Published in final edited form as:

J Comput Graph Stat. 2014 April ; 23(2): 439–459. doi:10.1080/10618600.2013.782818.

Sparse Covariance Matrix Estimation With Eigenvalue Constraints

Han LIU [Assistant Professor],

Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544 (hanliu@princeton.edu)

Lie WANG [Assistant Professor], and

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139 (liewang@math.mit.edu)

Tuo ZHAO

Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218 (tourzhao@jhu.edu)

Abstract

We propose a new approach for estimating high-dimensional, positive-definite covariance matrices. Our method extends the generalized thresholding operator by adding an explicit eigenvalue constraint. The estimated covariance matrix simultaneously achieves sparsity and positive definiteness. The estimator is rate optimal in the minimax sense and we develop an efficient iterative soft-thresholding and projection algorithm based on the alternating direction method of multipliers. Empirically, we conduct thorough numerical experiments on simulated datasets as well as real data examples to illustrate the usefulness of our method. Supplementary materials for the article are available online.

Keywords

Explicit eigenvalue constraint; High-dimensional data; Positive-definiteness guarantee

1. INTRODUCTION

We study the sparse covariance matrix estimation problem. Let $\mathbf{X} := (X_1, \dots, X_d)^T$ be a d -dimensional random vector with covariance matrix $\Sigma_{cov}^* := [(\Sigma_{cov}^*)_{jk}]_{1 \leq j, k \leq d}$, where $(\Sigma_{cov}^*)_{jk} := \mathbb{E}X_j X_k - \mathbb{E}X_j \mathbb{E}X_k$. Our goal is to estimate Σ_{cov}^* from n observational data points. The covariance matrix estimation problem plays an essential role in multivariate

© 2014 American Statistical Association, Institute of Mathematical Statistics, and Interface Foundation of North America

SUPPLEMENTARY MATERIALS

Appendices: Appendix A gives the proof of Propositions 3.1 and 3.6. Appendix B gives the proof of Lemmas 3.2. and 3.7. Appendix C gives the proof of Proposition 3.4. Appendix D gives the proof of Theorem 4.1. Appendix E gives the proof of Theorem 4.2. Appendix F gives the proof of Theorem 4.3.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/jcgs.

methods such as time series analysis (Box, Jenkins, and Reinsel 1994), spatial data analysis (Cressie 1992), and longitudinal data analysis (Searle, Casella, and McCulloch 1992). In this article, we are mainly interested in estimating covariance matrix in high-dimensional settings where $d \gg n$.

The problem of estimating high-dimensional covariance matrices has been extensively studied. By assuming the true covariance matrix is sparse, Wu and Pourahmadi (2003) analyzed longitudinal data where variables have a natural order and are weakly correlated when they are far apart. By using different forms of (weak) sparsity, researchers proposed many regularized estimators (Bickel and Levina 2004; Huang et al. 2006; Furrer and Bengtsson 2007; Levina, Rothman, and Zhu 2008; Bickel and Levina 2008b; Cai, Zhang, and Zhou 2010; Rothman, Levina, and Zhu 2010; Cai and Liu 2011). These estimators usually only require elementwise thresholding procedures (or combined with the Cholesky decomposition) and are computationally efficient. For other applications in which the variables do not have natural orders (e.g., genomic data, social science data, and economic data), Bickel and Levina (2008a) and Karoui (2008) showed that the thresholded sample covariance matrix also works well when the true covariance matrix is (weakly) sparse. Recently, Rothman, Levina, and Zhu (2009) established more comprehensive theoretical results for a generalized thresholding operator family including the soft-thresholding (Tibshirani 1996; Chen, Donoho, and Saunders 1998), hard-thresholding, Smoothly Clipped Absolute Deviation (SCAD)-thresholding (Fan and Li 2001), and adaptive thresholding operators (Zou 2006). These thresholding estimators, though simple, do not guarantee the positive definiteness of the estimated covariance matrix. Such indefiniteness may cause trouble in downstream analysis such as evaluating the predictive likelihood or linear discriminant analysis. The only exceptions are Lam and Fan (2009) and Rothman (2012), which both use a log-determinant function to explicitly enforce the positive definiteness of the estimated covariance matrix. The difference between them is that Rothman (2012) adopted a convex least-square formulation, while Lam and Fan (2009) adopted the penalized Gaussian likelihood approach, which results in a nonconvex formulation and there is no guarantee to find the global optimal solution in polynomial time.

In this article, we propose an alternative approach for estimating high-dimensional positive-definite covariance matrices. Our method can be viewed as an extension of the generalized thresholding operator (Rothman, Levina, and Zhu 2009). More specifically, we provide a convex formulation that secures the positive definiteness of the estimator by enforcing an eigenvalue constraint. Compared with the logarithmic barrier method of Rothman (2012), our approach explicitly constrains the smallest eigenvalue of the estimated covariance matrix and can be naturally extended to more flexible settings. Explicitly constraining the eigenvalues has its practical implications. For example, for certain downstream analysis of covariance matrix estimation, such as discriminant analysis and graphical estimation, we require the estimated covariance matrix to have a good condition number. Sometimes, we have prior knowledge about the smallest and largest eigenvalues and may want to incorporate such knowledge into the inference. Computationally, we adopt the alternating direction method of multipliers (ADMM) to develop an efficient ISP (Iterative Soft-thresholding and Projection) algorithm (Gabay and Mercier 1976; He and Yuan 2012). The

ISP algorithm iteratively conducts the soft-thresholding operation and eigenvalue decomposition. We prove that, with high probability, the obtained estimator is the same as the generalized thresholding estimator (when it is positive definite). This result implies that our procedure preserves the nice asymptotic properties of the generalized thresholding estimator. More precisely, let $\hat{\Sigma} \in \mathbb{R}^{d \times d}$ be our proposed covariance matrix estimator and $\|\hat{\Sigma}\|_2$ the spectral norm of $\hat{\Sigma}$, we have

$$\sup_{\Sigma_{cov}^* \in \mathcal{U}(q)} \mathbb{E} \|\hat{\Sigma} - \Sigma_{cov}^*\|_2 \asymp M_d \left(\sqrt{\frac{\log d}{n}} \right)^{1-q}, \quad (1.1)$$

where $\mathcal{U}$ denotes some parameter class (will be explained in Section 4) and M_d is some quantity that may scale with d . This rate of convergence is optimal in the minimax sense. Besides these nice asymptotic properties, our procedure has good finite sample performance. Empirically, we report thorough numerical experimental results on both simulated and real datasets to illustrate the usefulness of the proposed method.

While this article was under review, we learned that a similar positive-definite covariance estimator was independently proposed by Xue, Ma, and Zou (2012). There are several major differences between our estimator and theirs as follows:

1. Our estimator is different from that of Xue, Ma, and Zou (2012). Instead of estimating covariance matrix directly, we first estimate the correlation matrix and then combine the estimated correlation matrix with the marginal sample standard deviations to get the final covariance estimator. This same idea is also used in Lam and Fan (2009) and Rothman et al. (2008). The main advantage is that our estimator is adaptive to the marginal variability and more amendable to theoretical analysis. [In particular, we provide explicit rate of convergence for $\mathbb{E} \|\hat{\Sigma} - \Sigma_{cov}^*\|_2$, while Xue, Ma, and Zou (2012) only provided large probability bound for $\|\hat{\Sigma} - \Sigma_{cov}^*\|_{\text{F}}$.]
2. Besides the ℓ_1 penalty function, we also study other nonconvex penalty functions. An interesting fact we observed is that our proposed estimator can be combined with the nonconvex minimum concavity (MC+) penalty function without losing a global convex formulation. Therefore, we can simultaneously reduce the estimation bias and preserve the global convexity of the problem. In Xue, Ma, and Zou (2012), only the ℓ_1 penalty function is studied.
3. Xue, Ma, and Zou (2012) mainly focused on the optimization convergence property of the proposed ADMM-based computational algorithm. In contrast, we focus more on analyzing the statistical properties of the proposed method. We also illustrate the relationship between our proposed estimator and the generalized thresholding estimator. Such a relationship allows us to establish the minimax optimal rates of convergence of the resulting sparse covariance estimates under Frobenius and spectral norms.

