
Efficient Graph Laplacian Estimation by Proximal Newton

Yakov Medvedovsky

Electrical and Computer Engineering,
Ben-Gurion University.

Eran Treister

Computer Science Dept.,
Ben-Gurion University.

Tirza Routtenberg

Electrical and Computer Engineering,
Ben-Gurion University.

Abstract

The Laplacian-constrained Gaussian Markov Random Field (LGMRF) is a common multivariate statistical model for learning a weighted sparse dependency graph from given data. This graph learning problem can be formulated as a maximum likelihood estimation (MLE) of the precision matrix, subject to Laplacian structural constraints, with a sparsity-inducing penalty term. This paper aims to solve this learning problem accurately and efficiently. First, since the commonly used ℓ_1 -norm penalty is inappropriate in this setting and may lead to a complete graph, we employ the nonconvex min-max concave penalty (MCP), which promotes sparse solutions with lower estimation bias. Second, as opposed to existing first-order methods for this problem, we develop a second-order proximal Newton approach to obtain an efficient solver, utilizing several algorithmic features, such as using conjugate gradients, preconditioning, and splitting to active/free sets. Numerical experiments demonstrate the advantages of the proposed method in terms of both computational complexity and graph learning accuracy compared to existing methods.

1 INTRODUCTION

Graphs are fundamental mathematical structures used in various fields to represent data, signals, and processes. Weighted graphs naturally arise when processing networked data applications, such as computer, social, sensor, energy, transportation, and biological

networks (Newman, 2010), where the data is inherently related to a graph associated with the underlying network. As a result, graph learning is used in a wide range of applications, both for virtual, data-based networks and for physical networks of infrastructures; the latter include the brain (Vecchio et al., 2017), financial analysis (Giudici and Spelta, 2016), and electrical networks (Deka et al., 2018). In this context, a fundamental graph learning problem is the estimation of the Laplacian matrix, which plays a central role in spectral graph theory (Chung, 1997) and machine learning (Von Luxburg, 2007). Moreover, in the field of graph signal processing (GSP) (Shuman et al., 2013), the Laplacian matrix is used to extend basic signal processing operations, such as filtering (Milanfar, 2012), sampling (Anis et al., 2016), and signal recovery (Narang and Ortega, 2012; Kroizer et al., 2022). Therefore, accurate and efficient Laplacian learning is extremely valuable for data processing in networks.

Gaussian Markov Random Fields (GMRFs) (Rue and Held, 2005) are probabilistic graphical models that have been studied extensively over the last two decades (Banerjee et al., 2006; Hsieh et al., 2013, 2011). In particular, there has been a significant effort to develop approaches for learning sparse precision matrices under GMRF models, where this precision represents a graph topology. The graphical LASSO (GLASSO) (Friedman et al., 2008) is a widely-used approach to estimate the sparse precision matrix using ℓ_1 -norm regularization. Laplacian-constrained GMRF (LGMRF) models are a specific type of GMRF, in which the precision is a Laplacian matrix. Laplacian constraints imply that smooth signals, which have similar values for nodes connected by large weights, have a higher probability (Ying et al., 2020b). Smooth graph signals are common in various network applications, such as power systems (Dabush et al., 2023; Drayer and Routtenberg, 2020). Thus, LGMRF models have been widely studied in semi-supervised learning frameworks (Zhu et al., 2003), to analyze real-world datasets (Cardoso and Palomar, 2020), and in GSP (Dong et al., 2016; Egilmez et al., 2017).

The estimation of standard and Laplacian-constrained GMRF models are closely related. However, the extension of GLASSO and similar GMRF estimation algorithms to the case of LGMRF, which also considers the Laplacian constraints, is not straightforward. For example, since the Laplacian matrix is sparse in many applications, then, similar to the GLASSO approach, ℓ_1 -norm regularization methods were used to promote the Laplacian sparsity (Egilmez et al., 2017; Kumar et al., 2019; Liu et al., 2019; Zhao et al., 2019). However, it was recently shown that the ℓ_1 -norm is an inappropriate penalty for promoting the sparsity of the precision matrix under the LGMRF model, as it leads to an inaccurate recovery of the connectivity pattern of the graph (Ying et al., 2020a). Thus, to properly address the Laplacian estimation problem, other sparsity-promoting penalty terms should be used (Ying et al., 2020a,b). However, these existing methods are of first-order (i.e., gradient-based), while in general GMRF estimation literature (Oztoprak et al., 2012; Hsieh et al., 2013, 2011; Treister and Turek, 2014) it is known that second-order methods, i.e., proximal Newton (Lee et al., 2014), yield better performance in terms of run time. In such approaches, a quadratic approximation is applied on the smooth part, while the non-smooth regularization term remains intact. This exploits a connection between the gradient and Hessian of the MLE objective and hence is efficient for GMRF estimation. The work of Hassan-Moghaddam et al. (2016) employs a second-order approach for Laplacian-constrained topology identification. However, it is based on the convex ℓ_1 -norm penalty, which was found inappropriate for LGMRF (Ying et al., 2020a), and the edge weights are not constrained to be nonnegative, which may result in non-Laplacian solutions. Therefore, there is a need for second-order methods for Laplacian learning that utilize proper regularization (to yield accurate sparsity patterns) and are efficient in run time.

Contribution. In this paper, we present an efficient graph estimation method for the LGMRF model based on the proximal Newton approach. Our method involves solving a nonconvex penalized MLE problem with MCP penalty, instead of the commonly-used ℓ_1 -norm penalty. At each iteration, the smooth part of the objective is approximated by a second-order Taylor expansion around the current precision matrix, while the non-smooth penalty and Laplacian constraints remain intact. The resulting Newton problem is a penalized quadratic minimization under linear constraints to guarantee that the next iterate is a Laplacian matrix. To the best of our knowledge, the proposed method is the first proximal Newton method for the LGMRF model estimation. Furthermore, our framework includes two nontrivial algorithmic novelties in

this context. First, the inner, constrained, Newton problem is solved by a projected nonlinear conjugate gradient method. This yields a significant computational speedup over gradient-based first-order methods. Second, we introduce a diagonal preconditioner to improve the performance further. We show some theoretical results regarding the problem and the algorithms we propose. Furthermore, we demonstrate the effectiveness of the proposed method in learning sparse graphs through numerical experiments.

Organization and notation. The rest of this paper is organized as follows. In Section 2, we describe the considered Laplacian learning problem, simplify its formulation, and discuss the sparsity-promoting penalty function. In Section 3, we develop the proposed NewGLE method. Simulations are shown in Section 5 and the paper is concluded in Section 6.

In the rest of this paper we denote vectors and matrices in boldface lowercase and uppercase letters, respectively. The elements of the vectors $\mathbf{1}$ and $\mathbf{0}$ are ones and zeros, respectively. The notations $\|\cdot\|_{\mathbf{F}}$, $|\cdot|$, \otimes , $(\cdot)^T$, $(\cdot)^{-1}$, $(\cdot)^\dagger$, $|\cdot|_+$, and $\text{Tr}(\cdot)$ denote the Frobenius norm, determinant operator, Kronecker product, transpose, inverse, Moore-Penrose pseudo-inverse, pseudo-determinant, and trace, respectively. The sets \mathcal{S}^p and \mathcal{S}_+^p are of real symmetric and real positive semi-definite $p \times p$ matrices, respectively. For matrix $\mathbf{A} \in \mathbb{R}^{p \times p}$, $\nabla_{\mathbf{A}} f \in \mathbb{R}^{p \times p}$ and $\nabla_{\mathbf{A}}^2 f \in \mathbb{R}^{p^2 \times p^2}$ are the gradient and Hessian of the scalar function $f(\mathbf{A})$.

2 PROBLEM FORMULATION

In this section, we formulate the considered Laplacian learning problem. We start with definitions related to GSP. Then, we present the LGMRF and formulate the graph estimation problem under the LGMRF and MLE in Subsection 2.1. The sparsity-promoting penalty function is discussed in Subsection 2.2.

We consider an undirected, connected, and weighted graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \dots, p\}$ and $\mathcal{E} = \{1, \dots, m\}$ are the set of vertices and the set of edges, respectively. The nonnegative weighted adjacency matrix of the graph, \mathbf{W} , has the (i, j) -th element $w_{i,j} > 0$ if vertex i and j are connected, and zero otherwise. From spectral graph theory (Chung, 1997), the rank of the Laplacian matrix for a connected graph with p nodes is $p - 1$. Thus, the set of Laplacian matrices for connected graphs can be defined as (Ying et al., 2020b)

$$\mathcal{L} = \left\{ \mathbf{L} \in \mathcal{S}_+^p \mid \begin{array}{l} \mathbf{L}_{i,j} \leq 0, \forall i \neq j, i, j = 1, \dots, p \\ \mathbf{L}\mathbf{1} = \mathbf{0}, \text{rank}(\mathbf{L}) = p - 1 \end{array} \right\}. \quad (1)$$

2.1 GMRF and Constrained MLE

In this paper, we use the formulation of the graph estimation problem based on the probabilistic graphical model (Koller and Friedman, 2009; Banerjee et al., 2008). We assume that the data samples are obtained from a zero-mean Gaussian distribution parametrized by a positive semi-definite precision matrix, \mathbf{L} , i.e., $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{L}^\dagger)$, where $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. This defines an improper LGMRF model (Kumar et al., 2019) with parameters $(\mathbf{0}, \mathbf{L})$, where the precision matrix is $\mathbf{L} \in \mathcal{L}$ and \mathcal{L} is defined in (1).

Given n independent and identically distributed (i.i.d.) samples $\mathbf{x}_1, \dots, \mathbf{x}_n$ drawn from an LGMRF, our goal is to find the constrained MLE of \mathbf{L} based on these samples and under the Laplacian constraints. The MLE can be found by solving the following constrained minimization of the negative log-likelihood (Ying et al., 2021):

$$\hat{\mathbf{L}}_{MLE} = \underset{\mathbf{L} \in \mathcal{L}}{\operatorname{argmin}} \{ \operatorname{Tr}(\mathbf{L}\mathbf{S}) - \log |\mathbf{L}|_+ \}, \quad (2)$$

where $\mathbf{S} \triangleq \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^T$ is the empirical covariance matrix.

Sparsity plays an important role in high-dimensional learning, which helps avoid over-fitting and improves the identification of the relationships among data, especially where \mathbf{S} is low rank. In particular, the Laplacian is a sparse matrix in various applications, e.g. electrical networks (Halihal and Routtenberg, 2022; Grotas et al., 2019). A sparse graph estimation problem under the LGMRF model can be formulated by adding a sparse penalty function to the estimator in (2), which results in

$$\underset{\mathbf{L} \in \mathcal{L}}{\operatorname{minimize}} \quad \operatorname{Tr}(\mathbf{L}\mathbf{S}) - \log |\mathbf{L}|_+ + \rho(\mathbf{L}; \lambda), \quad (3)$$

where $\rho(\mathbf{L}; \lambda)$ is a general sparsity-promoting penalty function and λ is a tuning parameter.

The objective function in (3) (or in (2)) involves the pseudo-determinant term, since the target matrix \mathbf{L} is a singular matrix. As a result, the optimization problem becomes hard to solve (Holbrook, 2018). As a remedy to this problem, it is demonstrated in (Egilmez et al., 2017) that if $\mathbf{L} \in \mathcal{L}$, then we can use the relation

$$|\mathbf{L}|_+ = |\mathbf{L} + \mathbf{J}|, \quad (4)$$

where $\mathbf{J} = \frac{1}{p} \mathbf{1}\mathbf{1}^T$. By substituting (4) in (3), the problem is simplified. In addition, in the following theorem we show that the rank and the positive semi-definiteness constraints in (1) are redundant. The proof appears in Appendix 7.

Theorem 1. *The optimization problem*

$$\underset{\mathbf{L} \in \mathcal{F}}{\operatorname{minimize}} \quad \operatorname{Tr}(\mathbf{L}\mathbf{S}) - \log |\mathbf{L} + \mathbf{J}| + \rho(\mathbf{L}; \lambda), \quad (5)$$

where the feasible set is given by

$$\mathcal{F} = \left\{ \mathbf{L} \in \mathcal{S}^p \mid \begin{array}{l} \mathbf{L}_{i,j} \leq 0, \forall i \neq j, \quad i, j = 1, \dots, p \\ \mathbf{L}\mathbf{1} = \mathbf{0} \end{array} \right\}. \quad (6)$$

is equivalent to the optimization problem stated in (3).

