
Adaptive Estimation of Graphical Models under Total Positivity

Jiayi Ying^{1,2} José Vinícius de M. Cardoso¹ Daniel P. Palomar¹

Abstract

We consider the problem of estimating (diagonally dominant) M -matrices as precision matrices in Gaussian graphical models. These models exhibit intriguing properties, such as the existence of the maximum likelihood estimator with merely two observations for M -matrices (Lauritzen et al., 2019; Slawski & Hein, 2015) and even one observation for diagonally dominant M -matrices (Truell et al., 2021). We propose an adaptive multiple-stage estimation method that refines the estimate by solving a weighted ℓ_1 -regularized problem at each stage. Furthermore, we develop a unified framework based on the gradient projection method to solve the regularized problem, incorporating distinct projections to handle the constraints of M -matrices and diagonally dominant M -matrices. A theoretical analysis of the estimation error is provided. The proposed method outperforms state-of-the-art methods in precision matrix estimation and graph edge identification, as evidenced by synthetic and financial time-series data sets.

1. Introduction

Total positivity, as a strong form of positive dependence, has found its applications in various fields such as factor analysis in psychometrics (Lauritzen et al., 2019), taxonomic reasoning (Lake & Tenenbaum, 2010), graph signal processing (Egilmez et al., 2017), and financial markets (Agrawal et al., 2020). For instance, stock markets exemplify total positivity, as stocks generally exhibit positive dependence due to the influence of market factors (Agrawal et al., 2020). A distribution that satisfies total

positivity is also known as multivariate totally positive of order two (MTP₂). For a more in-depth understanding, we refer the reader to (Fallat et al., 2017). In the Gaussian context, the concept of MTP₂ becomes more straightforward: a Gaussian distribution is MTP₂ if and only if the precision matrix (*i.e.*, inverse covariance matrix) is an M -matrix (Karlin & Rinott, 1983). In this paper, our focus lies on estimating (diagonally dominant) M -matrices as precision matrices in Gaussian graphical models.

Estimation of precision matrices in general Gaussian graphical models have been extensively studied in the literature. A well-known method, known as graphical lasso (Banerjee et al., 2008; d’Aspremont et al., 2008), is formulated as an ℓ_1 -regularized Gaussian maximum likelihood estimation. Numerous extensions of graphical lasso have been explored (Friedman et al., 2008; Ravikumar et al., 2011; Lam & Fan, 2009; Fan et al., 2014; Loh & Wainwright, 2015; 2017; Sun et al., 2018; Hsieh et al., 2014). Other notable, though not exhaustive, works include graphical Dantzig (Yuan, 2010), CLIME (Cai et al., 2011), and TIGER (Liu & Wang, 2017). Recent research has illuminated the intriguing properties of incorporating additional M -matrix constraint, which proves to be advantageous in the realms of high-dimensional statistics and signal processing on graphs.

Recent studies (Lauritzen et al., 2019; Slawski & Hein, 2015; Lauritzen et al., 2021) show that the M -matrix constraint significantly reduces the sample size required for the maximum likelihood estimator (MLE) to exist. Specifically, Lauritzen et al. (2019); Slawski & Hein (2015) established that the MLE under the M -matrix constraint exists if the sample size $n \geq 2$, regardless of the underlying dimension p , substantially reducing the $n \geq p$ requirement for general Gaussian distributions. Soloff et al. (2020) revealed that the M -matrix constraint, serving as implicit regularization, is vital for achieving the minimax rate of $\sqrt{\frac{\log p}{n}}$; without this constraint, the rate cannot exceed $\sqrt{\frac{p}{n}}$. Remarkably, Lauritzen et al. (2021) proved that under the MTP₂ constraint, the MLE in binary distributions may exist with only $n = p$ observations, in contrast to the 2^p observations required in the absence of this constraint.

¹The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong ²HKUST Shenzhen-Hong Kong Collaborative Innovation Research Institute, Shenzhen, China. Correspondence to: Jiayi Ying <jx.ying@connect.ust.hk>.

The diagonally dominant M -matrix constraint is increasingly drawing interest in Gaussian models. In these distributions, diagonally dominant M -matrices, as precision matrices, exhibit a property called log- L^1 -concave (LLC) (Robeva et al., 2021) or strong MTP_2 (Röttger et al., 2021). This property is essential for analyzing positive dependence in multivariate Pareto distributions. A recent study (Truell et al., 2021) revealed that Gaussian models under LLC/strong MTP_2 act as a convex relaxation for Brownian motion tree models—a class of Gaussian models on trees—and showed that its MLE exists almost surely when $n = 1$. This finding was established by connecting Laplacian constrained Gaussian Markov random fields (LGMRF) (Ying et al., 2020b; Cardoso et al., 2021; 2022; Kumar et al., 2020) and leveraging results on MLE existence in LGMRF with a single observation (Ying et al., 2021a).

The estimation of MTP_2 graphical models has become a growing area of interest in the field of signal processing on graphs (Shuman et al., 2013; Ortega et al., 2018; Dong et al., 2019), which focuses on handling data defined on irregular graph domains. Precision matrices in MTP_2 graphical models belong to the class of generalized Laplacian matrices (Biyikoglu et al., 2007), each of which is a symmetric matrix with non-positive off-diagonal entries that are associated with a graph. The eigenvectors of a generalized Laplacian define a graph Fourier transform (Shuman et al., 2013), supported by the discrete nodal domain theorem (Davies et al., 2001) that an eigenvector corresponding to a larger eigenvalue is associated with a higher frequency. We note that such a spectral property does not hold for general positive definite matrices.

One approach to estimating MTP_2 graphical models is MLE (Lauritzen et al., 2019; Slawski & Hein, 2015), which implicitly promotes sparsity using the M -matrix constraint. However, it cannot adjust sparsity levels to specific values, a feature often desired in practice. The ℓ_1 -norm regularized MLE (Egilmez et al., 2017; Pavez & Ortega, 2016; Cai et al., 2021; Deng & So, 2020) offers improved sparsity control and can be solved using techniques like block coordinate descent (Egilmez et al., 2017; Pavez & Ortega, 2016), proximal point algorithm (Deng & So, 2020), and projected Newton-like methods (Cai et al., 2021). Notably, Pavez et al. (2018) found that some zero patterns in estimated matrix and graph components can be determined through thresholded sample covariance matrices. Estimating diagonally dominant M -matrices as precision matrices is explored in (Egilmez et al., 2017). Additionally, Wang et al. (2020) proposed a graph structure learning algorithm based on conditional independence testing, eliminating the need for adjustment of tuning parameters.

This paper focuses on estimating (diagonally dominant) M -matrices as precision matrices in MTP_2 Gaussian graphical models. The main contributions of this paper are threefold:

- We propose an adaptive multiple-stage estimation method that refines the estimate by solving a weighted ℓ_1 -regularized problem at each stage. Then we develop a unified framework based on the gradient projection method to solve the regularized problem, equipped with distinct projections to handle the constraints of M -matrices and diagonally dominant M -matrices.
- We provide a thorough theoretical analysis of the estimation error for our method, which comprises two components: optimization error and statistical error. The optimization error decays at a linear rate, highlighting the progressive refinement of the estimate across iterative stages, whereas the statistical error captures the intrinsic statistical rate.
- Experiments on synthetic and financial time-series data demonstrate that our method outperforms state-of-the-art methods, concurrently achieving lower precision matrix estimation error and higher graph edge selection accuracy. Additionally, we observe that incorporating M -matrix constraint enhances the robustness regarding the selection of the regularization parameter.

Lower and upper case bold letters denote vectors and matrices, respectively. $[p]$ denotes the set $\{1, \dots, p\}$. $\|\mathbf{x}\|_{\max} = \max_i |x_i|$. \mathbb{S}_{++}^p denotes the set of $p \times p$ positive definite matrices. $\lambda_{\max}(\mathbf{X})$ and $\lambda_{\min}(\mathbf{X})$ denote the maximum and minimum eigenvalues of \mathbf{X} , respectively. For matrix \mathbf{X} and set S , \mathbf{X}_S denotes a vector with dimension $|S|$, containing the entries of \mathbf{X} indexed by S .

2. Preliminaries and Problem Formulation

We first introduce preliminaries about MTP_2 Gaussian graphical models, then present the problem formulation.

2.1. MTP_2 Gaussian Graphical Models

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with the set of nodes \mathcal{V} and the set of edges \mathcal{E} . Associating a random vector $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}^*)$ with graph \mathcal{G} forms a Gaussian graphical model that satisfies the following properties:

$$\begin{aligned} \Theta_{ij}^* \neq 0 &\iff \{i, j\} \in \mathcal{E} \quad \forall i \neq j, \\ \Theta_{ij}^* = 0 &\iff x_i \perp\!\!\!\perp x_j \mid \mathbf{x}_{[p] \setminus \{i, j\}}, \end{aligned}$$

where $x_i \perp\!\!\!\perp x_j \mid \mathbf{x}_{[p] \setminus \{i, j\}}$ indicates that x_i is conditionally independent of x_j given the other random variables. Consequently, graph \mathcal{G} characterizes the sparsity pattern of the precision matrix $\boldsymbol{\Theta}^*$, where missing edges represent conditional independence between random variables.

We concentrate on estimating precision matrices for random variables following an MTP_2 Gaussian distribution. A multivariate Gaussian distribution with a positive definite covariance matrix Σ^* is MTP_2 if and only if its precision matrix Θ^* is a symmetric M -matrix (Karlin & Rinott, 1983), i.e., $\Theta_{ij}^* \leq 0$ for all $i \neq j$. This is equivalent to stating that all partial correlations are nonnegative, where the partial correlation between any two variables x_i and x_j conditioned on all other variables equals $-\Theta_{ij}^* / \sqrt{\Theta_{ii}^* \Theta_{jj}^*}$.

2.2. Problem Formulation

Consider a zero-mean random vector $\mathbf{x} = (x_1, \dots, x_p)$ following a Gaussian distribution $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma^*)$. Our goal is to estimate the precision matrix $\Theta^* := (\Sigma^*)^{-1}$, which is a (diagonally dominant) M -matrix, given n independent and identically distributed observations $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)} \in \mathbb{R}^p$.

To estimate the precision matrix under MTP_2 Gaussian distributions, we can solve the penalized Gaussian maximum likelihood estimation problem as follows:

$$\underset{\Theta \in \mathcal{S}}{\text{minimize}} \quad -\log \det(\Theta) + \text{tr}(\Theta \widehat{\Sigma}) + \sum_{i \neq j} h_\lambda(|\Theta_{ij}|), \quad (1)$$

where h_λ is a sparsity penalty function, such as the ℓ_1 -norm considered in (Egilemez et al., 2017), $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}^{(i)} (\mathbf{x}^{(i)})^\top$ is the sample covariance matrix, and \mathcal{S} is a feasible set. We consider two cases for \mathcal{S} :

$$\mathcal{M}^p = \{\Theta \in \mathbb{S}_{++}^p \mid \Theta_{ij} = \Theta_{ji} \leq 0, \forall i \neq j\}, \quad (2)$$

and

$$\mathcal{M}_D^p = \{\Theta \in \mathbb{S}_{++}^p \mid \Theta_{ij} = \Theta_{ji} \leq 0, \forall i \neq j, \Theta \mathbf{1} \geq \mathbf{0}\}, \quad (3)$$

where \mathcal{M}^p represents the set of all symmetric positive definite M -matrices, and \mathcal{M}_D^p adds constraint $\Theta \mathbf{1} \geq \mathbf{0}$ compared to \mathcal{M}^p , leading Θ to be a diagonally dominant M -matrix. We propose an adaptive method for estimating M -matrices and diagonally dominant M -matrices.

3. Proposed Method

We first propose an adaptive multiple-stage method, then develop a unified framework for solving each stage.

3.1. Adaptive Estimation

Our method comprises multiple stages, where each stage refines the previous stage's estimate by solving a weighted ℓ_1 -regularized problem. Specifically, at the k -th stage, we obtain the estimate $\widehat{\Theta}^{(k)}$ by solving the following problem:

$$\underset{\Theta \in \mathcal{S}}{\text{minimize}} \quad -\log \det(\Theta) + \text{tr}(\Theta \widehat{\Sigma}) + \sum_{i \neq j} \lambda_{ij}^{(k)} |\Theta_{ij}|, \quad (4)$$

where $\lambda_{ij}^{(k)} = p_\lambda(|\widehat{\Theta}_{ij}^{(k-1)}|)$ with p_λ the weight-updating function, and $\widehat{\Theta}^{(k-1)}$ is obtained at the $(k-1)$ -th stage.