Moreover, we should point out that the ISP algorithm was first developed in Bien and Tibshirani (2011) to solve a subproblem obtained by their iterative majorize/minimize procedure. Their subproblem is identical to our proposed convex program in the numerical form. More details about this point will be discussed in Section 3.1.

The rest of this article is organized as follows. In Section 2, we briefly review the covariance matrix estimation problem. In Section 3, we introduce our new method and present the ISP algorithm. In Section 4, we analyze the theoretical properties of the proposed estimator. Numerical simulations and real data experiments are reported in Section 5. We give brief conclusions and discussions in the last section. All the proofs are in Appendices.

2. NOTATION AND BACKGROUND

Let $\mathbf{A} = [\mathbf{A}_{jk}] \in \mathbb{R}^{d \times d}$, $\mathbf{B} = [\mathbf{B}_{jk}] \in \mathbb{R}^{d \times d}$ and $\mathbf{v} = (v_1, \dots, v_d)^T \in \mathbb{R}^d$. We denote by $\Lambda_{\min}(\mathbf{A})$ and $\Lambda_{\max}(\mathbf{A})$ the smallest and largest eigenvalues of $\mathbf{A}$. The inner product of $\mathbf{A}$ and $\mathbf{B}$ is defined as $\langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}^T \mathbf{B})$. We define the vector norms: $\|\mathbf{v}\|_1 = \sum_j |v_j|$, $\|\mathbf{v}\|_2^2 = \sum_j v_j^2$, $\|\mathbf{v}\|_\infty = \max_j |v_j|$. We also define matrix operator and elementwise norms: $\|\mathbf{A}\|_{1,\text{off}} = \sum_{j \neq k} |\mathbf{A}_{jk}|$, $\|\mathbf{A}\|_{\infty,\text{off}} = \max_{j \neq k} |\mathbf{A}_{jk}|$, $\|\mathbf{A}\|_2^2 = \Lambda_{\max}(\mathbf{A}^T \mathbf{A})$, $\|\mathbf{A}\|_F^2 = \sum_{j,k} \mathbf{A}_{jk}^2$. We use $\mathbf{S}_{\text{cov}}$ and $\mathbf{S}$ to denote the sample covariance and correlation matrices. The true covariance and correlation matrices are denoted by Σ_{cov}^* and Σ^* , respectively. We write $a_n \asymp b_n$ if there are positive constants c_1 and c_2 independent of n such that $c_1 b_n \leq a_n \leq c_2 b_n$.

2.1 Connection Between the Covariance and Correlation Estimation

A closely related problem is the estimation of correlation matrix, in which we are interested in estimating the correlation matrix $\Sigma^* := (\Theta^*)^{-1} \Sigma_{\text{cov}}^* (\Theta^*)^{-1}$, where $\Theta^* := \text{diag}(\theta_1^*, \dots, \theta_d^*)$ is a diagonal matrix, which with θ_j^* denotes the standard deviation of the j th variable. Using this decomposition, we can construct a covariance matrix estimator $\hat{\Sigma}_{\text{cov}}$ from $\hat{\Sigma}$:

$$\hat{\Sigma}_{\text{cov}} = \hat{\Theta} \hat{\Sigma} \hat{\Theta} \quad \text{with} \quad \hat{\Theta} = \text{diag}(\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_d), \quad (2.1)$$

where $\hat{\theta}_j$ is the sample standard deviation for the j th variable. Such a procedure is more amenable to theoretical analysis (more details can be found in Remark 4.4 in Section 4) and also makes the estimator adaptive to the marginal variability (more details can be found in Rothman et al. 2008; Lam and Fan 2009).

2.2 Soft-Thresholding Operator

Recall that $\mathbf{S}$ is the sample correlation matrix, the generalized thresholding operator (Rothman, Levina, and Zhu 2009) with ℓ_1 penalty solves the following optimization problem:

$$\hat{\Sigma}^{STO} = \underset{\Sigma}{\text{argmin}} \frac{1}{2} \|\mathbf{S} - \Sigma\|_F^2 + \lambda \|\Sigma\|_{1,\text{off}}. \quad (2.2)$$

This simple formulation has a closed-form solution:

$$\hat{\Sigma}_{jk}^{STO} = \begin{cases} \text{sign}(\mathbf{S}_{jk}) \cdot \max\{|\mathbf{S}_{jk}| - \lambda, 0\} & \text{if } j \neq k \\ \mathbf{S}_{jk} & \text{otherwise} \end{cases}, \quad (2.3)$$

which is well known as the univariate soft-thresholding operator (STO). Once the STO estimate $\hat{\Sigma}^{STO}$ is obtained, we plug it into (2.1) for estimating the covariance matrix $\hat{\Sigma}_{COV}$. Since (2.2) has no control of the positive definiteness, $\hat{\Sigma}^{STO}$ STO may be indefinite.

3. OUR PROPOSED METHOD

Our method, named Estimation of Covariance with Eigenvalue Constraints (EC2), extends the STO method by explicitly adding eigenvalue constraints. Recall that $\mathbf{S}$ is the sample correlation matrix, the EC2 estimator is defined as

$$\hat{\Sigma}^{EC2} := \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{S} - \Sigma\|_F^2 + \lambda \|\Sigma\|_{1,\text{off}} \quad \text{s.t.} \quad \tau \leq \Lambda_{\min}(\Sigma), \quad (3.1)$$

where $\lambda > 0$ is a regularization parameter, and $\tau > 0$ is a desired minimum eigenvalue lower bound of the estimator, which can also be chosen in a data-dependent way. The EC2 method simultaneously conducts sparse estimation and guarantees the positive definiteness of the solution. The equality constraint $\Sigma_{jj} = 1$ ensures all diagonal entries of $\hat{\Sigma}$ to be 1, since we are estimating the correlation matrix. Once $\hat{\Sigma}^{EC2}$ is obtained, we convert it to our covariance matrix estimator as (2.1). The next proposition shows that the formulation (3.1) is convex:

Proposition 3.1. The optimization problem in (3.1) is a convex program.

Proof. The proof of this proposition is presented in Appendix A.

We proceed to develop an efficient algorithm to solve (3.1) using the ADMM. The resulting algorithm is iterative and solves a closed-form soft-thresholding step and a spectral projection step within each iteration.

3.1 Iterative Soft-Thresholding and Projection Algorithm

We first reparameterize (3.1) by introducing an auxiliary variable Γ :

$$(\hat{\Sigma}, \hat{\Gamma}) = \underset{\Sigma_{jj}=1, \tau \leq \Lambda_{\min}(\Gamma)}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{S} - \Gamma\|_F^2 + \lambda \|\Sigma\|_{1,\text{off}} \quad \text{s.t.} \quad \Gamma = \Sigma. \quad (3.2)$$

This reparameterization decouples the computational dependency in the optimization problem. Therefore, a complicated problem decouples into multiple simpler subproblems, which can be solved in closed forms. More specifically, Equation (3.2) can be rewritten in the following equivalent saddle point problem

$$(\hat{\Sigma}, \hat{\Gamma}, \hat{\mathbf{U}}) = \underset{\Sigma_{jj}=1, \tau \leq \Lambda_{\min}(\Gamma)}{\operatorname{argmin}} \max_{\mathbf{U}} \frac{1}{2} \|\mathbf{S} - \Gamma\|_{\mathbf{F}}^2 + \lambda \|\Sigma\|_{1, \text{off}} + \langle \mathbf{U}, \Gamma - \Sigma \rangle + \frac{\rho}{2} \|\Gamma - \Sigma\|_{\mathbf{F}}^2, \quad (3.3)$$

where $\rho > 0$ is the penalty multiplier and $\mathbf{U} \in \mathbb{R}^{d \times d}$ is the Lagrange multiplier matrix. The ADMM works in an iterative fashion. Suppose we obtain the solution $\Sigma^{(t)}$, $\Gamma^{(t)}$, $\mathbf{U}^{(t)}$ at the t th iteration, the algorithm proceeds as follows:

Step 1. Update Σ by

$$\Sigma^{(t+1)} = \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \lambda \|\Sigma\|_{1, \text{off}} + \frac{\rho}{2} \left\| \frac{1}{\rho} \mathbf{U}^{(t)} + \Gamma^{(t)} - \Sigma \right\|_{\mathbf{F}}^2. \quad (3.4)$$

