\lambda_k(\mathbf{\Omega})^2/(\lambda_k(\mathbf{\Omega}) + (1/\sigma^2))$ is close to $\lambda_1(\mathbf{\Omega})$. The large norm of $\mathbf{L}_{\mathbf{\Omega}}$ implies that the spikiness condition is not well satisfied, thus the error bound of LVGGM is large. Note, however, that the first few eigenvectors of $\mathbf{\Omega}$ are the last few eigenvectors of $\mathbf{\Sigma}$. Our analysis shows that the error bound of the estimate of PCA+GGM is small when \mathbf{v} is aligned with the first few eigenvectors of $\mathbf{\Omega}$ and σ is large.

3.4 PCA+LVGGM

In this section we discuss the PCA+LVGGM method briefly. We use the same formulation as (8) to (10). We claim that PCA+LVGGM outperforms using PCA+GGM or LVGGM individually when $\mathbf{L}_{\mathbf{\Sigma}}$'s spectral norm is large compared to that of $\mathbf{L}'_{\mathbf{\Omega}}$ and $\mathbf{\Sigma}$, $\mathbf{L}_{\mathbf{\Sigma}}$'s vectors are not aligned with the first few eigenvectors of $\mathbf{\Sigma}$, and the spectral norm of $\mathbf{L}'_{\mathbf{\Omega}}$ is not significantly larger than that of $\mathbf{\Sigma}$. This is because based on Theorem 1 and 3, PCA+GGM is effective only when the spectral norm of the low-rank confounding is larger than that of the signal. PC-correction, however, can only effectively remove $\mathbf{L}_{\mathbf{\Sigma}}$ but not $\mathbf{L}'_{\mathbf{\Omega}}$ because the norm of $\mathbf{L}'_{\mathbf{\Omega}}$ is not significantly larger than that of $\mathbf{\Sigma}$. In contrast, LVGGM can estimate $\mathbf{L}'_{\mathbf{\Omega}}$ well, but not $\mathbf{L}_{\mathbf{\Sigma}}$ because it has a larger spectral norm and its eigenvectors might be aligned with the first few eigenvectors of $\mathbf{\Omega}$.

3.5 Tuning parameter selection

In both LVGGM and PCA+GGM there are crucial tuning parameters to select. For LVGGM, recall that λ controls the sparsity of $\mathbf{\Omega}$ and γ controls the rank of $\mathbf{L}_{\mathbf{\Omega}}$. Chandrasekaran et al. (2012) argues that λ should be proportional to $\sqrt{p/n}$, the rate in the convergence analysis, and choose γ among a range of values that makes the graph structure of $\hat{\mathbf{\Omega}}$ stable (see Chandrasekaran et al., 2011, for more detail).

When using PCA+GGM, we need to determine the rank first (i.e., how many principal components should be removed). Leek and Storey (2007); Leek et al. (2012) suggest using the `sva` function from `Bioconductor`, which is based on parallel analysis Horn (1965); Buja and Eyuboglu (1992); Lim and Jahng (2019). Parallel analysis compares the eigenvalues of the sample correlation matrix to the eigenvalues of a random correlation matrix for which no factors are assumed. Given the number of principal components to remove, we can use model selection tools such as AIC, BIC or cross-validation to choose the sparsity parameter in Glasso. One may also decide how many principal components to

remove by considering the number of top eigenvalues of the observed covariance matrix (see Section 3.1 and Section 3.2). Note that these rank selection approaches perform well when the low-rank confounding has large enough spectral norm compared to the norm of signal (more details on these conditions are discussed in Leek et al. (2012); Lim and Jahng (2019)). We will see in later sections that these conditions can be satisfied in many real-world applications. When the spectral norm of the latent confounder is small, approaches which do not account for confounding, such as Glasso and CLIME, are actually robust enough to perform well even when confounding exists. This is theoretically proved by Ren and Zhou (2012) and our simulations in next section also confirm this.

The PCA+LVGGM method has three tuning parameters: the rank of \mathbf{L}_Σ , γ and λ . To start, we first look at eigenvalues or use the `sva` package to determine the total rank of the low-rank component, $\mathbf{L}_\Sigma + \mathbf{L}'_\Omega$. We think it is natural to determine the rank of confounder first because we will see in later applications, we can have some domain knowledge on the ranks of confounders, e.g. in finance applications, some financial theory suggests the number of latent variables in the market. We then need to partition the total rank between \mathbf{L}_Σ and \mathbf{L}'_Ω . If we determine that $\text{rank}(\mathbf{L}_\Sigma + \mathbf{L}'_\Omega) = k$, we look for an eigengap in the first k eigenvalues and allocate the largest $m < k$ eigenvalues for PC-removal. Our experiments in Section 5 show that domain knowledge can be used to motivate the number of components for PC-removal. After removing the principal components, we choose γ in LVGGM so that \mathbf{L}'_Ω is approximately rank $k - m$. We observe that when running LVGGM, the rank won't change for a range of λ values when using a fixed γ . Thus, it suffices to fix γ first to control the rank, then determine λ to control the sparsity.

Practically, network estimation is often used to help exploratory data analysis and hypothesis generation. For these purposes, model selection methods such as AIC, BIC or cross-validation may tend to choose models that are too dense Danaher et al. (2014). This fact can also be observed by our experiments. Therefore, we recommend that model selection should be based on prior knowledge and practical purposes, such as network interpretability and stability, or identification of important edges with low false discovery rate Meinshausen and Bühlmann (2010). Thus, we recommend that the selection of tuning parameters should be driven by applications. For example, for biological applications, the model should be biologically plausible, sufficiently complex to include important information and sparse enough to be interpretable. In this context, a robustness analysis can be used to explore how edges change over a range of tuning parameters.

4 Simulations

In this section, numerical experiments illustrate the utility of each sparse plus low rank method. In Section 4.1 we illustrate the behavior of Glasso, LVGGM and PCA+GGM under different assumptions about rank-one confounding. In Section 4.2, we show the efficacy of PCA+LVGGM in a variety of simulation scenarios. In all experiments, we set $p = 100$ and use the scale-free and random networks from `huge.generator` function in R package `huge` Zhao et al. (2012). To generate random networks, each pair of off-diagonal elements are randomly set, while the graph is generated using B-A algorithm under scale-free structures Albert and Barabási (2002). Due to space limit, we only include results on the scale-free structure.

4.1 The efficacy of LVGGM and PCA+GGM

We compare the relative performance of PCA+GGM, LVGGM and Glasso in the presence of a rank-one confounder, \mathbf{L} . Guided by our analysis in Section 3, we show that the relationship between \mathbf{L} and the eigenstructure of Σ determines the performance of these three methods. We first generate the data with

$$\mathbf{X}_{obs}^{(i)} = \mathbf{X}^{(i)} + \mathbf{L}^{(i)}, \quad \mathbf{L}^{(i)} = \sigma \mathbf{V} \mathbf{Z}^{(i)}, \quad i \in \{1, \dots, n\},$$

where $\mathbf{X}^{(i)} \in \mathbb{R}^p$ is normally distributed with mean zero and covariance matrix Σ . $\mathbf{Z}^{(i)} \in \mathbb{R}^r$, the low-rank confounder, follows a normal distribution with mean zero and identity covariance matrix. $\mathbf{V} \in \mathbb{R}^{p \times r}$ is a non-random semi-orthogonal matrix satisfying $\mathbf{V}^T \mathbf{V} = \mathbf{I}$, and $\sigma \in \mathbb{R}$ represents the

magnitude of the confounder. Without loss of generality, we assume that $\mathbf{X}^{(i)}$ and $\mathbf{Z}^{(i)}$ are independent. We illustrate the performance of different methods under various choices for, \mathbf{V} , the eigenvectors of \mathbf{L} .