Algorithm 1 Solve Problem (5)

- 1: **Input:** Sample covariance matrix $\widehat{\Sigma}$, $\Lambda^{(k)}$, σ , α , β .
- 2: **for** $t = 0, 1, 2, \dots$ **do**
- 3: Compute $\nabla f(\Theta_t) = -\Theta_t^{-1} + \widehat{\Sigma} - \Lambda^{(k)}$;
- 4: $m \leftarrow 0$;
- 5: **repeat**
- 6: Update $\Theta_{t+1} = \mathcal{P}_{\mathcal{S}_t}(\Theta_t - \sigma \beta^m \nabla f(\Theta_t))$;
- 7: $m \leftarrow m + 1$;
- 8: **until** $\Theta_{t+1} \in \mathbb{S}_{++}^p$ and
- 9: $f(\Theta_{t+1}) \leq f(\Theta_t) - \alpha \sigma \beta^m \|G_{\frac{1}{\eta_t}}(\Theta_t)\|_{\mathbb{F}}^2$;
- 10: **end for**
- 11: **Output:** $\widehat{\Theta}^{(k)}$.

When the estimate $\widehat{\Theta}^{(k-1)}$ from the previous stage exhibits a large coefficient for its (i, j) entry, it is reasonable to assign a small $\lambda_{ij}^{(k)}$ in Problem (4). Therefore, p_λ should be monotonically non-increasing. In the initial stage, we set each $\lambda_{ij}^{(1)} = \lambda$, reducing Problem (4) to the well-studied ℓ_1 -regularized estimation (Egilemez et al., 2017; Ying et al., 2021b). Our method refines the ℓ_1 -norm estimate in subsequent stages, offering an improvement over the standard ℓ_1 -norm approach.

By choosing p_λ as the derivative of the sparsity penalty function h_λ in Problem (1), the sequence $\{\widehat{\Theta}^{(k)}\}_{k \geq 1}$ converges to a stationary point of Problem (1) with a nonconvex penalty. In this context, our method aligns with the local linear approximation method (Zou & Li, 2008) and the broader framework of multi-stage convex relaxation (Zhang, 2010b). Nonconvex approaches may encounter issues with sensitivity to the regularization parameter selection, as illustrated in Figure 6. However, we demonstrate that incorporating the M -matrix constraint considerably enhances the robustness regarding the selection of regularization parameter.

3.2. Gradient Projection Method

For each $\Theta \in \mathcal{S}$, where $\mathcal{S} = \mathcal{M}^p$ or \mathcal{M}_D^p , all off-diagonal entries are nonpositive, leading to the following representation of Problem (4):

$$\underset{\Theta \in \mathcal{S}}{\text{minimize}} \quad -\log \det(\Theta) + \text{tr}(\Theta \widehat{\Sigma}) - \sum_{i \neq j} \lambda_{ij}^{(k)} \Theta_{ij}. \quad (5)$$

We design a gradient projection method to solve Problem (5). Initially, we perform a gradient descent step, $\Theta_t - \eta_t \nabla f(\Theta_t)$, with ∇f representing the gradient of the objective function f . As this step may exceed the constraint set \mathcal{S} , we subsequently project it back onto \mathcal{S} . However, such projection onto \mathcal{S} does not exist due to its nonclosedness, resulting from the nonclosedness of \mathbb{S}_{++}^p .

Instead, we express the set \mathcal{S} as the intersection of a closed set \mathcal{S}_d and the open set \mathbb{S}_{++}^p :

$$\mathcal{S} = \mathcal{S}_d \cap \mathbb{S}_{++}^p. \quad (6)$$

We handle the constraints of \mathcal{S}_d and \mathbb{S}_{++}^p separately. The closed set \mathcal{S}_d constraint is manageable through projection, leading to the following projected gradient descent:

$$\Theta_{t+1} = \mathcal{P}_{\mathcal{S}_d}(\Theta_t - \eta_t \nabla f(\Theta_t)), \quad (7)$$

where $\mathcal{P}_{\mathcal{S}_d}$ denotes the projection onto the set \mathcal{S}_d with respect to the Frobenius norm.

The constraint of \mathbb{S}_{++} cannot be managed by projection due to its nonclosedness. One efficient approach to enforce positive definiteness is using a line search procedure (Hsieh et al., 2014). Specifically, at the t -th iteration, we try the step size $\eta_t \in \sigma \{\beta^0, \beta^1, \beta^2, \dots\}$, with $\beta \in (0, 1)$ and $\sigma > 0$, until we find the smallest $m \in \mathbb{N}$ such that the iterate Θ_{t+1} in (7), with $\eta_t = \sigma \beta^m$, ensures $\Theta_{t+1} \in \mathbb{S}_{++}^p$ and satisfies:

$$f(\Theta_{t+1}) \leq f(\Theta_t) - \alpha \eta_t \|G_{\frac{1}{\eta_t}}(\Theta_t)\|_{\mathbb{F}}^2, \quad (8)$$

where $\alpha \in (0, 1)$, and $G_{\frac{1}{\eta_t}}(\Theta_t) = \frac{1}{\eta_t}(\Theta_t - \Theta_{t+1})$. We present our algorithm in Algorithm 1, where $\Lambda^{(k)}$ contains all weights $\lambda_{ij}^{(k)}$, and diagonal entries are set to zero.

Theorem 3.1. *The sequence $\{\Theta_t\}_{t \geq 0}$ established by Algorithm 1 converges to the optimal solution of Problem (5).*

The proof is provided in Appendix B.2. It is important to note that the existence of the minimizer for Problem (5) is assumed throughout this paper, which can be guaranteed almost surely if the sample size $n \geq 2$ for the case of $\mathcal{S} = \mathcal{M}^p$ (Lauritzen et al., 2019; Slawski & Hein, 2015) and $n \geq 1$ for the case of $\mathcal{S} = \mathcal{M}_D^p$ (Truell et al., 2021).

Although our method needs to solve a sequence of log-determinant programs (5), numerical results indicate a rapid decrease in the number of iterations required for solving (5) as k increases (see Figure 8). This accelerated convergence can be attributed to the use of the estimate from the previous stage as an initial point, thereby providing a warm start for faster convergence.

We now present notable extensions of the proposed gradient projection algorithm. First, the algorithm can be readily extended to solve Problem (1) with a nonconvex penalty. Second, the algorithm can be further adapted to solve the following problem for the case of multivariate t -distribution:

$$\underset{\Theta \in \mathcal{S}}{\text{minimize}} -\log \det(\Theta) + \frac{p + \nu}{n} \sum_{i=1}^n \log \left(1 + \frac{1}{\nu} (\mathbf{x}^{(i)})^\top \Theta \mathbf{x}^{(i)} \right),$$

where ν denotes the number of degrees of freedom, and each $\mathbf{x}^{(i)}$ is an observation. This formulation can be employed

to learn positive partial correlation graphs. Notably, elliptical MTP₂ distributions are highly restrictive, and the total positivity property is not applicable to t -distributions (Rossell & Zwiernik, 2021). It is worth mentioning that, unlike block coordinate descent (Egilemez et al., 2017) and projected Newton-like methods (Cai et al., 2021), the proposed algorithm can handle these extensions, resulting in a more flexible and versatile approach.

3.3. Computation of Projections

We present the computation of projection $\mathcal{P}_{\mathcal{S}_d}$ in (7) for both cases of estimating M -matrices and diagonally dominant M -matrices, i.e., $\mathcal{S} = \mathcal{M}^p$ and \mathcal{M}_D^p in Problem (5).

3.3.1. ESTIMATION OF M-MATRICES

For the case of estimating an M -matrix as the precision matrix, the constraint set of Problem (5) is set as $\mathcal{S} = \mathcal{M}^p$, defined in (2). The closed set \mathcal{S}_d in (6) then becomes

$$\mathcal{S}_d = \{\Theta \in \mathbb{R}^{p \times p} \mid \Theta_{ij} = \Theta_{ji} \leq 0, \forall i \neq j\}, \quad (9)$$

which can be written as the intersection of two sets:

$$\mathcal{S}_d = \mathcal{S}_A \cap \mathcal{S}_B, \quad (10)$$

where $\mathcal{S}_A := \{\Theta \in \mathbb{R}^{p \times p} \mid \Theta_{ij} = \Theta_{ji}, \forall i \neq j\}$ and $\mathcal{S}_B := \{\Theta \in \mathbb{R}^{p \times p} \mid \Theta_{ij} \leq 0, \forall i \neq j\}$. If Θ_t is symmetric, then the update $\mathcal{P}_{\mathcal{S}_B}(\Theta_t - \eta_t \nabla f(\Theta_t))$ preserves symmetry. Thus, we only need to project $\Theta_t - \eta_t \nabla f(\Theta_t)$ onto \mathcal{S}_B . As a result, the iterate (7) can be simplified to

$$\Theta_{t+1} = \mathcal{P}_{\mathcal{S}_B}(\Theta_t - \eta_t \nabla f(\Theta_t)), \quad (11)$$

where the projection $\mathcal{P}_{\mathcal{S}_B}$ can be computed as follows:

$$[\mathcal{P}_{\mathcal{S}_B}(\mathbf{X})]_{ij} = \begin{cases} \min(X_{ij}, 0) & \text{if } i \neq j, \\ X_{ij} & \text{if } i = j. \end{cases} \quad (12)$$

By initializing with a symmetric Θ_0 , every point in the sequence $\{\Theta_t\}_{t \geq 0}$ generated by (11) maintains symmetry.

3.3.2. ESTIMATION OF DIAGONALLY DOMINANT M-MATRICES

For the case of estimating a diagonally dominant M -matrix as the precision matrix, we set $\mathcal{S} = \mathcal{M}_D^p$ in (5), with \mathcal{M}_D^p defined in (3). The closed set \mathcal{S}_d in (6) then becomes

$$\mathcal{S}_d = \{\Theta \in \mathbb{R}^{p \times p} \mid \Theta_{ij} = \Theta_{ji} \leq 0, \forall i \neq j, \Theta \mathbf{1} \geq \mathbf{0}\}, \quad (13)$$

which can be written as the intersection of two sets:

$$\mathcal{S}_d = \mathcal{S}_A \cap \mathcal{S}_C, \quad (14)$$

where $\mathcal{S}_C := \{\Theta \in \mathbb{R}^{p \times p} \mid \Theta \mathbf{1} \geq \mathbf{0}, \Theta_{ij} \leq 0, \forall i \neq j\}$.

The iterate (7) is then written as:

$$\Theta_{t+1} = \mathcal{P}_{\mathcal{S}_A \cap \mathcal{S}_C}(\Theta_t - \eta_t \nabla f(\Theta_t)). \quad (15)$$

In contrast to (11), iterate (15) cannot be simplified, as $\mathcal{P}_{\mathcal{S}_C}(\Theta_t - \eta_t \nabla f(\Theta_t))$ does not guarantee symmetry.

We now discuss computing the projection $\mathcal{P}_{\mathcal{S}_A \cap \mathcal{S}_C}$ in (15), defined as the minimizer of the following problem:

$$\mathcal{P}_{\mathcal{S}_A \cap \mathcal{S}_C}(\mathbf{Y}) := \arg \min_{\mathbf{X} \in \mathcal{S}_A \cap \mathcal{S}_C} \|\mathbf{X} - \mathbf{Y}\|_F^2. \quad (16)$$

Problem (16) is a projection problem of a given point \mathbf{Y} onto the intersection of the sets \mathcal{S}_A and \mathcal{S}_C . To solve Problem (16), we design an algorithm based on Dykstra's projection (Boyle & Dykstra, 1986), which aims to find the nearest point projection onto the intersection of closed convex sets. The algorithm is summarized in Algorithm 2, with \mathbf{Q}_k denoting the increment associated with the set \mathcal{S}_C . We do not need to introduce the increment associated with \mathcal{S}_A for convergence, since \mathcal{S}_A is a subspace.

Theorem 3.2. *The sequences $\{\mathbf{A}_k\}$ and $\{\mathbf{C}_k\}$ generated by Algorithm 2 converge to the minimizer of Problem (16).*

The convergence established in Theorem 3.2 is based on the convergence results of Dykstra's projection algorithm, as presented in (Boyle & Dykstra, 1986).