Let $\tilde{\Sigma} := \Gamma^{(t)} + \frac{1}{\rho} \mathbf{U}^{(t)}$, Equation (3.4) has the closed-form solution by soft-thresholding,

$$\Sigma_{jk}^{(t+1)} = \begin{cases} \operatorname{sign}(\tilde{\Sigma}_{jk}) (|\tilde{\Sigma}_{jk}| - \lambda/\rho) & \text{if } j \neq k \\ 1 & \text{otherwise} \end{cases}. \quad (3.5)$$

Step 2. Given $\Sigma^{(t+1)}$, we then update Γ by

$$\Gamma^{(t+1)} = \underset{\tau \leq \Lambda_{\min}(\Gamma)}{\operatorname{argmin}} \left\| \Gamma - \frac{\mathbf{S} + \rho \Sigma^{(t+1)} - \mathbf{U}^{(t)}}{(1+\rho)} \right\|_{\mathbf{F}}^2. \quad (3.6)$$

Equation (3.6) has a closed-form solution $\Gamma^{(t+1)} = \mathcal{P}_+ \left(\frac{\mathbf{S} + \rho \Sigma^{(t+1)} - \mathbf{U}^{(t)}}{(1+\rho)}, \tau \right)$, where $\mathcal{P}_+(\cdot, \cdot)$ is a spectral projection operator and is characterized by the following lemma.

Lemma 3.2. Suppose $\mathbf{A}$ has the spectral decomposition: $\mathbf{A} = \sum_{j=1}^d \hat{\delta}_j \mathbf{v}_j \mathbf{v}_j^T$, where $\hat{\delta}_j$'s are the eigenvalues and $\mathbf{v}_j$'s are the corresponding eigenvectors. Let $\tilde{\delta}_j = \max(\hat{\delta}_j, \tau)$, we have

$$\mathcal{P}_+(\mathbf{A}, \tau) = \sum_{j=1}^d \tilde{\delta}_j \mathbf{v}_j \mathbf{v}_j^T.$$

Proof. The proof of this lemma is presented in Appendix B.

Updating Γ requires an eigenvalue decomposition. The eigenvalue decomposition has been intensively studied in the literatures of numerical analysis and parallel computing. Many sophisticated parallel algorithms such as the Lanczos algorithm have been developed and the eigenvalue decomposition now can scale to large problems using general-purpose graphics processing unit (GP/GPU) or multicore techniques.

Step 3. Update $\mathbf{U}$ by

$$\mathbf{U}^{(t+1)} = \mathbf{U}^{(t)} + \rho \left(\mathbf{\Gamma}^{(t+1)} - \mathbf{\Sigma}^{(t+1)} \right). \quad (3.7)$$

We stop the iteration when the following convergence criterion is satisfied

$$\max \left\{ \frac{\|\mathbf{\Sigma}^{(t+1)} - \mathbf{\Sigma}^{(t)}\|_F^2}{\|\mathbf{\Sigma}^{(t)}\|_F^2}, \frac{\|\mathbf{\Sigma}^{(t+1)} - \mathbf{\Gamma}^{(t+1)}\|_F^2}{\|\mathbf{\Sigma}^{(t)}\|_F^2} \right\} \leq \epsilon, \quad (3.8)$$

where $\epsilon > 0$ is a precision tolerance parameter. In applications, we set $\epsilon = 10^{-6}$. Xue, Ma, and Zou (2012) also developed a similar computational algorithm, and our ISP algorithm has a different updating formula for $\mathbf{\Sigma}$ due to the equality constraint $\Sigma_{jj} = 1$ for the correlation matrix estimation. Moreover, as a reviewer pointed out, such an ADMM-type algorithm was first derived in Bien and Tibshirani (2011) for sparse covariance estimation problem. More specifically, they adopted the majorize/minimize procedure to relax the nonconvex formation. The resulting subproblem is identical to (3.1), and solved by the same computational algorithm. For more details, please refer to Appendix 3 in Bien and Tibshirani (2011).

Remark 3.3. The convergence rate of our ISP algorithm is established in He and Yuan (2012). It achieves an $O(1/t)$ rate of convergence in the sense of variational inequality, where t is the number of iterations. Please refer to He and Yuan (2012) for more technical details. A similar result is also established in Xue, Ma, and Zou (2012). Compared with existing analysis, a more interesting problem would be establishing a possible linear convergence rate since the least-square loss function in (3.1) is strongly convex. However, this is beyond the scope of this article and will be left for future investigation.

3.2 Relationship Between EC2 and Soft-Thresholding Operator

By comparing the STO estimator $\hat{\mathbf{\Sigma}}^{STO}$ and EC2 estimator $\hat{\mathbf{\Sigma}}^{EC2}$ in (2.2) and (3.1), we can see that $\hat{\mathbf{\Sigma}}^{STO} = \hat{\mathbf{\Sigma}}^{EC2}$ holds, if $\hat{\mathbf{\Sigma}}^{STO}$ falls in the feasible region of (3.1), that is,

$\tau \leq \Lambda_{\min} \left(\hat{\mathbf{\Sigma}}^{STO} \right)$. In the next section, we will show that this relationship holds with high probability under appropriate scaling. This result has important implications for both computation and theoretical analysis. Empirically, we use $\hat{\mathbf{\Sigma}}^{STO}$ as the initial estimator and if $\hat{\mathbf{\Sigma}}^{STO}$ is already feasible, then the algorithm stops immediately.

3.3 Covariance Estimation with Adaptive Penalty

As mentioned earlier, the generalized thresholding operator family contains many thresholding functions induced by other nonconvex penalty functions such as SCAD. Similarly, we can also combined our EC2 with other nonconvex regularization to reduce the estimation bias. Usually a nonconvex regularization makes the penalized least-square formulation nonconvex and there is no guarantee to obtain a global solution in polynomial time. To retain the global convexity of the formulation, we usually choose the adaptive penalty in Zou (2006), which results in the following optimization problem:

$$\hat{\Sigma}^{EC2} \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{S} - \Sigma\|_F^2 + \lambda \|\mathbf{W} \circ \Sigma\|_{1,\text{off}} \quad \text{s.t.} \quad \tau \leq \Lambda_{\min}(\Sigma), \quad (3.9)$$

where $\mathbf{W}$ represents the weight matrix and $\circ$ denotes the Hadamard product, that is,

$\mathbf{W} \circ \Sigma = [\mathbf{W}_{jk} \cdot \Sigma_{jk}]_{j,k=1}^d$. There are multiple choices for $\mathbf{W}$, and Zou (2006) suggested to use the reciprocal of the ordinary least-square estimate as the weight, that is, $\mathbf{W}_{jk} = (\mathbf{S}_{jk})^{-1}$ in our case. Some recent results in Xue and Zou (2012) also imply that ℓ_1 penalized estimator is potentially a better choice, that is, $\mathbf{W}_{jk} = \left(|\hat{\Sigma}_{jk}^{STO}| + 1/n \right)^{-1}$ in our case, where $1/n$ is to avoid the reciprocal of zero elements in the ℓ_1 penalized estimator. We will compare these two weight matrices in numerical experiments in Section 5.

Adding the weights does not change the problem formulation. Therefore, we can solve the problem in (3.9) using the same ISP algorithm with Step 1 replaced by the following Step 1a,

Step 1a. Update Σ by

$$\Sigma^{(t+1)} = \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \lambda \|\mathbf{W} \circ \Sigma\|_{1,\text{off}} + \frac{\rho}{2} \left\| \frac{1}{\rho} \mathbf{U}^{(t)} + \mathbf{\Gamma}^{(t)} - \Sigma \right\|_F^2. \quad (3.10)$$

Let $\tilde{\Sigma} := \mathbf{\Gamma}^{(t)} + \frac{1}{\rho} \mathbf{U}^{(t)}$, Equation (3.10) has the closed-form solution by soft-thresholding,

$$\Sigma_{jk}^{(t+1)} = \begin{cases} \operatorname{sign}(\tilde{\Sigma}_{jk}) \left(|\tilde{\Sigma}_{jk}| - \lambda \mathbf{W}_{jk} / \rho \right) & \text{if } j \neq k \\ 1 & \text{otherwise} \end{cases}, \quad (3.11)$$

and the same convergence guarantee holds for the EC2 with adaptive penalty.