2.2 Sparsity Promoting Penalty Functions

Recent works like (Egilmez et al., 2017) introduced the ℓ_1 -norm penalized MLE under the LGMRF model to estimate a sparse graph. The choice of ρ in (5) as $\|\mathbf{L}\|_{1,\text{off}}$, the absolute sum of all off-diagonal elements of \mathbf{L} , results in the same objective function as in the well-known GLASSO problem but where the feasible set is restricted to $\mathbf{L} \in \mathcal{F}$. However, unlike in GLASSO, under Laplacian constraints the ℓ_1 -norm is ineffective in promoting a sparse solution. In particular, it is shown by Ying et al. (2020a) that substituting $\rho(\mathbf{L}; \lambda) = \lambda \|\mathbf{L}\|_{1,\text{off}}$ in the objective of (5) and choosing the penalty parameter λ to be large enough leads to a fully-connected (complete) graph as the solution. Moreover, Ying et al. (2020a) show in their numerical experiments that the number of the estimated positive edges increases as λ increases. An intuitive explanation for this phenomenon is that the ℓ_1 -norm promotes sparsity uniformly on all off-diagonal entrances of the matrix. Hence, for $\lambda > 0$, minimizing the ℓ_1 -norm separately with a positive definite constraint, as in GLASSO, leads to a diagonal matrix with a positive diagonal. However, unlike a regular GMRF model, the Laplacian constraints, $\mathbf{L}\mathbf{1} = \mathbf{0}$ and $\mathbf{L}_{i,j} \leq 0$, for $i \neq j$, relate the off-diagonal entries with the diagonal entries of \mathbf{L} . Hence, minimizing the ℓ_1 -norm would lead to the zeros matrix. Therefore, as λ increases, the eigenvalues of \mathbf{L} get closer to zero, and the $-\log |\cdot|$ term in the objective function increases significantly. Since ℓ_1 -norm penalizes all off-diagonal elements in a uniform way, to prevent a significant increase of the objective function, the minimization process inserts relatively small (in absolute value) negative numbers, even for zero elements in the ground-truth Laplacian matrix. This way, both the ℓ_1 -norm and the expression $-\log |\cdot|$ in the optimization problem remain relatively small.

Alternative approaches using nonconvex penalties, such as clipped absolute deviation (SCAD) (Fan and Li, 2001) and the minimax concave penalty (MCP) (Zhang, 2010a), have been proposed to alleviate this issue. These penalty functions penalize the elements in the precision matrix more selectively than the ℓ_1 -norm, such that elements that receive significant values from the MLE will receive insignificant penalties, while elements closer to zero will receive a greater penalty. In

this paper, we use the MCP penalty function, which is applied on the off-diagonal entries of \mathbf{L} as follows:

$$\rho_{MCP}(\mathbf{L}; \lambda) = \sum_{i \neq j} MCP(\mathbf{L}_{ij}; \gamma, \lambda), \quad (7)$$

where the $MCP(\cdot)$ is a scalar function, defined as

$$MCP(x; \gamma, \lambda) = \begin{cases} \lambda|x| - \frac{x^2}{2\gamma} & \text{if } |x| \leq \gamma\lambda \\ \frac{1}{2}\gamma\lambda^2 & \text{if } |x| > \gamma\lambda \end{cases}. \quad (8)$$

The function is constant and has a derivative of 0 for values larger than $\gamma\lambda$ in magnitude, hence only elements smaller than $\gamma\lambda$ are affected by $\rho_{MCP}(\cdot)$.

To conclude, our goal in this paper is to develop an efficient algorithm for finding the constrained MLE of the Laplacian matrix, which is the precision matrix of the LGMRF model. This is done by solving the optimization problem in (5) with $\rho(\mathbf{L}; \lambda) = \rho_{MCP}(\mathbf{L}; \lambda)$, defined in (7). I.e., the problem we solve is given by:

$$\underset{\mathbf{L} \in \mathcal{F}}{\text{minimize}} \quad \text{Tr}(\mathbf{L}\mathbf{S}) - \log |\mathbf{L} + \mathbf{J}| + \rho_{MCP}(\mathbf{L}; \lambda), \quad (9)$$

where \mathcal{F} is defined in (6).

3 METHOD

In this section, we develop the proposed proximal Newton method for the graph Laplacian estimation, named NewGLE, that aims to solve (9) efficiently using a proximal Newton approach adopted to the problem. First, we define the constrained Newton optimization problem for the graph learning in Subsection 3.1, keeping the Laplacian constraints when finding the Newton direction. Then, in Subsection 3.2, we present the approach that maps a Laplacian matrix to a vector and simplifies the set of constraints. Next, we describe the idea of restricting the Newton direction into a “free set” in Subsection 3.3. This constitutes the NewGLE method, which is summarized in Algorithm 1. Following that, we present our inner solver for finding the Newton direction, which is composed of the nonlinear projected conjugate gradient that is used together with a diagonal preconditioner, in Subsection 3.4 and Appendix 9. Finally, we provide an algorithm (Algorithm 2 in the Appendix) for finding the Newton direction.

3.1 Proximal Newton for Graph Laplacian Estimation

In the “proximal Newton” approach, a quadratic approximation is applied on the smooth part of the objective function, while leaving the non-smooth term intact, in order to obtain the Newton descent direction. This approach is considered to be among the state-of-the-art methods for solving the GLASSO problem

(Rolfs et al., 2012; Hsieh et al., 2013, 2011; Mazumder and Hastie, 2012; Treister et al., 2016), which is highly related to our problem in (9). The advantage of this method lies in the treatment of the log det term (as we show next), which appears in both problems. However, the GLASSO methods cannot be applied to our graph learning problems because we aim to estimate a precision matrix that satisfies the Laplacian constraints (and is singular). In contrast, the learned precision matrix under the GMRF model in GLASSO is a general positive definite matrix. In addition, we consider the MCP penalty, while the above mentioned methods consider the ℓ_1 -norm penalty, which is not suitable for our case, as discussed in Subsection 2.2.

In our proximal Newton approach we design a constrained Newton problem, where the smooth part of the objective function in (9) is

$$f(\mathbf{L}) = \text{Tr}(\mathbf{L}\mathbf{S}) - \log |\mathbf{L} + \mathbf{J}|, \quad (10)$$

and the penalty function, $\rho_{MCP}(\mathbf{L}; \lambda)$, is the non-smooth term. At the t -th iteration of the proximal Newton approach, the smooth part of the objective is approximated by a second-order Taylor expansion around the current estimation $\mathbf{L}^{(t)}$. To this end, we use the gradient and Hessian of $f(\mathbf{L})$ that are given by (Hsieh et al., 2014; Chin et al., 2013)

$$\nabla_{\mathbf{L}} f = \mathbf{S} - \mathbf{Q}, \quad \nabla_{\mathbf{L}}^2 f = \mathbf{Q} \otimes \mathbf{Q}, \quad (11)$$

where \otimes is the Kronecker product and $\mathbf{Q} = (\mathbf{L} + \mathbf{J})^{-1}$.

The gradient in (11) already shows the main difficulty in solving the optimization problem in (9): it contains \mathbf{Q} , the inverse of the matrix $\mathbf{L} + \mathbf{J}$, which is expensive to compute. The advantage of the proximal Newton approach for this problem is the low overhead: by calculating \mathbf{Q} in $\nabla_{\mathbf{L}} f(\mathbf{L})$, we also get part of the Hessian computation at the same cost.

In the same spirit of the proximal Newton approach, in addition to the penalty term, we also keep the Laplacian constraint intact in the Newton problem. That is, at iteration t , the Newton direction $\Delta^{(t)}$ solves the constrained penalized quadratic minimization problem:

$$\begin{aligned} \underset{\Delta \in \mathcal{S}^p}{\text{minimize}} \quad & f(\mathbf{L}^{(t)}) + \text{Tr}(\Delta(\mathbf{S} - \mathbf{Q})) + \frac{1}{2} \text{Tr}(\Delta \mathbf{Q} \Delta \mathbf{Q}) \\ & + \frac{1}{2} \|\varepsilon \cdot \Delta\|_F^2 + \rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) \\ \text{s.t.} \quad & \mathbf{L}^{(t)} + \Delta \in \mathcal{F}, \end{aligned} \quad (12)$$

where $\mathbf{Q} = (\mathbf{L}^{(t)} + \mathbf{J})^{-1}$. It should be noted that for the sake of simplicity, we use the notation \mathbf{Q} instead of $\mathbf{Q}^{(t)}$, although \mathbf{Q} is updated with the iteration t . Note that the gradient and Hessian of $f(\cdot)$ in (11) at the iterate $\mathbf{L}^{(t)}$ are featured in the second and third

terms in (12), respectively. The first term of (12) is constant and can be ignored. Altogether, the objective function in (12) is quadratic with the MCP sparsity-promoting penalty. In addition, we have a constraint to guarantee that the next update of \mathbf{L} is feasible, i.e., $(\mathbf{L}^{(t)} + \Delta) \in \mathcal{F}$. The fourth term, with a small and symmetric weight matrix $\varepsilon \geq 0$ penalizes the squares of the off-diagonals of Δ . This term is used for the stabilization of the iterative solution of (12), and to make sure that the Hessian of (12) is positive definite, since ρ_{MCP} is concave—more details are given later on in the analysis in Appendix 10.4. The advantage of solving (12) with ρ_{MCP} intact, as opposed to majorizing ρ_{MCP} (e.g., using its gradient instead) is that the underlying gradient-based iterative solver of (12) has access to the curvature of ρ_{MCP} and to its behavior in the sub-sections in its definition (8).

Once the Newton problem is approximately solved, yielding the direction $\Delta^{(t)}$, it is added to $\mathbf{L}^{(t)}$ employing a linesearch procedure to sufficiently reduce the objective in (9). To this end, the updated iterate is

$$\mathbf{L}^{(t+1)} = \mathbf{L}^{(t)} + \alpha^{(t)} \Delta^{(t)}, \quad (13)$$

where the parameter $\alpha^{(t)}$ is obtained by a backtracking linesearch using Armijo’s rule (Armijo, 1966). We state the following theorem on such a proximal quadratic solution process, where we approximate the true Hessians by SPD matrices $\mathbf{M}^{(t)}$ whose eigenvalues are bounded by $\lambda_{min}^{\mathbf{M}}$ and $\lambda_{max}^{\mathbf{M}}$ from below and above, respectively. The proof appears in Appendix 10.

Theorem 2. *Let $\{\mathbf{L}^{(t)}\}$ be a series of points produced by a sequence of proximal quadratic minimizations, constrained by \mathcal{F} , with some SPD matrices $0 \prec \lambda_{min}^{\mathbf{M}} \mathbf{I} \preceq \mathbf{M}^{(t)} \preceq \lambda_{max}^{\mathbf{M}} \mathbf{I}$ as the approximate Hessian (as in (35)), followed by the linesearch (13), starting from $\mathbf{L}^{(0)} \in \mathcal{F}$. Then any limit point $\bar{\mathbf{L}}$ of $\{\mathbf{L}^{(t)}\}$ is a stationary point of (9).*

The theorem shows that if we choose $\mathbf{M}^{(t)}$ such that $\lambda_{min}^{\mathbf{M}} > \gamma^{-1}$, then the method monotonically converges to a stationary point of (9). Note that $\mathbf{M}^{(t)}$ are matrices we choose for the solution, and are not part of the problem, so this requirement can easily be fulfilled. In the next section, we simplify problem (12) to ease the treatment of the constraints.

3.2 Laplacian parameterization

To simplify the Laplacian structural constraints in (12), we use a linear operator that maps a vector $\mathbf{w} \in \mathbb{R}^{p(p-1)/2}$ to a matrix $\mathcal{P}(\mathbf{w}) \in \mathbb{R}^{p \times p}$.

Definition 1. *The linear operator $\mathcal{P} : \mathbb{R}^{p(p-1)/2} \mapsto$*

$\mathbb{R}^{p \times p}$ is defined as (Ying et al., 2020a)

$$\mathcal{P}(\mathbf{w})_{i,j} = \begin{cases} -w_k & i > j \\ \mathcal{P}(\mathbf{w})_{j,i} & i < j \\ -\sum_{j \neq i} \mathcal{P}(\mathbf{w})_{i,j} & i = j \end{cases}, \quad (14)$$

where $k = i - j + \frac{j-1}{2} (2p - j)$ for $i > j$.