Algorithm 2 Compute $\mathcal{P}_{\mathcal{S}_A \cap \mathcal{S}_C}(\mathbf{Y})$

- 1: **Input:** $\mathbf{C}_0 = \mathbf{Y}$, $\mathbf{Q}_0 = \mathbf{0}$, ϵ ;
 - 2: $k = 1$;
 - 3: **repeat**
 - 4: $\mathbf{A}_k = \mathcal{P}_{\mathcal{S}_A}(\mathbf{C}_{k-1})$;
 - 5: $\mathbf{C}_k = \mathcal{P}_{\mathcal{S}_C}(\mathbf{A}_k + \mathbf{Q}_{k-1})$;
 - 6: $\mathbf{Q}_k = \mathbf{A}_k - \mathbf{C}_k + \mathbf{Q}_{k-1}$;
 - 7: $k = k + 1$;
 - 8: **until** $\|\mathbf{A}_k - \mathbf{C}_k\|_F < \epsilon$
 - 9: **Output:** $\mathcal{P}_{\mathcal{S}_A \cap \mathcal{S}_C}(\mathbf{Y}) = \mathbf{A}_k$.
-

We now present the computation of $\mathcal{P}_{\mathcal{S}_A}$ and $\mathcal{P}_{\mathcal{S}_C}$ from Algorithm 2. The computation of $\mathcal{P}_{\mathcal{S}_A}$ is straightforward:

$$\mathcal{P}_{\mathcal{S}_A}(\mathbf{Y}) = (\mathbf{Y} + \mathbf{Y}^\top)/2.$$

Projection $\mathcal{P}_{\mathcal{S}_C}$ is defined as follows:

$$\mathcal{P}_{\mathcal{S}_C}(\mathbf{Y}) := \arg \min_{\mathbf{X} \in \mathcal{S}_C} \frac{1}{2} \|\mathbf{X} - \mathbf{Y}\|_F^2. \quad (17)$$

To compute $\mathcal{P}_{\mathcal{S}_C}$, we solve Problem (17) row by row. For the r -th row, we address the following problem:

$$\begin{aligned} & \underset{\mathbf{x} \in \mathbb{R}^p}{\text{minimize}} && \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|^2, \\ & \text{subject to} && \mathbf{x}^\top \mathbf{1} \geq 0, \mathbf{x}_{\setminus r} \leq \mathbf{0}, \end{aligned} \quad (18)$$

where $\mathbf{y} \in \mathbb{R}^p$ contains all entries of the r -th row of \mathbf{Y} , and $\mathbf{x}_{\setminus r} \in \mathbb{R}^{p-1}$ includes all entries of \mathbf{x} except the r -th one.

Proposition 3.3 below, proven in Appendix B.3, provides the optimal solution of Problem (18), which is a variant of the projection onto the simplex (Condat, 2016).

Proposition 3.3. *Let $\hat{\mathbf{x}}$ be the optimal solution of Problem (18), which can be obtained as follows:*

- If $y_r \geq -\sum_{i \in [p] \setminus r} \min(y_i, 0)$, then $\hat{x}_r = y_r$, and $\hat{x}_i = \min(y_i, 0)$ for $i \neq r$.
- If $y_r < -\sum_{i \in [p] \setminus r} \min(y_i, 0)$, then $\hat{x}_r = y_r + \rho$, and $\hat{x}_i = \min(y_i + \rho, 0)$ for $i \neq r$, where ρ satisfies $\rho + \sum_{i \in [p] \setminus r} \min(y_i + \rho, 0) + y_r = 0$.

Remark 3.4. Let $g(\rho) := \rho + \sum_{i \in [p] \setminus r} \min(y_i + \rho, 0) + y_r$. Note that there exists a unique solution to the piecewise linear equation $g(\rho) = 0$. On one hand, $g(0) < 0$, and $g(a) > 0$ for any $a > \max_i |y_i|$. Therefore, there exists at least one solution to $g(\rho) = 0$ since $g(\rho)$ is continuous. On the other hand, the solution is unique because $g(\rho)$ is a monotone function that is strictly increasing.

Let $\mathbf{y}_{\setminus r} \in \mathbb{R}^{p-1}$ represent the vector that contains all entries of \mathbf{y} except the r -th one. To solve the equation $g(\rho) = 0$, we follow the procedures outlined in (Palomar & Fonollosa, 2005), which consists of three steps:

1. Sort $\mathbf{y}_{\setminus r}$ and obtain the sorted version $\tilde{\mathbf{y}}_{\setminus r}$, where $[\tilde{\mathbf{y}}_{\setminus r}]_1 \leq [\tilde{\mathbf{y}}_{\setminus r}]_2 \leq \dots \leq [\tilde{\mathbf{y}}_{\setminus r}]_{p-1}$.
2. Find $M := \max_{1 \leq m \leq p-1} \left\{ m \mid \frac{y_r + \sum_{i=1}^m [\tilde{\mathbf{y}}_{\setminus r}]_i}{m+1} > [\tilde{\mathbf{y}}_{\setminus r}]_m \right\}$.
3. Compute $\rho = \frac{-y_r - \sum_{m=1}^M [\tilde{\mathbf{y}}_{\setminus r}]_m}{M+1}$.

We note that sorting the vector $\mathbf{y}_{\setminus r}$ is the most computationally demanding step in the above procedures, typically necessitating $\mathcal{O}(p \log p)$ operations.

4. Analysis of Estimation Error

In this section, we provide theoretical analysis of the estimation error for the proposed method.

4.1. Theoretical Results

Let Σ^* and Θ^* denote the underlying covariance and precision matrices, respectively, where $\Theta^* \in \mathcal{S}$ ($\mathcal{S} = \mathcal{M}^p$ or \mathcal{M}_D^p). Define S^* as the support set of the underlying precision matrix Θ^* , excluding the diagonal entries, *i.e.*,

$$S^* := \{(i, j) \in [p]^2 \mid \Theta_{ij}^* \neq 0, i \neq j\}. \quad (19)$$

We require some mild conditions for the weight-updating function p_λ in Assumption 4.1 and the underlying precision matrix Θ^* in Assumption 4.2.

Assumption 4.1. The function $p_\lambda : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfies the following conditions:

1. $p_\lambda(x)$ is monotonically nonincreasing, and $p_\lambda(0) = \lambda$;
2. There exists a $\gamma > 0$ such that $p_\lambda(x) = 0$ for $x \geq \gamma\lambda$, and $p_\lambda(c_0\lambda) \geq \frac{\sqrt{2}}{2}\lambda$, where $c_0 = 12\lambda_{\max}^2(\Theta^*)$ is a constant.

Assumption 4.2. The underlying precision matrix Θ^* has at most s nonzero entries, and the off-diagonal nonzero entries satisfy $\min_{(i,j) \in S^*} |\Theta_{ij}^*| \geq (c_0 + \gamma)\lambda$, with c_0 and γ defined in Assumption 4.1. There exists a constant δ such that $0 < \delta \leq \lambda_{\min}(\Theta^*) \leq \lambda_{\max}(\Theta^*) \leq 1/\delta < \infty$.

Remark 4.3. Assumption 4.1 fundamentally covers folded concave penalties, such as smooth clipped absolute deviation (SCAD) (Fan & Li, 2001) and minimax concave penalty (MCP) (Zhang, 2010a). In Assumption 4.2, the conditions on the underlying precision matrix are relatively mild, as the regularization parameter λ in our theorems takes the order of $\sqrt{\log p/n}$, which could be small when the sample size n increases. The assumption regarding minimal magnitude of signals is often utilized in nonconvex optimization analysis (Wang et al., 2016; Ying et al., 2020a).

The following theorem, proven in Appendix B.4, provides a non-asymptotic guarantee for estimation error.

Theorem 4.4. Under Assumptions 4.1 and 4.2, take the regularization parameter $\lambda = \sqrt{c_1\tau \log p/n}$ with $\tau > 2$. If the sample size satisfies $n \geq 8c_0c_1\tau s \log p$, then the sequence $\hat{\Theta}^{(k)}$ generated by the proposed method satisfies

$$\|\hat{\Theta}^{(k)} - \Theta^*\|_F \leq \underbrace{4c_0\|(\Sigma^* - \hat{\Sigma})_{S^* \cup I_p}\|}_{\text{Statistical error}} + \underbrace{\rho^k\|\hat{\Theta}^{(0)} - \Theta^*\|_F}_{\text{Optimization error}},$$

with probability at least $1 - 4/p^{\tau-2}$, where $\rho = \frac{2+\sqrt{2}}{4}$, $c_0 = 12\lambda_{\max}^2(\Theta^*)$, and $c_1 = (80 \max_i \Sigma_{ii}^*)^2$.

Theorem 4.4 establishes that the estimation error between the estimated precision matrix and the underlying precision matrix can be upper bounded by two terms: optimization error and statistical error. The optimization error decreases at a linear rate, indicating the progressive refinement of the estimate across iterative stages. In contrast, the statistical error remains independent of k and does not decrease during iterations, capturing the intrinsic statistical rate. Consequently, the estimation error is primarily dominated by the statistical error.

We note that the statements in Theorem 4.4 are applicable to estimate both M -matrices and diagonally dominant M -matrices, i.e., $S = \mathcal{M}^p$ or \mathcal{M}_D^p in Algorithm 1. The parameter τ is user-defined; a larger τ increases the probability of the claims being true but imposes a stricter requirement on the sample size.

Corollary 4.5. Under the same assumptions and conditions as stated in Theorem 4.4, the sequence $\hat{\Theta}^{(k)}$ generated by the proposed method satisfies

$$\|\hat{\Theta}^{(k)} - \Theta^*\|_F \leq \underbrace{c\sqrt{s \log p/n}}_{\text{Statistical error}} + \underbrace{\rho^k\|\hat{\Theta}^{(0)} - \Theta^*\|_F}_{\text{Optimization error}},$$

with probability at least $1 - 4/p^{\tau-2}$, where $c = 2c_0\sqrt{2c_1\tau}$.

Corollary 4.5 reveals that the statistical error follows the order of $\sqrt{s \log p/n}$, aligning with the minimax rate achieved in unconstrained graphical models (Ravikumar et al., 2011; Loh & Wainwright, 2015). However, the impact of the M -matrix constraint on the minimax rate for estimating sparse precision matrices remains uncertain. We highlight a recent study (Soloff et al., 2020) that investigated estimating precision matrices under the nonpositivity constraint without sparsity regularization. The minimax optimal rate under symmetrized Stein loss is established to be $\sqrt{\log p/n}$, while removing the nonpositivity constraint results in a reduced rate of $\sqrt{p/n}$. This finding suggests that the nonpositivity constraint affects the minimax rate, particularly in non-sparse situations.

5. Experimental Results

We conduct experiments on synthetic and financial time-series data. The code of our method is publicly available at <https://github.com/jxying/ddmtp2>.

State-of-the-art methods for comparison include GLASSO (Friedman et al., 2008), CLIME (Cai et al., 2011), GSCAD (Loh & Wainwright, 2015), GGL (Egilmez et al., 2017), DDGL (Egilmez et al., 2017), SLTP (Wang et al., 2020), and GOLAZO (Lauritzen & Zwiernik, 2022). GLASSO, CLIME and GSCAD focus on estimating precision matrices for general graphical models. The first two employ the ℓ_1 -penalty, while the latter utilizes the SCAD penalty. In contrast, GGL, DDGL, and SLTP are specifically designed for MTP_2 graphical models. Both GGL and DDGL use the ℓ_1 -norm regularization, aiming for estimating M -matrices and diagonally dominant M -matrices, respectively. Meanwhile, SLTP concentrates on learning graph structures rather than estimating precision matrices. Lastly, GOLAZO is tailored for estimating positive graphical models.




Figure 1: Illustration of graph structures: (a) grid, (b) line, (c) Barabasi-Albert model, and (d) stochastic block model.




Figure 2: Estimation errors and true positive rates versus false positive rates on the grid graph with $\Theta^* \in \mathcal{M}^p$. The number of nodes $p = 100$, and the sample size $n = 100$.




Figure 3: Estimation errors and true positive rates versus false positive rates on the line graph with $\Theta^* \in \mathcal{M}^p$. The number of nodes $p = 100$, and the sample size $n = 40$.

5.1. Synthetic Data

We examine four prevalent graph structures in synthetic experiments: grid, line, Barabasi-Albert model (Barabási & Albert, 1999), and stochastic block model (Holland et al., 1983), as depicted in Figure 1. Estimating these graph structures is essential in various applications, including social network analysis, image segmentation, and community detection. The graph structures are generated randomly, with each edge assigned a random weight. To assess the performance of precision matrix estimation, we employ estimation error as a metric. For evaluating graph edge selection accuracy, we utilize true/false positive rate and F-score. Further details about experimental settings, including performance measure definitions and synthetic data generation, can be found in Appendix A.

Figure 2 presents estimation errors and true positive rates of various methods as a function of false positive rates on a grid graph, given the number of nodes ($p = 100$), sample size ($n = 100$), and underlying precision matrix $\Theta^* \in \mathcal{M}^p$. The curve depicting true positive rates against false positive rates in Figure 2 is known as Receiver Operating Characteristic (ROC) curve, which is generated by setting a range of regularization parameter values from 10^{-3} to 1 for each method, except for SLTP. Instead, SLTP uses a parameter γ to control the sparsity of the learned graph, with values ranging from 0.75 to 0.98. The curves in both Figures 2 and 3 are averaged over 50 realizations.




Figure 4: Estimation errors and true positive rates versus false positive rates for various methods on Barabasi-Albert model with $\Theta^* \in \mathcal{M}_D^p$, where $p = 100$ and $n = 100$.




