Joint Estimation of Multiple High-dimensional Precision Matrices

T. Tony Cai,

Professor of Statistics, Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104

Hongzhe Li,

Professor of Biostatistics, University of Pennsylvania Perelman School of Medicine, Philadelphia, PA 19104

Weidong Liu, and

Professor, Department of Mathematics, Institute of Natural Sciences and MOE-LSC, Shanghai Jiao Tong University, Shanghai, China

Jichun Xie

Assistant Professor, Department of Biostatistics and Bioinformatics, Duke University School of Medicine, Durham, NC 27707

Abstract

Motivated by analysis of gene expression data measured in different tissues or disease states, we consider joint estimation of multiple precision matrices to effectively utilize the partially shared graphical structures of the corresponding graphs. The procedure is based on a weighted constrained ℓ_∞/ℓ_1 minimization, which can be effectively implemented by a second-order cone programming. Compared to separate estimation methods, the proposed joint estimation method leads to estimators converging to the true precision matrices faster. Under certain regularity conditions, the proposed procedure leads to an exact graph structure recovery with a probability tending to 1. Simulation studies show that the proposed joint estimation methods outperform other methods in graph structure recovery. The method is illustrated through an analysis of an ovarian cancer gene expression data. The results indicate that the patients with poor prognostic subtype lack some important links among the genes in the apoptosis pathway.

Keywords

Constrained optimization; Convergence rate; Graph recovery; Precision matrices; Second-order cone programming; Sparsity

1. INTRODUCTION

Gaussian graphical models provide a natural tool for modeling the conditional independence relationships among a set of random variables (Lauritzen (1996); Whittaker (1990)). They

have been successfully applied to infer relationships between genes at transcriptional level (Schäfer and Strimmer (2005); Li and Gui (2006); Li, Hsu, Peng et al. (2013)). Gaussian graphical models are tightly linked to precision matrices. Suppose $X = (X_1,\ldots, X_p)'$ follows a multivariate Gaussian distribution $N_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The precision matrix $\boldsymbol{\Omega} = \boldsymbol{\Sigma}^{-1}$ describes the graphical structure of its corresponding Gaussian graph. If the (i, j)-th entry of the precision matrix ω_{ij} is equal to zero, then X_i and X_j are independent conditioning on all other variables X_k, $k \neq i, j$. Correspondingly, no edge exists between X_i and variable X_j in the graphical structure of Gaussian graphical model. If $\omega_{ij} \neq 0$, then X_i and X_j are conditionally dependent and they are therefore connected in the graphical structure. Define the support of $\boldsymbol{\Omega}$ by $\mathscr{S} = \{(i, j) : \omega_{ij} \neq 0\}$, which is also the set of edges in the Gaussian graphical model. If the maximum degree of $\boldsymbol{\Omega}$, $\max_i \sum_{j=1}^{p} I(\omega_{ij} \neq 0)$ is relatively small, we call $\boldsymbol{\Omega}$ sparse. Since the expression variation of a gene can usually be explained by a small subset of other genes, the precision matrix for gene expression data of a set of genes is expected to be sparse.

Many methods for estimating high-dimensional precision matrix or its Gaussian graphical model have been developed in the past decade. Meinshausen and Bühlmann (2006) introduced a neighborhood selection approach by regressing all other variables on each variable with an ℓ_1 penalty. The method consistently estimates the set of non-zero elements of the precision matrix. Efficient algorithms for exact maximization of the ℓ_1-penalized log-likelihood have also been proposed. Yuan and Lin (2007), Banerjee, Ghaoui and d'Aspremont (2008) and Dahl, Vandenberghe and Roychowdhury (2008) adopted an interior point optimization method to solve this problem. Based on the work of Banerjee, Ghaoui and d'Aspremont (2008) and a block-wise coordinate descent algorithm, Friedman, Hastie and Tibshirani (2008) developed the graphical Lasso (GLASSO) for sparse precision matrix estimation; it is computationally efficient even when the dimension is greater than the sample size. Yuan (2010) developed a linear programming procedure and obtained oracle inequalities for the estimation error in term of matrix operator norm. Cai, Liu and Luo (2011) developed a constrained ℓ_1 minimization approach (CLIME) to estimate sparse precision matrix. All of these methods addressed the problem of estimating a single precision matrix or a single Gaussian graphical model.

In many applications, the problem of estimating multiple precision matrices arises when data are collected among multiple groups. For example, gene expression levels are often measured over multiple groups (tissues, environments, or subpopulations). Their precision matrices and the corresponding graphical structures imply gene regulatory mechanisms and are of great biological interest. Since the gene regulatory networks in different groups are often similar to each other, the graphical structures share many common edges. Estimating a single precision matrix group by group ignores the partial homogeneity in their graphical structures, which often leads to low power. To effectively utilize the shared graphical structures and to increase the estimation precision, it is important to estimate multiple precision matrices jointly.

Previous attempts to jointly estimate multiple precision matrices include Guo, Levina, Michailidis et al. (2011) and Danaher, Wang and Witten (2014). Guo, Levina, Michailidis et al. (2011) proposed a hierarchical penalized model to perserve the common graphical

structure while allowing differences across groups. Their method achieves Frobenius norm convergence when $p \log(p)/n$ goes to zero, where p is the number of variables, and n is the total sample size. Unfortunately, for genomic applications, the number of genes often exceeds the total sample size and, as a result, invalidates the theoretical justification in Guo, Levina, Michailidis et al. (2011). Danaher, Wang and Witten (2014) proposed two algorithms of joint graphical lasso (FGL and GGL) to estimate precision matrices that share common edges. Their approach is based upon maximizing a penalized log likelihood with a fused Lasso or group Lasso penalty. The paper did not provide any theoretical justification on the statistical convergence rate of their estimators.

In this paper, we propose a weighted constrained ℓ_∞/ℓ_1 minimization method to estimate K sparse precision matrices (MPE) jointly. Different from Guo, Levina, Michailidis et al. (2011) and Danaher, Wang and Witten (2014), the proposed estimators converge to the true precision matrices even when $p = O\{\exp(n^a)\}$, for some $0 < a < 1$. In addition, when K is sufficiently large, compared to the estimators from separate estimation methods, our proposed estimators converge to the true precision matrices (under the entry-wise ℓ_∞ norm loss) faster. An additional thresholding step on the estimators with a carefully chosen threshold yields thresholded estimators with additional theoretical properties. The thresholded estimators from our method converge to the true precision matrices under the matrix ℓ_1 norm. Finally, when the graphical structures across groups are the same, our method leads to the exact recovery of the graph structures with a probability tending to 1.

The rest of the paper is organized as follows. Section 2 presents the estimation method and the optimization algorithm. Theoretical properties of the proposed method and accuracy of the graph structure recovery are studied in Section 3. Section 4 investigates the numerical performance of the method through a simulation study. The proposed method is compared with other competing methods. The method is also illustrated by an analysis of an epithelial ovarian cancer gene expression study in Section 5. A brief discussion is given in Section 6 and proofs are presented in the Appendix.