3.4 EC2 With MC+ Penalty

Most nonconvex penalty functions result in nonconvex formulations of EC2, there is usually no guarantee to get a global solution in polynomial time. However, one exception is the minimax concave (MC+) penalty, it is possible to retain the global convexity with a suitable choice of tuning parameter. Recall the minimax concave penalty in Zhang (2010):

$$\begin{aligned} \lambda P(\Sigma, \lambda, \gamma) &= \lambda \sum_{k \neq j} \int_0^{|\Sigma_{jk}|} \left(1 - \frac{t}{\gamma \lambda} \right)_+^{dt} \\ &= \lambda \sum_{k \neq j} \left\{ \left(|\Sigma_{jk}| - \frac{|\Sigma_{jk}|^2}{2\lambda\gamma} \right) \mathbf{I}(|\Sigma_{jk}| < \lambda\gamma) + \frac{\lambda\gamma}{2} \mathbf{I}(|\Sigma_{jk}| \geq \lambda\gamma) \right\}, \end{aligned} \quad (3.12)$$

where $\gamma > 1$ is a prefixed parameter. By varying from $\gamma \rightarrow \infty$ to $\gamma \rightarrow 0^+$, we can obtain a continuum of penalties and threshold operators for every $\lambda > 0$. The EC2 with MC+ penalty is formulated as

$$\hat{\Sigma}^{EC2} = \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{S} - \Sigma\|_F^2 + \lambda P(\Sigma, \lambda, \gamma) \text{ s.t. } \tau \leq \Lambda_{\min}(\Sigma). \quad (3.13)$$

The next proposition shows that the problem in (3.13) is convex whenever the MC+ parameter $\gamma > 1$. Therefore, by applying the MC+ penalty function to EC2, we enjoy the bias reduction of the nonconvex penalty without losing global convexity.

Proposition 3.4. The optimization problem in (3.13) is convex when $\gamma > 1$.

Proof. The proof of this proposition is presented in Appendix C.

Remark 3.5. A heuristic approach to solve (3.13) is to modify Step 1 of the ISP algorithm described in Section 3.1. More specifically, we replace Step 1 with Step 1b described as follows.

Step 1b. Update current values of $\Gamma^{(t)}$ and $\mathbf{U}^{(t)}$, we update Σ by Let

$$\Sigma^{(t+1)} = \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \lambda P(\Sigma, \lambda, \gamma) + \frac{\rho}{2} \|\mathbf{U}^{(t)} + \Gamma^{(t)} - \Sigma\|_F^2. \quad (3.14)$$

Let $\tilde{\Sigma} := \Gamma^{(t)} + \mathbf{U}^{(t)}$, for $\gamma > 1$, Equation (3.10) has a closed-form solution, which is a piecewise-linear thresholding function,

$$\Sigma_{jk}^{(t+1)} = \begin{cases} 0 & \text{if } j \neq k \text{ and } |\tilde{\Sigma}_{jk}| \leq \frac{\lambda}{\rho} \\ \frac{\operatorname{sign}(\tilde{\Sigma}_{jk})(|\tilde{\Sigma}_{jk}| - \lambda/\rho)}{1 - 1/\gamma} & \text{if } j \neq k \text{ and } \frac{\lambda}{\rho} \leq |\tilde{\Sigma}_{jk}| \leq \frac{\lambda\gamma}{\rho} \\ \tilde{\Sigma}_{jk} & \text{if } j \neq k \text{ and } |\tilde{\Sigma}_{jk}| > \frac{\lambda\gamma}{\rho} \\ 1 & \text{otherwise} \end{cases}. \quad (3.15)$$

However, the ADMM cannot guarantee the convergence to the global solution of (3.13) since existing analysis relies on the convexity of both the loss and penalty functions. Our empirical studies show that such an ADMM often fails to converge. Therefore, instead of using this algorithm, we propose an alternative algorithm using the local linear approximation (LLA).

The LLA is also an iterative algorithm, and suppose at the t th iteration, we already have the current solution $\Sigma^{(t)}$. We consider the linearization of MC+ at $\Sigma = \Sigma^{(t)}$, which is defined as

$$\lambda L(\Sigma, \Sigma^{(t)}) = \lambda \sum_{k \neq j} \left\{ \left(1 - \frac{|\Sigma_{jk}^{(t)}|}{\gamma \lambda} \right)_+ \left(|\Sigma_{jk}| - |\Sigma_{jk}^{(t)}| \right) \right\}, \quad (3.16)$$

then instead of directly minimizing (3.13), we solve the following convex relaxation of

$$\hat{\Sigma}^{(t+1)} = \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{S} - \Sigma\|_F^2 + \lambda L(\Sigma, \Sigma^{(t)}) \text{ s.t. } \tau \leq \Lambda_{\min}(\Sigma). \quad (3.17)$$

The LLA is usually used to solve nonconvex minimization problem and has a provable convergence to the local optimum. Since the formulation of (3.13) is convex, a global solution can be obtained. Several illustrative examples of the MC+ and LLA are shown in Figure 1.

3.5 EC2 With Largest Eigenvalue Constraint

In this section, we extend the EC2 by also constraining the largest eigenvalue. In particular, we consider the following optimization problem:

$$\hat{\Sigma}^{EC2} := \underset{\Sigma_{jj}=1}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{S} - \Sigma\|_{\mathbb{F}}^2 + \lambda \|\Sigma\|_{1,\text{off}} \quad \text{s.t.} \quad \tau_1 \leq \Lambda_{\min}(\Sigma) \leq \Lambda_{\max}(\Sigma) \leq \tau_2. \quad (3.18)$$

The formulation in (3.18) is also convex:

Proposition 3.6. The optimization problem in (3.18) is a convex program.

Proof. The proof of this proposition is presented in Appendix A.

To derive the algorithm, we only need to modify Step 2 of the ISP algorithm described in Section 3.1. More specifically, we replace Step 2 with Step 2a described as follows.

Step 2a. Given $\Sigma^{(t+1)}$, we update Γ by

$$\Gamma^{(t+1)} = \underset{\tau_1 \leq \Lambda_{\min}(\Gamma) \leq \Lambda_{\max}(\Gamma) \leq \tau_2}{\operatorname{argmin}} \left\| \Gamma - \frac{\mathbf{S} + \rho \Sigma^{(t+1)} - \rho \mathbf{U}^{(t)}}{(1+\rho)} \right\|_{\mathbb{F}}^2. \quad (3.19)$$

Equation (3.19) also has a closed-form solution

$$\Gamma^{(t+1)} = \mathcal{P}'_+ \left(\frac{\mathbf{S} + \rho \Sigma^{(t+1)} - \rho \mathbf{U}^{(t)}}{(1+\rho)}, \tau_1, \tau_2 \right), \quad (3.20)$$

where $\mathcal{P}'_+(\cdot, \cdot, \cdot)$ is a spectral projection operator and is characterized by the following lemma.

Lemma 3.7. Suppose $\mathbf{A}$ has the spectral decomposition: $\mathbf{A} = \sum_{j=1}^d \hat{\delta}_j \mathbf{v}_j \mathbf{v}_j^T$, where $\hat{\delta}_j$'s are the eigenvalues and $\mathbf{v}_j$'s are the corresponding eigenvectors. Let

$\tilde{\delta}_j = \tau_1 \cdot \mathbf{I}(\hat{\delta}_j < \tau_1) + \hat{\delta}_j \cdot \mathbf{I}(\tau_1 \leq \hat{\delta}_j \leq \tau_2) + \tau_2 \cdot \mathbf{I}(\hat{\delta}_j > \tau_2)$, we have

$$\mathcal{P}'_+(\mathbf{A}, \tau_1, \tau_2) = \sum_{j=1}^d \tilde{\delta}_j \mathbf{v}_j \mathbf{v}_j^T.$$

Proof. The proof of this lemma is presented in Appendix B.

Therefore, using an iterative soft-thresholding and spectral thresholding procedure, we simultaneously control the smallest and largest eigenvalues of the estimated covariance matrix. By replacing Step 2 with Step 2a, the rate of convergence of the ISP algorithm

preserves. The same convergence guarantee also holds for the EC2 with the largest eigenvalue constraint.