We first set $r = 1$, $p = 100$ and $n = 200$. The largest eigenvalue of $\mathbf{\Sigma}$ is around 5. We use \mathbf{v}_i to denote the i -th eigenvector of $\mathbf{\Sigma}$. When examining the effect of σ , we choose the 95-th eigenvector of $\mathbf{\Sigma}$ as \mathbf{V} to ensure that \mathbf{V} is not closely aligned with the first few eigenvectors of $\mathbf{\Sigma}$. We then compare the cases with $\sigma^2 = 20$ and 3. Next, we examine the effect of eigenvectors. We fix σ^2 as 20, and use the i -th eigenvector of $\mathbf{\Sigma}$ as \mathbf{V} , where $i \in \{1, 60, 95\}$. Following previous notation, we use \mathbf{v}_1 , \mathbf{v}_{60} and \mathbf{v}_{95} as \mathbf{V} . 1 is chosen as the rank for PC-correction and LVGGM. We generate ROC curves (Hastie et al., 2009, 9.2.5) for each method based on 50 simulated samples and use the average to draw the ROC curves (Fig. 1). We truncate the ROC curves at FPR=0.2, since the estimates with large FPR are typically less useful in practice.




Figure 1: $n = 200$. The low-rank component has rank 1. σ^2 is the magnitude of the low-rank component, \mathbf{v}_i is the i -th eigenvector of $\mathbf{\Sigma}$. The first row illustrates the effect of σ : we use the 95-th eigenvector of $\mathbf{\Sigma}$ as \mathbf{V} , and set σ^2 to 20 and 3 from left to right. The second row illustrates the effect of \mathbf{V} : we fixed σ^2 as 20, and use the 60-th and first eigenvector of $\mathbf{\Sigma}$ as \mathbf{V} from left to right. PCA+GGM works the best when σ^2 is large and \mathbf{V} is not aligned with the first eigenvector of $\mathbf{\Sigma}$. LVGGM works the best when σ^2 is large and \mathbf{V} is not aligned with the last eigenvector of $\mathbf{\Sigma}$. When σ^2 is small, Glasso works as well as the other two.

From Fig. 1, we observe that when the confounder has large norm and its eigenvectors are not closely aligned with the first few eigenvectors of $\mathbf{\Sigma}$, PCA+GGM performs better than LVGGM and Glasso. LVGGM preforms the best when the confounder has large norm and its eigenvectors are not aligned with the last few eigenvectors of $\mathbf{\Sigma}$ (also the first few eigenvectors of $\mathbf{\Sigma}^{-1}$). When the low-rank component does not have a large norm, Glasso also performs well. This reaffirms the fact that Glasso can be robust enough to address the low-rank confounding with small norm.

4.2 The efficacy of PCA+LVGGM

In this section, we use multiple examples to demonstrate the efficacy of the PCA+LVGGM. We introduce corruption of the signal with two low-rank confounders. The data is generated as follows:

$$\mathbf{X}_{obs}^{(i)} = \mathbf{X}^{(i)} + \mathbf{V}_1 \mathbf{D}_1 \mathbf{Z}_1^{(i)} + \mathbf{V}_2 \mathbf{D}_2 \mathbf{Z}_2^{(i)}, \quad i \in \{1, \dots, n\},$$

where $\mathbf{X}_i \in \mathbb{R}^p$ is normally distributed with mean zero and covariance matrix Σ . $\mathbf{Z}_1^{(i)} \in \mathbb{R}^{d_1}$, corresponding to the first source of low-rank confounder, has a normal distribution with mean zero and covariance matrix \mathbf{I}_{d_1} . $\mathbf{V}_1 \in \mathbb{R}^{p \times d_1}$ is a non-random, semi-orthogonal matrix satisfying $\mathbf{V}_1^T \mathbf{V}_1 = \mathbf{I}$, and $\mathbf{D}_1 \in \mathbb{R}^{d_1 \times d_1}$ is a diagonal matrix, measuring the magnitude of the first confounder. Similarly, $\mathbf{Z}_2^{(i)} \in \mathbb{R}^{d_2}$, corresponding to the second source of low-rank confounder, has normal distribution with mean zero and covariance matrix \mathbf{I}_{d_2} . $\mathbf{V}_2 \in \mathbb{R}^{p \times d_2}$ is a semi-orthogonal matrix satisfying $\mathbf{V}_2^T \mathbf{V}_2 = \mathbf{I}$, and $\mathbf{D}_2 \in \mathbb{R}^{d_2 \times d_2}$ is a diagonal matrix, measuring the magnitude of the second low-rank confounder. Without loss of generality, we assume that $\mathbf{X}^{(i)}$, $\mathbf{Z}_1^{(i)}$ and $\mathbf{Z}_2^{(i)}$ are three pairwise independent vectors. Hence the observed covariance matrix is

$$\text{Cov}(\mathbf{X}_{obs}) = \Sigma_{obs} = \Sigma + \mathbf{V}_1 \mathbf{D}_1^2 \mathbf{V}_1^T + \mathbf{V}_2 \mathbf{D}_2^2 \mathbf{V}_2^T = \Sigma + \mathbf{L}_1 + \mathbf{L}_2. \quad (19)$$

Our first simulation setup (case 1) shows an ideal case for PCA+LVGGM, meaning that PCA+LVGGM method performs much better than using PCA+GGM, LVGGM, or Glasso. Let $d_1 = d_2 = 3$. We set $p = 100$ and $n = 100$. In our first example, the columns of \mathbf{V}_1 and \mathbf{V}_2 come from the eigenvectors of Σ . We expect that PC-correction removes \mathbf{L}_2 , so we set the diagonal elements of \mathbf{D}_2^2 all 50, and use the last 3 eigenvectors of Σ as \mathbf{V}_2 . This can guarantee that PC-correction performs much better than LVGGM and Glasso when removing \mathbf{L}_2 . Then we use LVGGM to estimate \mathbf{L}_1 , so we need a moderately large magnitude. We set all diagonal elements of \mathbf{D}_1^2 to 20, and use the first 3 eigenvectors of Σ as \mathbf{V}_1 . This ensures that LVGGM performs better than PC-correction and Glasso when estimating \mathbf{L}_1 .

Using the `sva` package, we estimate the rank of $\mathbf{L}_1 + \mathbf{L}_2$ to be 6. Then we look at the eigenvalues of the observed sample covariance matrix and we can see the first 3 eigenvalues are much larger than the 4-th to 6-th eigenvalues (shown in the top row of Fig. 2). We therefore allocate 3 to PC-correction, and $6 - 3 = 3$ to LVGGM. We also try allocating 1 to PC-correction and 5 to LVGGM. Then we compare more approaches, including using PC-correction individually by removing only 3 principal components or 6 principal components, using LVGGM with rank 6 for the low-rank component as well as the uncorrected approach Glasso. We still use 50 datasets and draw the ROC curve for the averages with varying sparsity parameters λ . The ROC for the scale-free example is in the bottom row of Fig. 2. We also include the AUC (area under the ROC curve) for each method. We compare PCA+LVGGM with rank 3 in PC-correction with other methods. For each data set, we calculate $(\text{AUC of PCA+LVGGM}) / (\text{AUC of one other method})$, then compute the sample mean and sample standard deviation of that ratio over 50 data sets to compare the average performance and the variance of different methods. The results for the scale-free graph are shown in the first column of Table 4.2. We can see that PCA+LVGGM with rank 3 for PC-correction and 3 for LVGGM do perform much better than other methods for both graph structures, indicating that if the assumptions are satisfied, our method and parameter tuning procedure are useful.