2. METHODOLOGY

The following notations are used in the paper. For a vector $\mathbf{a} = (a_1,\ldots, a_p)^T \in \mathbb{R}^p$, define

$|\mathbf{a}|_1 = \sum_{j=1}^{p} |a_j|$ and $|\mathbf{a}|_2 = \left(\sum_{j=1}^{p} a_j^2\right)^{1/2}$. The vector $\mathbf{a}_{-i}$ is the vector of a without the i-th entry. The support of a is defined as $\mathrm{supp}(\mathbf{a}) = \{i : a_i \neq 0\}$. For a matrix $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{p \times q}$, its entrywise ℓ_∞ norm is denoted by $|\mathbf{A}|_\infty = \max_{i,j} |a_{ij}|$. Its matrix ℓ_1-norm is denoted by

$\|A\|_1 = \max_{1 \le j \le q} \sum_{i=1}^{p} |a_{ij}|$, and its spectral norm is denoted by $\|A\|_2$. The sub-matrix $\mathbf{A}_{-i,-i}$ is the matrix of $\mathbf{A}$ without the i-th row and the i-th column. Denote by $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ the largest and smallest eigenvalues of A, respectively. For two sequences of real numbers $\{a_n\}$ and $\{b_n\}$, write $a_n = O(b_n)$ if there exists a constant C such that $|a_n| \le C|b_n|$ holds for all sufficiently large n, write $a_n = o(b_n)$ if $\lim_{n\to\infty} a_n/b_n = 0$. If $a_n = O(b_n)$ and $b_n = O(a_n)$, then $a_n \asymp b_n$.

2.1 The Joint Estimation Method

We introduce an estimation method to jointly estimate K precision matrices with partial homogeneity in their graphical structures. The method uses a constrained ℓ_1 minimization approach, that has been successfully applied to high dimensional regression problems (Donoho, Elad and Temlyakov (2006); Candés and Tao (2007)) and signal precision matrix estimation problem (Cai, Liu and Luo (2011)) to recover the sparse vector or matrix.

For $1 \leq k \leq K$, let $X^{(k)} \sim N(\boldsymbol{\mu}^{(k)}, \boldsymbol{\Sigma}^{(k)})$ be a p-dimensional random vector. The precision matrix of $X^{(k)}$, denoted by $\boldsymbol{\Omega}^{(k)}=(\omega_{ij}^{(k)})$, is the inverse of the covariance matrix $\boldsymbol{\Sigma}^{(k)}$. Suppose there are n_k identically and independently distributed random samples of $X^{(k)}$:$\{X_j^{(k)}, 1 \leq j \leq n_k\}$. The sample covariance matrix for the k-th group is

$$\hat{\sum}^{(k)}=\frac{1}{n_k}\sum_{j=1}^{n_k}(X_j^{(k)} - \overline{X}^{(k)})(X_j^{(k)} - \overline{X}^{(k)})',$$

where $\overline{X}^{(k)}=\sum_{j=1}^{n_k} X_j^{(k)}/n_k$. Denote by $n = \Sigma_k\, n_k$ the total sample size and $w_k = n_k/n$ the weight of the k-th group. We estimate $\boldsymbol{\Omega}^{(k)}=(\omega_{ij}^{(k)})$ for $k = 1, \dots, K$ by the constrained optimization

$$\min_{\boldsymbol{\Omega}_1^{(k)} \in \mathbb{R}^{p \times p}, 1 \leq k \leq K} \left(\max_{1 \leq k \leq K} |\boldsymbol{\Omega}_1^{(k)}|_1 \right),$$
$$\text{subject to} \max_{i,j} \left\{ \sum_{k=1}^{K} w_k |(\hat{\sum}^{(k)}\boldsymbol{\Omega}_1^{(k)} - \mathbf{I})_{ij}|^2 \right\}^{1/2} \leq \lambda_n, \qquad (1)$$

where $\lambda_n = C(\log\, p/n)^{1/2}$ is a tuning parameter. The ℓ_∞/ℓ_1 objective function is used to encourage the sparsity of all K precision matrices. The constraint is imposed on the maximum of the element-wise group ℓ_2 norm to encourage the groups to share a common graphical structure.

Denote by $\hat{\boldsymbol{\Omega}}_1^{(k)} (1 \leq k \leq K)$ the solution to (1). They are not symmetric in general. To make the solution symmetric, the estimator $\hat{\boldsymbol{\Omega}}^{(k)}=(\hat{\omega}_{ij}^{(k)})$ is constructed by comparing $\hat{\omega}_{1ij}^{(k)}$ and $\hat{\omega}_{1ji}^{(k)}$ and assigning the one with a smaller magnitude at both entries,

$$\hat{\omega}_{ij}^{(k)}=\hat{\omega}_{ji}^{(k)}:=\hat{\omega}_{1ij}^{(k)}I(|\hat{\omega}_{1ij}^{(k)}| \leq |\hat{\omega}_{1ji}^{(k)}|)+\hat{\omega}_{1ji}^{(k)}I(|\hat{\omega}_{1ij}^{(k)}|>|\hat{\omega}_{1ji}^{(k)}|).$$

This symmetrizing procedure is not ad-hoc. The procedure assures that the final estimator $\hat{\boldsymbol{\Omega}}^{(k)}$ achieves the same entry-wise ℓ_∞ estimation error as $\hat{\boldsymbol{\Omega}}_1^{(k)}$. The details are discussed in Section 3.

2.2 Computational algorithm

The convex optimization problem (1) involves estimating K $p \times p$ precision matrices. To reduce the computation complexity, it can be further decomposed into p sub-problems that involve estimating K $p \times 1$ sparse vectors:

$$
\min_{\beta_j^{(k)} \in \mathbb{R}^p, 1 \le k \le K} \left(\max_k |\beta_1^{(k)}|_1 \right),
$$
$$
\text{subject to} \max_i \left\{ \sum_{k=1}^K w_k |(\hat{\sum}^{(k)} \beta_j^{(k)} - e_j)_i|^2 \right\}^{1/2} \le \lambda_n \tag{2}
$$

for $1 \le j \le p$, where $e_j \in \mathbb{R}^p$ is the unit vector with j-th element 1 and other elements 0. A lemma shows that solving (2) is equivalent to solving (1).

Lemma 1—Suppose $\hat{\Omega}_1^{(k)}$ is the solution to (1) and $\hat{B}^{(k)} := (\hat{\beta}_1^{(k)}, \dots, \hat{\beta}_p^{(k)})$, where $\hat{\beta}_j^{(k)}$ is the solution to (2). Then $\hat{\Omega}_1^{(k)} = \hat{\mathbf{B}}^{(k)}$ for $1 \le k \le K$.

Problem (2) can be solve by a second-order cone programming. The existing packages to solve (2) include the SDTP3 and the SeDuMi package in Matlab, and the CLSOCP package in R. CLSOCP uses a one-step smoothing Newton method of Liang, He and Hu (2009). This algorithm has good precision but works relatively slowly for high dimensional problems. SeDuMi and SDTP3 adopt the primal-dual infeasible-interior point algorithm (Newsterov and Todd (1998)). The most time-consuming part of the algorithm is to solve the Schur complement equation, which involves Cholesky factorization. The sparsity and the size of the Schur complement matrix are two factors that affect efficiency. SDTP3 is able to divide a high dimensional optimization problem into sparse blocks and uses the sparse solver for Cholesky factorizations. It is therefore faster than SeDuMi in solving (2). In this paper, we used the SDTP3 package for all the computations. For a problem with $p = 200$, $n_k = 150$ and $K = 3$, it takes a dual-core 2.7 GHz Intel Core i7 laptop approximately 11 minutes to solve (1).