4. STATISTICAL PROPERTIES

In this section, we analyze the theoretical properties of the EC2 estimator. In Section 3, we introduced several EC2 estimators: EC2 with ℓ_1 penalty, EC2 with MC+ penalty, EC2 with adaptive penalty, and EC2 with largest eigenvalue constraint. All these estimators have similar statistical properties. For conciseness, we focus on analyzing the EC2 estimator with ℓ_1 penalty as defined in (3.1) and (2.1) in this section. The same results hold for the other variants of the EC2 estimators by following the same argument.

Since EC2 first estimates the correlation matrix, we proceed to derive the asymptotic error bound for $\hat{\Sigma}^{EC2}$. Let $0 < q < 1$, we consider the following class of sparse correlation matrices:

$$\mathcal{M}(q, M_d, \delta) := \left\{ \Sigma : \max_{1 \leq j \leq d} \sum_{k \neq j} |\Sigma_{jk}|^q \leq M_d \text{ and } \Sigma_{jj} = 1 \text{ for all } j, \Lambda_{\min}(\Sigma) \geq \delta \right\}.$$

We also define a class of covariance matrices:

$$\mathcal{U}(\kappa, q, M_d, \delta) := \left\{ \Sigma : \max_j \Sigma_{jj} \leq \kappa \text{ and } \Theta^{-1} \Sigma \Theta^{-1} \in \mathcal{M}(q, M_d, \delta) \right\}, \quad (4.1)$$

where $\Theta^{-1} = \text{diag}(\sqrt{\Sigma_{11}}, \dots, \sqrt{\Sigma_{dd}})$. The definition of this class is similar to the “universal thresholding class” defined by Bickel and Levina (2008b) and Rothman, Levina, and Zhu (2009). The main difference is that instead of assuming each column of the covariance matrix belongs to an ℓ_q -ball, we impose such conditions on the correlation matrix. Such a class is more suitable for our setting since the EC2 first estimates the correlation matrix.

Here, we list all the required assumptions:

(A1) Each marginal distribution of $\mathbf{X} := (X_1, \dots, X_d)$ is sub-Gaussian, that is, there exists a finite constant K , such that

$$\max_{1 \leq j \leq d} \mathbb{E} \exp(t(X_j - \mathbb{E} X_j)) \leq \exp\left(\frac{K^2 t^2}{2}\right) \text{ for all } t. \quad (4.2)$$

(A2) The true covariance matrix $\Sigma_{cov}^* \in \mathcal{U}(\kappa, q, M_d, \delta_{\min})$ for some $\delta_{\min} > 0$.

(A3) The dimensionality d and sample size n satisfy $\limsup_{n \rightarrow \infty} (M_d)^{\frac{2}{1-q}} \frac{\log d}{n} = 0$.

These conditions are mild. Similar assumptions were used in Rothman, Levina, and Zhu (2009) to analyze the asymptotic properties of the STO estimator. Though this article

focuses on the sub-Gaussian family (as shown in Condition A1), the methodology and theory can be further extended to the semiparametric Gaussian copula family (Liu, Lafferty, and Wasserman 2009; Liu et al. 2012), which further benefits Gaussian copula-based statistical methods such as Gaussian copula discriminant analysis and Gaussian copula component analysis in Han, Zhao, and Liu (2013) and Han and Liu (2012).

In the rest of this section, we provide both asymptotic and nonasymptotic analysis. For the asymptotic result, we show that, with high probability, the EC2 estimator in (3.1) is the same as the STO estimator in (2.2). This result suggests that the EC2 estimator preserves all the asymptotic properties of the STO estimator and is thus rate optimal in the minimax sense. However, such an asymptotic result holds only for large enough sample size n , which may not be satisfied in the finite sample settings. To gain more insights, we also provide a nonasymptotic result of the EC2 estimator. Our nonasymptotic analysis achieves the same rate of convergence as that in Rothman (2012), which is optimal in terms of Frobenius norm error. To our best knowledge, EC2 is the first estimator that simultaneously achieves sparse and positive-definite estimates and owes the optimal rate of convergence in terms of the spectral norm error. The optimal rate of convergence under the spectral norm has not been established in Rothman (2012), Lam and Fan (2009), and Bien and Tibshirani (2011). It is still an open problem in the literature on whether their estimators can achieve the optimal rate of convergence under the spectral norm.

4.1 Asymptotic Results

We first present the asymptotic analysis, which suggests that, for large enough sample size n , the EC2 estimator is the same as the STO estimator. In this section, we always consider the high-dimensional settings where $d > n$. The analysis of the low-dimensional case when $d < n$ is trivial.

Theorem 4.1. Let $\hat{\Sigma}^{EC2}$ be defined as in (3.1). Under assumptions (A1) to (A3), there exist universal constants c_0 and c_1 , such that, by taking $\lambda = c_0 \sqrt{\log d/n}$ and $\tau = \delta_{\min}/2$,

whenever $n \geq \left(\frac{2c_1 M_d}{\delta_{\min}} \right)^{\frac{2}{1-q}} \cdot \log d$, we have $\mathbb{P} \left(\hat{\Sigma}^{EC2} = \hat{\Sigma}^{STO} \right) \geq 1 - \frac{1}{d^5}$.

Proof. The proof of this theorem is presented in Appendix D.

Theorem 4.1 shows that the solution to STO is also the solution to EC2 with high probability. Therefore, our ISP algorithm always first checks the minimum eigenvalue of the STO estimator. The algorithm proceeds only when the feasibility condition is not satisfied. Empirically, in many situations, the STO meets the requirement in Theorem 4.1. In these cases, the computation of the EC2 estimator is as efficient as the simple STO estimator. Also, in applications, we do not require exact knowledge of $\delta_{\min}$, a lower bound is enough. In fact, in later numerical experiment section, we propose a fully data-dependent approach to select τ .

The next theorem bounds the error term between the EC2 covariance estimate $\hat{\Sigma}_{cov}^{EC2}$ and the true covariance Σ_{cov}^* .

Theorem 4.2. Let $\hat{\Sigma}_{cov}^{EC2}$ be defined as in (2.1). Under assumptions (A1) to (A3) and $d > n$, there exist universal constants c_0 and c_1 , such that, by taking $\lambda = c_0 \sqrt{\log d/n}$ and $\tau = \delta_{\min}/2$, whenever $n \geq \left(\frac{2c_1 M_d}{\delta_{\min}}\right)^{\frac{2}{1-q}} \cdot \log d$, we have

$$\sup_{\Sigma_{cov}^* \in \mathcal{U}(\kappa, q, M_d, \delta_{\min})} \mathbb{E} \|\hat{\Sigma}_{cov}^{EC2} - \Sigma_{cov}^*\|_2 \leq c_1 \cdot M_d \left(\frac{\log d}{n}\right)^{\frac{1-q}{2}}. \quad (4.3)$$

Proof. The proof of this theorem is presented in Appendix E.

As has been proved in Cai and Zhou (2012), the rate of convergence in (4.3) is minimax optimal over the class $\mathcal{U}(\kappa, q, M_d, \delta_{\min})$. Therefore, we have

$$\sup_{\Sigma_{cov}^* \in \mathcal{U}(\kappa, q, M_d, \delta_{\min})} \mathbb{E} \|\hat{\Sigma}_{cov}^{EC2} - \Sigma_{cov}^*\|_2 \asymp M_d \left(\frac{\log d}{n}\right)^{\frac{1-q}{2}}. \quad (4.4)$$

Theorem 4.2 shows that the EC2 estimator is asymptotically rate optimal. The main catch of this result is that it requires the condition

$$n \geq \left(\frac{2c_1 M_d}{\delta_{\min}}\right)^{\frac{2}{1-q}} \cdot \log d.$$

Under Assumption (A3), this would not be a problem in the asymptotic sense. However, in finite sample settings, this may no longer be true. In the next section, we provide a nonasymptotic result of the EC2 estimator. This result is comparable to the theoretical result as in Rothman (2012).

4.2 Nonasymptotic Results

In this section, we derive the nonasymptotic error bound to further analyze the performance of EC2. To make our results comparable to that in Rothman (2012), we analyze the error rate in Frobenius norm and only consider the case that $\Sigma_{cov}^* \in \mathcal{U}(\kappa, 0, M_d, \delta_{\min})$. Our results show that EC2 achieves the optimal rate in Frobenius norm.