The operator \mathcal{P} defines a linear mapping from a *non-negative* weight vector $\mathbf{w} \geq 0$ to a Laplacian matrix $\mathcal{P}(\mathbf{w}) \in \mathcal{F}$, defined in (6). Using this parametrization, we rewrite the optimization problem in (12) such that the decision variable is a vector instead of a matrix. In particular, in Appendix 8, we show that (12) is equivalent to the following simpler vector-parameterized Newton problem that we solve

$$\begin{aligned} \underset{\delta \in \mathbb{R}^{p(p-1)/2}}{\text{minimize}} \quad & f_N(\delta) \triangleq \text{Tr}(\mathcal{P}(\delta)(\mathbf{S} - \mathbf{Q})) \\ & + \frac{1}{2} \text{Tr}(\mathcal{P}(\delta) \mathbf{Q} \mathcal{P}(\delta) \mathbf{Q}) + \|\tilde{\varepsilon} \odot \delta\|_2^2 \quad (15) \\ & + 2\tilde{\rho}_{MCP}(\mathbf{w} + \delta; \lambda) \\ \text{s.t.} \quad & \delta \geq -\mathbf{w}, \end{aligned}$$

where $\tilde{\varepsilon} \in \mathcal{R}^{p(p-1)/2}$ is the vectorized upper triangle of ε (without the diagonal) and $\tilde{\rho}_{MCP}$ is given by

$$\tilde{\rho}_{MCP}(\mathbf{w}; \lambda) \triangleq \sum_{k=1}^{p(p-1)/2} MCP(w_k; \gamma, \lambda) \quad (16)$$

and $MCP(x; \gamma, \lambda)$ is defined in (8). The only remaining constraint in (15) is an inequality constraint, which simplifies the projection onto the feasible space.

3.3 Restricting the updates to free sets

To ease the minimization of (9) we limit the minimization of the Newton problem to a “free set” of variables, while keeping the rest as zeros. This idea was suggested in (Hsieh et al., 2011) for the GLASSO problem that includes the ℓ_1 -norm penalty. The free set of a matrix \mathbf{L} is defined as

$$\mathbf{Free}(\mathbf{L}) = \left\{ k : \mathbf{L}_{i,j} \neq 0 \vee [\mathbf{S} - \mathbf{Q}]_{i,j} > \lambda \right\}, \quad (17)$$

where \vee denotes the logistic OR operator and the relation between (i, j) and k appears in Definition (1). This set comes from the subgradient of the ℓ_1 -norm term, in addition to the non-positivity constraint—the zero elements that are not in $\mathbf{Free}(\mathbf{L})$ will remain zero after a (projected) coordinate descent update. Because the subgradient of the MCP penalty is identical to that of ℓ_1 -norm at zero, the free set is defined in the same way as for the ℓ_1 -norm. As $\mathbf{L}^{(t)}$ approaches a solution $\bar{\mathbf{L}}$, $\mathbf{Free}(\mathbf{L}^{(t)})$ approaches $\{(i, j) : \bar{\mathbf{L}}_{i,j} = 0\}$. Restricting (15) to the free set variables improves the

Hessian’s condition number and decreases the number of iterations needed to solve (15).

So far, we have discussed the outer iteration of the proposed method, which is summarized in Algorithm 1.

Algorithm 1: NewGLE

Input: Empirical covariance matrix \mathbf{S}

Result: Estimated Laplacian matrix $\hat{\mathbf{L}}$

Initialization: Set $\mathbf{L}^{(0)} \in \mathcal{L}$

while *Stopping criterion is not achieved* **do**

- Compute $\mathbf{Q} = (\mathbf{L}^{(t)} + \mathbf{J})^{-1}$.
- Define **Free** $(\mathbf{L}^{(t)})$ as in (17).
- Find the Newton direction $\delta^{(t)}$ by solving (15) restricted to **Free** $(\mathbf{L}^{(t)})$. % by Alg. 2
- Update: $\mathbf{L}^{(t+1)} = \mathbf{L}^{(t)} + \alpha \mathcal{P}(\delta^{(t)})$, where α is achieved by linesearch.

end

3.4 Solution of the Newton problem

To get the Newton direction for the t -th iteration, we approximately solve the quadratic, linearly constrained and MCP-regularized problem in (15) using an iterative method. We use the nonlinear conjugate gradient (NLCG) method and modify it for our case. Specifically, we use the variant by Dai and Yuan (1999), and since (15) is linearly constrained we add projections between the iterations. Moreover, to speed up the convergence, we use NLCG together with a diagonal preconditioner, using the diagonal of the Hessian, as suggested in (Zibulevsky and Elad, 2010) for the LASSO problem.

The projected and preconditioned NLCG algorithm is the inner and most computationally expensive part of Algorithm 1. The algorithm itself and its details are given in Appendix 9. A key element of the approach is the diagonal preconditioner \mathbf{D} , which is applied by replacing the gradient $(\nabla_{\delta} f_N)$ in the search direction by $\mathbf{D}^{-1} \cdot \nabla_{\delta} f_N$. The matrix \mathbf{D} is defined as the diagonal of the Hessian of $f_N(\delta)$, i.e.

$$\mathbf{D}_{k,k} = \left([\mathbf{Q}]_{i,i} + [\mathbf{Q}]_{j,j} - 2[\mathbf{Q}]_{i,j} \right)^2 + \tilde{\varepsilon}_k^2, \quad (18)$$

where the relation between (i, j) and k appears in Definition 1. For the full derivation of the Hessian diagonal, please refer to Appendix 9. Note that the operator \mathbf{D} is diagonal, so calculating its inverse is trivial.

4 MORE RELATED WORKS

GMRF estimation with ℓ_1 regularization. The ℓ_1 regularized MLE under the GMRF model has been

extensively studied recently (see (d’Aspremont et al., 2008; Friedman et al., 2008; Rolfs et al., 2012; Hsieh et al., 2013, 2011; Treister et al., 2016; Shalom et al., 2023) and references therein). Some of these methods include a proximal Newton approach. However, these methods cannot be directly applied to the graph learning problem, as the goal here is to learn a Laplacian precision matrix as opposed to a general positive definite matrix.

The ℓ_1 regularization bias and its alternatives.

One drawback of the mentioned approaches is that the ℓ_1 penalty causes an estimation bias, which is relevant to several variable selection problems, e.g., least squares minimization and precision estimation under ℓ_1 priors. To reduce the ℓ_1 bias, two-stage methods, nonconvex penalties, and hard-thresholding approaches have been introduced for least squares variable selection (Zhang, 2010a; Breheny and Huang, 2011; Loh and Wainwright, 2013) and GMRF estimation (Chen et al., 2018; Lam and Fan, 2009; Shen et al., 2012; Finder et al., 2022). Several of these methods use nonconvex penalties, such as SCAD, MCP, and capped ℓ_1 -norm (Zhang, 2010b). Different algorithms have been proposed to solve the resulting nonconvex problems, e.g., minimization-maximization (Hunter and Li, 2005), local quadratic and linear approximations (Fan and Li, 2001; Zou and Li, 2008).

LGMRF estimation. The recent works (Egilmez et al., 2017; Kumar et al., 2019; Liu et al., 2019; Zhao et al., 2019) investigate the ℓ_1 -norm penalized MLE under the LGMRF model to estimate a graph Laplacian from data. However, it has been shown recently that imposing an ℓ_1 -norm penalty to the Laplacian-constrained MLE produces an unexpected behavior: the number of nonzero graph weights grows as the regularization parameter increases (Ying et al., 2020a).

To alleviate this issue, nonconvex penalties, such as the MCP, have been proposed together with LGMRF estimation. These methods are first-order (i.e., gradient-based) methods that do not take Hessian information into account, and, thus, are expected to converge slower than second-order methods. For example, the work of Koyakumar et al. (2023) suggests adding Tikhonov regularization to convexify the cost function and use the primal-dual splitting method to solve the optimization problem. However, as the Tikhonov regularization parameter increases, the bias of the estimation increases as well. Another approach in (Ying et al., 2020a,b; Vieyra, 2022) uses a nonconvex estimation method by solving a sequence of weighted ℓ_1 -norm penalized subproblems, according to a majorization-minimization framework. In each majorization step, the ℓ_1 -norm weights are updated, and then in the minimization step, the ℓ_1 -penalized function is minimized

by the projected gradient descent (PGD) method. The approach in (Ying et al., 2021; Tugnait, 2021) requires solving two subproblems: The first is an initial estimator of the precision matrix, e.g., the MLE. Then, the estimation is used to define the ℓ_1 weights for the second subproblem. Both stages are solved by PGD.

All the abovementioned methods are gradient-based methods that do not consider Hessian information. On the other hand, the experience with the closely related GMRF estimation problem suggests that proximal Newton approaches would result in superior computational efficiency compared to those methods.

5 EXPERIMENTAL RESULTS

In this section, we present experimental results for the evaluation of the proposed method. We define the experimental setup in Subsection 5.1. Then, we present the results in terms of graph learning performance and computational efficiency in Subsections 5.2 and 5.3, respectively. Our MATLAB code is publicly available on GitHub at <https://github.com/BGUCompSci/GraphLaplacianEstimationProxNewton>.

5.1 Experimental setup

We compare the performance of the following methods:

- MLE, defined in (2), where the objective function does not include any sparsity-promoting penalty.
- Nonconvex Graph Learning (NGL) with MCP (Ying et al., 2020b), which solves a sequence of weighted ℓ_1 -norm penalized subproblems via the majorization-minimization framework.
- Adaptive Laplacian-constrained Precision matrix Estimation (ALPE) (Ying et al., 2021), which is based on a weighted ℓ_1 -norm-regularized MLE. The weights are precalculated using the MLE.
- PGD for solving (9) using the Projected Gradient Descent method. We developed this method to demonstrate the advantages of using the second-order approximation.
- The proposed method (NewGLE) for the MCP-regularized MLE, which is summarized in Algorithm 1 and Algorithm 2 (in Appendix 9). We used $\gamma = 1.01$ and $\varepsilon = 0$, and to start the solution we applied a few PGD iterations, which were taken into account in the results.

The regularization parameter, λ , is fine-tuned for each method for the best performance. In addition, all algorithms use the convergence criterion $\frac{\|\mathbf{L}^{(t+1)} - \mathbf{L}^{(t)}\|_F}{\|\mathbf{L}^{(t+1)}\|_F} \leq \epsilon$ with the tolerance $\epsilon = 10^{-4}$. All the experiments were conducted on a machine with 2 Intel Xeon E5-2660 2.0GHz processors with 28 cores and 512GB RAM.

We use the relative error (RE) and F-score (FS) to evaluate the performance of the algorithms, where $\text{RE} = \frac{\|\bar{\mathbf{L}} - \mathbf{L}^*\|_F}{\|\mathbf{L}^*\|_F}$ in which $\bar{\mathbf{L}}$ and \mathbf{L}^* denote the estimated and true precision matrices, respectively, and $\text{FS} = \frac{2\text{tp}}{2\text{tp} + \text{fp} + \text{fn}}$, where the true-positive (tp), false-positive (fp), and false-negative (fn) detection of graph edges are calculated by comparing the supports of $\bar{\mathbf{L}}$ and \mathbf{L}^* . The F-score takes values in $[0, 1]$, where 1 indicates perfect recovery of the support.

We generated several datasets based on different graph-based models. To create the datasets, we first created a graph, and then its associated Laplacian matrix, \mathbf{L}^* , is used to generate independent data samples from the Gaussian distribution $\mathcal{N}(\mathbf{0}, (\mathbf{L}^*)^\dagger)$. The graph connectivity is based on two options: 1) a random planar graph consisting of $p = 1,000$ nodes (Treister et al., 2016), and 2) a random Barabasi-Albert graph (Zadorozhnyi and Yudin, 2012) of degree 2 with $p = 100$ nodes. The edge weights for both graph models are uniformly sampled from the range $[0.5, 2]$. The curves in Figs. 1 and 5 are the results of an average of 10 Monte Carlo realizations, and Fig. 2 shows the results of an average of 50 realizations. The standard deviation of all plots in these figures is less than 0.002, which is also approximately the gap between the two leading methods.