2.3 Tuning Parameter Selection

Choosing the tuning parameters in regularized estimation is in general a difficult problem. For linear regression models, Chen and Chen (2008); Wang, Li and Leng (2009); Wang and Zhu (2011) studied how to consistently choose the tuning parameters when $p = O(n^a)$ for some $a > 0$. Recently, Fan and Tang (2013) proposed a general information criterion (GIC) for choosing the tuning parameter for estimating the generalized linear model in ultra-high dimensional settings, $p = O(\exp(n^a))$ for some $a > 0$. The GIC criterion adopts a novel penalty on the degree freedom of the model so that it consistently chooses the proper tuning parameter under mild conditions. Unfortunately, the Gaussian graphical model is different from the generalized linear model, and therefore the justification of GIC does not apply to our problem. We propose a tuning parameter selection method based on BIC and a re-fitted precision matrix on the restricted model.

Since $\boldsymbol{X}_1^{(k)}(1 \leq l \leq n_k)$ follows a multivariate Gaussian distribution $N_p(\boldsymbol{\mu}^{(k)}, \boldsymbol{\Sigma}^{(k)})$, we have

$$X_{l,i}^{(k)} = \alpha_i^{(k)} + \boldsymbol{X}_{l,-i}^{(k)'} \boldsymbol{\beta}_i^{(k)} + \varepsilon_{l,i}^{(k)}, \quad (3)$$

where $\varepsilon_{l,i}^{(k)} \sim N(0, 1/\omega_{ij}^{(k)})$. The regression coefficients satisfy (Anderson (2003))

$$\alpha_i^{(k)} = \mu_i^{(k)} - \sum\nolimits_{i,-i}^{(k)} \left(\sum\nolimits_{-i,-i}^{(k)}\right)^{-1} \boldsymbol{\mu}_{-i}^{(k)}, \quad \boldsymbol{\beta}_i^{(k)} = -\left(\omega_{ii}^{(k)}\right)^{-1} \boldsymbol{\Omega}_{-i,i}^{(k)}. \quad (4)$$

Based on these results, we propose a tuning parameter selection procedure:

1. For a given λ, calculate the estimator $\hat{\boldsymbol{\Omega}}^{(k)}$. Based on the support of $\hat{\boldsymbol{\Omega}}^{(k)}$, use least squares and neighborhood selection to re-fit the precision matrix estimator $\hat{\boldsymbol{\Omega}}_2^{(k)}$.

2. Define $\mathscr{S}_i^{(k)} = \{j : \hat{\omega}_{ij}^{(k)} \neq 0, j \neq i\}$, the set of non-zero off-diagonal elements of the i-th row of $\hat{\boldsymbol{\Omega}}^{(k)}$.

3. If $\mathrm{card}\,(\mathscr{S}_i^{(k)}) \geq n_k$, let the i-th column of $\hat{\boldsymbol{\Omega}}_2^{(k)}$ equal to the i-th column of $\hat{\boldsymbol{\Omega}}_{\cdot j}^{(k)}, \hat{\boldsymbol{\Omega}}_{2,\cdot i}^{(k)} = \hat{\boldsymbol{\Omega}}_{\cdot i}^{(k)}$. If $\mathrm{card}\,(\mathscr{S}_i^{(k)}) < n_k$, fit the regression model

    $$X_{l,i}^{(k)} = \alpha_i^{(k)} + \sum_{j \in \mathscr{S}_i^{(k)}} X_{l,j}^{(k)'} \beta_{ij}^{(k)} + \varepsilon_{l,i}^{(k)}. \quad (5)$$

    If $\mathscr{S}_i^{(k)}$ equals the true support $\mathscr{S}_{0,i}^{(k)} = \{j : \omega_{0,ij}^{(k)} \neq 0, j \neq i\}$, $\beta_{ij}^{(k)} = -\omega_{0,ij}^{(k)}/\omega_{0,ii}^{(k)}$ and $\mathrm{Var}\,(\varepsilon_{l,i}^{(k)}) = 1/\omega_{0,ii}^{(k)}$. Suppose the solution to (5) is $\hat{\beta}_{ij}^{(k)}$. Define

    $$\hat{\varepsilon}_{l,i}^{(k)} = (X_{l,i}^{(k)} - \overline{X}_i^{(k)}) - \sum_{j \in \mathscr{S}_i^{(k)}} (X_{l,j}^{(k)} - \overline{X}_j^{(k)})' \hat{\beta}_{ij}^{(k)}.$$

    After fitting Model (5), let $\hat{\omega}_{2,ii}^{(k)} = n_k / \sum_{l=1}^{n_k} \left(\varepsilon_{l,i}^{(k)}\right)^2$, and $\hat{\omega}_{2,ij}^{(k)} = -\hat{\beta}_{ij}^{(k)} \hat{\omega}_{2,ii}^{(k)}$.

4. Repeat Step 3 for $i = 1,\ldots, p$ and $k = 1,\ldots, K$. The resulting matrices $\hat{\boldsymbol{\Omega}}_2^{(k)}$, $k = 1$, $\ldots, K$ are not symmetric. We symmetrize $\hat{\boldsymbol{\Omega}}_2^{(k)}$ by the same procedure:

    $$\hat{\omega}_{3,ij}^{(k)} = \hat{\omega}_{3,ji}^{(k)} := \hat{\omega}_{2,ij}^{(k)} I(|\hat{\omega}_{2,ij}^{(k)}| \leq |\hat{\omega}_{2,ji}^{(k)}|) + \hat{\omega}_{2,ji}^{(k)} I(|\hat{\omega}_{2,ij}^{(k)}| > |\hat{\omega}_{2,ji}^{(k)}|).$$

We use $\hat{\boldsymbol{\Omega}}_3^{(k)} = (\hat{\omega}_{3,ij}^{(k)})$, $k = 1,\ldots, K$ as the estimators corresponding to the tuning parameter λ. The optimal tuning parameter can be selected by Bayesian information criterion (BIC),

$$\text{BIC}(\lambda) = \sum_{k=1}^{K} \left\{ n_k \text{tr}\left(\hat{\sum}^{(k)} \hat{\boldsymbol{\Omega}}_3^{(k)} \right) - n_k \log(\det \hat{\boldsymbol{\Omega}}_3^{(k)}) + \log(n_k) s_k \right\}, \tag{6}$$

where $s_k = \mathsf{Card}\ \{(i,j) : \hat{\omega}_{i,j} \neq 0, 1 \le i < j \le p\}$. We obtain the solution to our method over a wide range of tuning parameters and choose $\hat{\lambda}_n$ that minimizes $\text{BIC}(\lambda)$.

Since the refitted estimator can potentially reduce some bias introduced in the optimization due to the penalty term, using it in BIC (6) improves the tuning parameter selection in the numerical studies. However, using $\hat{\boldsymbol{\Omega}}_3^{(k)}$ as an estimator of $\boldsymbol{\Omega}^{(l)}$ is not recommended because when $\mathsf{Card}\ (\mathscr{S}_i^{(k)})$ is large, the re-fitted estimator might lead to overfitting. Overfitting does not severely affect the tuning parameter selection, because BIC puts penalties on complicated models that are less likely to be chosen.

3. THEORETICAL PROPERTIES

3.1 Estimation Error Bound

We investigate the properties of the proposed estimator by considering the convergence rates of $\hat{\boldsymbol{\Omega}}^{(k)} - \boldsymbol{\Omega}^{(k)}$, including estimation error bounds and graph structure recovery. We assume the following conditions:

(C1)   There exists some constant $a > 0$, such that

$$\log p = o\left(\frac{n}{K^{2a}(\log n)^2} \right), \quad \text{and}\quad \max(K, K^{4-a}\log K) = o(\log p).$$

(C2) $\sup_{1 \le k \le K} \{\lambda_{\max}(\boldsymbol{\Omega}^{(k)})/\lambda_{\min}(\boldsymbol{\Omega}^{(k)})\} \le M_0$ for some bounded constant $M_0 > 0$.