Finally, we try setups that are more similar to real world data. We still use (19) to generate the data and set $p = 100$ and $n = 100$. Differently from previous settings, we now use some randomly generated eigenvectors as columns of \mathbf{V}_1 and \mathbf{V}_2 . We look at the distribution of eigenvalues of gene co-expression and stock return data covariance matrices, and try to make simulation settings similar to those examples. We run two setups - the first is called a large-magnitude case (case 2), with \mathbf{D}_1^2 a diagonal matrix with diagonal elements (7, 6, 6) and \mathbf{D}_2^2 a diagonal matrix with diagonal elements (20, 10, 10). The second setup is referred to a moderately large magnitude case (case 3), in which the low-rank component has the same eigenvectors as the large-magnitude case, but the elements of \mathbf{D}_1 and \mathbf{D}_2 become smaller, with diagonal elements of \mathbf{D}_1^2 (3, 3, 3) and diagonal elements of \mathbf{D}_2^2 (10, 8, 6).

Using the `sva` package, we estimate the rank of $\mathbf{L}_1 + \mathbf{L}_2$ to be 6 for both case 2 and case 3. We observe that the first 3 eigenvalues are larger than the rest, so we allocate 3 to PC-correction and use $6 - 3 = 3$

as the rank for the low-rank component for LVGGM. We also try allocating 1 to PC-correction and 5 to LVGGM, using PC-correction by removing only 3 PC-components and 6 PC-components, using LVGGM with rank 6 and using Glasso. Again, we run over 50 datasets and include ROC curves and AUC tables.

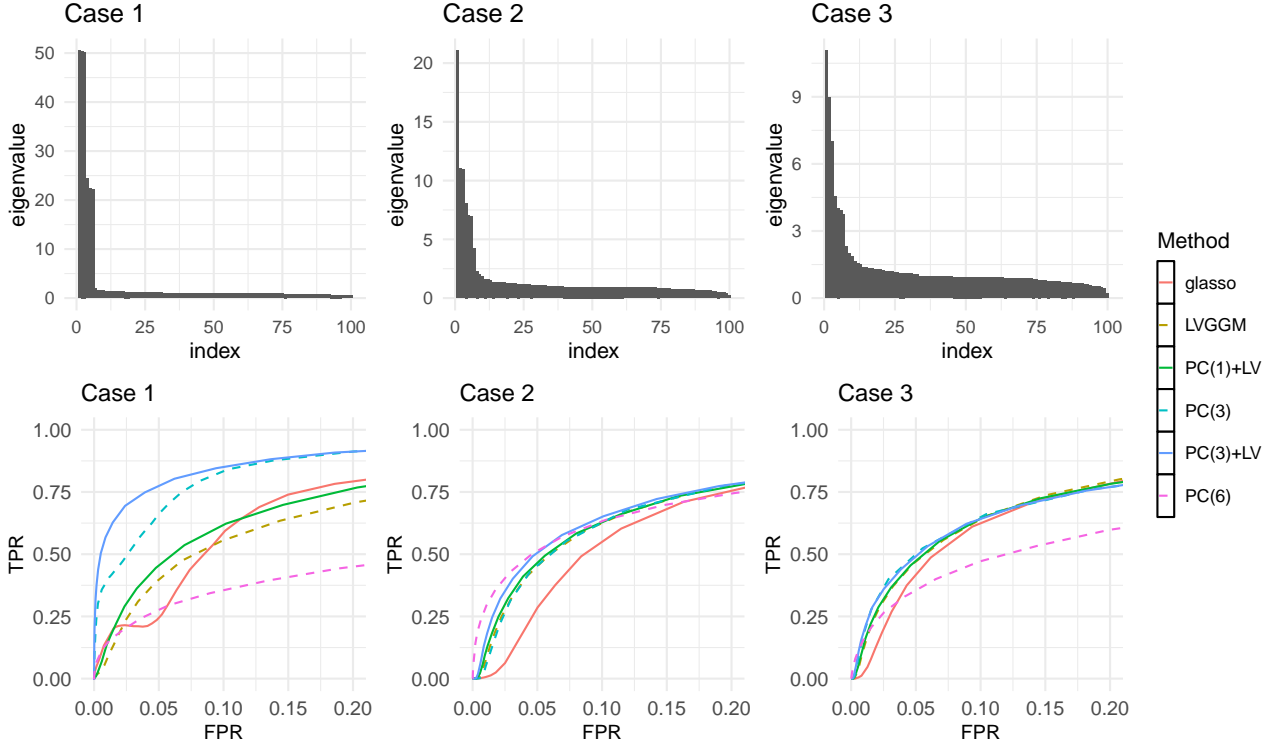


Figure 2: We use the scale-free structure when generating graphs. The first row shows the eigenvalues of Σ_{obs} under 3 setups, and the second row shows the corresponding ROC curves with different methods. PC(k) means that we use k as the rank in PC-correction and PC(k)+LV means that we use PCA+LVGGM with k as the rank for PC-correction.

Method	Case 1	Case 2	Case 3
PCA(3)+LVGGM	1	1	1
Glasso	1.58(0.073)	1.25(0.080)	1.07(0.048)
LVGGM	1.64(0.088)	1.06(0.044)	1.01(0.036)
PCA(Full)	2.46(0.22)	1.01(0.12)	1.36(0.14)
PCA(3)	1.08(0.017)	1.05(0.027)	0.99(0.018)
PCA(1)+LVGGM	1.47(0.11)	1.04(0.038)	1.01(0.035)

Table 1: We use the scale-free structure when generating graphs. We compute the ratio of AUC between PCA+LVGGM with rank 3 in PC-correction and other methods, using PCA+LVGGM as the numerator. The table shows the sample mean and sample standard deviations of that ratio (in the parenthesis) over 50 data sets. In case 3, the magnitude of the confounding is not as large as other cases, so PC-correction with rank 3 has the best performance.

From Fig. 2 and Table 4.2, we can see that other approaches considered hardly outperform PCA+LVGGM. Actually, using PCA+GGM or LVGGM can be viewed as a special case of the PCA+LVGMM methods. To see that, we can have LVGGM from PCA+LVGGM by allocating a rank of 0 to PC-correction. From the simulation and real data examples, we observe that using PCA+GGM with higher ranks often removes some useful information, resulting in more false negatives. On the other hand, if the effect of multiple confounders exists in the data that are not well represented by the first few principal components, using PCA+GGM alone might not be enough to remove the additional sources of noise. Note that LVGGM may not be enough to remove the confounding with large norm, leading to spurious

connections between nodes. In this case, we would suggest PCA+LVGGM as a default setting and a starting point for problems with low-rank confounding. We can adjust different rank allocations based on the specific problems and goals of interest.