Figure 5: Estimation errors and true positive rates versus false positive rates for different methods on the stochastic block model with $\Theta^* \in \mathcal{M}_D^p$, where $p = 100$ and $n = 100$.

Figure 2 shows that the proposed method attains a low estimation error and a high true positive rate simultaneously while maintaining a small false positive rate. In comparison, GLASSO and CLIME exhibit the best performance in estimation error and true positive rate within a region of high false positive rates. A higher true positive rate coupled with a lower false positive rate signifies better performance in identifying underlying graph edges. Our method surpasses GSCAD in achieving a smaller estimation error and a higher true positive rate, as we incorporate more prior knowledge of the constraint set \mathcal{M}^p . Although GGL also employs \mathcal{M}^p , our method achieves a better performance because it refines the GGL estimate in subsequent stages. The comparison with SLTP regarding estimation errors is excluded. It is noteworthy that GGL, SLTP, and our method all possess an upper bound on false positive rates.

Figure 3 displays the estimation errors and true positive rates of different methods as a function of false positive rates on a line graph. Our proposed method consistently achieves smaller estimation errors and higher true positive rates compared to state-of-the-art methods in the region of low false positive rates. This region is of primary interest, as an ideal estimate should exhibit low estimation error while maintaining high true positive rate and low false positive rate. Consequently, our method is capable of simultaneously achieving a reduced precision matrix estimation error and enhanced graph edge selection accuracy.




Figure 6: Estimation errors and F-scores as a function of varying regularization parameter values on a line graph.

Figure 4 presents comparisons on the Barabasi-Albert model with the underlying precision matrix $\Theta^* \in \mathcal{M}_D^p$. Our method consistently achieves lower estimation errors and higher true positive rates compared to state-of-the-art methods, highlighting its superior performance in both precision matrix estimation and graph edge identification. SLTP attains a higher true positive rate than all other methods when the false positive rate is below 0.05; however, its highest true positive rate is still lower than that of our method. Similar observations can be made for the stochastic block model, as displayed in Figure 5. The curves in Figures 4 and 5 are averaged over 50 realizations.

Figure 6 illustrates the robustness of GSCAD and our method regarding regularization parameter choices in terms of estimation error and F-score. Experiments are conducted on a line graph with $p = 100$ and $n = 40$. Figure 6 reveals that both estimation errors and F-scores of our method are more stable across varying regularization parameter values than those of GSCAD. Notably, the estimation error of GSCAD grows rapidly as the regularization parameter decreases, while that of our method remains steady. Thus, our method exhibits less sensitivity to the regularization parameter selection than GSCAD.

The key difference between our method and GSCAD is the former imposes an additional M -matrix constraint, while the latter does not. Thus, our method searches for a solution within a smaller space than GSCAD, excluding some stationary points that yield large estimation errors. This is why our method is more robust to regularization parameter selection. In fact, an extremely small regularization parameter can render the optimization problem for GSCAD ill-defined, meaning the minimum is not finite.

Table 1 displays the estimation errors for various methods under varying sample size ratios n/p , maintaining a low false positive rate ($\text{FPR} < 0.05$). The proposed method consistently achieves lower estimation errors compared to state-of-the-art methods across a range of sample size ratios. Notably, the superior performance of our method becomes more evident in cases with smaller sample size ratios, which are indicative of high-dimensional scenarios.

n/p	GLASSO	CLIME	GSCAD	GGL	Proposed
1	0.421	0.472	0.328	0.403	0.267
2	0.366	0.423	0.212	0.344	0.171
5	0.358	0.242	0.110	0.297	0.100

Table 1: Estimation errors for various methods across varying sample size ratios, with a consistently low false positive rate ($\text{FPR} < 0.05$).

5.2. Financial Time-series Data

In this subsection, we present comparisons of various methods on the financial time-series data, where the MTP_2 assumption has been well-justified due to the market factor causing positive dependence among stocks (Agrawal et al., 2020; Cardoso et al., 2020). We collect data from 204 stocks composing the S&P 500 index for the period between January 5th, 2004, and December 30th, 2006, resulting in 503 observations per stock (*i.e.*, $p = 204$ and $n = 503$). We construct the log-returns data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ as follows:

$$X_{i,j} = \log P_{i,j} - \log P_{i-1,j}, \quad (20)$$

where $P_{i,j}$ represents the closing price of the j -th stock on the i -th day. We then use the sample correlation matrix as input for each method. The 204 stocks are categorized into five sectors according to the Global Industry Classification Standard (GICS) system: Consumer Staples, Consumer Discretionary, Industrials, Energy, and Information Technology. We compare the proposed method with GLASSO, GSCAD, GGL, and GOLAZO on this financial time-series data.

In contrast to synthetic data, we cannot evaluate estimation error for each method, as the underlying precision matrix is unavailable. However, we expect stocks within the same sector to be somewhat interconnected. Consequently, we use the *modularity* (Newman, 2006) to assess the estimation performance. Define a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, where \mathcal{V} is the vertex set, \mathcal{E} is the edge set, and \mathbf{A} is the adjacency matrix. The *modularity* of graph \mathcal{G} is given by:

$$Q := \frac{1}{2|\mathcal{E}|} \sum_{i,j \in |\mathcal{V}|} \left(A_{ij} - \frac{d_i d_j}{2|\mathcal{E}|} \right) \delta(c_i, c_j),$$

where d_i represents the degree of the i -th node, c_i denotes the type of the i -th node, and $\delta(\cdot, \cdot)$ is the Kronecker delta function with $\delta(a, b) = 1$ if $a = b$ and 0 otherwise. A stock graph with high *modularity* exhibits dense connections between stocks within the same sector and sparse connections between stocks from different sectors. Therefore, a higher *modularity* value indicates a better representation of the actual stock network.




Figure 7: Stock graphs characterized by the precision matrices estimated via (a) GLASSO, (b) GSCAD, (c) GGL, (d) GOLAZO, and (e) the proposed method. The *modularity* values for GLASSO, GSCAD, GGL, GOLAZO, and our method are 0.3534, 0.3433, 0.3530, 0.3578, and 0.4978, respectively. We fine-tune the regularization parameter for each method based on the *modularity* value while allowing only a few isolated nodes.

Table 2: *Modularity* values of S&P 500 stock graphs learned via various methods using a rolling-window approach, covering the period from January 3rd, 2014, to October 31st, 2017.

Methods	Jan. 2016	Apr. 2016	Jul. 2016	Oct. 2016	Jan. 2017	Apr. 2017	Jul. 2017	Oct. 2017
GLASSO	0.4038	0.4072	0.4115	0.3956	0.3982	0.3910	0.3911	0.4184
GSCAD	0.3880	0.3850	0.3935	0.3803	0.3848	0.3771	0.3822	0.3860
GGL	0.4088	0.4118	0.4193	0.3991	0.3927	0.3896	0.3845	0.4156
GOLAZO	0.4068	0.4098	0.4173	0.4007	0.4029	0.3945	0.3953	0.4199
Proposed	0.5441	0.5289	0.5298	0.5139	0.5138	0.5118	0.5219	0.5019

Figure 7 shows stock graphs based on precision matrices estimated using various methods. It is observed that the performance of the proposed method stands out, as most connections are between nodes within the same sector, and only a few connections (gray-colored edges) are between nodes from distinct sectors, which are often spurious from a practical perspective. We further compare the *modularity* value for each method. The *modularity* values for GLASSO, GSCAD, GGL, GOLAZO, and our method are 0.3534, 0.3433, 0.3530, 0.3578, and 0.4978, respectively, indicating that our method outperforms state-of-the-art methods in representing the stock network.

In the subsequent experiment, we collect data from 82 stocks in the S&P 500 index, representing three sectors: Communication Services, Utilities, and Real Estate. The data spans a timeframe from January 3rd, 2014, to October 31st, 2017, resulting in 945 observations.

Table 2 presents *modularity* values of graphs learned using a rolling-window approach. We select a window length of 504 days (roughly two years of stock market days) and shift it by 63 days (approximately three months of stock market days). Higher *modularity* signifies a more accurate representation of the actual stock network. Table 2 demonstrates that the proposed method consistently attains the highest *modularity* values throughout the entire period.

6. Conclusions

In this paper, we have proposed an adaptive multiple-stage method to estimate both M -matrices and diagonally dominant M -matrices as precision matrices in MTP_2 Gaussian graphical models. The proposed method starts with an initial estimate obtained via an ℓ_1 -regularized maximum likelihood estimation and refines it in subsequent stages by solving a sequence of adaptive ℓ_1 -regularized problems. We have provided theoretical analysis for estimation error. Experiments on both synthetic and real-world data have shown that our method outperforms state-of-the-art methods in estimating precision matrices and identifying graph edges. We have also demonstrated that incorporating the M -matrix constraint improves robustness regarding regularization parameter selection.

Acknowledgements

This work was supported by the Hong Kong Research Grants Council GRF 16207820, 16310620, and 16306821, the Hong Kong Innovation and Technology Fund (ITF) MHP/009/20, and the Project of Hetao Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone under Grant HZQB-KCZYB-2020083. We would also like to thank the anonymous ICML reviewers for their valuable feedback on the manuscript.

References

- Agrawal, R., Roy, U., and Uhler, C. Covariance Matrix Estimation under Total Positivity for Portfolio Selection. *Journal of Financial Econometrics*, 20(2):367–389, 2020.
- Banerjee, O., Ghaoui, L. E., and d’Aspremont, A. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516, 2008.
- Barabási, A.-L. and Albert, R. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- Beck, A. *Introduction to nonlinear optimization: Theory, algorithms, and applications with MATLAB*. SIAM, 2014.
- Biyikoglu, T., Leydold, J., and Stadler, P. F. *Laplacian eigenvectors of graphs: Perron-Frobenius and Faber-Krahn type theorems*. Springer, 2007.
- Boyle, J. P. and Dykstra, R. L. A method for finding projections onto the intersection of convex sets in Hilbert spaces. In *Advances in order restricted statistical inference*, pp. 28–47. Springer, 1986.
- Cai, J.-F., Cardoso, J. V. D. M., Palomar, D. P., and Ying, J. Fast projected newton-like method for precision matrix estimation under total positivity. *arXiv preprint arXiv:2112.01939*, 2021.
- Cai, T., Liu, W., and Luo, X. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106(494):594–607, 2011.
- Cardoso, J. V. d. M., Ying, J., and Palomar, D. P. Algorithms for learning graphs in financial markets. *arXiv preprint arXiv:2012.15410*, 2020.
- Cardoso, J. V. D. M., Ying, J., and Palomar, D. P. Graphical models in heavy-tailed markets. In *Advances in Neural Information Processing Systems*, 2021.
- Cardoso, J. V. D. M., Ying, J., and Palomar, D. P. Learning bipartite graphs: Heavy tails and multiple components. In *Advances in Neural Information Processing Systems*, 2022.
- Condat, L. Fast projection onto the simplex and the l_1 ball. *Mathematical Programming*, 158(1):575–585, 2016.
- d’Aspremont, A., Banerjee, O., and El Ghaoui, L. First-order methods for sparse covariance selection. *SIAM Journal on Matrix Analysis and Applications*, 30(1):56–66, 2008.
- Davies, E. B., Leydold, J., and Stadler, P. F. Discrete nodal domain theorems. *Linear Algebra and its Applications*, 336(1):51–60, 2001.
- Deng, Z. and So, A. M.-C. A fast proximal point algorithm for generalized graph Laplacian learning. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5425–5429, 2020.
- Dong, X., Thanou, D., Rabbat, M., and Frossard, P. Learning graphs from data: A signal representation perspective. *IEEE Signal Processing Magazine*, 36(3):44–63, 2019.
- Drazin, M. P. and Haynsworth, E. V. Criteria for the reality of matrix eigenvalues. *Mathematische Zeitschrift*, 78(1):449–452, 1962.
- Egilmez, H. E., Pavez, E., and Ortega, A. Graph learning from data under Laplacian and structural constraints. *IEEE Journal of Selected Topics in Signal Processing*, 11(6):825–841, 2017.
- Fallat, S., Lauritzen, S., Sadeghi, K., Uhler, C., Wermuth, N., and Zwiernik, P. Total positivity in Markov structures. *The Annals of Statistics*, 45(3):1152–1184, 2017.
- Fan, J. and Li, R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- Fan, J., Xue, L., and Zou, H. Strong oracle optimality of folded concave penalized estimation. *Annals of statistics*, 42(3):819, 2014.
- Fan, J., Liu, H., Sun, Q., and Zhang, T. I-lamm for sparse learning: Simultaneous control of algorithmic complexity and statistical error. *Annals of statistics*, 46(2):814, 2018.
- Friedman, J., Hastie, T., and Tibshirani, R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.
- Hsieh, C.-J., Sustik, M. A., Dhillon, I. S., and Ravikumar, P. QUIC: quadratic approximation for sparse inverse covariance estimation. *Journal of Machine Learning Research*, 15(1):2911–2947, 2014.
- Karlin, S. and Rinott, Y. M-matrices as covariance matrices of multinormal distributions. *Linear algebra and its applications*, 52:419–438, 1983.