(C3)   If $e_{ij} = I(i = j)$, $Z_{ij}^{(k)} = (\boldsymbol{X}^{(k)}\boldsymbol{X}^{(k)'}\boldsymbol{\Omega}^{(k)})_{ij} - e_{ij}$, $Z_{ij} = (Z_{ij}^{(1)}, \ldots, Z_{ij}^{(K)})'$ and the largest eigenvalue of $\text{Cov}(Z_{ij})$ is $\lambda_{\max,ij}$, $\sup_{ij} \lambda_{\max,ij} \le M_1$.

(C4) $n_1 \asymp n_2 \asymp \cdots \asymp n_K \asymp n/K$.

Condition (C1) allows p to grow exponentially fast as n. It also allows the number of groups K to grow slowly with p and n. For example, when $\log p = O(n^r)$ and $K = O(n^b)$ for $r + 2b < 1$ and $3b < r$, (C1) holds. Condition (C3) allows $\boldsymbol{X}^{(k)}$ to be dependent across groups. When $\boldsymbol{X}^{(k)}$ are independent, $\text{Cov}(Y_{ij}) = \boldsymbol{I}_K$, and thus $\max_{ij} \lambda_{\max,ij} = 1$.

Let $M_n = \sup_{1 \le k \le K} \max_j \sum_{i=1}^{p} |\omega_{ij}^{(k)}| = \sup_{1 \le k \le K} \|\boldsymbol{\Omega}^{(k)}\|_1$ be the maximum matrix ℓ_1 norms of the K matrices. A theorem establishes the convergence rate of the precision matrix estimates under the element-wise ℓ_∞ norm.

Theorem 1—Let $\lambda_n = C_0 (\log p/n)^{1/2}$ for some constant $C_0 > \sqrt{2M_0+2}$. If (*C1*)–(*C4*) hold,

$$\sup_{i,j}\left\{\sum_{k=1}^{K} w_k |(\hat{\boldsymbol{\Omega}}^{(k)} - \boldsymbol{\Omega}^{(k)})_{i,j}|^2\right\}^{1/2} \leq C_1 M_n \left(\frac{\log K \cdot \log p}{n}\right)^{1/2} \tag{7}$$

with a high probability converging to *1* and $C_1 = 2C_0$.

> Remark 1: The value of C_0 depends on M_0. In practice, M_0 is often unknown. However, we can use the tuning parameter selection method, such as BIC in (6), to choose λ_n. The details are discussed in Section 2.3.

> Remark 2: Theorem 1 (and Theorems 2 and 3) does not require the true precision matrices $\boldsymbol{\Omega}^{(k)}$ to have identical graphical structures. Both the values and locations of non-zero entries can differ across $\boldsymbol{\Omega}^{(k)}$, $k = 1,\dots, K$.

Define

$$\mathscr{U}(M) = \Big\{(\boldsymbol{\Omega}^{(1)},\dots,\boldsymbol{\Omega}^{(K)}):\lambda_{\max}(\boldsymbol{\Omega}^{(k)})/\lambda_{\min}(\boldsymbol{\Omega}^{(k)}) \leq M,$$
$$\|\boldsymbol{\Omega}^{(k)}\|_1 ≤ M_n = o\{(n/\log p)^{1/2}\}, k=1,\dots,K\Big\}.$$

Proposition 1—Let $\tilde{\mathscr{U}}$ be the set of estimators $(\tilde{\boldsymbol{\Omega}}^{(1)},\dots,\tilde{\boldsymbol{\Omega}}^{(K)})$, where $\tilde{\boldsymbol{\Omega}}^{(k)}$ only depends on the k-th sample $\left\{\boldsymbol{X}_j^{(k)}; 1 \leq j \leq n_k\right\}$. Assume the samples are independent across K groups. Under (*C4*), there exists a constant $a > 0$, such that

$$\inf_{i,j}\ \inf_{(\tilde{\boldsymbol{\Omega}}^{(1)},\dots,\tilde{\boldsymbol{\Omega}}^{(K)})\in\tilde{\mathscr{U}}}\ \sup_{(\boldsymbol{\Omega}^{(1)},\dots,\boldsymbol{\Omega}^{(K)})\in\mathscr{U}} P\left[\left\{\sum_{k=1}^{K} w_k |(\tilde{\boldsymbol{\Omega}}^{(k)} - \tilde{\boldsymbol{\Omega}}^{(k)})_{i,j}|^2\right\}^{1/2} \geq \alpha M_n \left(\frac{K \cdot \log p}{n}\right)^{1/2}\right] \geq \alpha^K,$$

for sufficiently large n.

Theorem 1 shows that the convergence rate of $\left\{\sum_{k=1}^{K} w_k |(\tilde{\boldsymbol{\Omega}}^{(k)} - \boldsymbol{\Omega}^{(k)})_{i,j}|^2\right\}^{1/2}$ from our joint estimation method is less than or equal to $C_1 M_n (\log K \log p/n)^{1/2}$ with a probability tending to 1. When K is bounded, Proposition 1 shows that with a non-vanishing probability a^K, the minimax convergence rate of any separate estimation method is at least $aM_n(K \log p/n)^{1/2}$. For bounded but sufficiently large K, $C_1 (\log K)^{1/2} \quad aK^{1/2}$. Therefore, the

convergence upper bound $C_1 M_n \left(\frac{\log K \cdot \log p}{n}\right)^{1/2}$ in (7) for the joint estimation method is

less than the convergence lower bound $\alpha M_n \left(\frac{K \cdot \log p}{n}\right)^{1/2}$ in Proposition 1 for the separate estimation method. In other words, compared to separate estimation methods, our joint estimation method leads to estimators with a faster convergence.

An additional thresholding step on the estimators with a careful chosen threshold leads to new estimators, that converge to the true precision matrices under the matrix operator norm.

Define the thresholded estimator $\breve{\boldsymbol{\Omega}}^{(k)} = (\breve{\omega}_{ij}^{(k)})$ as follows,

$$
\breve{\omega}_{ij}^{(k)} = \hat{\omega}_{ij}^{(k)} I \left\{ \left(\sum_{k=1}^{K} w_k (\hat{\omega}_{ij}^{(k)})^2 \right)^{1/2} > C_1 M_n \left(\frac{\log K \cdot \log p}{n} \right)^{1/2} \right\}
$$

with C_1 the constant defined in (7). Let $\mathscr{S}_j^{(k)} = \{(i,j) : \omega_{ij}^{(k)} \neq 0, i < j\}$ and $\mathscr{S}_j = \cup_{k=1}^{k} S_j^{(k)}$. Define $s_0(p) = \max_{1 \leq j \leq p} \text{Card}(\mathscr{S}_j)$ as the union sparsity.

Theorem 2—Suppose that (*C1*)–(*C4*) hold. Then with a high probability converging to 1,

$$
\max_j \sum_{i=1}^{p} \left\{ \sum_{k=1}^{K} w_k (\breve{\boldsymbol{\Omega}}^{(k)} - \boldsymbol{\Omega}^{(k)})_{ij}^2 \right\}^{1/2} \leq C_1 M_n s_0(p) \left(\frac{\log K \cdot \log p}{n} \right)^{1/2}. \tag{8}
$$

The convergence rates of $\breve{\boldsymbol{\Omega}}^{(k)}$ depend on the union sparsity level $s_0(p)$. When the precision matrices share the same graphical structure, $s_0(p) = \max_{1 \leq j \leq p} \text{Card}(\mathscr{S}_j^{(k)})$ for all $k = 1$, ..., K. When the number of shared edges in the graphical structures increases, the union sparsity $s_0(p)$ decreases, and consequently $\breve{\boldsymbol{\Omega}}^{(k)}$ converges to $\boldsymbol{\Omega}^{(k)}$ faster.