5.2 Results - graph learning performance

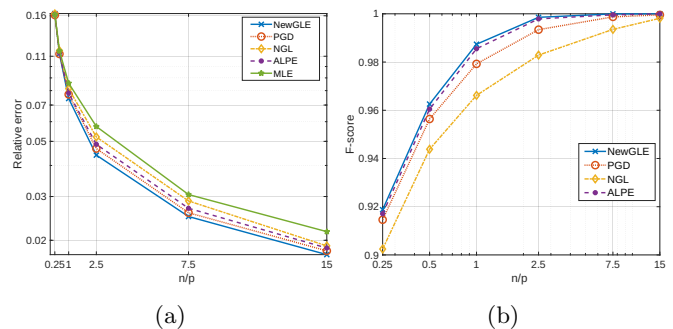


Figure 1: Performance comparisons under (a) RE, (b) F-score with different sample size ratios n/p for planar graphs with 1,000 nodes.

Figures 1 and 2 show the graph learning performance of different methods for learning the topology of planar graphs with 1,000 nodes (Fig. 1) and of Barabasi-Albert graphs of degree two with 100 nodes (Fig. 2). The performance is presented in terms of RE and F-score versus the sample size ratio n/p , where n is the number of data samples used to calculate the sample covariance matrix, \mathbf{S} , and p is the number of nodes.

Figures 1(a) and 2(a) show that the REs of all the estimators decrease as the sample size ratio, n/p , in-

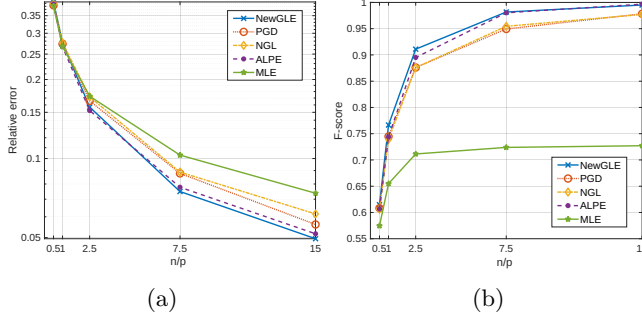


Figure 2: Performance comparisons in terms of (a) RE, and (b) F-score with different sample size ratios n/p for Barabasi-Albert graph of degree 2 with 100 nodes.

creases. Figures 1(b) and 2(b) show that the F-score of all the methods increases as the sample size ratio n/p increases. Figures 1 and 2 show that the proposed NewGLE method outperforms all compared methods (ALPE, NGL, PGD, and MLE) in both the RE and F-score senses. It should be noted that in 1(b) the comparison with the MLE has been removed since it has a significantly lower F-score than the other methods and in order to highlight the differences between the other methods. As shown in Figs. 1(b) and 2(b) for a small sample size ratio, the methods do not provide a perfect F-score. This is because there may not be enough data samples to recover the graph connectivity effectively. Still, the proposed NewGLE method consistently outperforms all compared methods in the F-score sense, especially when the sample size ratio is small. This outcome is significant in large-scale problems since, in these cases, the sample size ratio is usually small.

5.3 Results - Computational Costs

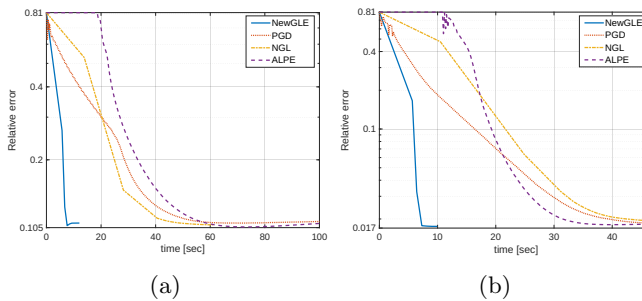


Figure 3: Run-time performance comparisons under RE with a sample size ratio (a) $n/p = 0.5$ and (b) $n/p = 15$ for random planar graphs with 1,000 nodes.

Figures 3-5 compare the computational efficiency of the proposed NewGLE method with the other methods in terms of run time. These experiments were conducted on relatively large planar graphs of $p = 1,000$ nodes so that the measuring times less depend on mi-

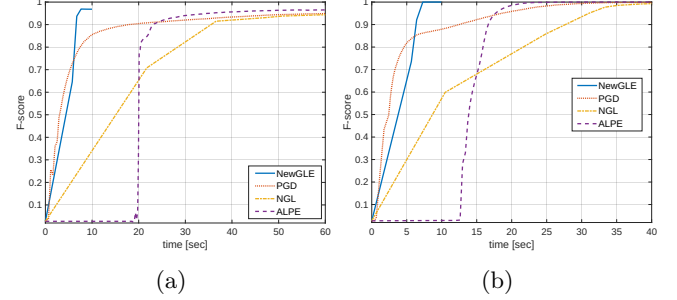


Figure 4: Run-time performance comparisons under F-score with a sample size ratio (a) $n/p = 0.5$ and (b) $n/p = 15$ for random planar graphs with 1,000 nodes.

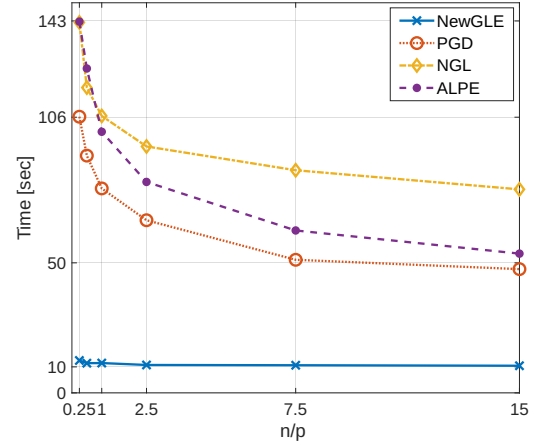


Figure 5: Average convergence times over sample size ratio n/p for random planar graphs with 1,000 nodes.

nor implementation details. Figure 3 shows the convergence of the compared methods in terms of RE for two sample size ratios: (a) $n/p = 0.5$, and (b) $n/p = 15$. Notably, the proposed method converges significantly faster than the other methods in both cases, while the difference between the methods is more significant for the case when the sample size ratio is smaller, $n/p = 0.5$.

Figure 4 shows the F-score of the compared methods during two realizations: (a) with sample size ratio $n/p = 0.5$, and (b) with sample size ratio $n/p = 15$. Notably, the proposed method reached the highest F-score and got it significantly faster than the run time taken for the other methods to achieve their highest F-score values. In addition, it can be seen that in the case of $n/p = 15$, all the methods succeeded in estimating the connectivity of the graph, and the NewGLE, PGD, and ALPE methods managed to estimate it perfectly (i.e., F-score = 1). In addition, it can be seen that the advantage of the proposed method is more notable for a smaller sample size ratio, i.e., $n/p = 0.5$. Additionally, the results presented in Figs. 3 and 4 reveal that

while the RE and F-score performance of the ALPE method, as demonstrated in Subsection (5.2), is commendable and comparable to the proposed method, the implementation of the ALPE method necessitates a significant investment in computational resources to effectively learn a graph, as measured by RE and F-score metrics. The reason for this is that the ALPE method requires an estimation of the weights, which are used for sparsity-promoting regularization.

Finally, Fig. 5 presents the average convergence time for each method over the sample size ratio n/p . It shows that the average convergence time of all estimators decreases as the sample size ratio n/p increases. In addition, it emphasizes the significant advantage of the proposed method in terms of computational efficiency, especially when there are few samples and the sample size ratio is small (i.e., $n/p \leq 1$), which is typically the case in large-scale problems. More experiments for real data are presented in Fig. 7 in the Appendix, demonstrating the computational efficiency of NewGLE.

5.4 Ablation Study

In this section, we provide some intuition on different features of our method. Specifically, we demonstrate the influence of the following three choices:

1. The free set mechanism.
2. The nonlinear conjugate gradients method for the constrained Newton problem.
3. The influence of the ϵ parameters.
4. The influence of the diagonal preconditioner.

We show the performance using the random planar graph example with $p = 1,000$ nodes and $n = 500$ samples. Each run includes the NewGLE method without one of the algorithmic features. The regularization parameter for the experiment is $\lambda = 0.25$, which is the best performing choice in terms of the error and F-score for this case. Figure 6 shows a typical instance of this example, demonstrating the behavior of NewGLE without each of the features. All the results are comparable in their accuracy (error and F-scores). It is clear that the most important feature in our study is the usage of nonlinear projected conjugate gradients for solving the inner Newton problem. Without it (i.e., using preconditioned gradient descent), the method is much slower. The second most important feature is the diagonal preconditioner, which approximately halves the cost in this example, and is significant also in other cases in our experience. The use of the free set imposes only a minor improvement in this example and is the less dominant feature. The use of ϵ slows




Figure 6: Ablation study of computational times versus algorithmic choices.

down the algorithm a bit, but it is important for theoretical stability, according to our theory in Theorem 2.

6 CONCLUSION

In this paper, we developed a proximal Newton method for minimizing a regularized maximum likelihood objective function for estimating a sparse graph Laplacian precision matrix from data distributed according to an LGMRF model. Our method is based on a second-order approximation of the smooth part of the objective function, while keeping the regularization and constraints intact. Methods of this type have been proven to be efficient for estimating standard GMRF models using ℓ_1 -norm penalty, while in our case, the estimation is also subject to Laplacian constraint. Furthermore, we used the nonconvex MCP function to promote the sparsity of the resulting graph without introducing a significant bias. Beyond the quadratic approximation itself, we used three algorithmic features to treat the Newton problem at each iteration efficiently: (1) we used a free/active set splitting, (2) we used a nonlinear conjugate gradient method, and (3) we used diagonal preconditioning. This resulted in an accurate and efficient estimation of the sparse weighted graph.

Numerical studies demonstrated the effectiveness of the proposed algorithm in terms of accuracy and computational complexity. The proposed method outperformed recent methods in terms of convergence time, while achieving similar or better F-scores and relative errors, particularly for problems with a small sample size ratio (i.e., $n/p < 1$). The plausible performance of the proposed method is particularly important in high-dimensional problems, where computational efficiency is crucial, and sample scarcity is a common issue.

ACKNOWLEDGEMENTS

This research was supported by the Israel Ministry of National Infrastructure, Energy, and Water Resources and by the ISRAEL SCIENCE FOUNDATION (grant No. 1148/22). This work was also supported in part by the Israeli Council for Higher Education (CHE) via the Data Science Research Center, Ben-Gurion University of the Negev, Israel.