5 Applications

5.1 Gene co-expression networks

Our first application is to reanalyze the gene co-expression networks originally analyzed by Parsana et al. (2019). The goal of gene co-expression network analysis is to identify transcriptional patterns indicating functional and regulatory relationships between genes. In biology, it is of great interest to infer the structure of these networks; however, the construction of such networks from data is challenging, since the data is usually corrupted by technical and unwanted biological variability known to confound expression data. The influence of such artifacts can often introduce spurious correlation between genes; if we apply sparse precision matrix inference directly without addressing confounding, we may obtain a graph including many false positive edges. Parsana et al. (2019) uses PCA+GGM to estimate this network and shows that PC-correction can be an effective way to control the false discovery rate of the network. In practice, however, some effects of confounding may not be represented in the top few principal components. This motivates the more flexible PCA+LVGGM approach. The PC-correction effectively removes high variance confounding, and then LVGGM subsequently accounts for any remaining low-rank confounding. We consider gene expression data from 3 diverse tissues: blood, lung and tibial nerve, with sample sizes between 300 to 400 each. 1000 genes are chosen from each tissue. More detail about the source of the data and pre-processing steps are introduced in Appendix D.

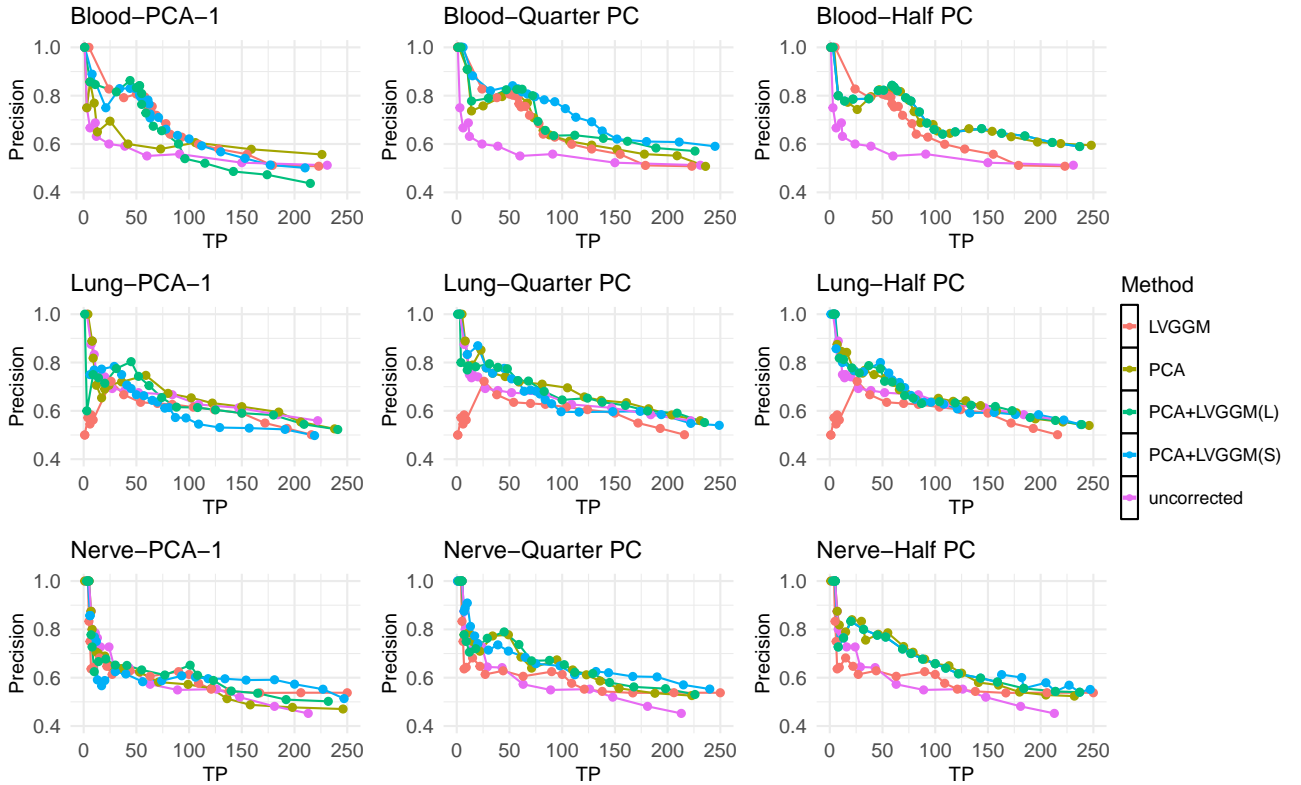


Figure 3: Precision-recall plots for gene expression data. TP represents number of true positives. PCA+LVGGM(L) means larger γ in LVGGM and PCA+LVGGM(S) means small γ in LVGGM. We can see that PCA+LVGGM performs the best or equivalently well compared to other approaches for almost all 3 tissues.

We observe that all of the sample covariance matrices are approximately low-rank by looking at the

eigenvalues of the covariance matrices of genes, indicating the potential existence of high variance confounding. Then we use `sva` package to estimate the rank for PC-correction and call this the full `sva` rank correction. Parsana et al. (2019) suggests that the rank estimated by `sva` might be so large that some useful network signal is removed. To reduce the effect of over-correction, we apply the PC-correction with half and one quarter of the `sva` rank, which we refer to as half `sva` rank correction and quarter `sva` rank correction, respectively. For many tissues, the first eigenvalue is much larger than the rest, this motivates us to try rank-1 PC-correction. We include the results with half `sva` rank, quarter `sva` rank and rank 1 PC-corrections in Fig. 3. After running the above PC-corrections to remove high-variance confounding, we run LVGGM as an additional step to further estimate and remove the low-rank noise with moderate variance. We use two different values as the γ parameters in LVGGM. Larger γ leads to removing lower-rank confounding and smaller γ leads to remove higher-rank confounding. We show the results for both choices of γ . We use different λ to control sparsity of the estimated graph. Specifically, following Parsana et al. (2019), we use 50 values of λ between 0.3 and 1. We draw Fig. 3 similar to the precision recall plot. The y-axis represents the precision (True Positives/(True Positives + False Positives)), and the x-axis is the number of true positives. We can see that PCA+LVGGM can yield better or equivalently good results compared to other methods, indicating that it can be useful to run LVGGM after the PC-correction when estimating gene co-expression networks.

5.2 Stock return data

In finance, the Capital Asset Pricing Model (CAPM) states that there is a widespread market factor which dominates the movement of all stock prices. Empirical evidence for the market trend can be found in the first principal component of the stock data, which is dense and has approximately equal loadings across all stocks (Fig. 5, left). In fact, the first few eigenvalues of the stock correlation matrix are significantly larger than the rest Fama and French (2004), which suggests that only a few latent factors are mainly driving stock correlations.

In this section, we posit that the conditional dependence structure after accounting for these latent effects is more likely to reflect direct relationships between companies aside from the market and, perhaps, sector trends. Our interest is in recovering the undirected graphical model (conditional dependence) structure between stock returns after controlling for potential low rank confounders.

We compare networks inferred by PCA+LVGGM, PCA+GGM, LVGGM and Glasso by analyzing monthly returns of component stocks in S&P 100 index between 2008 and 2019 Hayden et al. (2015). The 49 chosen companies are in 6 sectors: technology (10 companies), finance (11), energy (7), health (8), capital goods (7) and non-cyclical stocks (6). For PCA+GGM, we remove the first eigenvector which corresponds to the overall market trend. For the other latent variable methods we use the `sva` package to identify a plausible rank. For PCA+LVGGM, we remove the first principal component corresponding to the overall market trend and use LVGGM to estimate remaining latent confounders and the graph. Fig. 4 shows the networks obtained by each approach.