The EC2 solves a convex optimization problem, which results in an M -estimator. Therefore, we can apply similar technique as in Negahban et al. (2012) to derive the following nonasymptotic error bound.

Theorem 4.3. Let $\hat{\Sigma}_{cov}^{EC2}$ be defined as in (3.1). We assume (A1) and (A3) hold. For (A2), we assume $\Sigma_{cov}^* \in \mathcal{U}(\kappa, 0, M_d, \delta_{\min})$. We denote N_d to be the number of nonzero off-diagonal elements in Σ_{cov}^* . Then there exist universal constants c_2 and c_3 such that, by taking $\lambda = c_2 \sqrt{\log d/n}$ and $\tau = \delta_{\min}$, we have

$$\mathbb{P} \left(\|\hat{\Sigma}^{EC2} - \Sigma^*\|_F \leq c_3 \sqrt{\frac{N_d \log d}{n}} \right) \geq 1 - \frac{1}{d^5}. \quad (4.5)$$

Proof. The proof of this theorem is presented in Appendix F.

Remark 4.4. If we directly exploit EC2 to estimate Σ_{cov}^* , we do not have $(\Sigma_{cov}^*)_{jj} = (\mathbf{S}_{cov})_{jj}$ in (F.5). The proof will end up with a rate $O_p \left(\sqrt{(d+N_d) \log d/n} \right)$ under the Frobenius norm. The extra term $d \log d$ on the numerator reflects the effort for estimating the diagonal elements.

The rate of convergence under the Frobenius norm in Theorem 4.3 matches the minimax lower bound for correlation matrix estimation and is the same as the one obtained by Rothman (2012). Such a result leads to an error bound under the spectral norm as in the following corollary.

Corollary 4.5. Under the same conditions as in Theorem 4.3 and $d > n$, we have

$$\sup_{\Sigma_{cov}^* \in \mathcal{U}(\kappa, 0, M_d, \delta_{min})} \mathbb{E} \|\hat{\Sigma}_{cov}^{EC2} - \Sigma_{cov}^*\|_2 \asymp \sqrt{\frac{N_d \log d}{n}}. \quad (4.6)$$

Proof. The proof follows from the same argument as Theorem 4.2, combined with the trivial bound

$$\|\hat{\Sigma}^{EC2} - \Sigma^*\|_2 \leq \|\hat{\Sigma}^{EC2} - \Sigma^*\|_F. \quad (4.7)$$

The detailed proof is omitted for conciseness.

5. NUMERICAL EXPERIMENTS

We compare different variants of EC2 with the estimators proposed by Rothman, Levina, and Zhu (2009) and Rothman (2012). More specifically, we consider nine methods:

- HTO: generalized thresholding with L_0 penalty (or the hard-thresholding estimator).
- STO: generalized thresholding with ℓ_1 penalty (or the soft-thresholding estimator).
- ATO-S: adaptive STO using the sample correlation matrix as the weight matrix.
- ATO-G: adaptive STO using the STO estimator as the weight matrix.
- MCTO: generalized thresholding estimator using the MC+ penalty.
- EC2-L1: EC2 with ℓ_1 penalty.
- AEC2-S: adaptive EC2 using the sample correlation matrix as the weight matrix.

- AEC2-G: adaptive EC2 using the STO estimator as the weight matrix.
- EC2-M: EC2 with MC+ penalty.
- Log-Det: the estimator proposed in Rothman (2012). (The R package is available on <http://cran.r-project.org/web/packages/PDSCE/index.html>).

We select the tuning parameters of these methods in a data-dependent way as described in the next section.

5.1 Tuning Parameter Selection

We choose the tuning parameter by subsampling (Bickel and Levina 2008a). More specifically, we randomly split the dataset K times. For splitting, the data are divided into a validation set of size n_2 and a training set of size $n_1 = n - n_2$. Theoretically, we can choose $n_2 \asymp n/\log n$. Practically, we choose $n_1 = 3/4n$. Then we select the tuning parameter $\gamma = 1.01$ for EC2-M, which makes the performance of EC2-M close to HTO. We choose λ for STO and MCTO as

$$\hat{\lambda} = \underset{\lambda}{\operatorname{argmin}} \sum_{k=1}^K \|\hat{\Sigma}_{\lambda}^{k,n_1} - \mathbf{S}^{k,n_2}\|_{\mathbb{F}}^2, \quad (5.1)$$

where $\hat{\Sigma}_{\lambda}^{k,n_1}$ is the STO or MCTO correlation estimator, with λ computed with the training set of the k th split and $\mathbf{S}^{k,n_2}$ is the sample correlation of the validation set of the k th split. For EC2-L1 and EC2-M, we choose λ as the same as STO and MCTO, respectively, and then tune τ as

$$\hat{\tau} = \underset{\tau}{\operatorname{argmin}} \sum_{k=1}^K \|\hat{\Sigma}_{\tau}^{k,n_1} - \mathbf{S}^{k,n_2}\|_{\mathbb{F}}^2, \quad (5.2)$$

where $\hat{\Sigma}_{\tau}^{k,n_1}$ is the EC2-L1 or EC-M estimator using a similar notion in (5.1). For ATO, we tune λ over the grid using the same criterion as (5.1). For AEC2, we choose the same τ as EC2 and the same λ as ATO. For Log-Det, we take $\tau = 10^{-4}$ as suggested in Rothman (2012) and tune λ over the grid using the same criterion as (5.1).

5.2 Simulated Datasets

We use three models for the population covariance matrix. The heatmaps of these three covariance models with $d = 200$ are shown in Figure 2:

1. Toeplitz matrix: $\Sigma_{jk}^* = 0.75^{|j-k|}$. We know that for $d = 100, 200$, or 400 , $\Lambda_{\min}(\Sigma^*) \approx 0.143$.
2. Block matrix: The indices $1, \dots, d$ are evenly divided into 10 groups with $\Sigma_{jk}^* = 0.8$ if k and j ($k \neq j$) belong to the same group and 0 otherwise. We know that for $d = 100, 200$, or 400 , $\Lambda_{\min}(\Sigma^*) = 0.2$.

3. Banded matrix: $\Sigma_{jk}^* = 1 - \frac{|j-k|}{20}$ if $|j-k| \leq 20$ and 0 otherwise. We know that $\Lambda_{\min}(\Sigma^*) \approx 0.034, 0.010, 0.002$ for $d = 100, 200, 400$, respectively.

Using these covariance models, we generate $n = 200$ independent data points from a d -dimensional Gaussian random variable with mean $\mathbf{0}$. The resulting covariance estimates are compared with the population covariance matrix using the spectral and Frobenius norms as evaluation metrics. We repeat this procedure 200 times and summarize the averaged estimation errors in Tables 1-3 with the corresponding standard errors in the parentheses. The histograms of the minimum eigenvalues of the HTO, STO, ATO-G, ATO-S, and MCTO estimators are provided in Figures 3-5.

From these results, we see that STO, EC2-L1, and Log-Det achieve similar estimation performance in both Frobenius norm and spectral norm errors. In particular, EC2-L1 and Log-Det slightly outperforms STO in Frobenius norm errors, while STO has a slightly better spectral norm performance. For both block covariance and banded covariance models, EC2-L1 always guarantees the positive definiteness of the estimated covariance matrices, in contrast, STO almost never delivers positive-definite estimates (especially for $d = 400$). This is similar to the results obtained by Rothman (2012).

Since the adaptive penalty reduces the estimation bias, ATO-S and ATO-G perform better than STO. ATO-G performs better than ATO-S. Similarly, both AEC2-S and AEC2-G outperform EC2-L1 with AEC2-G performs better than AEC2-S. Comparing AEC2-G with ATO-G, we see that AEC2-G outperforms ATO-G in most experimental settings. This result suggests that by using the eigenvalue constraint, AEC2-G further improves the estimation performance.

HTO has a comparable performance to AEC2-G. For the Toeplitz covariance model, AEC2-G clearly outperforms HTO. For the block covariance model, HTO outperforms AEC2-G. For the banded covariance model, these two methods behave similarly: HTO slightly outperforms AEC2-G in spectral norm error, while AEC2-G is slightly superior to HTO in Frobenius norm error. The main reason that HTO achieves good estimation performance for the block covariance model is because the hard-thresholded covariance matrix is always positive definite in this setting. The performance of HTO decreases on both Toeplitz and banded covariance models since the hard-thresholded covariance matrices are never positive definite in these two settings. This again confirms the fact that using the eigenvalue constraints improves the estimation performance of an estimator.