- Kumar, S., Ying, J., Cardoso, J. V. d. M., and Palomar, D. P. A unified framework for structured graph learning via spectral constraints. *Journal of Machine Learning Research*, 21(22):1–60, 2020.
- Lake, B. and Tenenbaum, J. Discovering structure by learning sparse graphs. In *Proceedings of the Annual Cognitive Science Conference*, pp. 778–783, 2010.
- Lam, C. and Fan, J. Sparsistency and rates of convergence in large covariance matrix estimation. *The Annals of Statistics*, 37(6B):4254, 2009.
- Lauritzen, S. and Zwiernik, P. Locally associated graphical models and mixed convex exponential families. *The Annals of Statistics*, 50(5):3009–3038, 2022.
- Lauritzen, S., Uhler, C., and Zwiernik, P. Maximum likelihood estimation in Gaussian models under total positivity. *The Annals of Statistics*, 47(4):1835–1863, 2019.
- Lauritzen, S., Uhler, C., and Zwiernik, P. Total positivity in exponential families with application to binary variables. *The Annals of Statistics*, 49(3):1436–1459, 2021.
- Liu, H. and Wang, L. Tiger: A tuning-insensitive approach for optimally estimating Gaussian graphical models. *Electronic Journal of Statistics*, 11(1):241–294, 2017.
- Loh, P.-L. and Wainwright, M. J. Regularized M-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *Journal of Machine Learning Research*, 16(1):559–616, 2015.
- Loh, P.-L. and Wainwright, M. J. Support recovery without incoherence: A case for nonconvex regularization. *Annals of Statistics*, 45(6):2455–2482, 2017.
- Newman, M. E. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.
- Ortega, A., Frossard, P., Kovačević, J., Moura, J. M., and Vandergheynst, P. Graph signal processing: Overview, challenges, and applications. *Proceedings of the IEEE*, 106(5):808–828, 2018.
- Palomar, D. P. and Fongollosa, J. R. Practical algorithms for a family of waterfilling solutions. *IEEE transactions on Signal Processing*, 53(2):686–695, 2005.
- Pavez, E. and Ortega, A. Generalized Laplacian precision matrix estimation for graph signal processing. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 6350–6354, 2016.
- Pavez, E., Egilmez, H. E., and Ortega, A. Learning graphs with monotone topology properties and multiple connected components. *IEEE Transactions on Signal Processing*, 66(9):2399–2413, 2018.
- Ravikumar, P., Wainwright, M. J., Raskutti, G., and Yu, B. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- Robeva, E., Sturmfels, B., Tran, N., and Uhler, C. Maximum likelihood estimation for totally positive log-concave densities. *Scandinavian Journal of Statistics*, 48(3):817–844, 2021.
- Rossell, D. and Zwiernik, P. Dependence in elliptical partial correlation graphs. *Electronic Journal of Statistics*, 15(2):4236–4263, 2021.
- Röttger, F., Engelke, S., and Zwiernik, P. Total positivity in multivariate extremes. *arXiv preprint arXiv:2112.14727*, 2021.
- Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013.
- Slawski, M. and Hein, M. Estimation of positive definite M-matrices and structure learning for attractive Gaussian Markov random fields. *Linear Algebra and its Applications*, 473:145–179, 2015.
- Soloff, J. A., Guntuboyina, A., and Jordan, M. I. Covariance estimation with nonnegative partial correlations. *arXiv preprint arXiv:2007.15252*, 2020.
- Sun, Q., Tan, K. M., Liu, H., and Zhang, T. Graphical nonconvex optimization via an adaptive convex relaxation. In *International Conference on Machine Learning*, pp. 4810–4817, 2018.
- Truell, M., Hütter, J.-C., Squires, C., Zwiernik, P., and Uhler, C. Maximum likelihood estimation for Brownian motion tree models based on one sample. *arXiv preprint arXiv:2112.00816*, 2021.
- Wang, L., Ren, X., and Gu, Q. Precision matrix estimation in high dimensional Gaussian graphical models with faster rates. In *Artificial Intelligence and Statistics*, pp. 177–185, 2016.
- Wang, Y., Roy, U., and Uhler, C. Learning high-dimensional Gaussian graphical models under total positivity without adjustment of tuning parameters. In *International*

- Conference on Artificial Intelligence and Statistics*, pp. 2698–2708, 2020.
- Xiao, L. and Zhang, T. A proximal-gradient homotopy method for the sparse least-squares problem. *SIAM Journal on Optimization*, 23(2):1062–1091, 2013.
- Ying, J., Cardoso, J. V. d. M., and Palomar, D. P. Does the ℓ_1 -norm learn a sparse graph under Laplacian constrained graphical models? *arXiv preprint arXiv:2006.14925*, 2020a.
- Ying, J., de Miranda Cardoso, J. V., and Palomar, D. P. Nonconvex sparse graph learning under Laplacian constrained graphical model. In *Advances in Neural Information Processing Systems*, 2020b.
- Ying, J., Cardoso, J. V. D. M., and Palomar, D. Minimax estimation of Laplacian constrained precision matrices. In *International Conference on Artificial Intelligence and Statistics*, pp. 3736–3744, 2021a.
- Ying, J., Cardoso, J. V. D. M., and Palomar, D. P. A fast algorithm for graph learning under attractive Gaussian Markov random fields. In *Asilomar Conference on Signals, Systems, and Computers*, pp. 1520–1524, 2021b.
- Yuan, M. High dimensional inverse covariance matrix estimation via linear programming. *The Journal of Machine Learning Research*, 11:2261–2286, 2010.
- Zhang, C.-H. Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics*, 38(2): 894–942, 2010a.
- Zhang, T. Analysis of multi-stage convex relaxation for sparse regularization. *Journal of Machine Learning Research*, 11:1081–1107, 2010b.
- Zou, H. and Li, R. One-step sparse estimates in nonconcave penalized likelihood models. *Annals of Statistics*, 36(4): 1509, 2008.

A. Experimental Settings and Additional Results

We first provide a more in-depth description of experimental settings in Appendix A.1. Next, we present additional experimental results in Appendix A.2.

A.1. Experimental Settings

To generate synthetic data, we first construct a graph structure randomly, then associate each edge with a weight that is uniformly sampled from $U(2, 5)$. Then we obtain a weighted adjacency matrix \mathbf{A} , where A_{ij} denotes the graph weight between node i and node j .

To generate an underlying precision matrix $\Theta^* \in \mathcal{M}^p$, we adopt the procedures used in (Slawski & Hein, 2015; Wang et al., 2020). We first set

$$\Theta = \delta \mathbf{I} - \mathbf{A}, \quad \text{with } \delta = 1.05 \lambda_{\max}(\mathbf{A}), \quad (21)$$

where $\lambda_{\max}(\mathbf{A})$ denotes the largest eigenvalue of \mathbf{A} . We set $\Theta^* = \mathbf{E}\Theta\mathbf{E}$, where \mathbf{E} is a diagonal matrix chosen such that the covariance matrix $(\Theta^*)^{-1}$ has unit diagonal elements. The matrix Θ^* as generated above must be an M -matrix, but usually not diagonally dominant. Given a precision matrix Θ^* , we generate n independent and identically distributed observations $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)} \sim \mathcal{N}(\mathbf{0}, (\Theta^*)^{-1})$. The sample covariance matrix is constructed by $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}^{(i)} (\mathbf{x}^{(i)})^\top$.

We also conduct experiments for estimating diagonally dominant M -matrices. To generate an underlying precision matrix $\Theta^* \in \mathcal{M}_D^p$, we construct $\Theta^* = \mathbf{D} - \mathbf{A} + \mathbf{V}$, where \mathbf{D} and \mathbf{V} are diagonal matrices with each $D_{ii} = \sum_{j=1}^p A_{ij}$ and each V_{ii} uniformly sampled from $U(0, 1)$.

To evaluate the estimation performance across different methods, we compute the estimation error as $\|\hat{\Theta} - \Theta^*\|_{\text{F}} / \|\Theta^*\|_{\text{F}}$, where $\hat{\Theta}$ and Θ^* denote the estimated and underlying precision matrices, respectively. The true positive rate (TPR) and false positive rate (FPR) are defined as

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad \text{and} \quad \text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}},$$

where TP, FP, TN and FN denote the number of true positives, false positives, true negatives, and false negatives, respectively. The true positive rate denotes the proportion of correctly identified edges in the graph, and the false positive rate denotes the proportion of incorrectly identified edges in the graph. An estimate with a higher true positive rate and a lower false positive rate indicates a better performance of identifying the underlying graph edges.

The F-score of a graph is defined as

$$\text{F-score} := \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}. \quad (22)$$

The F-score takes values in $[0, 1]$, and 1 indicates perfect structure recovery. Our method uses the weight-updating function corresponding to the SCAD penalty.

A.2. Additional Experimental Results

We perform experiments to demonstrate the empirical convergence of the proposed gradient projection algorithm, as illustrated in Figure 8. We observe that the number of iterations required for convergence descends rapidly with increasing stage k , and empirically the algorithm enjoys a linear convergence rate.

When estimating diagonally dominant M -matrices, our proposed algorithm involves double loops of iterations to solve the subproblem at each stage, potentially leading to computational inefficiency due to the use of Dykstra's algorithm for handling the diagonally dominant M -matrix (DDM-matrix) constraint. Table 3 displays the running time of each algorithm for the case of estimating DDM-matrices. We observe that our algorithm is less computationally efficient in this scenario, primarily due to managing multiple constraints to ensure a DDM-matrix solution. To enhance our algorithm's efficiency, we could compute an approximate solution when addressing subproblems in earlier stages, a technique employed in previous studies (Xiao & Zhang, 2013; Fan et al., 2018). Investigating the development of numerical algorithms for handling diagonally dominant M -matrix constraints more efficiently would be a valuable future research direction.




Figure 8: An empirical convergence result of the gradient projection algorithm at different stages. At the k -th stage with $k = 1, 2, \dots, 8$, we plot the relative error $\|\Theta_t - \widehat{\Theta}^{(k)}\|_F / \|\widehat{\Theta}^{(k)}\|_F$ as a function of the iteration t , where $\{\Theta_t\}_{t \geq 1}$ is established by the gradient projection algorithm, and $\widehat{\Theta}^{(k)}$ is the optimal solution at the k -th stage.

Table 3: Comparison of running times for methods in the case of estimating diagonally dominant M -matrices.

Methods	SLTP	GLASSO	CLIME	GSCAD	DDGL	Proposed
Running time (s)	9.48	0.05	4.06	0.14	1.17	7.85

It is essential to emphasize that our proposed algorithm is the sole method effectively addressing the Gaussian maximum likelihood estimation problem under the DDM-matrix constraints. In contrast, GLASSO, CLIME, and GSCAD tackle unconstrained problems and cannot guarantee a DDM-matrix solution. Such a solution provides desirable spectral properties for performing the graph Fourier transform and is crucial for analyzing positive dependence in multivariate Pareto distributions. Furthermore, the popular Block Coordinate Descent (BCD) algorithm cannot ensure convergence to the global minimizer, as demonstrated in Table 4.

We compare the optimality gap for solutions obtained by the BCD algorithm and our proposed algorithm. The optimality gap is evaluated using two approaches:

$$\Delta_{\Theta} = \|\widehat{\Theta} - \Theta^*\|_F / \|\Theta^*\|_F, \quad \text{and} \quad \Delta_f = \hat{f} - f^*,$$

where $\widehat{\Theta}$ is the output from BCD or our algorithm, Θ^* is the global minimizer obtained through CVX, and \hat{f} and f^* represent the objective function values at $\widehat{\Theta}$ and Θ^* , respectively. Experiments are conducted on the Barabasi-Albert model with 50 nodes ($p = 50$) and the number of samples n ranging from 50 to 5000. For both algorithms, the stopping criterion is met when successive iterations satisfy the condition $\|\Theta_{t+1} - \Theta_t\|_F / \|\Theta_t\|_F < 10^{-12}$.

Table 4: Optimality gap comparison for algorithms in the case of estimating diagonally dominant M -matrices.