For each method, the sparsity-inducing tuning parameter was chosen to minimize negative log-likelihood using a 6-fold cross-validation procedure, and the number of low rank components are chosen manually. Specifically, in cross-validation, we use negative log-likelihood to measure the out-of-sample error and choose the parameters that minimize the average out-of sample error over 6 validation sets. We observe that when using LVGGM, allocating rank 1 or 2 to the low-rank component won't make the estimates very different from Glasso, while allocating ranks higher than 6 to the low-rank component leads to higher out-of-sample error, so 5, the rank picked by `sva`, is among the best choices. For PCA-based methods, removing more than 1 principal components leads to higher out-of-sample error. As expected, the Glasso result is denser than the networks learned with sparse plus low rank methodology with PCA+LVGGM yielding the sparsest network.

For LVGGM, we note that the method effectively controls for sector effect but is less effective in controlling for the effect of the overall market trend. Let $\hat{\Sigma}_{obs}$ be the empirical observed covariance matrix and \hat{v}_i be its i -th eigenvector. We have the following observations: first, the first principal




Figure 4: Stock connections between 2008 and 2019 learned by different methods. The following sectors are included: Tech, Finance, Energy, Health, Capital goods and Non-cyclical (NC) from left to right.

component is closely aligned with the overall market trend, because the absolute value of the inner product between the first eigenvector of $\hat{\Sigma}_{obs}$ and the normalized “all ones” vector is 0.98. Second, the observed empirical covariance matrix has an approximately low-rank structure, because the first eigenvalue of $\hat{\Sigma}_{obs}$ is 18.25 and the second is 3.5 and all other eigenvalues are close or smaller than 1. Third, LVGGM does not capture the full effect of the market trend. Let \hat{L}'_{Ω} be the estimate of L'_{Ω} in (9). When we apply LVGGM on $\hat{\Sigma}_{obs}$, the inner product between the first eigenvector of \hat{L}'_{Ω} and \hat{v}_1 is close to 1 but the first eigenvalue of \hat{L}'_{Ω} is only 0.55, much smaller than the first eigenvalue of $\hat{\Sigma}_{obs}$.

We argue that PCA+LVGGM is the most appropriate method for this application because it appropriately controls for both market and sector effects. Let \hat{L}_{Σ} and \hat{L}'_{Ω} be the estimates of the low-rank components defined in (10). For PCA+LVGGM, we remove L_{Σ} by removing the first eigenvector of Σ_{obs} , then run LVGGM to estimate L_{Ω} and Ω . We claim that PCA+LVGGM can remove the confounding effect fully in the market trend direction, as well as the remaining confounding effect in other directions. To see that, first, \hat{v}_1 is removed in PC-correction. Second, the inner product between the first eigenvector of \hat{L}'_{Ω} and \hat{v}_2 , the second eigenvector of $\hat{\Sigma}_{obs}$, is 0.99. The first eigenvalue of \hat{L}'_{Ω} is 0.4 and the second eigenvalue of $\hat{\Sigma}_{obs}$ is 3.5. This shows that when applying LVGGM, only part of the information in the direction of \hat{v}_2 has been removed. We know that the direction of \hat{v}_1 reflects the market trend, but \hat{v}_2 might include both true graph information and some latent confounding effect, hence using LVGGM might be a good choice for capturing the confounding effect in the direction of \hat{v}_2 . Overall PCA+LVGGM, therefore, might be a better choice than LVGGM and the PCA-based method.

Fig. 5 shows heat maps of \hat{L}_{Σ} obtained via PCA and \hat{L}'_{Ω} obtained with LVGGM (rank 5). As expected, the elements of \hat{L}_{Σ} are roughly equal in magnitude, reflect the market trend and the large first eigenvalue of $\hat{\Sigma}_{obs}$. In contrast, \hat{L}'_{Ω} shows a block-diagonal structure and its elements have smaller magnitudes, which suggests that LVGGM does not adequately account for the overall the market trend.

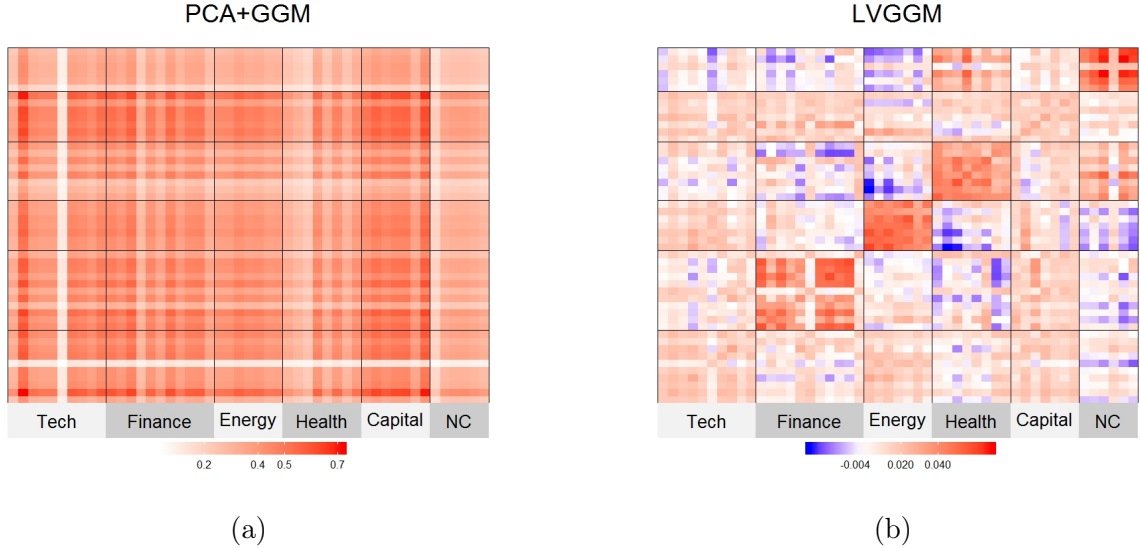


Figure 5: (a) Rank one approximation to $\hat{\mathbf{L}}_{\Sigma}$ obtained with PCA. (b) $\hat{\mathbf{L}}'_{\Omega}$ obtained with LVGGM with rank 5. The rank-one approximation to $\hat{\mathbf{L}}_{\Sigma}$ is close to a constant matrix. In contrast, $\hat{\mathbf{L}}'_{\Omega}$ reflects sector effects but does not reflect the strong effect due to overall market trends.

On the other hand, the block diagonal structure of $\hat{\mathbf{L}}'_{\Omega}$ reflects inferred sector effects. PCA+GGM is most effective at reducing confounding from overall market trends and LVGGM is more effective at accounting for remaining confounding, such as the sector effect. Therefore, PCA+LVGGM, which combines the benefits of PCA and LVGGM is arguably the most appropriate choice for addressing the latent confounding in this context.