Let $\hat{a}_{ij} = \sqrt{\sum_{k=1}^{K} w_k (\breve{\boldsymbol{\Omega}}^{(k)} - \boldsymbol{\Omega}^{(k)})_{ij}^2}$. Matrix $\hat{\mathbf{A}} = (\hat{a}_{ij})_{p \times p}$ measures the overall estimation errors among the entries of $\breve{\boldsymbol{\Omega}}^{(k)}$ and $\boldsymbol{\Omega}^{(k)}$ for $k = 1, ..., K$.

Corollary 1—Suppose that the conditions in Theorem 2 hold. Then with a high probability converging to 1,

$$
\|\hat{\mathbf{A}}\|_2 \leq \|\hat{\mathbf{A}}\|_1 \leq C_1 M_n s_0(p) \left(\frac{\log K \cdot \log p}{n} \right)^{1/2}.
$$

3.2 Graphical Structure Recovery

Theoretical analysis for graphical structure recovery is very complicated when the graphical structures of the precision matrices are different across the K groups since the results depend on the structures of the shared edges. Here we focus on the case in which the K precision matrices have a common support. Let $\mathscr{S}_k = \{(i,j) : \omega_{ij}^{(k)} \neq 0\}$ be the support for the k-th precision matrix, and the common support be $\mathscr{S} = \cap_k \mathscr{S}_k$. When $\mathscr{S}_1 = \ldots = \mathscr{S}_K$, $\mathscr{S} = \mathscr{S}_k$, $k = 1, ..., K$, by Theorem 1 we estimate $\mathscr{S}$ by

$$\hat{\mathscr{S}} = \left[(i,j) : \left\{ \sum_{k=1}^{K} w_k \left(\hat{\omega}_{ij}^{(k)} \right)^2 \right\}^{1/2} > C_1 M_n \left(\frac{\log K \cdot \log p}{n} \right)^{1/2} \right],$$

where C_1 is a constant given in Theorem 1. Let

$$\theta_n = \min_{(i,j) \in \mathscr{S}} \left\{ \sum_{k=1}^{K} w_k \left(\omega_{ij}^{(k)} \right)^2 \right\}^{1/2}.$$

We have a result on support recovery.

Theorem 3—Suppose that the conditions in Theorem 1 hold. Assume that

$$\theta_n > 2 C_1 M_n \left(\frac{\log K \cdot \log p}{n} \right)^{1/2}. \tag{9}$$

Then $\hat{\mathscr{S}} = \mathscr{S}$ with a high probability converging to 1.

When the graphical structures are the same across all K groups, the lower bound condition (9) is weaker than the lower bound condition needed for graphical structure recovery by separate estimation methods. Based on Proposition 1 and its proof, to fully recover the shared graphical structure by separate estimation methods, a necessary condition is

$$\inf_{(i,j) \in \mathscr{S}} \inf_k |\omega_{ij}^{(k)}| > 2\alpha M_n (K \log p / n)^{1/2}.$$

When K is sufficiently large, this condition is stronger than (9).

4. SIMULATION STUDIES

4.1 Data generation

We evaluated the numerical performance of the proposed method and other competitive methods, including the separate precision matrix estimation procedures proposed by Friedman, Hastie and Tibshirani (2008) and Cai, Liu and Luo (2011) and the joint estimation method proposed by Guo, Levina, Michailidis et al. (2011) and Danaher, Wang and Witten (2014). The separate precision matrix estimation methods were applied to each group, and therefore ignored the partial homogeneity in graphical structures among groups. In all numerical studies, we set $p = 200$, $K = 3$ and $(n_1, n_2, n_3) = (80, 120, 150)$. The simulated observations were generated in each group independently from a multivariate Gaussian distribution $N\{0, (\Omega^{(k)})^{-1}\}$, where $\Omega^{(k)}$ is the precision matrix in the k-th group. For each model, 100 replications were performed.

We present results for two different types of graphical models: the Erdös and Rényi (ER) model (Erdös and Rényi (1960)) and the Watts-Strogatz (WS) model (Watts and Strogatz (1998)). For the ER model, the graph contains p vertices and each pair of vertices are connected with a probability 0.05. For the WS model, first a ring lattice of p vertex is created; one vertex is connected with its neighbors within order distance of 15, and then the edges of the lattice are rewired uniformly and randomly with a probability 0.01. These graph models have several topological properties such as sparcity and a "small world" property often observed in true biological gene regulatory networks. See Fell and Wagner (2000); Jeong, Tombor, Albert et al. (2000); Vendrascolo, Dokholyan, Paci et al. (2002); Greene and Higman (2003).

Based on the ER model or the WS model, a common graph structure was generated. Let M be the number of edges in the common graph structure. Then, $\lfloor \rho M \rfloor$ random edges were added to the common graph structure to generate graph structures for each group, where the parameter ρ was the ratio between the number of group-specific edges and common edges. We considered $\rho = 0$, 1/4, and 1. The first setting represents the scenario where the precision matrices in all groups share the same graph structure. After the graph structure of each group was determined, the values of non-zero off-diagonal entries were generated independently as uniform in $[-1, -0.5] \cup [0.5, 1]$. The diagonal values were assigned to a constant so that the condition number of each precision matrix was equal to p.

4.2 Simulation results

Each method was evaluated for a range of tuning parameters under each model. The optimal tuning parameter was chosen by BIC (6). Several measures are used to compare the performances of these estimators. The estimation error was evaluated in terms of average matrix ℓ_1 norm, ℓ_2 norm (spectral norm), and Frobenius norm:

$$L_1 = \frac{1}{K} \sum_{k=1}^{K} \|\hat{\mathbf{\Omega}}^{(k)} - \mathbf{\Omega}_0^{(k)}\|_1,$$
$$L_2 = \frac{1}{K} \sum_{k=1}^{K} \|\hat{\mathbf{\Omega}}^{(k)} - \mathbf{\Omega}_0^{(k)}\|_2,$$
$$L_F = \frac{1}{K} \sum_{k=1}^{K} \|\hat{\mathbf{\Omega}}^{(k)} - \mathbf{\Omega}_0^{(k)}\|_F.$$

The graph structure recovery results were evaluated by average sensitivity (SEN), specificity (SPE), and Matthews correlation coefficient (MCC). For a true precision matrix $\mathbf{\Omega}_0 = (\omega_{0,ij})$ with support set $\mathscr{S}_0 = \{(i,j) : \omega_{0,ij} \neq 0 \text{ and } i \neq j\}$, suppose its estimator $\hat{\mathbf{\Omega}}$ has the support set $\hat{\mathscr{S}}$. Then the measures with respect to $\mathbf{\Omega}_0$ and $\hat{\mathbf{\Omega}}$ are defined as follows:

$$\text{SPE} = \frac{\text{TN}}{\text{TN+FP}}, \quad \text{SEN} = \frac{\text{TP}}{\text{TP+FN}},$$
$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\{(\text{TP+FP})(\text{TP+FN})(\text{TN+FP})(\text{FP+FN})\}^{1/2}}.$$

Here, TP, TN, FP, and FN are the numbers of true positives, true negatives, false positives and false negatives:

$$\text{TP} = \#\{(i,j) : (i,j) \in \mathscr{S}_0 \cap \hat{\mathscr{S}}\}, \quad \text{TN} = \#\{(i,j) : (i,j) \in \mathscr{S}_0^C \cap \hat{\mathscr{S}}^C\},$$
$$\text{FP} = \#\{(i,j) : (i,j) \in \mathscr{S}_0^C \cap \hat{\mathscr{S}}\}, \quad \text{FN} = \#\{(i,j) : (i,j) \in \mathscr{S}_0 \cap \hat{\mathscr{S}}^C\}.$$

We compared $\hat{\boldsymbol{\Omega}}^{(k)}$ and $\boldsymbol{\Omega}_0^{(k)}$ and report the average sensitivities (SEN), specificities (SPE), and Matthews correlation coefficient (MCC) among K groups.