Algorithms	BCD		Proposed	
$n = 50$	$\Delta_{\Theta} = 0.5295,$	$\Delta_f = 0.0738$	$\Delta_{\Theta} = 2.81 \times 10^{-6},$	$\Delta_f = 1.49 \times 10^{-12}$
$n = 500$	$\Delta_{\Theta} = 0.0765,$	$\Delta_f = 0.0022$	$\Delta_{\Theta} = 1.09 \times 10^{-5},$	$\Delta_f = 2.68 \times 10^{-12}$
$n = 5000$	$\Delta_{\Theta} = 0.0166,$	$\Delta_f = 0.0004$	$\Delta_{\Theta} = 1.39 \times 10^{-5},$	$\Delta_f = 2.77 \times 10^{-12}$

Table 4 reveals that the popular BCD algorithm exhibits a substantial optimality gap, indicating its failure to converge to the global minimizer. In contrast, our proposed algorithm effectively converges to the minimizer. We note that the BCD algorithm can perform well under the M -matrix constraint without diagonal dominance.

B. Proofs

We first present two lemmas in Appendix B.1. Then we provide the proofs of Theorem 3.1 in Appendix B.2, Proposition 3.3 in Appendix B.3, and Theorem 4.4 and Corollary 4.5 in Appendix B.4.

We begin by introducing the notations used in proofs. Define

$$L_\alpha(\Theta) := \{(i, j) \mid |\Theta_{ij}| \geq \alpha, i \neq j\}, \quad (23)$$

where $\alpha = c_0\lambda$ with c_0 defined in Assumption 4.1. Then we define the set $T_\alpha(\Theta)$ as follows:

$$T_\alpha(\Theta) := L_\alpha(\Theta) \cup S^*, \quad (24)$$

where S^* is defined in (19). Define

$$G_{\lambda, \hat{\Sigma}}(\tilde{\Theta}) := \arg \min_{\Theta \in \mathcal{M}^p} -\log \det(\Theta) + \text{tr}(\Theta \hat{\Sigma}) + \sum_{i \neq j} p_\lambda(|\tilde{\Theta}_{ij}|) |\Theta_{ij}|.$$

For ease of presentation, by a slight abuse of notation, we write $p_\lambda(|\Theta_{S^*}|)$ for $(p_\lambda(|\Theta_{ij}|))_{(i,j) \in S^*}$.

B.1. Lemmas

Lemma B.1. *Let $g(\Theta) = -\log \det(\Theta)$. Define a local region of Θ^* by*

$$\mathcal{B}(\Theta^*, r) = \{\Theta \in \mathcal{M}^p \mid \|\Theta - \Theta^*\|_F \leq r\}.$$

Then for any $\Theta_1, \Theta_2 \in \mathcal{B}(\Theta^; \lambda_{\max}(\Theta^*))$, we have*

$$\langle \nabla g(\Theta_1) - \nabla g(\Theta_2), \Theta_1 - \Theta_2 \rangle \geq \frac{1}{4} \lambda_{\max}^{-2}(\Theta^*) \|\Theta_1 - \Theta_2\|_F^2.$$

Proof. For any $\Theta_1, \Theta_2 \in \mathcal{B}(\Theta^*; \lambda_{\max}(\Theta^*))$, by Mean Value Theorem, one obtains

$$g(\Theta_2) = g(\Theta_1) + \langle \nabla g(\Theta_1), \Theta_2 - \Theta_1 \rangle + \frac{1}{2} \text{vec}(\Theta_2 - \Theta_1)^\top \nabla^2 g(\Theta_t) \text{vec}(\Theta_2 - \Theta_1),$$

where $\Theta_t = t\Theta_2 + (1-t)\Theta_1$ with $t \in [0, 1]$. Note that $\Theta_t \in \mathcal{B}(\Theta^*; r)$. One has

$$\lambda_{\min}(\nabla^2 g(\Theta_t)) = \lambda_{\min}((\Theta_t \otimes \Theta_t)^{-1}) = \lambda_{\max}^{-2}(\Theta_t),$$

where the second equality follows from the property that the eigenvalues of $\mathbf{A} \otimes \mathbf{B}$ are $\lambda_i \mu_j$, where $\lambda_1, \dots, \lambda_p$ and μ_1, \dots, μ_p are the eigenvalues of $\mathbf{A} \in \mathbb{R}^{p \times p}$ and $\mathbf{B} \in \mathbb{R}^{p \times p}$, respectively. Following from the Weyl's inequality, one obtains

$$\lambda_{\max}(\Theta_t) \leq \lambda_{\max}(\Theta_t - \Theta^*) + \lambda_{\max}(\Theta^*) \leq r + \lambda_{\max}(\Theta^*).$$

Therefore, We have

$$g(\Theta_2) \geq g(\Theta_1) + \langle \nabla g(\Theta_1), \Theta_2 - \Theta_1 \rangle + \frac{1}{2} (r + \lambda_{\max}(\Theta^*))^{-2} \|\Theta_2 - \Theta_1\|_F^2, \quad (25)$$

and

$$g(\Theta_1) \geq g(\Theta_2) + \langle \nabla g(\Theta_2), \Theta_1 - \Theta_2 \rangle + \frac{1}{2} (r + \lambda_{\max}(\Theta^*))^{-2} \|\Theta_2 - \Theta_1\|_F^2. \quad (26)$$

Setting $r = \lambda_{\max}(\Theta^*)$ and combining (25) and (26), we obtain

$$\langle \nabla g(\Theta_1) - \nabla g(\Theta_2), \Theta_1 - \Theta_2 \rangle \geq \frac{1}{4} \lambda_{\max}^{-2}(\Theta^*) \|\Theta_1 - \Theta_2\|_F^2.$$

completing the proof. \square

Lemma B.2. Suppose the event $\|\Sigma^* - \widehat{\Sigma}\|_{\max} \leq \frac{\sqrt{2}}{2}\lambda$ holds and the sample size satisfies $n \geq 8c_0c_1\tau s \log p$, and take $\lambda = \sqrt{c_1\tau \log p/n}$. Under Assumptions 4.1 and 4.2, if $|T_\alpha(\widehat{\Theta})| \leq 2s - p$ holds for some $\widehat{\Theta}$, then

$$\|G_{\lambda, \widehat{\Sigma}}(\widehat{\Theta}) - \Theta^*\|_{\mathbb{F}} \leq \frac{c_0}{2} (\|p_\lambda(|\widehat{\Theta}_{S^*}|)\| + \|(\Sigma^* - \widehat{\Sigma})_{T_\alpha(\widehat{\Theta}) \cup I_p}\|),$$

and

$$|T_\alpha(G_{\lambda, \widehat{\Sigma}}(\widehat{\Theta}))| \leq 2s - p. \quad (27)$$

Proof. Drawing inspiration from (Sun et al., 2018), our proof employs techniques that establish upper-bounded estimation error in regularized maximum likelihood estimation problems using Karush–Kuhn–Tucker (KKT) conditions. Notably, our formulation incorporates multiple constraints, contrasting the unconstrained problem in (Sun et al., 2018), which leads to distinct KKT conditions and error-bounding techniques.

Due to the sign constraint, $G_{\lambda, \widehat{\Sigma}}(\widehat{\Theta})$ can be expressed as the minimizer of the following problem:

$$\underset{\Theta \in \mathcal{M}^p}{\text{minimize}} -\log \det(\Theta) + \text{tr}(\Theta \widehat{\Sigma}) - \sum_{i \neq j} W_{ij} \Theta_{ij}, \quad (28)$$

where $W_{ij} = p_\lambda(|\widehat{\Theta}_{ij}|)$ if $i \neq j$, and zero otherwise. Define a local region

$$\mathcal{B}(\Theta^*, \lambda_{\max}(\Theta^*)) = \{\Theta \in \mathcal{M}^p \mid \|\Theta - \Theta^*\|_{\mathbb{F}} \leq \lambda_{\max}(\Theta^*)\}.$$

For ease of presentation, we denote $G_{\lambda, \widehat{\Sigma}}(\widehat{\Theta})$ by $\widehat{\Theta}$. We first prove that $\widehat{\Theta} \in \mathcal{B}(\Theta^*, \lambda_{\max}(\Theta^*))$ under event \mathcal{J} defined in (53). We construct an intermediate estimate:

$$\Theta_t := \Theta^* + t(\widehat{\Theta} - \Theta^*), \quad (29)$$

where t is taken such that $\|\Theta_t - \Theta^*\|_{\mathbb{F}} = \lambda_{\max}(\Theta^*)$ if $\|\widehat{\Theta} - \Theta^*\|_{\mathbb{F}} > \lambda_{\max}(\Theta^*)$, and $t = 1$ otherwise. Hence $\|\Theta_t - \Theta^*\|_{\mathbb{F}} \leq \lambda_{\max}(\Theta^*)$ holds for any $t \in [0, 1]$. Then we apply Lemma B.1 with $\Theta_1 = \Theta_t$, $\Theta_2 = \Theta^*$, and get

$$\frac{c_0}{2} t \langle (\Theta^*)^{-1} - \Theta_t^{-1}, \widehat{\Theta} - \Theta^* \rangle \geq \|\Theta_t - \Theta^*\|_{\mathbb{F}}^2. \quad (30)$$

Construct a function $q(t) := -\log \det(\Theta^* + t(\widehat{\Theta} - \Theta^*)) + t \langle (\Theta^*)^{-1}, \widehat{\Theta} - \Theta^* \rangle$ and $t \in [0, 1]$. One has

$$q'(t) = \langle (\Theta^*)^{-1} - \Theta_t^{-1}, \widehat{\Theta} - \Theta^* \rangle, \quad (31)$$

Let $\mathbf{A} = \Theta_t^{-1}$, and $\mathbf{B} = \widehat{\Theta} - \Theta^*$. Then one also has

$$q''(t) = \text{tr}(\mathbf{A}\mathbf{B}\mathbf{A}\mathbf{B}),$$

Next, we show that $q''(t) \geq 0$. Let $\mathbf{C} = \mathbf{A}\mathbf{B}$. Following from Theorem 1 in (Drazin & Haynsworth, 1962), all eigenvalues of a matrix \mathbf{X} are real if there exists a symmetric and positive definite matrix \mathbf{Y} such that $\mathbf{X}\mathbf{Y}$ are symmetric. Since $\mathbf{C}\mathbf{A}$ is symmetric and \mathbf{A} symmetric and positive definite, all eigenvalues of \mathbf{C} are real. Suppose $\lambda_1, \dots, \lambda_p$ are the eigenvalues of \mathbf{C} . Then one obtains $q''(t) = \sum_{i=1}^p \lambda_i^2 \geq 0$, implying that $q'(t)$ is monotonically non-decreasing. Then one obtains

$$\langle (\Theta^*)^{-1} - \Theta_t^{-1}, \widehat{\Theta} - \Theta^* \rangle \geq q'(t) \geq \frac{2}{c_0 t} \|\Theta_t - \Theta^*\|_{\mathbb{F}}^2. \quad (32)$$

where the first inequality follows from $q'(1) \geq q'(t)$ and $t \leq 1$, and the second inequality follows from (30).

Following from (32), we have

$$\|\Theta_t - \Theta^*\|_{\mathbb{F}}^2 \leq \frac{c_0}{2} t \langle (\Theta^*)^{-1} - \Theta_t^{-1}, \widehat{\Theta} - \Theta^* \rangle. \quad (33)$$

We present the Lagrangian of the optimization (28) as follows:

$$L(\Theta, \Gamma) = -\log \det(\Theta) + \text{tr}(\Theta \widehat{\Sigma}) - \sum_{i \neq j} W_{ij} \Theta_{ij} + \langle \Gamma, \Theta \rangle,$$

where Γ is a KKT multiplier with $\Gamma_{ii} = 0$ for $i \in [p]$. Let $(\hat{\Theta}, \hat{\Gamma})$ be the primal and dual optimal point, which must satisfy the KKT conditions as follows:

$$-\hat{\Theta}^{-1} + \hat{\Sigma} - \mathbf{W} + \hat{\Gamma} = \mathbf{0}; \quad (34)$$

$$\hat{\Theta}_{ij}\hat{\Gamma}_{ij} = 0, \hat{\Theta}_{ij} \leq 0, \hat{\Gamma}_{ij} \geq 0, \forall i \neq j; \quad (35)$$

$$\hat{\Gamma}_{ii} = 0, \text{ for } i \in [p]; \quad (36)$$

According to (34), one has

$$\langle -\hat{\Theta}^{-1} + \hat{\Sigma}, \hat{\Theta} - \Theta^* \rangle = \langle \mathbf{W} - \hat{\Gamma}, \hat{\Theta} - \Theta^* \rangle. \quad (37)$$

Plugging (37) into (33) yields

$$\|\Theta_t - \Theta^*\|_F^2 = \frac{c_0}{2} t \left(\underbrace{\langle \hat{\Gamma}, \Theta^* - \hat{\Theta} \rangle}_{\text{term I}} + \underbrace{\langle \mathbf{W}, \hat{\Theta} - \Theta^* \rangle}_{\text{term II}} + \underbrace{\langle \Sigma^* - \hat{\Sigma}, \hat{\Theta} - \Theta^* \rangle}_{\text{term III}} \right). \quad (38)$$

The term I in (38) can be bounded by

$$\langle \hat{\Gamma}, \Theta^* - \hat{\Theta} \rangle = \sum_{i \neq j} \hat{\Gamma}_{ij} \Theta_{ij}^* \leq 0.$$

where the equality follows from (35) and (36), and the inequality follows from $\hat{\Gamma}_{ij} \geq 0$ in (35) and the fact that $\Theta_{ij}^* \leq 0$ for any $i \neq j$.