6 Conclusion

We have studied the problem of estimating the graph structure under Gaussian graphical models when the data is corrupted by latent confounders. We compare two popular methods PCA+GGM and LVGGM. One of our contributions is to show the connection and difference between these two approaches, both theoretically and empirically. Based on that, we propose a new method, PCA+LVGGM. The effectiveness of this method is supported by theoretical analysis, simulations and applications. Actually, our analysis provides guidance on when to use which approach to estimate GMM in the presence of latent confounders. We believe that this guidance can help researchers in many fields, such as finance and biology, when there exist problems of graph estimation with confounders.

There are several future directions. First, we can extend the current framework to other distributions, such as the transelliptical distributions Liu et al. (2012) and the Ising model Ravikumar et al. (2010); Nussbaum and Giesen (2019). Secondly, we can consider more structures - for example, we are interested in what will happen if the principal components of the observed data are sparse. Given the distributions of eigenvalues or eigenvectors of Σ , we can provide more precise bounds. Deriving sharper bounds to more general settings is a useful extension. We can also extend LVGGM based on our observations in this work. For instance, it is interesting to consider replacing the nuclear norm penalty in LVGGM with some unbiased regularization such as SCAD Fan and Li (2001) and MCP Zhang (2010) to avoid shrinking too much on large eigenvalues. Other future works include applying our current methods and analysis in more applications, such as the functional magnetic resonance imaging (fMRI) in neuroscience.

Acknowledgments

The authors thank Christos Thrampoulidis and Megan Elcheikhali for the helpful discussion. The authors thank anonymous reviewers for their valuable feedback.

Appendix A Proofs

Proof outline: To prove Theorems 1 and 2, we first write $\hat{\Sigma} - \Sigma$ as below:

$$\hat{\Sigma} - \Sigma = (\hat{\Sigma}_{obs} - \hat{\lambda}_1 \hat{\theta}_1 \hat{\theta}_1^T) - (\Sigma_{obs} - \lambda_1 \theta_1 \theta_1^T + \lambda_k \theta_k \theta_k^T) = (\hat{\Sigma}_{obs} - \Sigma_{obs}) + (\lambda_1 \theta_1 \theta_1^T - \hat{\lambda}_1 \hat{\theta}_1 \hat{\theta}_1^T) - \lambda_k \theta_k \theta_k^T, \quad (20)$$

where λ_k is the k -th eigenvalue of Σ_{obs} , and θ_k is the k -th eigenvector of Σ_{obs} . To bound $\hat{\Sigma} - \Sigma$, we need to bound the norms of $\Sigma_{obs} - \hat{\Sigma}_{obs}$, $\lambda_1 - \hat{\lambda}_1$ and $\theta_1 - \hat{\theta}_1$, and the lemmas below show these bounds. \square

Lemma 1. *Assuming that Σ_{obs}^{-1} satisfies the maximum and minimum eigenvalue condition in (14), then*

$$P(\|\hat{\Sigma}_{obs} - \Sigma_{obs}\|_{\infty} \geq t) \leq C_1 p^{-1}, \quad t = C_2 \sqrt{\frac{\ln p}{n}},$$

where C_1 and C_2 depends on the eigenvalue bound M in (14).

Proof: The proof follows (Bickel et al., 2008b, Lemma A.3). \square

The next two lemmas provide bounds for $|\lambda_1 - \hat{\lambda}_1|$.

Lemma 2. *Under the assumptions of Theorem 1, we have*

$$|\lambda_1 - \hat{\lambda}_1| \leq C_1 \lambda_1 \sqrt{\frac{p}{n}},$$

with probability at least $1 - 2e^{-p/C_2}$ for some constants C_i 's > 1 .

Proof: By Weyl's lemma Horn and Johnson (2012)

$$\max_{j=1, \dots, p} |\lambda_j(\Sigma_{obs}) - \lambda_j(\hat{\Sigma}_{obs})| \leq \|\hat{\Sigma}_{obs} - \Sigma_{obs}\|_2.$$

The bound on $\|\hat{\Sigma}_{obs} - \Sigma_{obs}\|_2$ is then obtained from (Wainwright, 2019, Theorem 6.5). \square

Following (Koltchinskii and Lounici, 2017, Theorem 5), a tighter bound can be obtained using the effective rank $r(\Sigma_{obs})$ defined in (15).

Lemma 3. *Under the assumptions of Theorem 2, we have*

$$|\lambda_1 - \hat{\lambda}_1| \leq \|\hat{\Sigma}_{obs} - \Sigma_{obs}\|_2 \leq C_1 \left(\sqrt{\frac{p\lambda_1}{n}} \vee \frac{p}{n} \right),$$

with probability at least $1 - e^{-p/C_2}$ for some constants C_i 's > 1 .

Proof: Theorem 5 in Koltchinskii and Lounici (2017) shows that

$$|\lambda_1 - \hat{\lambda}_1| \leq \|\hat{\Sigma}_{obs} - \Sigma_{obs}\|_2 \leq C \lambda_1 \left(\sqrt{\frac{r(\Sigma_{obs})}{n}} \vee \frac{r(\Sigma_{obs})}{n} \right).$$

The bound above is at the same order as $\sqrt{(p\lambda_1)/n} \vee (p/n)$ because $r(\Sigma_{obs}) \asymp (p/\lambda_1)$. \square

Then the two lemmas below provide bounds for $\|\boldsymbol{\theta}_1 - \hat{\boldsymbol{\theta}}_1\|_\infty$.

Lemma 4 (adapted from Wainwright 2019, Corollary 8.7). *Under the assumptions of Theorem 1, suppose $n \geq p$ and $\|\boldsymbol{\Sigma}\|_2 \sqrt{(\nu+1)/\nu^2} \sqrt{p/n} \leq 1/128$, then*

$$\|\boldsymbol{\theta}_1 - \hat{\boldsymbol{\theta}}_1\|_2 \leq C_1 \sqrt{\frac{\nu+1}{\nu^2}} \sqrt{\frac{p}{n}},$$

with probability at least $1 - C_2 e^{-p/C_3}$ for some C_i 's > 1 , where $\nu = \lambda_1(\boldsymbol{\Sigma}_{obs}) - \lambda_2(\boldsymbol{\Sigma}_{obs})$.

The lemma below shows a tighter bound with a large eigengap $\nu = \lambda_1(\boldsymbol{\Sigma}_{obs}) - \lambda_2(\boldsymbol{\Sigma}_{obs})$ following (Fan et al., 2018, Section 3.1).

Lemma 5. *Under the assumptions of Theorem 2,*

$$\|\boldsymbol{\theta}_1 - \hat{\boldsymbol{\theta}}_1\|_\infty \asymp O_p\left(\frac{p \vee \lambda_1}{\nu p} \sqrt{\frac{\ln p}{n}}\right) \asymp O_p\left(\frac{\sqrt{p}}{\lambda_1} \sqrt{\frac{\ln p}{n}}\right),$$

with probability at least $1 - C/p$ for some constant C .