The comparisons of the results for the four graphical models are shown in Tables 1 and 2. It shows that when $\rho = 0$, the true graph structures are the same across all three groups, joint estimation methods perform much better than the separate estimation methods. As ρ increases, the structures across different groups become more different, and the joint estimation methods gradually lose their advantages. Our method has the best performance in graph structure recovery among all the methods. Even when $\rho = 1$, it still performs significantly better than the separate estimation methods. Our method has the best performance in graph structure recovery and it achieves the highest Matthews correlation coefficient. Its estimation error measured under the matrix ℓ_1, ℓ_2, and the Frobenius norm are comparable to other joint estimation methods.

Since the tuning parameter selection may affect the performance of the methods, we plot in Figure 1 the receiver operating characteristic (ROC) curves averaged over 100 repetitions with the false positive rate controlled under 10%. The methods proposed by Danaher, Wang and Witten (2014) have two tuning parameters. For each sparsity tuning parameter, we first chose an optimal similarity tuning parameter from a grid of candidates by BIC and then plotted the ROC curves based on a sequence of sparsity tuning parameters and their corresponding optimal similarity tuning parameters. Because these methods involve choosing two tuning parameters, they are slower than our method in computation (The computation of FGL and GGL is based on the R package "JGL" contributed by Danaher (2013)). Figure 1 shows that our method consistently outperforms the other methods in graph structure recovery.

5. EPITHELIAL OVARIAN CANCER DATA ANALYSIS

Epithelial ovarian cancer is a molecularly diverse cancer that lacks effective personalized therapy. Tothill, Tinker, George et al. (2008) identified six molecular subtypes of ovarian cancer, labeled as C1–C6, where the C1 subtype was characterized by significant differential expressions of genes associated with stromal and immune cell types. The patients in C1 subtype group have shown to have a lower survival rate compared to the patients from other subtypes. The data set includes RNA expression data collected from $n = 78$ patients of C1 subtype and $n = 113$ patients from the other subtypes. We are interested to see how the wiring (conditional dependency) of the genes at the transcription levels differs among molecular subgroups of ovarian cancer. We focus on the apoptosis pathway from the KEGG database (Orgata, Goto, Sato et al. (1999); Kanehisa, Goto, Sato et al. (2012)) to see whether the genes related to this pathway ($p = 87$) are differentially wired between the C1 and other subtypes.

To stabilize the graph structure selection, we bootstraped the samples 100 times within each group. At each time, I_{ik} was sampled uniformly taking values in $i = \{1,\ldots, n_k\}$, with $k = 1, 2$. Let $\tilde{\mathbf{x}}_i^{(k)} = \mathbf{x}_{I_{ik}}^{(k)}$, where $\mathbf{x}_{I_{ik}}^{(k)}$ is the p-dimensional gene expression data for the I_{ik}-th patient in the kth subtype group. The bootstrap sample is $\tilde{\mathbf{X}}^{(k)} = (\tilde{\mathbf{x}}_1^{(k)}, \ldots, \tilde{\mathbf{x}}_{n_k}^{(k)})$, with $k = 1, 2$. We then applied our proposed method and its competitors to each of the bootstrapped samples to obtain the estimators of the two precision matrices $\hat{\boldsymbol{\Omega}}^{(k)}$, $k = 1, 2$. The supports of the estimators were recorded so that $\tilde{\boldsymbol{\Omega}}^{(k)} = (I(\hat{\omega}_{ij}^{(k)} \neq 0))$. We then added $\hat{\boldsymbol{\Omega}}^{(k)}$ up for all bootstrap samples and got the final frequency of each edge being selected. Those edges that were selected in more than 50 of the 100 bootstrap samples were finally selected as important edges. This type of bootstrap aggregation methods has been commonly used in recovering the sparse graphical structures (Meinshausen and Bühlmann (2010); Li, Hsu, Peng et al. (2013)), which often leads to better selection stability for sparse precision matrix.

Table 3 lists the number of edges selected by the bootstrap aggregation of our proposed method and its competitors. The separate estimation methods (CLIME and GLASSO) resulted in graphs that share fewer edges in the precision matrices of the two cancer subtype groups. JEMGM resulted in most shared edges, followed by GGL and our method (MPE). Overall, FGL and GGL selected a lot more linked genes than other methods. Figure 2 shows the Gaussian graphs estimated by these six methods. FGL, GGL and MPE selected more unique edges among the gene expression levels for the C2–C6 subtype cancer than those for the C1 subtype. This suggests that the patients with poor prognostic subtype (C1) lack some important links among the Apoptosis genes.

We further defined the nodes with degrees equal or larger than five based on the union of the estimated graphs of two subtypes as the hub nodes. FGL and GGL yielded estimators with most of the hub nodes completely unlinked in the estimated graph for C1 cancer subtype. The estimators by MPE had several edges between the hub nodes shared by both subtype groups, while also displaying some links unique to each group. The hub nodes identified by MPE were FASLG, CASP10, CSF2RB, IL1B, MYD88, NFKB1, NFKBIA, PIK3CA, IKBKG, and PIK3R5. Among these, CASP10, PIK3CA, IL1B, and NFKb1 have been implicated in ovarian cancer risk or progression. In particular, PIK3CA has been implicated as an oncogene in ovarian cancer (Shayesteh, Lu, Kuo et al. (1999)), indicating the importance of these hub genes in ovarian cancer progression.

6. DISCUSSION

It is of interest to discuss the connection and difference between the problem considered in this paper and the problem of estimating matrix graphical models (Leng and Tang (2012); Yin and Li (2012); Zhou (2014)). Matrix graphical models consider the random matrix variate $\mathbf{X}$ following the distribution $MN_{p \times q}(\mathbf{M}; \mathbf{U}, \mathbf{V})$ with the probability density function (pdf)

$$p(\mathbf{X}|\mathbf{M}, \mathbf{U}, \mathbf{V}) = (2\pi)^{-pq/2}|\mathbf{U}^{-1}|^{q/2}|\mathbf{V}^{-1}|^{q/2}\exp\{-\mathrm{t\,r}\,[(\mathbf{X} - \mathbf{M})'\mathbf{V}^{-1}(\mathbf{X} - \mathbf{M})'\mathbf{U}^{-1}/2]\}.$$

Here $\mathbf{M} \in \mathbb{R}^{p \times q}$ is the mean matrix, $\mathbf{U} \in \mathbb{R}^{p \times p}$ is the row covariance matrix, and $\mathbf{V} \in \mathbb{R}^{q \times q}$ is the column covariance matrix. Thus, each row and each column of $\mathbf{X}$ share the same covariance matrices $\mathbf{U}$ and $\mathbf{V}$. This distribution implies that vec($\mathbf{X}$) follows a vector multivariate Gaussian distribution $N_{pq}(\text{vec}(\mathbf{M}), \mathbf{U} \otimes \mathbf{V})$, where "vec" is the vectorization operator and "$\otimes$" is the Kronecker product.