EC2-M and MCTO achieve good estimation performance in many settings. Especially for both block and banded covariance models, EC2-M achieves the best estimation performance in both Frobenius and spectral norm errors, and MCTO falls slightly behind the EC2-M due to the indefiniteness.

In summary, by adding the eigenvalue constraints we can obtain equal or better estimation than the naive thresholding. Also, by examining the histograms for minimum eigenvalues of the HTO, STO, ATO-G, ATO-S, and MCTO estimates in Figures 3-5, we see that all these

thresholded estimators suffer different extents of indefiniteness, especially for the banded setting, no positive-definite output is found over all replications.

5.3 Real Dataset

To illustrate the efficacy of the proposed methods, we adopt the Parkinson's disease dataset from the UCI machine learning data repository, which is the same as Rothman (2012). It contains 195 speech signals, of which 147 are from cases and 48 are from controls. Each signal is converted into 22 numerical variables. We plug the estimated covariance matrix into quadratic discriminant analysis (QDA), and the performance of the proposed covariance estimators are evaluated based on the classification error rates. Similar to the simulated datasets, we randomly partition the data 10 times into 65 training samples (49 from cases and 16 from controls) and 130 testing cases (98 from cases and 32 from controls).

As shown in Table 4, for all 200 partitions, HTO, ATO-G, ATO-S, MCTO are not applicable to QDA due to the indefiniteness. The STO estimator is positive definite and applicable to QDA for 136 out of 200 random partitions. The average classification error rate of all the 136 applicable STO estimates is 22.91%. While all EC2 estimators are guaranteed to be positive definite and work well for QDA. In terms of classification error, EC2-M achieves the best performance with an error rate 21.45%. The performance of AEC2 is close to EC2-M. This result again confirms that the EC2 estimators achieve better estimation performance than the simple thresholded estimators.

6. CONCLUSIONS

We propose a novel approach named Estimation of Covariance with Eigenvalue Constraints (EC2) for estimating high-dimensional sparse covariance matrices with explicit eigenvalue constraints. EC2 is computationally tractable and theoretically justifiable. In particular, EC2 achieves the optimal rates of convergence in terms of both spectral norm and Frobenius norm errors. Practically, we adopt a data-dependent approach for tuning parameter selection and numerical experiments on both simulated and real datasets are used to illustrate the usefulness of our method.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

We thank the editor, associate editor, and two referees for their helpful comments. Han Liu and Tuo Zhao are supported by NSF grant III-1116730, and Lie Wang is supported by NSF grant DMS-1005539.

REFERENCES

- Bickel P, Levina E. Some Theory for Fisher's Linear Discriminant Function, 'Naive Bayes,' and Some Alternatives When There are Many More Variables Than Observations. *Bernoulli*. 2004; 10:989–1010. 439.
- Bickel P, Levina E. Covariance Regularization by Thresholding. *The Annals of Statistics*. 2008a; 36:2577–2604. 440,452.

- Bickel P, Levina E. Regularized Estimation of Large Covariance Matrices. *The Annals of Statistics*. 2008b; 36:199–227. 439,449.
- Bien J, Tibshirani R. Sparse Estimation of a Covariance Matrix. *Biometrika*. 2011; 98:807–820. 441,444,450. [PubMed: 23049130]
- Box, G.; Jenkins, G.; Reinsel, G. *Time Series Analysis*. 3rd ed.. Prentice Hall; 1994. 439
- Cai T, Liu W. Adaptive Thresholding for Sparse Covariance Matrix Estimation. *Journal of the American Statistical Association*. 2011; 106:672–684. 440.
- Cai T, Zhang C, Zhou H. Optimal Rates of Convergence for Covariance Matrix Estimation. *The Annals of Statistics*. 2010; 38:2118–2144. 440.
- Cai T, Zhou H. Optimal Rates of Convergence for Sparse Covariance Matrix Estimation. *The Annals of Statistics*. 2012; 40:2389–2420. 450.
- Chen S, Donoho D, Saunders M. Atomic Decomposition by Basis Pursuit. *SIAM Journal on Scientific Computing*. 1998; 20:33–61. 440.
- Cressie N. *Statistics for Spatial Data*. Terra Nova. 1992; 4:613–617. 439.
- Fan J, Li R. Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties. *Journal of the American Statistical Association*. 2001; 96:1348–1360. 440.
- Furrer R, Bengtsson T. Estimation of High-Dimensional Prior and Posterior Covariance Matrices in Kalman Filter Variants. *Journal of Multivariate Analysis*. 2007; 98:227–255. 439.
- Gabay D, Mercier B. A Dual Algorithm for the Solution of Nonlinear Variational Problems via Finite Element Approximation. *Computers & Mathematics with Applications*. 1976; 2:17–40. 440.
- Han F, Liu H. Semiparametric Principal Component Analysis. *Advances in Neural Information Processing Systems*. 2012:171–179. 449.
- Han F, Zhao T, Liu H. Coda: High Dimensional Copula Discriminant Analysis. *Journal of Machine Learning Research*. 2013; 14:629–671. 449.
- He, B.; Yuan, X. Technical Report. Nanjing University; 2012. On Non-Ergodic Convergence Rate of Douglas-Rachford Alternating Direction Method of Multipliers. 440,444
- Huang J, Liu N, Pourahmadi M, Liu L. Covariance Matrix Selection and Estimation via Penalised Normal Likelihood. *Biometrika*. 2006; 93:85–98. 439.
- Karoui N. Operator Norm Consistent Estimation of Large Dimensional Sparse Covariance Matrices. *The Annals of Statistics*. 2008; 36:2717–2756. 440.
- Lam C, Fan J. Sparsistency and Rates of Convergence in Large Covariance Matrix Estimation. *The Annals of Statistics*. 2009; 37:42–54. 440,441,442,450.
- Levina E, Rothman A, Zhu J. Sparse Estimation of Large Covariance Matrices via a Nested Lasso Penalty. *The Annals of Applied Statistics*. 2008; 2:245–263. 439.
- Liu H, Han F, Yuan M, Lafferty J, Wasserman L. High Dimensional Semiparametric Gaussian Copula Graphical Models. *The Annals of Statistics*. 2012; 40:2293–2326. 449.
- Liu H, Lafferty J, Wasserman L. The Nonparanormal: Semiparametric Estimation of High Dimensional Undirected Graphs. *Journal of Machine Learning Research*. 2009; 10:2295–2328. 449.
- Negahban S, Ravikumar P, Wainwright M, Yu B. A Unified Framework for High-Dimensional Analysis of m -Estimators With Decomposable Regularizers. *Statistical Science*. 2012; 27:538–557. 451.
- Rothman A. Positive Definite Estimators of Large Covariance Matrices. *Biometrika*. 2012; 99:733–740. 440,449,451,452,453,454,457.
- Rothman A, Bickel P, Levina E, Zhu J. Sparse Permutation Invariant Covariance Estimation. *Electronic Journal of Statistics*. 2008; 2:494–515. 441,442.
- Rothman A, Levina E, Zhu J. Generalized Thresholding of Large Covariance Matrices. *Journal of the American Statistical Association*. 2009; 104:177–186. 440,442,449,452.
- Rothman A, Levina E, Zhu J. A New Approach to Cholesky-based Covariance Regularization in High Dimensions. *Biometrika*. 2010; 97:539–550. 440.
- Searle, S.; Casella, G.; McCulloch, C. *Variance Components*. Vol. 419. Wiley Online Library; 1992. 439

- Tibshirani R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society, Series B*. 1996; 58:267–288. 440.
- Wu W, Pourahmadi M. Nonparametric Estimation of Large Covariance Matrices of Longitudinal Data. *Biometrika*. 2003; 90:831–844. 439.
- Xue L, Ma S, Zou H. Positive Definite l_1 Penalized Estimation of Large Covariance Matrices. *Journal of the American Statistical Association*. 2012; 107:1480–1491. 441,444.
- Xue L, Zou H. Regularized Rank-Based Estimation of High-Dimensional Nonparanormal Graphical Models. *The Annals of Statistics*. 2012; 40:2541–2571. 445.
- Zhang C. Nearly Unbiased Variable Selection Under Minimax Concave Penalty. *The Annals of Statistics*. 2010; 38:894–942. 446.
- Zou H. The Adaptive Lasso and its Oracle Properties. *Journal of the American Statistical Association*. 2006; 101:1418–1429. 440,445.