To bound term II, we separate the support of \mathbf{W} off the diagonal into two parts, S^* and T^* which is defined as follows:

$$T^* = \{(i, j) \mid \Theta_{ij}^* = 0, i \neq j\}.$$

Let $I_p := \{(i, i) \mid i \in [p]\}$. Define

$$\bar{T}_\alpha(\tilde{\Theta}) = T_\alpha(\tilde{\Theta}) \cup I_p \quad \text{and} \quad \bar{T}_\alpha^c(\tilde{\Theta}) = T_\alpha^c(\tilde{\Theta}) \setminus I_p. \quad (39)$$

Since $|T_\alpha(\tilde{\Theta})| \leq 2s - p$, $|\bar{T}_\alpha(\tilde{\Theta})| \leq 2s$. Then one has

$$\begin{aligned} \langle \mathbf{W}, \hat{\Theta} - \Theta^* \rangle &= \langle \mathbf{W}_{S^*}, (\hat{\Theta} - \Theta^*)_{S^*} \rangle + \langle \mathbf{W}_{T^*}, \hat{\Theta} - \Theta^* \rangle \\ &\leq \|\mathbf{W}_{S^*}\| \|(\hat{\Theta} - \Theta^*)_{S^*}\| + \langle \mathbf{W}_{\bar{T}_\alpha^c(\tilde{\Theta})}, \hat{\Theta}_{\bar{T}_\alpha^c(\tilde{\Theta})} \rangle, \end{aligned}$$

where the inequality follows from the fact that $W_{ij} \geq 0$, $\hat{\Theta}_{ij} \leq 0$ for $i \neq j$, and $\bar{T}_\alpha^c(\tilde{\Theta}) \subseteq T^*$.

For term III, we separate the support of $(\Sigma^* - \hat{\Sigma})$ into parts, $\bar{T}_\alpha(\tilde{\Theta})$ and $\bar{T}_\alpha^c(\tilde{\Theta})$. Then one has

$$\langle \Sigma^* - \hat{\Sigma}, \hat{\Theta} - \Theta^* \rangle \leq \langle (\Sigma^* - \hat{\Sigma})_{\bar{T}_\alpha(\tilde{\Theta})}, \hat{\Theta}_{\bar{T}_\alpha(\tilde{\Theta})} \rangle + \|(\Sigma^* - \hat{\Sigma})_{\bar{T}_\alpha^c(\tilde{\Theta})}\| \|(\hat{\Theta} - \Theta^*)_{\bar{T}_\alpha^c(\tilde{\Theta})}\|.$$

We note that, for any $(i, j) \in \bar{T}_\alpha^c(\tilde{\Theta})$, $\hat{\Theta}_{ij} \leq 0$ and $W_{ij} \geq \frac{\sqrt{2}}{2}\lambda$, following from the definition of T_α in (24) and Assumption 4.1. Then under the event that $\|\Sigma^* - \hat{\Sigma}\|_{\max} \leq \frac{\sqrt{2}}{2}\lambda$, we have

$$\langle (\Sigma^* - \hat{\Sigma} + \mathbf{W})_{\bar{T}_\alpha^c(\tilde{\Theta})}, \hat{\Theta}_{\bar{T}_\alpha^c(\tilde{\Theta})} \rangle \leq 0. \quad (40)$$

By bounding term I, term II, and term III in (38), together with (40), we obtain

$$\|\Theta_t - \Theta^*\|_F^2 \leq \frac{c_0}{2} t \left(\|\mathbf{W}_{S^*}\| \|(\hat{\Theta} - \Theta^*)_{S^*}\| + \|(\Sigma^* - \hat{\Sigma})_{\bar{T}_\alpha(\tilde{\Theta})}\| \|(\hat{\Theta} - \Theta^*)_{\bar{T}_\alpha(\tilde{\Theta})}\| \right). \quad (41)$$

Through the definition of Θ_t in (29), one has $\|\Theta_t - \Theta^*\|_F = t \|\hat{\Theta} - \Theta^*\|_F$. Then (41) can be written as follows:

$$\|\Theta_t - \Theta^*\|_F \leq \frac{c_0}{2} \left(\|\mathbf{W}_{S^*}\| + \|(\Sigma^* - \hat{\Sigma})_{\bar{T}_\alpha(\tilde{\Theta})}\| \right). \quad (42)$$

Recall that $W_{ij} = p_\lambda(|\tilde{\Theta}_{ij}|) \leq \lambda$ for any $i \neq j$ according to Assumption 4.1, and $|S^*| \leq s - p$. Thus one has

$$\|\mathbf{W}_{S^*}\| \leq \sqrt{s}\lambda.$$

Under the event \mathcal{J} , one can bound

$$\|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{\bar{T}_\alpha(\tilde{\Theta})}\| \leq |\bar{T}_\alpha(\tilde{\Theta})|^{\frac{1}{2}} \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{\bar{T}_\alpha(\tilde{\Theta})}\|_{\max} \leq \sqrt{s}\lambda, \quad (43)$$

where the second inequality follows from $|\bar{T}_\alpha(\tilde{\Theta})| \leq 2s$.

Plugging (42) and (43) into (41) yields

$$\|\boldsymbol{\Theta}_t - \boldsymbol{\Theta}^*\|_{\text{F}} \leq c_0 \sqrt{s}\lambda \leq \lambda_{\max}(\boldsymbol{\Theta}^*),$$

which indicates that $t = 1$ in (29), i.e., $\boldsymbol{\Theta}_t = \widehat{\boldsymbol{\Theta}}$. The last inequality is established by plugging $\lambda = \sqrt{c_1 \tau \log p/n}$ with $n \geq 8c_0 c_1 \tau s \log p$. Therefore, we obtain

$$\|\widehat{\boldsymbol{\Theta}} - \boldsymbol{\Theta}^*\|_{\text{F}} \leq c_0 \sqrt{s}\lambda. \quad (44)$$

Putting $\mathbf{W}_{S^*} = p_\lambda(|\tilde{\Theta}_{S^*}|)$ and $\boldsymbol{\Theta}_t = G_{\lambda, \widehat{\boldsymbol{\Sigma}}}(\tilde{\Theta})$ into (42) gets

$$\|G_{\lambda, \widehat{\boldsymbol{\Sigma}}}(\tilde{\Theta}) - \boldsymbol{\Theta}^*\|_{\text{F}} \leq \frac{c_0}{2} (\|p_\lambda(|\tilde{\Theta}_{S^*}|)\| + \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{\bar{T}_\alpha(\tilde{\Theta})}\|).$$

To establish $|T_\alpha(\widehat{\boldsymbol{\Theta}})| \leq 2s - p$, we separate the set $T_\alpha(\widehat{\boldsymbol{\Theta}})$ into two parts, S^* and $L_\alpha(\widehat{\boldsymbol{\Theta}}) \setminus S^*$. For any $(i, j) \in L_\alpha(\widehat{\boldsymbol{\Theta}}) \setminus S^*$, one has $|\widehat{\Theta}_{ij}| \geq \alpha$, and further obtains

$$|L_\alpha(\widehat{\boldsymbol{\Theta}}) \setminus S^*| \leq \frac{\|\widehat{\boldsymbol{\Theta}}_{L_\alpha(\widehat{\boldsymbol{\Theta}}) \setminus S^*}\|^2}{\alpha^2} \leq \frac{\|\widehat{\boldsymbol{\Theta}} - \boldsymbol{\Theta}^*\|_{\text{F}}^2}{\alpha^2} \leq s,$$

where the last inequality follows from (44). Then we have

$$|T_\alpha(\widehat{\boldsymbol{\Theta}})| = |S^* \cup L_\alpha(\widehat{\boldsymbol{\Theta}}) \setminus S^*| = |S^*| + |L_\alpha(\widehat{\boldsymbol{\Theta}}) \setminus S^*| \leq 2s - p,$$

completing the proof. \square

B.2. Proof of Theorem 3.1

Proof. Our convergence analysis largely builds upon standard techniques for the projected gradient descent, as provided in (Beck, 2014). The main difference lies in the open constraint set in our formulation, which leads to an additional requirement regarding the positive definiteness of the matrix during stepsize selection, extending beyond the backtracking condition. Consequently, it is crucial to incorporate certain adaptations in comparison to the standard methods employed for analyzing the convergence of projected gradient descent.

Given $\boldsymbol{\Theta}^o \in \mathcal{S}$, define the lower level set of the objective f of Problem (5) as follows:

$$L_f := \{\boldsymbol{\Theta} \in \mathcal{S} \mid f(\boldsymbol{\Theta}) \leq f(\boldsymbol{\Theta}^o)\}. \quad (45)$$

Following from (Cai et al., 2021, Lemma 3.5), there exists $m > 0$ such that $\boldsymbol{\Theta} \succeq m\mathbf{I}$ holds for any $\boldsymbol{\Theta} \in L_f$. The gradient of f is Lipschitz continuous with parameter $L = m^{-2}$ over L_f , since the Hessian is $\boldsymbol{\Theta}^{-1} \otimes \boldsymbol{\Theta}^{-1}$.

Suppose that the initial point of the sequence $\boldsymbol{\Theta}_0 \in L_f$. To guarantee this, we may consider a $\boldsymbol{\Theta}^o$ with $f(\boldsymbol{\Theta}^o)$ sufficiently large. Let $\boldsymbol{\Theta}_t(\eta) = \mathcal{P}_{\mathcal{S}_d}(\boldsymbol{\Theta}_t - \eta \nabla f(\boldsymbol{\Theta}_t))$. At the t -th iteration, if $\boldsymbol{\Theta}_t \in L_f$, then there must exist $\gamma > 0$ such that for any $\eta \in (0, \gamma)$, $\boldsymbol{\Theta}_t(\eta)$ satisfies that $\boldsymbol{\Theta}_t(\eta) \in \mathbb{S}_{++}^p$. This is because $\boldsymbol{\Theta}_t$ is an interior point of \mathbb{S}_{++}^p , indicating that there exists $r > 0$, such that for any $\boldsymbol{\Theta}$ satisfying $\|\boldsymbol{\Theta} - \boldsymbol{\Theta}_t\|_{\text{F}} < r$, $\boldsymbol{\Theta} \in \mathbb{S}_{++}^p$. Taking $\eta < \frac{r}{\|\nabla f(\boldsymbol{\Theta}_t)\|_{\text{F}}}$, one has,

$$\|\boldsymbol{\Theta}_t(\eta) - \boldsymbol{\Theta}_t\|_{\text{F}} = \|\mathcal{P}_{\mathcal{S}_d}(\boldsymbol{\Theta}_t - \eta \nabla f(\boldsymbol{\Theta}_t)) - \mathcal{P}_{\mathcal{S}_d}(\boldsymbol{\Theta}_t)\|_{\text{F}} \leq \eta \|\nabla f(\boldsymbol{\Theta}_t)\|_{\text{F}} < r,$$

establishing that $\Theta_t(\eta) \in \mathcal{S}_{++}^p$. The projection $\mathcal{P}_{\mathcal{S}_d}$ ensures $\Theta_t(\eta) \in \mathcal{S}_d$. Totally, $\Theta_t(\eta) \in \mathcal{S}$.