Our model assumes $X^{(k)}$ follows the vector multivariate Gaussian distribution $N_p(\boldsymbol{\mu}, \mathbf{\Sigma}^{(k)})$. If $n = n_1 = \ldots = n_K$, $\mathbf{X} = (X^{(1)}, \ldots, X^{(K)}) \in \mathbb{R}^{p \times K}$ is a random matrix variate. However, each column of $\mathbf{X}$ has its own covariance matrix $\mathbf{\Sigma}^{(k)}$, and each row of $\mathbf{X}$ may also have its own. Therefore, the covariance matrix of vec($\mathbf{X}$) cannot be expressed as the Kronecker product of two positive definite matrices. In general, the degree of freedom of $\text{Cov}(\text{vec}(\mathbf{X}))$ is larger than that of $\mathbf{U} \otimes \mathbf{V}$ mentioned above.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

Anderson, TW. An introduction to multivariate statistical analysis. Wiley-Interscience; 2003.

Banerjee O, Ghaoui LE, d'Aspremont A. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. J Machine Learning Research. 2008; 9:485–516.

Cai T, Liu W, Luo X. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. Journal of American Statistical Association. 2011; 106:594–607.

Candés E, Tao T. The dantzig selector: Statistical estimation when p is much larger than n. The Annals of Statistics. 2007; 35:2313–2351.

Chen J, Chen Z. Extended bayesian information criterion for model selection with large model space. Biometrika. 2008; 95:232–253.

Dahl J, Vandenberghe L, Roychowdhury V. Covariance selection for non-chordal graphs via chordal embedding. Optimization Methods and Software. 2008; 23:501–420.

Danaher P. JGL: Performs the Joint Graphical Lasso for sparse inverse covariance estimation on multiple classes. R package version 2.3. 2013

Danaher P, Wang P, Witten D. The joint graphical lasso for inverse covariance estimation across multiple classes. Journal of the Royal Statistical Society, Series B. 2014:76–2.

Donoho D, Elad M, Temlyakov V. Stable recovery of sparse overcomplete representations in the presence of noise. IEEE Transactions on Information Theory. 2006; 52:6–18.

Erdös P, Rényi A. On the evolution of random graphs. Publications of the Mathematical Institute of the Hungarian Academy of Sciences. 1960; 5:17–61.

Fan Y, Tang C. Tuning parameter selection in high dimensional penalized likelihood. Journal of the Royal Statistical Society. 2013; 75:531–552.

Fell D, Wagner A. The small world of metabolism. Nature Biotechnology. 2000:18–11.

Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. Biostatistics. 2008; 9:432–441. [PubMed: 18079126]

Greene L, Higman V. Uncovering network systems within protein structures. Journal of Molecular Biology. 2003; 334:781–791. [PubMed: 14636602]

Guo J, Levina E, Michailidis G, Zhu J. Joint estimation of mulitple graphical models. Biometrika. 2011; 98:1–15. [PubMed: 23049124]

Jeong H, Tombor B, Albert R, Oltvai Z, Barabási A. The large-scale organization of metabolic networks. Nature. 2000:407–6804.

Kanehisa M, Goto S, Sato Y, Furumichi M, Tanabe M. KEGG for integration and interpretation of large-scale molecular data sets. Nucleic Acids Research. 2012; 40:D109–D114. [PubMed: 22080510]

Lauritzen, SL. Graphical Models. Clarendon Press; Oxford: 1996.

Leng C, Tang C. Sparse matrix graphical models. Journal of American Statistical Association. 2012:107–499.

Li H, Gui J. Gradient directed regularization for sparse gaussian concentration graphs, with applications to inference of genetic networks. Biostatistics. 2006; 7:302–317. [PubMed: 16326758]

Li S, Hsu L, Peng J, Wang P. Bootstrap inference for network construction with an application to a breast cancer microarray study. The Annals of Applied Statistics. 2013:7–1.

Liang F, He G, Hu Y. A new smoothing Newton-type method for second-order cone programming problems. Applied Mathematics and Compuation. 2009; 215:1020–1029.

Meinshausen N, Bühlmann P. High-dimensional graphs and variable selection with the lasso. Annals of Statistics. 2006; 34:1436–1462.

Meinshausen N, Bühlmann P. Stability selection. Journal of Royal Statistics Society, Series B. 2010; 72:417–473.

Newsterov Y, Todd M. Primal-dual interior-point methods for self-scaled cones. SIAM Journal on Optimization. 1998; 8:324–364.

Orgata H, Goto S, Sato K, Fujibuchi W, Bono H, Kanehisa M. KEGG: Kyoto encyclopedia of genes and genomes. Nucleic Acids Research. 1999; 27:29–34. [PubMed: 9847135]

Schäfer J, Strimmer K. An empirical bayes approach to inferring large-scale gene association networks. Bioinformatics. 2005; 21:754–764. [PubMed: 15479708]

Shayesteh L, Lu Y, Kuo W, Baldocchi R, Godfrey T, Collins C, Pinkel D, Powell B, Mills G, Gray J. Pik3ca is implicated as an oncogene in ovarian cancer. Nature Genetics. 1999; 21:99–102. [PubMed: 9916799]

Tothill R, Tinker A, George J, Brown R, Fox S, Lade S, Johnson D, Trivett J, Etemadmoghadam D, Locandro B, Traficante N, Fereday S, Hung J, Chiew Y, Haviv I, Group A. O. C. S. Gertig D, deFazio A, Bowtell D. Novel molecular subtypes of serous and endometrioid ovarian cancer linked to clinical outcome. Clinical Cancer Research. 2008; 14:5198–5208. [PubMed: 18698038]

Vendrascolo M, Dokholyan N, Paci E, Karplus M. Small-world view of the amino acids that play a key role in protein folding. Physical Review E. 2002; 65

Wang H, Li B, Leng C. Shrinkage tuning parameter selection with a diverging number of parameters. Journal of Royal Statistical Society, Series B. 2009; 71:671–683.

Wang T, Zhu L. Consistent tuning parameter selection in high dimensional sparse linear regression. Journal of Multivariate Analysis. 2011; 102:1141–1151.

Watts D, Strogatz S. Collective dynamics of "small world" networks. Nature. 1998; 393:440–442. [PubMed: 9623998]

Whittaker, J. Graphical Models in Applied Multivariate Analysis. Wiley; 1990.

Yin J, Li H. Model selection and estimation in the matrix normal graphical model. Journal of Multivariate Analysis. 2012; 107:119–140. [PubMed: 22368309]

Yuan M. Sparse inverse covariance matrix estimation via linear programming. Journal of Machine Learning Research. 2010; 11:2261–2286.

Yuan M, Lin Y. Model selection and estimation in the gaussian graphical model. Biometrika. 2007; 94:19–35.

Zhou S. Gemini: Graph estimation with matrix variate normal instances. Annals of Statistics. 2014; 42:532–562.

Figure 1.
Receiver operator characteristic curves for graph structure recovery for the simulated Erdős and Rényi graphs (the first row), and the Watts-Strogatz graphs (the second row). The x-axis and y-axis of each panel are average false positive rate and average sensitivity across $K = 3$ groups. Black solid line: CLIME; black dot-dashed line: GLASSO; black long-dashed line: JEMGM; grey solid line: FGL; grey dot-dashed line: GGL; grey long-dashed line: MPE.