Before moving to the proofs of Theorem 1 and 2, we first show an upper bound for $\|\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_\infty$ using (20).

$$\|\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_\infty \leq \|\boldsymbol{\Sigma}_{obs} - \hat{\boldsymbol{\Sigma}}_{obs}\|_\infty + \|\lambda_1 \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T - \hat{\lambda}_1 \hat{\boldsymbol{\theta}}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty + \|\lambda_k \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T\|_\infty. \quad (21)$$

The term $\|\lambda_1 \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T - \hat{\lambda}_1 \hat{\boldsymbol{\theta}}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty$ can be expressed as:

$$\begin{aligned} \|\lambda_1 \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T - \hat{\lambda}_1 \hat{\boldsymbol{\theta}}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty &= \|\lambda_1 \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T - \lambda_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T + \lambda_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T - \hat{\lambda}_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T + \hat{\lambda}_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T - \hat{\lambda}_1 \hat{\boldsymbol{\theta}}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty \\ &\leq \|\lambda_1 \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T - \lambda_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty + \|\lambda_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T - \hat{\lambda}_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty + \|\hat{\lambda}_1 \boldsymbol{\theta}_1 \hat{\boldsymbol{\theta}}_1^T - \hat{\lambda}_1 \hat{\boldsymbol{\theta}}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty \\ &\leq |\lambda_1| \|\boldsymbol{\theta}_1\|_\infty \|\boldsymbol{\theta}_1 - \hat{\boldsymbol{\theta}}_1\|_\infty + |\lambda_1 - \hat{\lambda}_1| \|\boldsymbol{\theta}_1\|_\infty \|\hat{\boldsymbol{\theta}}_1\|_\infty + |\hat{\lambda}_1| \|\hat{\boldsymbol{\theta}}_1\|_\infty \|\boldsymbol{\theta}_1 - \hat{\boldsymbol{\theta}}_1\|_\infty. \end{aligned} \quad (22)$$

We can then use the bounds for $|\lambda_1 - \hat{\lambda}_1|$ and $\|\boldsymbol{\theta}_1 - \hat{\boldsymbol{\theta}}_1\|_\infty$ in previous lemmas to bound $\|\lambda_1 \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T - \hat{\lambda}_1 \hat{\boldsymbol{\theta}}_1 \hat{\boldsymbol{\theta}}_1^T\|_\infty$.

Proofs of Theorems 1 and 2: Now we are ready to prove Theorem 1. We first plug in the bounds in Lemmas 1, 2 and 4 to (22). Since $\boldsymbol{\theta}_1$ is the eigenvector of a matrix, $\|\boldsymbol{\theta}_1\|_\infty \leq 1$. Combining these two results completes the proof.

To prove Theorem 2, we need to plug in the bounds in Lemmas 1, 3 and 5 to (22). Given Lemmas 3 and 5, we can bound the second term in (22),

$$|\lambda_1 - \hat{\lambda}_1| \|\boldsymbol{\theta}_1\|_\infty \|\hat{\boldsymbol{\theta}}_1\|_\infty = O_p\left(\sqrt{\frac{\lambda_1}{np}} \vee \frac{1}{n}\right) = O_p\left(\sqrt{\frac{\ln p}{n}}\right).$$

Similarly, the first and third terms in (22) have the convergence rate of

$$O_p\left(\lambda_1 \|\boldsymbol{\theta}_1\|_\infty \left(\frac{\sqrt{p}}{\lambda_1} \sqrt{\frac{\ln p}{n}}\right)\right) = O_p\left(\sqrt{\frac{\ln p}{n}}\right).$$

Then with incoherent $\|\boldsymbol{\theta}_1\|_\infty = O(1/\sqrt{p})$, the last term in (21) $\|\lambda_k \boldsymbol{\theta}_1 \boldsymbol{\theta}_1^T\|_\infty$ has the convergence rate of $O_p(1/p)$. \square

Proof of Theorem 3: The proof follows the proof of (Cai et al., 2011, Theorem 6). First we know,

$$\|\hat{\Sigma}\Omega - I\|_\infty = \|(\hat{\Sigma} - \Sigma)\Omega\|_\infty \leq \|\hat{\Sigma} - \Sigma\|_\infty \|\Omega\|_{L_1}.$$

Then we have,

$$\|\hat{\Sigma}(\Omega - \hat{\Omega}_1)\|_\infty = \|\hat{\Sigma}\Omega - I + I - \hat{\Sigma}\hat{\Omega}_1\|_\infty \leq \|\hat{\Sigma}\Omega - I\|_\infty + \|I - \hat{\Sigma}\hat{\Omega}_1\|_\infty + \|\hat{\Sigma} - \Sigma\|_\infty \|\Omega\|_{L_1} + \lambda_n.$$

We know,

$$\|\Omega - \hat{\Omega}_1\|_\infty = \|\Omega\Sigma(\Omega - \hat{\Omega}_1)\|_\infty \leq \|\Sigma(\Omega - \hat{\Omega}_1)\|_\infty \|\Omega\|_{L_1}.$$

To bound the terms above, we need,

$$\|\Sigma(\Omega - \hat{\Omega}_1)\|_\infty \leq \|\hat{\Sigma}(\Omega - \hat{\Omega}_1)\|_\infty + \|\hat{\Sigma} - \Sigma\|_\infty \|\Omega\|_{L_1}.$$

We know $\|\Omega\|_{L_1} \leq M_0$ from (17) and combining the relations above with the result of Theorem 1 or Theorem 2, with the choice of λ_n specified in Theorem 3, we can obtain the bound for $\hat{\Omega}_1$. The bound of the same order can be obtained for $\hat{\Omega}$, the symmetric version of $\hat{\Omega}_1$. \square

Appendix B Generalization of Section 3

The analysis in Section 3 assumes that the low-rank confounder is independent of \mathbf{X} and the eigenvector of the covariance of the low-rank confounding is one of the eigenvectors of Σ , the covariance of \mathbf{X} . Those two assumptions can be extended to the more general setups. In equation (11), when \mathbf{X} and \mathbf{Z} are not independent, the covariance matrix for \mathbf{X}_{obs} becomes

$$\Sigma_{obs} = \Sigma + \sigma Cov(\mathbf{X}, \mathbf{Z})\mathbf{v}^T + \sigma\mathbf{v}Cov(\mathbf{X}, \mathbf{Z})^T + \sigma^2\mathbf{v}\mathbf{v}^T,$$

where $Cov(\mathbf{X}, \mathbf{Z})$ is a p -dimensional column vector. We can see that $\sigma Cov(\mathbf{X}, \mathbf{Z})\mathbf{v}^T + \sigma\mathbf{v}Cov(\mathbf{X}, \mathbf{Z})^T + \sigma^2\mathbf{v}\mathbf{v}^T$ has rank at most 3, hence Σ_{obs} can still be expressed as the sum of Σ and a low-rank matrix. Here, to ensure that the confounding can be identified in PCA-based approach, we assume that both σ and σ^2 are large compared to the eigenvalues of Σ . Then, our analysis in Section 3 can still be applied here, but the eigenvectors of the low-rank matrix are not necessarily the eigenvectors of Σ .

Appendix C Eigenvalues of sparse graphs




Figure 6: The distribution of the top 25 eigenvalues created by scale-free, random and cluster graphs using `huge` package with $p = 100$ and $n = 10000$.

Fig. 6 shows the first 25 eigenvalues of sparse graphs. We use `huge` package Zhao et al. (2012) to generate the sparse graphs with three different structures: scale-free, random and clustered. We set $p = 100, n = 10000$ and assume default for all other parameters (see Zhao et al. (2012) for more detail). We generate 10 realizations for each graph structure and show the distribution of the first 25 eigenvalues of Σ in Fig. 6. We notice that the top eigenvalues are typically larger than the rest, especially for the scale-free graphs.