References

- Anis, A., Gadde, A., and Ortega, A. (2016). Efficient sampling set selection for bandlimited graph signals using graph spectral proxies. *IEEE Transactions on Signal Processing*, 64(14):3775–3789.
- Armijo, L. (1966). Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific Journal of mathematics*, 16(1):1–3.
- Banerjee, O., El Ghaoui, L., and d’Aspremont, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *The Journal of Machine Learning Research*, 9:485–516.
- Banerjee, O., Ghaoui, L. E., d’Aspremont, A., and Natsoulis, G. (2006). Convex optimization techniques for fitting sparse gaussian graphical models. In *Proceedings of the 23rd international conference on Machine learning*, pages 89–96.
- Breheny, P. and Huang, J. (2011). Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *The annals of applied statistics*, 5(1):232.
- Cardoso, J. V. d. M. and Palomar, D. P. (2020). Learning undirected graphs in financial markets. In *Asilomar Conference on Signals, Systems, and Computers*, pages 741–745.
- Chen, J., Xu, P., Wang, L., Ma, J., and Gu, Q. (2018). Covariate adjusted precision matrix estimation via nonconvex optimization. In *International Conference on Machine Learning*, pages 922–931. PMLR.
- Chin, G. M., Nocedal, J., Olsen, P. A., and Rennie, S. J. (2013). Second order methods for optimizing convex matrix functions and sparse covariance clustering. *IEEE Transactions on Audio, Speech, and Language Processing*, 21(11):2244–2254.
- Chung, F. R. (1997). *Spectral graph theory*, volume 92. American Mathematical Soc.
- Dabush, L., Kroizer, A., and Routtenberg, T. (2023). State estimation in unobservable power systems via graph signal processing tools. *Sensors*, 23(3).
- Dai, Y.-H. and Yuan, Y. (1999). A nonlinear conjugate gradient method with a strong global convergence property. *SIAM Journal on optimization*, 10(1):177–182.
- d’Aspremont, A., Banerjee, O., and El Ghaoui, L. (2008). First-order methods for sparse covariance selection. *SIAM Journal on Matrix Analysis and Applications*, 30(1):56–66.
- Deka, D., Backhaus, S., and Chertkov, M. (2018). Structure learning in power distribution networks. *IEEE Trans. Control of Network Systems*, 5(3):1061–1074.
- Dong, X., Thanou, D., Frossard, P., and Vandergheynst, P. (2016). Learning Laplacian matrix in smooth graph signal representations. *IEEE Transactions on Signal Processing*, 64(23):6160–6173.
- Drayer, E. and Routtenberg, T. (2020). Detection of false data injection attacks in smart grids based on graph signal processing. *IEEE Systems Journal*, 14(2):1886–1896.
- Egilmez, H. E., Pavez, E., and Ortega, A. (2017). Graph learning from data under Laplacian and structural constraints. *IEEE Journal of Selected Topics in Signal Processing*, 11(6):825–841.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360.
- Finder, S. E., Treister, E., and Freifeld, O. (2022). Effective learning of a GMRF mixture model. *IEEE Access*, 10:7289–7299.
- Friedman, J., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441.
- Giudici, P. and Spelta, A. (2016). Graphical network models for international financial flows. *Journal of Business & Economic Statistics*, 34(1):128–138.
- Grotas, S., Yakoby, Y., Gera, I., and Routtenberg, T. (2019). Power systems topology and state estimation by graph blind source separation. *IEEE Transactions on Signal Processing*, 67(8):2036–2051.
- Halihal, M. and Routtenberg, T. (2022). Estimation of the admittance matrix in power systems under Laplacian and physical constraints. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5972–5976.
- Hassan-Moghaddam, S., Dhingra, N. K., and Jovanović, M. R. (2016). Topology identification of undirected consensus networks via sparse inverse covariance estimation. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 4624–4629.
- Holbrook, A. (2018). Differentiating the pseudo determinant. *Linear Algebra and its Applications*, 548:293–304.

- Horn, R. A. and Johnson, C. R. (2012). *Matrix analysis*. Cambridge university press.
- Hsieh, C.-J., Dhillon, I., Ravikumar, P., and Sustik, M. (2011). Sparse inverse covariance matrix estimation using quadratic approximation. *Neural information processing systems*, 24.
- Hsieh, C.-J., Sustik, M. A., Dhillon, I. S., and Ravikumar, P. (2014). QUIC: quadratic approximation for sparse inverse covariance estimation. *Journal of Machine Learning Research*, 15(1):2911–2947.
- Hsieh, C.-J., Sustik, M. A., Dhillon, I. S., Ravikumar, P. K., and Poldrack, R. (2013). BIG & QUIC: Sparse inverse covariance estimation for a million variables. *Advances in neural information processing systems*, 26.
- Hunter, D. R. and Li, R. (2005). Variable selection using MM algorithms. *Annals of statistics*, 33(4):1617.
- Koller, D. and Friedman, N. (2009). *Probabilistic graphical models: principles and techniques*. MIT press.
- Koyakumar, T., Yukawa, M., Pavez, E., and Ortega, A. (2023). Learning sparse graph with min-max concave penalty under Gaussian Markov random fields. *IEICE Transaction on Fundamentals of Electronics, Communications and Computer Sciences*, 106(1):23–34.
- Kroizer, A., Routtenberg, T., and Eldar, Y. C. (2022). Bayesian estimation of graph signals. *IEEE Transactions on Signal Processing*, 70:2207–2223.
- Kumar, S., Ying, J., de Miranda Cardoso, J. V., and Palomar, D. (2019). Structured graph learning via Laplacian spectral constraints. *Neural information processing systems*, 32.
- Lam, C. and Fan, J. (2009). Sparsistency and rates of convergence in large covariance matrix estimation. *The Annals of statistics*, 37(6B):4254–4278.
- Lee, J. D., Sun, Y., and Saunders, M. A. (2014). Proximal Newton-type methods for minimizing composite functions. *SIAM Journal on Optimization*, 24(3):1420–1443.
- Li, L. and Toh, K.-C. (2010). An inexact interior point method for l_1 -regularized sparse covariance selection. *Mathematical Programming Computation*, 2:291–315.
- Liu, T., Hoang, M. T., Yang, Y., and Pesavento, M. (2019). A block coordinate descent algorithm for sparse gaussian graphical model inference with Laplacian constraints. In *International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 236–240.
- Loh, P.-L. and Wainwright, M. J. (2013). Regularized M-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *Advances in Neural Information Processing Systems*, 26.
- Mazumder, R. and Hastie, T. (2012). Exact covariance thresholding into connected components for large-scale graphical lasso. *The Journal of Machine Learning Research*, 13(1):781–794.
- Milanfar, P. (2012). A tour of modern image filtering: New insights and methods, both practical and theoretical. *IEEE signal processing magazine*, 30(1):106–128.
- Narang, S. K. and Ortega, A. (2012). Perfect reconstruction two-channel wavelet filter banks for graph structured data. *IEEE Transactions on Signal Processing*, 60(6):2786–2799.
- Newman, M. (2010). *Networks: An Introduction*. Oxford University Press, Inc., New York, NY, USA.
- Oztoprak, F., Nocedal, J., Rennie, S., and Olsen, P. A. (2012). Newton-like methods for sparse inverse covariance estimation. *Neural information processing systems*, 25.
- Rolfs, B., Rajaratnam, B., Guillot, D., Wong, I., and Maleki, A. (2012). Iterative thresholding algorithm for sparse inverse covariance estimation. *Advances in Neural Information Processing Systems*, 25.
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. Chapman and Hall/CRC.
- Shalom, G., Treister, E., and Yavneh, I. (2023). pISTA: preconditioned iterative soft thresholding algorithm for graphical Lasso. *SIAM Journal on Scientific Computing*, 0(0):S445–S466.
- Shen, X., Pan, W., and Zhu, Y. (2012). Likelihood-based selection and sharp parameter estimation. *Journal of the American Statistical Association*, 107(497):223–232.
- Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. (2013). 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.
- Treister, E. and Turek, J. S. (2014). A block-coordinate descent approach for large-scale sparse inverse covariance estimation. In *Advances in neural information processing systems*, pages 927–935.
- Treister, E., Turek, J. S., and Yavneh, I. (2016). A multilevel framework for sparse optimization with application to inverse covariance estimation and logistic regression. *SIAM Journal on Scientific Computing*, 38(5):S566–S592.

- Tugnait, J. K. (2021). Sparse graph learning under Laplacian-related constraints. *IEEE Access*, 9:151067–151079.
- Vecchio, F., Miraglia, F., and Rossini, P. M. (2017). Connectome: Graph theory application in functional brain network architecture. *Clinical neurophysiology practice*, 2:206–213.
- Vieyra, M. V. (2022). Robust estimation of Laplacian constrained Gaussian graphical models with trimmed non-convex regularization. In *ICML 2022-Workshop on Principles of Distribution Shift*.
- Von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416.
- Ying, J., Cardoso, J. V. d. M., and Palomar, D. P. (2020a). Does the ℓ_1 -norm learn a sparse graph under Laplacian constrained graphical models? *arXiv preprint arXiv:2006.14925*.
- Ying, J., de M. Cardoso, J. V., and Palomar, D. (2020b). Nonconvex sparse graph learning under Laplacian constrained graphical model. *Advances in Neural Information Processing Systems*, 33:7101–7113.
- Ying, J., de Miranda Cardoso, J. V., and Palomar, D. (2021). Minimax estimation of Laplacian constrained precision matrices. In *International Conference on Artificial Intelligence and Statistics*, pages 3736–3744. PMLR.
- Zadorozhnyi, V. N. and Yudin, E. B. (2012). Structural properties of the scale-free barabasi-albert graph. *Automation and Remote Control*, 73:702–716.
- Zhang, C.-H. (2010a). Nearly unbiased variable selection under minimax concave penalty. *The Annals of statistics*, 38(2):894–942.
- Zhang, T. (2010b). Analysis of multi-stage convex relaxation for sparse regularization. *Journal of Machine Learning Research*, 11(3).
- Zhao, L., Wang, Y., Kumar, S., and Palomar, D. P. (2019). Optimization algorithms for graph Laplacian estimation via ADMM and MM. *IEEE Transactions on Signal Processing*, 67(16):4231–4244.
- Zhu, X., Lafferty, J., and Ghahramani, Z. (2003). Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In *ICML 2003 workshop on the continuum from labeled to unlabeled data in machine learning and data mining*, volume 3.
- Zibulevsky, M. and Elad, M. (2010). L1-L2 optimization in signal and image processing. *IEEE Signal Processing Magazine*, 27(3):76–88.
- Zou, H. and Li, R. (2008). One-step sparse estimates in nonconcave penalized likelihood models. *The Annals of statistics*, 36(4):1509–1533.

Checklist

- For all models and algorithms presented, check if you include:
 - A clear description of the mathematical setting, assumptions, algorithm, and/or model. Yes.
 - An analysis of the properties and complexity (time, space, sample size) of any algorithm. We provide an analysis of the convergence guarantees.
 - (Optional) Anonymized source code, with specification of all dependencies, including external libraries. Yes.
- For any theoretical claim, check if you include:
 - Statements of the full set of assumptions of all theoretical results. Yes.
 - Complete proofs of all theoretical results. Yes.
 - Clear explanations of any assumptions. Yes.
- For all figures and tables that present empirical results, check if you include:
 - The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). Yes. Instructions to reproduce the experiments are in the paper. Our MATLAB code is publicly available on Github.
 - All the training details (e.g., data splits, hyperparameters, how they were chosen). Hyperparameters are given in the paper, including instructions on the data generation.
 - A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). Yes, although the variance is quite low and hence omitted.
 - A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). Yes.
- If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - Citations of the creator If your work uses existing assets. Yes.

- (b) The license information of the assets, if applicable. Not Applicable.
 - (c) New assets either in the supplemental material or as a URL, if applicable. Not Applicable.
 - (d) Information about consent from data providers/curators. Not Applicable.
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. Not Applicable.
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. Not Applicable.
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. Not Applicable.
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. Not Applicable.

Efficient Graph Laplacian Estimation by Proximal Newton: Supplementary Materials

7 PROOF OF THEOREM 1

In this section we prove Theorem 1 (which is also repeated below). We note that this property was used in the works (Egilmez et al., 2017; Ying et al., 2020b) without explicit proof, and here we provide it.

Theorem 1 *The optimization problem*

$$\underset{\mathbf{L} \in \mathcal{F}}{\text{minimize}} \quad \text{Tr}(\mathbf{L}\mathbf{S}) - \log |\mathbf{L} + \mathbf{J}| + \rho(\mathbf{L}; \lambda),$$

where the feasible set is given by

$$\mathcal{F} = \left\{ \mathbf{L} \in \mathcal{S}^p \mid \mathbf{L}_{i,j} \leq 0, \forall i \neq j, i, j = 1, \dots, p, \mathbf{L}\mathbf{1} = \mathbf{0} \right\}.$$

is equivalent to the optimization problem stated in (3).

Proof. Since the objective functions in (3) and in (5) are identical (through (4)), it remains to prove that the feasible sets (solution spaces) of the two problems are equal under this objective function. It should be noted that there are two distinctions between the feasible sets \mathcal{L} and \mathcal{F} in problems (3) and (5), respectively. The first is that \mathcal{L} includes a positive semi-definite (PSD) constraint (i.e., $\mathbf{L} \in \mathcal{S}_+^p$), and the second is that \mathcal{L} includes a rank constraint (i.e., $\text{rank}(\mathbf{L}) = p - 1$).

To address the PSD constraint in problem (3), we observe that the constraints $\mathbf{L}\mathbf{1} = \mathbf{0}$, and $[\mathbf{L}]_{i,j} \leq 0$ for $i \neq j$ are present. Therefore, it follows that every matrix $\mathbf{L} \in \mathcal{F}$ is a diagonally dominant matrix with positive diagonal entries. Hence, requiring symmetry is sufficient to ensure that the minimizer of (5) is PSD as well (Horn and Johnson, 2012).