Figure 2.
Estimated Gaussian Graphs by the proposed method and its competitors. The dashed edges are links unique to the precision matrix estimator of the C1 subtype, the dotted edges are unique to that of other subtypes, and the solid edges are shared by both estimators. The size of node is a linear function of its degree. Upper left panel: CLIME; upper middle panel: GLASSO; upper right panel: JEMGM; bottom left panel: FGL; bottom middle panel: GGL; bottom right panel: MPE.

Table 1

Simulation results for data generated based on the Erdös and Rényi graph with different ratios of the number of individual-specific edges to the number of shared edges. Results are based on 100 replications. The numbers in the brackets are standard errors. CLIME: method of Cai, Liu and Luo (2011); GLASSO: graphical Lasso; JEMGM: method of Guo, Levina, Michailidis et al. (2011); fgland GGL: methods of Danaher, Wang and Witten (2014); MPE: proposed method.

Model(ρ)	Method	Performance					
		L_1	L_2	L_F	SEN	SPE	MCC
ER(0)	CLIME	19.62(0.32)	7.72(0.08)	49.48(0.83)	0.25(0.01)	0.99(0.00)	0.34(0.01)
	GLASSO	20.57(0.38)	7.91(0.12)	46.81(0.68)	0.03(0.02)	1.00(0.00)	0.11(0.02)
	JEMGM	16.21(0.47)	6.82(0.05)	39.24(0.20)	0.31(0.01)	0.99(0.00)	0.46(0.01)
	FGL	20.26(0.72)	7.89(0.09)	46.89(0.62)	0.05(0.02)	1.00(0.00)	0.12(0.02)
	GGL	20.15(0.83)	7.87(0.10)	46.42(0.64)	0.05(0.03)	1.00(0.00)	0.13(0.02)
	MPE	18.33(0.38)	7.32(0.10)	48.27(0.98)	0.47(0.02)	1.00(0.00)	0.63(0.01)
ER(1/4)	CLIME	20.60(0.27)	8.62(0.09)	57.13(0.86)	0.20(0.02)	0.98(0.00)	0.27(0.01)
	GLASSO	21.48(0.56)	8.22(0.19)	51.82(1.47)	0.17(0.05)	0.97(0.01)	0.19(0.03)
	JEMGM	19.17(0.42)	7.65(0.06)	47.07(0.21)	0.25(0.01)	0.99(0.00)	0.37(0.01)
	FGL	21.72(0.59)	8.51(0.05)	54.21(0.39)	0.11(0.01)	0.98(0.00)	0.16(0.01)
	GGL	21.28(0.36)	8.56(0.09)	54.37(0.65)	0.10(0.03)	0.99(0.00)	0.16(0.03)
	MPE	19.47(0.29)	8.36(0.10)	55.93(0.88)	0.32(0.01)	0.99(0.00)	0.49(0.01)
ER(1)	CLIME	28.73(0.30)	10.86(0.17)	72.98(1.68)	0.11(0.03)	0.98(0.01)	0.17(0.01)
	GLASSO	35.47(48.52)	11.40(10.62)	74.19(75.65)	0.11(0.08)	0.97(0.08)	0.13(0.02)
	JEMGM	28.25(0.62)	9.91(0.07)	62.74(0.19)	0.12(0.00)	0.99(0.00)	0.22(0.01)
	FGL	29.96(0.56)	10.56(0.08)	68.89(0.67)	0.05(0.02)	0.99(0.00)	0.09(0.01)
	GGL	30.25(0.66)	10.58(0.06)	68.71(0.51)	0.05(0.01)	0.99(0.00)	0.09(0.01)
	MPE	28.43(0.32)	10.60(0.09)	70.83(0.85)	0.19(0.01)	0.99(0.00)	0.34(0.01)

Table 2

Simulation results for data generated based on the Watt-Strogatz graph with different ratios of the number of individual-specific edges to the number of shared edges. Results are based on 100 replications. The numbers in the brackets are standard errors. CLIME: method of Cai, Liu and Luo (2011); GLASSO: graphical Lasso; JEMGM: method of Guo, Levina, Michailidis et al. (2011); fgland GGL: methods of Danaher, Wang and Witten (2014); MPE: proposed method.

Model(ρ)	Method	Performance					
		L_1	L_2	L_F	SEN	SPE	MCC
WS(0)	CLIME	29.80(0.23)	13.05(0.19)	87.71(1.96)	0.11(0.03)	0.99(0.00)	0.25(0.02)
	GLASSO	29.55(0.25)	12.35(0.08)	79.20(0.47)	0.08(0.01)	0.99(0.00)	0.20(0.01)
	JEMGM	29.59(0.46)	11.89(0.27)	74.32(2.50)	0.11(0.04)	1.00(0.00)	0.27(0.03)
	FGL	29.45(0.33)	12.43(0.13)	80.10(1.22)	0.10(0.02)	0.99(0.00)	0.23(0.03)
	GGL	29.65(0.23)	12.65(0.11)	80.53(0.67)	0.08(0.02)	1.00(0.00)	0.22(0.02)
	MPE	28.99(0.22)	12.72(0.12)	84.35(1.11)	0.16(0.01)	1.00(0.00)	0.34(0.01)
WS(1/4)	CLIME	42.70(0.35)	14.80(0.19)	102.03(2.08)	0.08(0.02)	0.98(0.01)	0.15(0.01)
	GLASSO	42.79(0.58)	14.25(0.07)	95.38(0.49)	0.04(0.00)	0.99(0.00)	0.09(0.00)
	JEMGM	43.06(0.60)	13.44(0.17)	84.35(1.75)	0.12(0.02)	0.98(0.00)	0.21(0.01)
	FGL	43.33(0.59)	14.31(0.06)	96.24(0.49)	0.06(0.01)	0.98(0.00)	0.11(0.01)
	GGL	42.84(0.54)	14.30(0.11)	95.95(0.95)	0.05(0.01)	0.99(0.00)	0.11(0.01)
	MPE	42.15(0.39)	14.60(0.11)	100.13(1.16)	0.10(0.00)	0.99(0.00)	0.23(0.01)
WS(1)	CLIME	63.27(0.35)	18.64(0.19)	128.83(1.94)	0.04(0.01)	0.99(0.00)	0.08(0.01)
	GLASSO	62.94(0.31)	17.85(0.11)	120.11(1.01)	0.02(0.01)	0.99(0.00)	0.04(0.01)
	JEMGM	63.67(0.85)	16.82(0.23)	107.53(2.20)	0.07(0.01)	0.98(0.00)	0.12(0.01)
	FGL	63.11(0.28)	17.86(0.07)	120.82(0.66)	0.02(0.00)	0.99(0.00)	0.04(0.00)
	GGL	63.13(0.35)	17.87(0.12)	120.31(1.05)	0.03(0.01)	0.99(0.01)	0.04(0.01)
	MPE	62.89(2.57)	18.18(0.47)	123.98(2.71)	0.08(0.09)	0.98(0.09)	0.15(0.01)

Table 3

Number of edges selected by the proposed method and its competitors. "C1 unique" counts the number of edges that only appear in the precision matrix of the gene expression levels in C1 cancer subtype; "Other unique" counts the number of edges that only appear in C2–C6 cancer subtypes; and "Common" counts the number of edges shared by both precision matrices.

Method	C1 unique	Other unique	Common
CLIME	40	43	20
GLASSO	11	11	7
JEMGM	23	22	77
FGL	8	112	23
GGL	14	148	44
MPE	13	38	42