Figure 1.

Illustration of the minimax concave penalty function and its local linear approximations. In (a), we illustrate the MC+ penalty with different γ 's. In (b)–(d), the LLA approximations of different MC+ penalties are plotted.

Figure 2.
Heatmaps of three simulated covariance matrices for $d = 200$.

Figure 3.
The histograms of the minimum eigenvalues for the Toeplitz setting.

Figure 4.
The histograms of the minimum eigenvalues for the block setting.

Figure 5.
The histograms of the minimum eigenvalues for the banded setting.

Table 1

Quantitative comparison among nine different methods for the Toeplitz setting

d	HTO	STO	ATO-G	ATO-S	MCTO	EC2-L1	AEC2-G	AEC2-S	EC2-M	Log-Det
$\ \cdot\ _F$	100	4.9884 (0.2365)	4.8897 (0.3145)	4.5851 (0.2588)	4.5836 (0.2632)	4.9873 (0.2523)	4.8871 (0.3154)	4.4231 (0.2592)	4.4660 (0.2657)	4.8884 (0.3077)
	200	7.5706 (0.2147)	7.7281 (0.2982)	7.0908 (0.2515)	7.5720 (0.2389)	7.5826 (0.2486)	7.7254 (0.2992)	6.6514 (0.2346)	7.1443 (0.2386)	7.7528 (0.3025)
	400	11.292 (0.2305)	11.946 (0.3293)	10.779 (0.2462)	13.170 (0.2631)	11.262 (0.2453)	11.945 (0.3299)	9.8182 (0.2386)	11.748 (0.2487)	11.901 (0.3301)
$\ \cdot\ _2$	100	1.8999 (0.2452)	2.3238 (0.2321)	1.8246 (0.2301)	1.8762 (0.2310)	1.8957 (0.2543)	2.3248 (0.2318)	1.8487 (0.2370)	1.8946 (0.2329)	2.3808 (0.2333)
	200	2.1781 (0.1920)	2.7299 (0.1686)	2.0674 (0.1995)	2.1844 (0.1931)	2.1775 (0.1569)	2.7306 (0.1683)	2.1221 (0.1939)	2.2052 (0.1891)	2.7352 (0.1678)
	400	2.4055 (0.1886)	3.0652 (0.1488)	2.2739 (0.1749)	2.6387 (0.1521)	2.3408 (0.1924)	3.0654 (0.1488)	2.3680 (0.1799)	2.5857 (0.1511)	3.0593 (0.1501)
P. D.	100	0/200	200/200	12/200	0/200	0/200	200/200	200/200	200/200	200/200
	200	0/200	200/200	0/200	0/200	0/200	200/200	200/200	200/200	200/200
	400	0/200	200/200	0/200	0/200	0/200	200/200	200/200	200/200	200/200

Table 2

Quantitative comparison among nine different methods for the block setting

d	HTO	STO	ATO-G	ATO-S	MCTO	EC2-L1	AEC2-G	AEC2-S	EC2-M	Log-Det	
$\ \cdot\ _F$	100	2.8596 (0.5048)	4.2816 (0.5283)	3.5540 (0.5561)	3.5612 (0.6116)	2.8512 (0.5733)	4.2713 (0.5318)	3.3175 (0.5623)	2.6381 (0.5439)	4.2735 (0.5251)	
	200	5.7245 (1.0861)	8.4954 (1.0666)	7.0914 (1.0696)	7.1843 (1.1559)	5.9215 (1.1145)	8.4488 (1.0813)	6.1891 (1.0291)	5.4628 (1.0063)	8.4821 (1.0774)	
	400	11.5622 (1.8796)	17.3455 (2.1898)	14.6285 (2.1387)	14.8318 (2.4461)	11.085 (1.8598)	17.2004 (2.2313)	12.8227 (2.0607)	13.6331 (2.2910)	10.254 (1.8910)	17.331 (2.1517)
$\ \cdot\ _2$	100	1.6136 (0.4352)	2.1871 (0.3730)	1.8590 (0.4381)	1.8723 (0.4235)	1.5957 (0.4191)	2.1904 (0.3734)	1.7173 (0.4151)	1.4968 (0.4065)	2.1911 (0.3901)	
	200	3.1632 (0.8901)	4.4013 (0.6786)	3.5856 (0.8145)	3.6534 (0.8067)	3.2964 (0.8873)	4.4140 (0.6811)	3.3320 (0.7236)	3.5143 (0.7233)	4.5012 (0.6978)	
	400	6.4316 (1.5812)	8.8578 (1.3325)	7.4399 (1.6952)	7.5318 (1.7206)	6.0839 (1.6128)	8.8864 (1.3349)	6.8903 (1.5598)	7.2267 (1.5976)	5.7409 (1.5215)	8.8488 (1.3035)
P. D.	100	198/200	200/200	7/200	98/200	200/200	200/200	200/200	200/200	200/200	200/200
	200	200/200	179/200	0/200	2/200	200/200	200/200	200/200	200/200	200/200	200/200
	400	200/200	0/200	0/200	0/200	200/200	200/200	200/200	200/200	200/200	200/200

Table 3

Quantitative comparison among nine different methods for the banded setting

d	HTO	STO	ATO-G	ATO-S	MCTO	EC2-L1	AEC2-G	AEC2-S	EC2-M	Log-Det
$\ \cdot\ _F$	100	5.8906 (1.0158)	6.2322 (1.0459)	5.6524 (1.0382)	5.7007 (1.0497)	5.9849 (1.0566)	6.2309 (1.0463)	5.5816 (1.0401)	5.6515 (1.0559)	5.4667 (1.1640)
	200	9.0050 (0.7644)	10.2387 (1.0703)	8.9650 (0.9514)	8.9833 (0.9171)	8.8706 (0.8509)	10.2301 (1.0734)	8.7203 (0.9565)	8.8152 (0.9314)	7.9104 (1.0280)
	400	13.5656 (0.8763)	16.1438 (1.0779)	13.9411 (1.1702)	14.8573 (1.1959)	13.620 (0.8347)	16.1224 (1.0814)	13.2626 (1.1416)	14.3028 (1.1758)	12.014 (0.9258)
$\ \cdot\ _2$	100	3.3106 (1.0020)	4.0640 (0.9660)	3.4677 (0.9929)	3.5170 (0.9852)	3.2976 (1.0265)	4.0643 (0.9661)	3.4701 (0.9830)	3.5184 (0.9832)	3.2817 (1.0044)
	200	3.8574 (0.7660)	5.2307 (0.8045)	4.2794 (0.8243)	4.3174 (0.7712)	3.8045 (0.7759)	5.2328 (0.8047)	4.2751 (0.8023)	4.3199 (0.7752)	3.7769 (0.8422)
	400	4.3551 (0.7719)	6.1805 (0.6044)	4.9620 (0.8302)	5.3328 (0.8063)	4.3890 (0.8208)	6.1846 (0.6043)	4.9356 (0.7235)	5.2862 (0.7599)	4.2781 (0.7020)
P. D.	100	0/200	0/200	0/200	0/200	0/200	200/200	200/200	200/200	200/200
	200	0/200	0/200	0/200	0/200	0/200	200/200	200/200	200/200	200/200
	400	0/200	0/200	0/200	0/200	0/200	200/200	200/200	200/200	200/200

Table 4

Classification error rates among nine different estimators when plugged into QDA

	HTO	STO	ATO-G	ATO-S	MCTO	EC2-L1	AEC2-G	AEC2-S	EC2-M	Log-Det
Test. Err.	N.A. (N.A.)	0.2291 (0.0266)	N.A. (N.A.)	N.A. (N.A.)	N.A. (N.A.)	0.2252 (0.0179)	0.2182 (0.0199)	0.2198 (0.0202)	0.2145 (0.0225)	22.56 (0.200)
P. D.	0/200	136/200	0/200	0/200	0/200	200/200	200/200	200/200	200/200	200/200