Next, we show that for any graph Laplacian $\mathbf{L} \in \mathcal{F}$, adding the constraint $\text{rank}(\mathbf{L}) = p - 1$ is either equivalent to $\text{rank}(\mathbf{L} + \mathbf{J}) = p$, or that the matrix $(\mathbf{L} + \mathbf{J})$ is non-singular under the considered objective function. To this end, we note that any $\mathbf{L} \in \mathcal{F}$ admits an eigenvalue decomposition $\mathbf{U}\Lambda_L\mathbf{U}^T$, where \mathbf{U} is $p \times p$ matrix of the eigenvectors of \mathbf{L} , and Λ_L is a diagonal eigenvalue matrix as follows

$$\Lambda_L = \text{diag} \left([0, \lambda_2, \dots, \lambda_p]^T \right), \tag{19}$$

where $0, \lambda_2, \dots, \lambda_p$ are the eigenvalues of \mathbf{L} . The rank of \mathbf{L} is the number of positive eigenvalues. In addition, the first eigenvector of \mathbf{L} is $\frac{1}{\sqrt{p}}\mathbf{1}$ and it corresponds to the eigenvalue $\lambda_1 = 0$. On the other hand, an eigendecomposition of \mathbf{J} can be $\mathbf{U}\Lambda_J\mathbf{U}^T$, where

$$\Lambda_J = \text{diag} \left([1, 0, \dots, 0]^T \right) \tag{20}$$

and the eigenvector of \mathbf{J} that is associated with the single nonzero eigenvalue is $\frac{1}{\sqrt{p}}\mathbf{1}$ —the same one as in \mathbf{L} . Hence, for the same eigenvector matrix \mathbf{U} , the eigenvalues of the matrix $\mathbf{L} + \mathbf{J}$ are $1, \lambda_2, \dots, \lambda_p$, and $\text{rank}(\mathbf{L} + \mathbf{J}) = \text{rank}(\mathbf{L}) + 1$. Hence, to conclude, if $\mathbf{L} \in \mathcal{F}$ and $(\mathbf{L} + \mathbf{J})$ is a non-singular matrix, then $\text{rank}(\mathbf{L}) = p - 1$.

Now we come to the last part of the proof. Assume by contradiction that the minimizer of (5) satisfies that the matrix $(\mathbf{L} + \mathbf{J})$ is singular, then the expression $-\log(|\mathbf{L} + \mathbf{J}|)$ in the objective approaches infinity. Thus, any other matrix $\mathbf{L} \in \mathcal{F}$ for which the matrix $(\mathbf{L} + \mathbf{J})$ is a non-singular matrix will yield a lower (finite) objective value in (5). Hence it can be concluded that the solution of (5) satisfies that the $(\mathbf{L} + \mathbf{J})$ matrix is non-singular, i.e., $\text{rank}(\mathbf{L} + \mathbf{J}) = p$ and hence $\text{rank}(\mathbf{L}) = p - 1$. \square

8 DERIVATION OF THE PARAMETERIZED NEWTON PROBLEM

First, we simplify the writing constraints of (12). We assume that the previous iterate $\mathbf{L}^{(t)}$ satisfies the constraints, and by definition $\mathbf{L}^{(t+1)}$ will satisfy them as well. To achieve that we also choose $\mathbf{L}^{(0)} \in \mathcal{F}$. Based on the definition of \mathcal{F} in (6), under the assumption that $\mathbf{L}^{(t)}$ satisfies the constraints, in order to obtain an update $\mathbf{L}^{(t+1)}$ that satisfy these constraints, we require the following conditions on the variable Δ :

- $\mathbf{L}^{(t)} + \Delta \in \mathcal{S}^p$. Hence, we require that $\Delta \in \mathcal{S}^p$.
- The off-diagonal elements should satisfy $[\mathbf{L}^{(t)} + \Delta]_{i,j} \leq 0, \forall i \neq j$. Hence, we require that $\Delta_{i,j} \leq -\mathbf{L}_{i,j}^{(t)}$.
- $(\mathbf{L}^{(t)} + \Delta) \mathbf{1} = \mathbf{0}$. Assuming that $\mathbf{L}^{(t)} \mathbf{1} = \mathbf{0}$, then we only need to require $\Delta \mathbf{1} = \mathbf{0}$.

By substituting these conditions in (12) and removing the constant term in the objective, we can rewrite the Newton problem from (12) as follows:

$$\begin{aligned} \underset{\Delta \in \mathcal{S}^p}{\text{minimize}} \quad & \text{Tr}(\Delta(\mathbf{S} - \mathbf{Q})) + \frac{1}{2} \text{Tr}(\Delta \mathbf{Q} \Delta \mathbf{Q}) + \frac{1}{2} \|\varepsilon \odot \Delta\|_F^2 + \rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) \\ \text{s.t.} \quad & \Delta_{i,j} \leq -\mathbf{L}_{i,j}^{(t)}, \forall i \neq j \\ & \Delta \mathbf{1} = \mathbf{0}. \end{aligned} \tag{21}$$

Next, we further simplify the problem by replacing the symmetric unknown matrix $\Delta \in \mathcal{S}^p$ with a vector $\delta \in \mathbb{R}^{p(p-1)/2}$ by using Definition 1 of the linear operator $\mathcal{P} : \mathbb{R}^{p(p-1)/2}$. According to this definition, given any $\mathbf{w} \geq 0$, $\mathcal{P}(\mathbf{w})$ satisfies $\sum_j \mathcal{P}(\mathbf{w})_{i,j} = 0$ for any $i \in \{1, \dots, p\}$, and $\mathcal{P}(\mathbf{w})_{i,j} = \mathcal{P}(\mathbf{w})_{j,i}$ for any $i \neq j$. In other words, $\mathcal{P}(\mathbf{w})$ is a Laplacian matrix for any \mathbf{w} which has nonnegative values. The simple example below illustrates the definition of \mathcal{P} . Let $\mathbf{w} = [w_1, w_2, w_3] \in \mathbb{R}^3$ be a nonnegative vector. Then

$$\mathcal{P}(\mathbf{w}) = \begin{bmatrix} \sum_{i=1,2} w_i & -w_1 & -w_2 \\ -w_1 & \sum_{i=1,3} w_i & -w_3 \\ -w_2 & -w_3 & \sum_{i=2,3} w_i \end{bmatrix}. \tag{22}$$

Since the Laplacian is a symmetric matrix, we can rewrite the sparse promoting term in (21) as:

$$\begin{aligned} \rho_{MCP}(\mathcal{P}(\mathbf{w}); \lambda) &\stackrel{(a)}{=} \sum_{i \neq j} MCP(\mathbf{L}_{ij}; \gamma, \lambda) \\ &\stackrel{(b)}{=} 2 \sum_{i > j} MCP(\mathbf{L}_{ij}; \gamma, \lambda) \\ &\stackrel{(c)}{=} 2 \sum_{k=1}^{p(p-1)/2} MCP(-w_k; \gamma, \lambda) \\ &\stackrel{(d)}{=} 2\tilde{\rho}_{mcp}(\mathbf{w}; \lambda), \end{aligned} \tag{23}$$

where $\mathbf{L} = \mathcal{P}(\mathbf{w})$. In (23), (a) is obtained by substituting (7), (b) is obtained from the symmetry of \mathbf{L} (i.e., since $\mathbf{L} \in \mathcal{S}^p$), and (c) is obtained by substituting (14). Finally, (d) is obtained since $\tilde{\rho}_{MCP}(\mathbf{w}; \lambda)$ defined in (16) is an even function. Noting that for $\mathbf{w}, \delta \in \mathbb{R}^{p(p-1)/2}$, such that $\mathbf{L} = \mathcal{P}(\mathbf{w})$ and $\Delta = \mathcal{P}(\delta)$, we can write:

- $\rho_{mcp}(\mathbf{L} + \Delta; \lambda) \stackrel{(23)}{=} 2\tilde{\rho}_{mcp}(\mathbf{w} + \delta; \lambda)$.
- $\Delta \in \left\{ \Delta \in \mathcal{S}^p \mid \Delta \mathbf{1} = \mathbf{0}, \Delta_{i,j} \leq -\mathbf{L}_{i,j}^{(k)}, \forall i \neq j \right\} \stackrel{(14)}{\iff} \delta \in \left\{ \delta \in \mathbb{R}^{p(p-1)/2} \mid \delta \geq -\mathbf{w} \right\}$.
- $\frac{1}{2} \|\varepsilon \odot \Delta\|_F^2 = \|\tilde{\varepsilon} \odot \delta\|_2^2$, which holds trivially because the diagonal of ε is 0 and it is symmetric.

Hence, (21), which is equivalent to (12), is also equivalent to (15), which is the problem we solve in practice in the inner Newton problem.

9 NONLINEAR PRECONDITIONED CONJUGATE GRADIENTS

Generally, iterates of projection methods are obtained by

$$\delta^{(t+1)} = \Pi_U \left\{ \delta^{(t)} + \zeta \mathbf{d}^{(t)} \right\}, \quad (24)$$

where $\delta^{(t)}$ is the current iterate, $\mathbf{d}^{(t)}$ is the descent (search) direction, $\zeta > 0$ is a step-size parameter, U is some feasible set, and $\Pi_U \{ \mathbf{v} \}$ is a projection of \mathbf{v} onto U . That is, using the ℓ_2 norm we have $\Pi \{ \mathbf{v} \} = \arg \min_{\mathbf{x}} \| \mathbf{x} - \mathbf{v} \|_2$. Since there is only a nonnegativity constraint in (15), the projection here is simple to calculate and is given by:

$$\Pi_{\mathbf{v} + \mathbf{w} \geq \mathbf{0}} \{ \mathbf{v} \} = \max \{ \mathbf{v}, -\mathbf{w} \}. \quad (25)$$

The step-size $\zeta > 0$ is computed by a linesearch procedure, and the descent direction $\mathbf{d}^{(t)}$ is given by

$$\mathbf{d}^{(t)} = \begin{cases} -\nabla_{\delta} f_N(\delta^{(t)}) & \text{if } t = 0 \\ -\nabla_{\delta} f_N(\delta^{(t)}) + \beta_t \mathbf{d}^{(t-1)} & \text{if } t > 0 \end{cases}, \quad (26)$$

where β_t is a scalar and $f_N(\delta^{(t)})$ is the objective defined in (15) at the point $\delta^{(t)}$. Since 1952, there have been many formulas for the scalar β_t , and in this paper, we use the DY method (Dai and Yuan, 1999), given by

$$\beta_t^{DY} = \max \left\{ 0, \frac{\| \nabla_{\delta} f_N^{(t)} \|_2^2}{\left(\nabla_{\delta} f_N^{(t)} - \nabla_{\delta} f_N^{(t-1)} \right)^T \mathbf{d}^{(t-1)}} \right\}. \quad (27)$$

One remarkable property of the DY method is that it produces a descent direction at each iteration and converges globally for convex problems in the case where linesearch is used (Dai and Yuan, 1999).

To apply the projected NLCG, the gradient of $f_N(\delta)$ (defined in (15)) is needed and is given by

$$\nabla_{\delta} f_N = \mathcal{P}^* (\nabla_{\Delta} f_N) = \mathcal{P}^* (\mathbf{S} - \mathbf{Q} + \mathbf{Q} \mathcal{P}(\delta) \mathbf{Q}) + 2\tilde{\varepsilon}^2 \odot \delta + 2\nabla_{\delta} \tilde{\rho}_{mcp}(\mathbf{w} + \delta; \lambda), \quad (28)$$

where $\mathcal{P}(\delta) = \Delta$, and \mathcal{P}^* is the adjoint of \mathcal{P} , that is

$$[\mathcal{P}^*(\mathbf{Y})]_k \triangleq Y_{ii} + Y_{jj} - Y_{ij} - Y_{ji} \quad (29)$$

for k as in Definition 1. Here we see the advantage of the proximal Newton approach. In (28) the gradient of $\tilde{\rho}$ is computed at $\mathbf{w} + \delta$, instead of \mathbf{w} alone if we were to naively majorize $\tilde{\rho}$ in the outer Newton problem. So, the inner solver is able to adapt to the curvature of $\tilde{\rho}$ and to the different sections in its definition.