The line search condition (8) ensures that $f(\Theta_t(\eta)) \leq f(\Theta_t)$, implying that $\Theta_t(\eta) \in L_f$. Then one has

$$f(\Theta_t(\eta)) \leq f(\Theta_t) + \langle \nabla f(\Theta_t), \Theta_t(\eta) - \Theta_t \rangle + \frac{L}{2} \|\Theta_t(\eta) - \Theta_t\|_{\mathbb{F}}^2. \quad (46)$$

Following from projection theorem (Beck, 2014, Theorem 9.8), one has

$$\langle \Theta_t - \eta \nabla f(\Theta_t) - \Theta_t(\eta), \Theta_t - \Theta_t(\eta) \rangle. \quad (47)$$

Together with (46), one has

$$f(\Theta_t(\eta)) \leq f(\Theta_t) - \eta \left(1 - \frac{\eta L}{2}\right) \|G_{\frac{1}{\eta}}(\Theta_t)\|_{\mathbb{F}}^2. \quad (48)$$

Taking $\eta \leq \frac{2(1-\alpha)}{L}$ leads to $1 - \frac{\eta L}{2} \geq \alpha$. Therefore, for any $\eta \leq \min\left(\frac{2(1-\alpha)}{L}, \gamma\right)$, $\Theta_t(\eta)$ can simultaneously satisfy $\Theta_t(\eta) \in \mathcal{S}$ and the line search condition (8). Thus, the step size in line search has a lower bound $\eta_t \geq \min\left(\frac{2(1-\alpha)\beta}{L}, \gamma, \sigma\right)$, and $\Theta_{t+1} \in L_f$. By induction, $\Theta_t \in L_f$ for any $t \geq 0$. Sequence $\{f(\Theta_t)\}$ is monotonically decreasing, and $f(\Theta_{t+1}) < f(\Theta_t)$ until $G_{\frac{1}{\eta_t}}(\Theta_t) = \mathbf{0}$, indicating that Θ_t is a stationary point (Beck, 2014, Theorem 9.10). Since Problem (5) is a strictly convex problem, the stationary point is the unique minimizer. Therefore, the proposed algorithm converges to the optimal solution. \square

B.3. Proof of Proposition 3.3

Proof. The Lagrangian of the optimization (18) is

$$L(\mathbf{x}, \boldsymbol{\mu}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|^2 - \boldsymbol{\mu}_r \mathbf{x}^\top \mathbf{1} + \langle \boldsymbol{\mu}_{\setminus r}, \mathbf{x}_{\setminus r} \rangle,$$

where $\boldsymbol{\mu}$ is a KKT multiplier. Let $(\hat{\mathbf{x}}, \hat{\boldsymbol{\mu}})$ be the primal and dual optimal point. Then $(\hat{\mathbf{x}}, \hat{\boldsymbol{\mu}})$ must satisfy the KKT system:

$$\hat{x}_i - y_i - \hat{\mu}_r + \hat{\mu}_i = 0, \quad \hat{\mu}_i \hat{x}_i = 0, \quad \text{for } i \neq r; \quad (49)$$

$$\hat{x}_i - y_i - \hat{\mu}_i = 0, \quad \hat{\mu}_i \hat{\mathbf{x}}^\top \mathbf{1} = 0, \quad \text{for } i = r; \quad (50)$$

$$\hat{\mu}_1, \dots, \hat{\mu}_p \geq 0; \quad (51)$$

Therefore, for any $i \neq r$, it holds that $\hat{x}_i = y_i + \hat{\mu}_r - \hat{\mu}_i$, $\hat{\mu}_i \hat{x}_i = 0$, $\hat{\mu}_i \geq 0$ and $\hat{x}_i \leq 0$. As a result, we obtain the following results:

- If $y_i + \hat{\mu}_r \leq 0$, then $\hat{\mu}_i = 0$, indicating that $\hat{x}_i = y_i + \hat{\mu}_i$.
- If $y_i + \hat{\mu}_r > 0$, then $\hat{\mu}_i = y_i + \hat{\mu}_r$, indicating that $\hat{x}_i = 0$.

Overall, we obtain that

$$\hat{x}_i = \min(y_i + \hat{\mu}_r, 0), \quad \forall i \neq r. \quad (52)$$

On the other hand, \hat{x}_r and $\hat{\mu}_r$ satisfy that $\hat{x}_r = y_r + \hat{\mu}_r$, $\hat{\mu}_r \hat{\mathbf{x}}^\top \mathbf{1} = 0$, $\hat{\mu}_r \geq 0$, and $\hat{\mathbf{x}}^\top \mathbf{1} \geq 0$. Then we have the following results:

- If $y_r \geq -\sum_{i \in [p] \setminus r} \min(y_i, 0)$, then $\hat{\mu}_r = 0$, indicating that $\hat{x}_r = y_r$. Following from (52), $\hat{x}_i = \min(y_i, 0)$ for any $i \neq r$.
- If $y_r < -\sum_{i \in [p] \setminus r} \min(y_i, 0)$, then $\hat{\mu}_r \neq 0$. This is because $\hat{\mu}_r = 0$ will result in $\hat{\mathbf{x}}^\top \mathbf{1} < 0$, which does not fulfill the KKT system. Together with the KKT condition that $\hat{\mu}_r \hat{\mathbf{x}}^\top \mathbf{1} = 0$, one has $\hat{\mathbf{x}}^\top \mathbf{1} = 0$. Therefore, one obtains that $\hat{x}_r = y_r + \hat{\mu}_r$, where $\hat{\mu}_r$ satisfies $\hat{\mu}_r + y_r + \sum_{i \in [p] \setminus r} \min(y_i + \hat{\mu}_r, 0) = 0$.

Then the ρ in Proposition 3.3 is the dual optimal point $\hat{\mu}_r$ above. \square

B.4. Proofs of Theorem 4.4 and Corollary 4.5

Proof. We prove the theorem under the case of estimating M -matrices, i.e., $\mathcal{S} = \mathcal{M}^p$. The proof can be directly applied to the case of estimating diagonally dominant M -matrices, i.e., $\mathcal{S} = \mathcal{M}_D^p$, since \mathcal{M}_D^p is a subset of \mathcal{M}^p . Provided that the following event holds:

$$\mathcal{J} := \left\{ \|\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}}\|_{\max} \leq \frac{\sqrt{2}}{2}\lambda \right\}. \quad (53)$$

Recall that $\widehat{\boldsymbol{\Theta}}^{(k)} = G_{\lambda, \widehat{\boldsymbol{\Sigma}}}(\widehat{\boldsymbol{\Theta}}^{(k-1)})$ for $k \geq 2$. According to Assumption 4.1 that $p_\lambda(0) = \lambda$, we can rewrite $\widehat{\boldsymbol{\Theta}}^{(1)}$ as $\widehat{\boldsymbol{\Theta}}^{(1)} = G_{\lambda, \widehat{\boldsymbol{\Sigma}}}(\widehat{\boldsymbol{\Theta}}^{(0)})$, where $\widehat{\boldsymbol{\Theta}}_{ij}^{(0)} = 0$ for any $i \neq j$. Thus, $|T_\alpha(\widehat{\boldsymbol{\Theta}}^{(0)})| \leq 2s$. Following from Lemma B.2, we can further obtain that $|T_\alpha(\widehat{\boldsymbol{\Theta}}^{(1)})| \leq 2s$. To simplify notation, we denote $T_\alpha(\widehat{\boldsymbol{\Theta}}^{(k)})$ by T_α^k for short. By induction, we obtain that, for any $k \geq 1$,

$$|T_\alpha(\widehat{\boldsymbol{\Theta}}^{(k)})| \leq 2s, \quad (54)$$

and

$$\|\widehat{\boldsymbol{\Theta}}^{(k)} - \boldsymbol{\Theta}^*\|_{\text{F}} \leq \frac{c_0}{2} (\delta_1^{k-1} + \delta_2^{k-1}), \quad (55)$$

where $\delta_1^{k-1} = \|p_\lambda(|\widehat{\boldsymbol{\Theta}}_{S^*}^{(k-1)}|)\|$ and $\delta_2^{k-1} = \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{T_\alpha^{k-1} \cup I_p}\|$.

In what follows, we show that the term δ_1^{k-1} in (55) can be bounded in terms of $\|\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*\|$. For any $(i, j) \in S^*$ and $\boldsymbol{\Theta} \in \mathbb{R}^{p \times p}$, if $|\Theta_{ij}^* - \Theta_{ij}| \geq \alpha$, then one has

$$0 \leq p_\lambda(|\Theta_{ij}|) \leq \lambda \leq \frac{\lambda}{\alpha} |\Theta_{ij}^* - \Theta_{ij}|,$$

where the first two inequalities follows from Assumption 4.1. If $|\Theta_{ij}^* - \Theta_{ij}| \leq \alpha$, then one has

$$0 \leq p_\lambda(|\Theta_{ij}|) \leq p_\lambda(|\Theta_{ij}^*| - \alpha) = 0, \quad (56)$$

where the second inequality follows from Assumption 4.1 that p_λ is monotonically non-increasing, and the equality is established because $\min_{(i,j) \in S^*} |\Theta_{ij}^*| - \alpha \geq \gamma\lambda$ and $p_\lambda(x) = 0$ for any $x \geq \gamma\lambda$ following from Assumptions 4.1 and 4.2. As a result, one can obtain

$$\delta_1^{k-1} \leq \frac{\lambda}{\alpha} \|(\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*)_{S^*}\| \leq \frac{\lambda}{\alpha} \|\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*\|_{\text{F}}. \quad (57)$$

To bound δ_2^{k-1} , we separate the set T_α^{k-1} into two parts, S^* and $L_\alpha^{k-1} \setminus S^*$, where L_α^k denotes the set $L_\alpha(\widehat{\boldsymbol{\Theta}}^{(k)})$. The term $\|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{L_\alpha^{k-1} \setminus S^*}\|$ can be bounded in terms of $\|\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*\|_{\text{F}}$ as follows:

$$\|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{L_\alpha^{k-1} \setminus S^*}\| \leq \sqrt{|L_\alpha^{k-1} \setminus S^*|} \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{L_\alpha^{k-1} \setminus S^*}\|_{\max} \leq \frac{\sqrt{2}\lambda}{2\alpha} \|\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*\|_{\text{F}},$$

where the last second equality follows from

$$\sqrt{|L_\alpha^{k-1} \setminus S^*|} \leq \frac{\|\widehat{\boldsymbol{\Theta}}_{L_\alpha^{k-1} \setminus S^*}^{(k-1)}\|}{\alpha} \leq \frac{\|\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*\|_{\text{F}}}{\alpha}, \quad (58)$$

where the first inequality follows from the definition of L_α^{k-1} . Thus one has

$$\delta_2^{k-1} \leq \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{S^* \cup I_p}\| + \frac{\sqrt{2}\lambda}{2\alpha} \|\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*\|_{\text{F}}. \quad (59)$$

Substituting (57) and (59) into (55) yields

$$\|\widehat{\boldsymbol{\Theta}}^{(k)} - \boldsymbol{\Theta}^*\|_{\text{F}} \leq \frac{c_0}{2} \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{S^* \cup I_p}\| + \rho \|\widehat{\boldsymbol{\Theta}}^{(k-1)} - \boldsymbol{\Theta}^*\|_{\text{F}},$$

where $\rho = \frac{2+\sqrt{2}}{4}$. By induction, if $d_k \leq a_0 + \rho d_{k-1}$ for any $k \geq 1$ with $\rho \in [0, 1)$, then

$$d_k \leq \frac{1 - \rho^k}{1 - \rho} a_0 + \rho^k d_0. \quad (60)$$

Applying (60) with $a_0 = \frac{c_0}{2} \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{S^* \cup I_p}\|$, and $d_k = \|\widehat{\boldsymbol{\Theta}}^{(k)} - \boldsymbol{\Theta}^*\|_{\text{F}}$ yields

$$\|\widehat{\boldsymbol{\Theta}}^{(k)} - \boldsymbol{\Theta}^*\|_{\text{F}} \leq 4c_0 \|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{S^* \cup I_p}\| + \rho^k \|\widehat{\boldsymbol{\Theta}}^{(0)} - \boldsymbol{\Theta}^*\|_{\text{F}}.$$

Next, we establish Corollary 4.5. Under the event \mathcal{J} , one has

$$\|(\boldsymbol{\Sigma}^* - \widehat{\boldsymbol{\Sigma}})_{S^* \cup I_p}\| \leq \frac{\sqrt{2}}{2} \lambda \sqrt{s}, \quad (61)$$

where the inequality follows from $|S^* \cup I_p| \leq s$. As a result, we plug $\lambda = \sqrt{c_1 \tau \log p/n}$ into (61) and obtain

$$\|\widehat{\boldsymbol{\Theta}}^{(k)} - \boldsymbol{\Theta}^*\|_{\text{F}} \leq c \sqrt{s \log p/n} + \rho^k \|\widehat{\boldsymbol{\Theta}}^{(0)} - \boldsymbol{\Theta}^*\|_{\text{F}}, \quad (62)$$

where $c = 2c_0 \sqrt{2c_1 \tau}$.

Finally, we calculate the probability that event \mathcal{J} holds. We apply (Ravikumar et al., 2011, Lemma 1) and union sum bound, and obtain

$$\mathbb{P} \left[\|\boldsymbol{\Theta}^* - \widehat{\boldsymbol{\Sigma}}\|_{\max} \geq \frac{\sqrt{2}}{2} \lambda \right] \leq 4p^2 \exp \left(-\frac{n\lambda^2}{(80 \max_i \Sigma_{ii}^*)^2} \right).$$

Take $\lambda = \sqrt{c_1 \tau \log p/n}$ with $\tau > 2$ and $c_1 = (80 \max_i \Sigma_{ii}^*)^2$. By calculation, we obtain that

$$\mathbb{P} \left[\|\boldsymbol{\Theta}^* - \widehat{\boldsymbol{\Sigma}}\|_{\max} \leq \frac{\sqrt{2}}{2} \lambda \right] \geq 1 - 4/p^{\tau-2}, \quad (63)$$

completing the proof. \square