Appendix D Gene co-expression networks data

Now we briefly introduce the data and the pre-processing procedure of gene co-expression networks in Section 5.1. More detail can be found from Parsana et al. (2019). We use the RNA-Seq data from the Genotype-Tissue Expression (GTEx) project v6p release. We consider three diverse tissues with sample sizes between 300 to 400 each: blood, lung and tibial nerve. We first filter the non-overlapping protein genes and perform a log transformation with base 2 to scale the data following (Parsana et al., 2019, Appendix 2.4). Since the underlying true network structure is generally unknown, we obtain the interaction information from some canonical pathway databases including KEGG, Biocarta, Reactome and Pathway Interaction Database. To make better use of those information, we pick 1000 high-variance genes which are included in all these databases, thus $p = 1000$ in this example.

References

- A. Agarwal, S. Negahban, M. J. Wainwright, et al. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics*, 40(2):1171–1197, 2012.
- R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.
- D. M. Barch, G. C. Burgess, M. P. Harms, S. E. Petersen, B. L. Schlaggar, M. Corbetta, M. F. Glasser, S. Curtiss, S. Dixit, C. Feldt, et al. Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage*, 80:169–189, 2013.
- D. J. Bartholomew, M. Knott, and I. Moustaki. *Latent variable models and factor analysis: A unified approach*. John Wiley & Sons, 2011.
- P. J. Bickel, E. Levina, et al. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008a.
- P. J. Bickel, E. Levina, et al. Regularized estimation of large covariance matrices. *The Annals of Statistics*, 36(1):199–227, 2008b.
- A. Buja and N. Eyuboglu. Remarks on parallel analysis. *Multivariate behavioral research*, 27(4):509–540, 1992.
- T. Cai, W. Liu, and X. Luo. A constrained l_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106(494):594–607, 2011.
- T. T. Cai, Z. Ren, H. H. Zhou, et al. Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation. *Electronic Journal of Statistics*, 10(1):1–59, 2016.
- E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.
- V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2):572–596, 2011.
- V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky. Latent variable graphical model selection via convex optimization. *The Annals of Statistics*, pages 1935–1967, 2012.
- M. J. Choi, V. Y. Tan, A. Anandkumar, and A. S. Willsky. Learning latent tree graphical models. *Journal of Machine Learning Research*, 12(May):1771–1812, 2011.
- P. Danaher, P. Wang, and D. M. Witten. The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(2):373–397, 2014.
- E. F. Fama and K. R. French. The capital asset pricing model: Theory and evidence. *Journal of economic perspectives*, 18(3):25–46, 2004.

- J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- J. Fan, W. Wang, and Y. Zhong. An l_∞ eigenvector perturbation bound and its application to robust covariance estimation. *Journal of Machine Learning Research*, 18(207):1–42, 2018.
- M. D. Fox and M. E. Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature reviews neuroscience*, 8(9):700–711, 2007.
- S. Freytag, J. Gagnon-Bartsch, T. P. Speed, and M. Bahlo. Systematic noise degrades gene co-expression signals but can be corrected. *BMC bioinformatics*, 16(1):309, 2015.
- J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- J. A. Gagnon-Bartsch, L. Jacob, and T. P. Speed. Removing unwanted variation from high dimensional data with negative controls. *Berkeley: Tech Reports from Dep Stat Univ California*, pages 1–112, 2013.
- S. Geng, M. Kolar, and O. Koyejo. Joint nonparametric precision matrix estimation with confounding. *arXiv preprint arXiv:1810.07147*, 2018.
- T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
- L. X. Hayden, R. Chachra, A. A. Alemi, P. H. Ginsparg, and J. P. Sethna. Canonical sectors and evolution of firms in the us stock markets. *arXiv preprint arXiv:1503.06205*, 2015.
- J. L. Horn. A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30(2):179–185, 1965.
- R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge university press, 2012.
- L. Jacob, J. A. Gagnon-Bartsch, and T. P. Speed. Correcting gene expression data when neither the unwanted variation nor the factor of interest are observed. *Biostatistics*, 17(1):16–28, 2016.
- I. Jolliffe. Principal component analysis. *Encyclopedia of statistics in behavioral science*, 2005.
- K. Khare, S.-Y. Oh, and B. Rajaratnam. A convex pseudolikelihood framework for high dimensional partial correlation estimation with convergence guarantees. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(4):803–825, 2015.
- V. Koltchinskii and K. Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23(1):110–133, 2017.
- C. Lam and J. Fan. Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of statistics*, 37(6B):4254, 2009.
- S. L. Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.
- J. T. Leek and J. D. Storey. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS genetics*, 3(9), 2007.
- J. T. Leek, W. E. Johnson, H. S. Parker, A. E. Jaffe, and J. D. Storey. The sva package for removing batch effects and other unwanted variation in high-throughput experiments. *Bioinformatics*, 28(6):882–883, 2012.
- S. Lim and S. Jahng. Determining the number of factors using parallel analysis and its recent variants. *Psychological methods*, 24(4):452, 2019.
- H. Liu, F. Han, and C.-h. Zhang. Transelliptical graphical models. In *Advances in neural information processing systems*, pages 800–808, 2012.

- N. Meinshausen and P. Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- N. Meinshausen, P. Bühlmann, et al. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3):1436–1462, 2006.
- Z. Meng, B. Eriksson, and A. Hero. Learning latent variable gaussian graphical models. In *International Conference on Machine Learning*, pages 1269–1277, 2014.
- F. Nussbaum and J. Giesen. Ising models with latent conditional gaussian variables. *arXiv preprint arXiv:1901.09712*, 2019.
- P. Parsana, C. Ruberman, A. E. Jaffe, M. C. Schatz, A. Battle, and J. T. Leek. Addressing confounding artifacts in reconstruction of gene co-expression networks. *Genome biology*, 20(1):94, 2019.
- J. Peng, P. Wang, N. Zhou, and J. Zhu. Partial correlation estimation by joint sparse regression models. *Journal of the American Statistical Association*, 104(486):735–746, 2009.
- T. Price, C.-Y. Wee, W. Gao, and D. Shen. Multiple-network classification of childhood autism using functional connectivity dynamics. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 177–184. Springer, 2014.
- P. Ravikumar, M. J. Wainwright, J. D. Lafferty, et al. High-dimensional ising model selection using ℓ_1 -regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.
- P. Ravikumar, M. J. Wainwright, G. Raskutti, B. Yu, et al. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- Z. Ren and H. H. Zhou. Discussion: Latent variable graphical model selection via convex optimization1. *The Annals of Statistics*, 40(4):1989–1996, 2012.
- A. J. Rothman, P. J. Bickel, E. Levina, J. Zhu, et al. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.
- O. Stegle, C. Lippert, J. M. Mooij, N. D. Lawrence, and K. Borgwardt. Efficient inference in matrix-variate gaussian models with iid observation noise. In *Advances in neural information processing systems*, pages 630–638, 2011.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- K. R. Van Dijk, M. R. Sabuncu, and R. L. Buckner. The influence of head motion on intrinsic functional connectivity mri. *Neuroimage*, 59(1):431–438, 2012.
- M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- M. Yuan and Y. Lin. Model selection and estimation in the gaussian graphical model. *Biometrika*, 94(1):19–35, 2007.
- C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.
- T. Zhao, H. Liu, K. Roeder, J. Lafferty, and L. Wasserman. The huge package for high-dimensional undirected graph estimation in r. *The Journal of Machine Learning Research*, 13:1059–1062, 2012.