9.1 Diagonal Preconditioning

The NLCG method that we use in this work is essentially an accelerated gradient-based method. To further accelerate it with rather minimal computational effort, a diagonal approximation to the Hessian can be used, acting as a preconditioner in the NLCG method. The aim of the diagonal preconditioner is to properly scale the gradient and speed up the rate of convergence of the iterative inner method for the Newton direction. In this paper, we suggest using the diagonal of the Hessian with respect to δ as a preconditioner with NLCG. Below we show the derivation of this preconditioner.

Our diagonal preconditioner matrix, \mathbf{D} , is defined as the diagonal of the Hessian of f_N . To ease the derivation, we ignore the ε term of (15) in this section. Following the Hessian of f at (11) and considering the linear operator \mathcal{P} as a matrix in $\mathbb{R}^{p^2 \times p(p-1)/2}$ (i.e., ignoring the vector-to-matrix reshaping in its definition), the Hessian of f_N is

$$\mathbf{H}_N = \mathcal{P}^T (\mathbf{Q} \otimes \mathbf{Q}) \mathcal{P} \in \mathbb{R}^{p(p-1)/2 \times p(p-1)/2}. \quad (30)$$

To get its (k, k) diagonal entry for our preconditioner, we multiply \mathbf{H}_N by a unit vector from both sides:

$$[\mathbf{D}]_{k,k} = [\mathbf{H}_N]_{k,k} = \mathbf{e}_k^T \mathbf{H}_N \mathbf{e}_k, \quad (31)$$

where $\mathbf{e}_k \in \mathbb{R}^{p(p-1)/2}$ is the unit vector of zeros, except the k -th entry, which equals 1. Equivalently to (31), we can write

$$[\mathbf{H}_N]_{k,k} = \text{Tr}(\mathcal{P}(\mathbf{e}_k)\mathbf{Q}\mathcal{P}(\mathbf{e}_k)\mathbf{Q}), \quad (32)$$

where the matrix $\mathcal{P}(\mathbf{e}_k) \in \mathbb{R}^{p \times p}$ is a matrix that is essentially the reshape of the k -th column of the operator \mathcal{P} , if we consider it as a matrix. If we look at the example of (22), then taking $k = 1$ corresponds to

$$\mathcal{P}(\mathbf{e}_1) = \begin{bmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Note that $k = 1$ in (22) influences only the 2×2 sub-matrix of the indices $(1, 2)$ in $\mathcal{P}(\mathbf{e}_1)$. More generally, $\mathcal{P}(\mathbf{e}_k)$ results in a matrix of zeros, except a 2×2 nonzero submatrix of the indices (i, j) where k corresponds to the pair (i, j) according to Definition 1. Hence, multiplying $\mathcal{P}(\mathbf{e}_k)\mathbf{Q}$, will result in a matrix of zeros except the i -th and j -th rows, which are equal to $\mathbf{q}_i - \mathbf{q}_j$ and $\mathbf{q}_j - \mathbf{q}_i$, respectively, where \mathbf{q}_i is the i -th column/row of (the symmetric) \mathbf{Q} . According to (32) we now have

$$\begin{aligned} [\mathbf{H}_N]_{k,k} &= \text{Tr}((\mathcal{P}(\mathbf{e}_k)\mathbf{Q})^2) \\ &= \text{Tr}\left(\begin{bmatrix} [\mathbf{Q}]_{i,i} - [\mathbf{Q}]_{j,i} & [\mathbf{Q}]_{i,j} - [\mathbf{Q}]_{j,j} \\ [\mathbf{Q}]_{j,i} - [\mathbf{Q}]_{i,i} & [\mathbf{Q}]_{j,j} - [\mathbf{Q}]_{j,i} \end{bmatrix}^2\right) \\ &= \left([\mathbf{Q}]_{i,i} + [\mathbf{Q}]_{j,j} - 2[\mathbf{Q}]_{i,j}\right)^2 = [\mathbf{D}]_{k,k}, \end{aligned} \quad (33)$$

where (i, j) corresponds to the index k according to Definition 1. In the last equality we used the symmetry of \mathbf{Q} , i.e., $[\mathbf{Q}]_{j,i} = [\mathbf{Q}]_{i,j}$. Below, Algorithm 2 summarizes the projected and preconditioned NLCG method for solving (12). Note that the linesearch in the last bullet of the algorithm is applied before the projection, hence taking a smaller steplength ξ means we effectively enlarge the Hessian approximation at each step.

Algorithm 2: Projected & preconditioned NLCG

Input: $\mathbf{L}^{(t)} = \mathcal{P}(\mathbf{w})$, \mathbf{Q} , $\text{Free}(\mathbf{L}^{(t)})$, λ

Result: Newton direction $\Delta^{(t)}$

Initialization: Set $\delta^{(0)} = \mathbf{0}$

while *Stopping criterion is not achieved* **do**

- Compute the preconditioner \mathbf{D} as in (33).
- Compute the gradient $\nabla_{\delta} f_N$ as in (28).
- Apply the preconditioner $\mathbf{g} = (\mathbf{D}^{-1} \cdot \nabla_{\delta} f_N)$.
- Zero non-free elements: $g_k \leftarrow 0 \forall k \in \text{Free}(\mathbf{L}^{(t)})$.
- Compute \mathbf{d} as in (26), with \mathbf{g} in the role of $\nabla_{\delta} f_N$.
- Update: $\delta^{(t+1)} = \mathbf{max}(\delta^{(t)} + \zeta \mathbf{d}, -\mathbf{w})$, where ζ achieved by linesearch.

end

10 CONVERGENCE GUARANTEES: PROOF OF THEOREM 2

In this section, we provide theoretical proof for the convergence of our method to a stationary point (Theorem 2). We note that the MCP penalty is nonconvex; hence one typically gets to a local minimum using a monotone series of iterations. Some results in our convergence proof follow (Hsieh et al., 2014) (e.g., Lemma 1 and Lemma 3). However, in the work (Hsieh et al., 2014), the ℓ_1 norm is used as a penalty, and the proof heavily relies on its convexity. Here, we had to adapt the proofs for these Lemmas to the concave MCP penalty, which is far from trivial and has not been done before, to the best of our knowledge.

This section is organized as follows. In subsection 10.1 we define the considered setup of proximal methods for MCP regularized penalties. In subsection 10.2 we provide some spectral bounds on our problem, which are needed for the proofs that follow. In subsection 10.3 we provide a lemma proving that our method is a fixed point iteration. In subsection 10.4 we prove that the objective value sufficiently decreases between the iterations, which is the main point needed for the convergence proof. Lastly, in subsection 10.5 we finalize the convergence proof given the previous auxiliary lemmas.

10.1 The Proximal Methods Setup

Let us write the problem (9) as follows

$$\underset{\mathbf{L} \in \mathcal{F}}{\text{minimize}} \quad F(\mathbf{L}) = f(\mathbf{L}) + \rho_{MCP}(\mathbf{L}; \lambda), \quad (34)$$

where f is the smooth part of the objective, defined in (10), \mathcal{F} and $\rho_{MCP}(\mathbf{L}; \lambda)$ are defined in (6) and (7), respectively. In the setup of general proximal methods, at each iteration t , we minimize the penalized and constrained quadratic objective, as follows:

$$\begin{aligned} \Delta^{(t)} = \underset{\Delta \in \mathcal{S}^p}{\text{argmin}} \quad & G(\Delta) = f(\mathbf{L}^{(t)}) + \langle \nabla f(\mathbf{L}^{(t)}), \Delta \rangle + \frac{1}{2} \langle \Delta, \Delta \rangle_{\mathbf{M}^{(t)}} + \rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) \\ \text{s.t.} \quad & (\mathbf{L}^{(t)} + \Delta) \in \mathcal{F}, \end{aligned} \quad (35)$$

where $\mathbf{M}^{(t)}$ is some positive definite matrix that is an approximation of the true Hessian, or the true Hessian itself. In particular, Proximal Gradient Descent uses $\mathbf{M} = \mathbf{I}$, while proximal Newton will have $\mathbf{M} = \nabla^2 f(\mathbf{L}) + \text{diag}(\text{vec}(\varepsilon))$, as in (12) and Algorithm 1. Once the quadratic problem (35) is solved, we apply linesearch as in (13), finding a step size $\alpha^{(t)}$ that ensures a reduction in the objective F . We now prove the convergence of this framework to our Laplacian estimation problem.

Optimality Conditions We can state that by the definition of \mathcal{F} our optimality condition is given by

$$\begin{cases} 0 \in \frac{\partial F(\mathbf{L}^*)}{\partial \mathbf{L}_{i,j}^*} + \mu_{i,j}, & i > j \\ \mathbf{L}_{j,i}^* = \mathbf{L}_{i,j}^* \geq 0, & i \neq j \\ \mathbf{L}_{i,i}^* = \sum_{j \neq i} -\mathbf{L}_{i,j}^* & i = 1, \dots, p \end{cases}, \quad (36)$$

where $\frac{\partial F(\mathbf{L}^*)}{\partial \mathbf{L}_{i,j}} = [\nabla f]_{i,j} + \frac{\partial MCP(\mathbf{L}_{i,j}^*; \gamma, \lambda)}{\partial \mathbf{L}_{i,j}}$ is the sub-differential of F w.r.t $\mathbf{L}_{i,j}$, $\mu_{i,j} \geq 0$ is an inequality Lagrange multiplier for the constraint $\mathbf{L}_{i,j} \leq 0$, and the sub-differential of the MCP function in (8) is given by

$$\frac{\partial MCP(x; \gamma, \lambda)}{\partial x} = \begin{cases} \in [-\lambda, \lambda] & x = 0 \\ \lambda \cdot \text{sign}(x) - \frac{x}{\gamma} & 0 < |x| \leq \gamma\lambda \\ 0 & |x| > \gamma\lambda \end{cases}. \quad (37)$$

If the conditions in (36) are fulfilled, then \mathbf{L}^* is a stationary point of (9). Since the MCP penalty is nonconvex, we expect (9) to have multiple minima, and we show that Algorithm 1 converges to one of them. We initialize our algorithm with $\mathbf{L}^{(0)} \in \mathcal{F}$ and keep $\mathbf{L}^{(t)} \in \mathcal{F}$ in all iterations.

10.2 Spectral Bounds

Let us first note that the $-\log \det(\mathbf{L} + \mathbf{J})$ term in (9) acts as a barrier function. That is, if $\mathbf{L} + \mathbf{J}$ gets closer to being singular, then $\det(\mathbf{L} + \mathbf{J}) \rightarrow 0$ and therefore, $F(\mathbf{L}) \rightarrow \infty$ following the log term. Therefore, we can say

that as long as $F(\mathbf{L}^{(t)})$ is bounded, there exists some $\beta_1 > 0$ for which $\mathbf{L}^{(t)} + \mathbf{J} \succeq \beta_1 \mathbf{I}$ for all t . Furthermore, by (Ying et al., 2021), we have that for every Laplacian matrix \mathbf{L}

$$\langle \mathbf{S}, \mathbf{L} \rangle = \langle \mathbf{v}, \mathbf{w} \rangle, \quad (38)$$

where $\mathbf{L} = \mathcal{P}\mathbf{w}$, $\mathbf{w} \geq 0$, and

$$\mathbf{v}_k = [\mathcal{P}^* \mathbf{S}]_k = \left[\mathcal{P}^* \frac{1}{n} \sum_{q=1}^n \mathbf{x}_q \mathbf{x}_q^T \right]_k = \frac{1}{n} \sum_{q=1}^n ([\mathbf{x}_q]_i - [\mathbf{x}_q]_i)^2 > 0 \quad (39)$$

holds with a very high probability, and \mathcal{P}^* is the adjoint of \mathcal{P} in (29). Following that, we have

$$\langle \mathbf{v}, \mathbf{w} \rangle \geq \min_j \{\mathbf{v}_j\} \sum_q \mathbf{w}_q = \frac{1}{2} \min_j \{\mathbf{v}_j\} \text{Tr}(\mathbf{L}). \quad (40)$$

Now, for any $\mathbf{L} \in \mathcal{F}$ such that $F(\mathbf{L}) \leq F(\mathbf{L}^{(0)})$ we have

$$\text{Tr}(\mathbf{L}\mathbf{S}) - \log |\mathbf{L}|_+ + \rho(\mathbf{L}; \lambda) \leq F(\mathbf{L}^{(0)}),$$

and since $\rho(\mathbf{L}; \lambda) \geq 0$ we can write:

$$\text{Tr}(\mathbf{L}\mathbf{S}) \leq F(\mathbf{L}^{(0)}) + \log |\mathbf{L}|_+.$$

Furthermore, using (40) and $\log |\mathbf{L}|_+ \leq (p-1) \log(\lambda_{\max}(\mathbf{L}))$, we get

$$\min_k \{\mathbf{v}_k\} \frac{1}{2} \lambda_{\max}(\mathbf{L}) \leq \min_k \{\mathbf{v}_k\} \frac{1}{2} \text{Tr}(\mathbf{L}) \leq F(\mathbf{L}^{(0)}) + (p-1) \log(\lambda_{\max}(\mathbf{L})). \quad (41)$$

Now, since in (41) the left-hand-side grows much faster than the right-hand-side, then there is an upper-bound on $\lambda_{\max}(\mathbf{L})$, depending on \mathbf{S} and $\mathbf{L}^{(0)}$, assuming (39). Hence, we can conclude that if the iterations $\{\mathbf{L}^{(t)}\}$ are monotonically decreasing starting from $\mathbf{L}^{(0)} \in \mathcal{F}$, there exist upper and lower bounds

$$\beta_1 \mathbf{I} \preceq \mathbf{L}^{(t)} + \mathbf{J} \preceq \beta_2 \mathbf{I}. \quad (42)$$

This also means that the level set $\mathcal{R} = \{\mathbf{L} \in \mathcal{F} : F(\mathbf{L}) \leq F(\mathbf{L}^{(0)})\}$ is compact, and the Hessian is bounded, i.e., $\nabla^2 f(\mathbf{L}) = (\mathbf{L} + \mathbf{J})^{-1} \otimes (\mathbf{L} + \mathbf{J})^{-1} \preceq \frac{1}{\beta_1^2} \mathbf{I} = \theta \mathbf{I}$, where $\theta > 0$. The last upper bound comes from the properties of eigenvalues of Kronecker products.

We summarize the conclusions above in the following lemma:

Lemma 1. *Assume that the iterations $\mathbf{L}^{(t)}$ are monotonically decreasing with respect to the objective in (34), starting from a feasible $\mathbf{L}^{(0)}$. Then there exist constants β_1, β_2 such that $\beta_1 \mathbf{I} \preceq \mathbf{L}^{(t)} + \mathbf{J} \preceq \beta_2 \mathbf{I}$, and the level-set $\mathcal{R} = \{\mathbf{L} \in \mathcal{F} : F(\mathbf{L}) \leq F(\mathbf{L}^{(0)})\}$ is compact. Also, the Hessian is bounded: $\nabla^2 f(\mathbf{L}) \preceq \frac{1}{\beta_1^2} \mathbf{I} = \theta \mathbf{I}$*

10.3 Fixed Point Iteration

The following lemma shows that if $\mathbf{L}^{(t)}$ is feasible, and the quadratic minimization (35) ends with $\Delta^{(t)} = 0$, then $\mathbf{L}^{(t)}$ is a minimum point of (34).

Lemma 2. *Assume that $\mathbf{L}^{(t)} \in \mathcal{F}$. If $\Delta^{(t)} = 0$ is a minimizer of $G(\Delta)$ in (35), then $\mathbf{L}^{(t+1)} = \mathbf{L}^{(t)}$ is a minimizer of $F(\mathbf{L})$ in (34).*

Proof. Since $\Delta^{(t)}$ is a minimizer of (35), the following optimality conditions are held

$$\begin{cases} 0 \in \frac{\partial G(\Delta^{(t)})}{\partial \Delta_{i,j}} + \mu_{i,j}, & i > j \\ \mathbf{L}_{j,i}^{(t)} + \Delta_{j,i}^{(t)} = \mathbf{L}_{i,j}^{(t)} + \Delta_{i,j}^{(t)} \geq 0, & i \neq j \\ \mathbf{L}_{i,i}^{(t)} + \Delta_{i,i}^{(t)} = \sum_{j \neq i} -(\mathbf{L}_{i,j}^{(t)} + \Delta_{i,j}^{(t)}) & i = 1, \dots, p \end{cases}, \quad (43)$$

where $\frac{\partial G(\Delta)}{\partial \Delta_{i,j}} = [\nabla f]_{i,j} + [\langle \mathbf{M}\Delta \rangle]_{i,j} + \frac{\partial MCP(\mathbf{L}_{i,j}^{(t)} + \Delta_{i,j}; \gamma, \lambda)}{\partial \Delta_{i,j}}$, and $\langle \mathbf{M}\Delta \rangle$ denotes a matrix vector multiplication of \mathbf{M} with the vectorized Δ , i.e., $\text{mat}(\mathbf{M} \cdot \text{vec}(\Delta))$. Placing $\Delta^{(t)} = 0$ in (43) yields the same conditions as in (36) regardless of the specific choice of the matrix $\mathbf{M} \succ 0$, hence $\mathbf{L}^{(t)}$ is a stationary point of (34). \square

10.4 Sufficient Decrease in Objective

We now show that we have a decrease in the objective F following the minimization in (35) and the linesearch in (13). This lemma is the key to our convergence proof.

Lemma 3. *Assume that the Hessian is bounded, $\nabla^2 f \preceq \theta \mathbf{I}$ for $\theta > 0$. Also assume that the iterates $\{\mathbf{L}^{(t)}\}$ are defined by (35) followed by a linesearch (13), starting from $\mathbf{L}^{(0)} \in \mathcal{F}$, with $\mathbf{M}^{(t)}$ satisfying $\mathbf{M}^{(t)} \succeq \lambda_{\min}^{\mathbf{M}} \mathbf{I} \succ \gamma^{-1} \mathbf{I}$, where $\lambda_{\min}^{\mathbf{M}}$ is a constant and γ is the MCP parameter. Then*

$$F(\mathbf{L}^{(t+1)}) - F(\mathbf{L}^{(t)}) \leq -c \cdot \|\Delta^{(t)}\|_F^2, \quad (44)$$

where $c > 0$. Furthermore, the linesearch parameter $\alpha^{(t)}$ in (13) can be chosen to be bounded away from zero, i.e., $\alpha^{(t)} \geq \alpha_{\min} > 0$.

Proof. For $\Delta^{(t)}$, the solution of (35), and some linesearch parameter $0 < \alpha < 1$ we have:

$$\begin{aligned} & \langle \nabla f(\mathbf{L}^{(t)}), \Delta \rangle + \frac{1}{2} \langle \Delta^{(t)}, \Delta^{(t)} \rangle_{\mathbf{M}^{(t)}} + \rho_{MCP}(\mathbf{L}^{(t)} + \Delta^{(t)}; \lambda) \\ & \leq \langle \nabla f(\mathbf{L}^{(t)}), \alpha \Delta^{(t)} \rangle + \frac{1}{2} \langle \alpha \Delta^{(t)}, \alpha \Delta^{(t)} \rangle_{\mathbf{M}^{(t)}} + \rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta^{(t)}; \lambda). \end{aligned} \quad (45)$$

That is because the objective in (35) is lower for $\Delta^{(t)}$ than for $\alpha \Delta^{(t)}$ (because $\Delta^{(t)}$ is assumed to be the minimum). Also, since \mathcal{F} is convex, then $\mathbf{L}^{(t)} + \alpha \Delta \in \mathcal{F}$. Thus, we have

$$(1 - \alpha) \langle \nabla f(\mathbf{L}^{(t)}), \Delta \rangle + (1 - \alpha^2) \frac{1}{2} \langle \Delta, \Delta \rangle_{\mathbf{M}^{(t)}} + \rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) - \rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta; \lambda) \leq 0. \quad (46)$$

After dividing by $(1 - \alpha) > 0$ we get:

$$\langle \nabla f(\mathbf{L}^{(t)}), \Delta \rangle \leq -(1 + \alpha) \frac{1}{2} \langle \Delta, \Delta \rangle_{\mathbf{M}^{(t)}} - \frac{1}{1 - \alpha} \left(\rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) - \rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta; \lambda) \right). \quad (47)$$

Now, by the definition of MCP in (8), for every x, y such that $\text{sign}(x) = \text{sign}(x + y)$ the function is twice differentiable in the interval $(x, x + y)$, and we can write

$$MCP(x + y) = MCP(x) + MCP'(x)y - \frac{1}{2\gamma}(\omega y)^2, \quad (48)$$

where $0 \leq \omega \leq 1$. Given that the feasible \mathcal{F} states that $\mathbf{L}^{(t)} \leq 0$ and $\mathbf{L}^{(t)} + \Delta \leq 0$ at all times, we can write

$$\rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) = \rho_{MCP}(\mathbf{L}^{(t)}; \lambda) + \langle \nabla \rho_{MCP}, \Delta \rangle - \frac{1}{2\gamma} \langle \Omega \odot \Delta, \Omega \odot \Delta \rangle \quad (49)$$

where $0 \leq \Omega \leq 1$ is a matrix of the ω values as in (48), and \odot is an element-wise product. Similarly, let

$$\rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta; \lambda) = \rho_{MCP}(\mathbf{L}^{(t)}; \lambda) + \langle \nabla \rho_{MCP}, \alpha \Delta \rangle - \frac{1}{2\gamma} \langle \Omega_\alpha \odot \alpha \Delta, \Omega_\alpha \odot \alpha \Delta \rangle, \quad (50)$$

where $0 \leq \Omega_\alpha \leq 1$. To show a decrease in the function F following the solution of (35) and a linesearch, we write

$$F(\mathbf{L}^{(t)} + \alpha \Delta^{(t)}) - F(\mathbf{L}^{(t)}) = f(\mathbf{L}^{(t)} + \alpha \Delta^{(t)}) + \rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta; \lambda) - f(\mathbf{L}^{(t)}) - \rho_{MCP}(\mathbf{L}^{(t)}; \lambda) \quad (51)$$

$$= \langle \nabla f(\mathbf{L}^{(t)}), \alpha \Delta^{(t)} \rangle + \frac{1}{2} \langle \alpha \Delta^{(t)}, \alpha \Delta^{(t)} \rangle_{\nabla^2 f(\mathbf{L}^{(t)})} + O(\|\Delta\|^3) \quad (52)$$

$$\begin{aligned} & + \rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta; \lambda) - \rho_{MCP}(\mathbf{L}^{(t)}; \lambda) \\ & \leq -\frac{\alpha(1 + \alpha)}{2} \langle \Delta, \Delta \rangle_{\mathbf{M}^{(t)}} + \frac{1}{2} \langle \alpha \Delta^{(t)}, \alpha \Delta^{(t)} \rangle_{\nabla^2 f(\mathbf{L}^{(t)})} + O(\|\Delta\|^3) \\ & \quad - \frac{\alpha}{1 - \alpha} \left(\rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) - \rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta; \lambda) \right) \\ & \quad + \rho_{MCP}(\mathbf{L}^{(t)} + \alpha \Delta; \lambda) - \rho_{MCP}(\mathbf{L}^{(t)}; \lambda). \end{aligned} \quad (53)$$

Equation (52) is a simple Taylor expansion, and the inequality in (53) stems from (47). At this point, we will focus on the last two lines of (53):

$$-\frac{\alpha}{1-\alpha} \left(\rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) - \rho_{MCP}(\mathbf{L}^{(t)} + \alpha\Delta; \lambda) \right) + \rho_{MCP}(\mathbf{L}^{(t)} + \alpha\Delta; \lambda) - \rho_{MCP}(\mathbf{L}^{(t)}; \lambda) \quad (54)$$

$$= \frac{1}{1-\alpha} \left((1-\alpha+\alpha)\rho_{MCP}(\mathbf{L}^{(t)} + \alpha\Delta; \lambda) - \alpha\rho_{MCP}(\mathbf{L}^{(t)} + \Delta; \lambda) \right) - \rho_{MCP}(\mathbf{L}^{(t)}; \lambda) \quad (55)$$

$$= \frac{1}{1-\alpha} \left((1-\alpha)\rho_{MCP}(\mathbf{L}^{(t)}; \lambda) - \frac{1}{2\gamma} \langle \alpha^2 \Omega_\alpha^2 - \alpha \Omega^2 \odot \Delta, \Delta \rangle \right) - \rho_{MCP}(\mathbf{L}^{(t)}; \lambda) \quad (56)$$

$$= \frac{1}{1-\alpha} \left(-\frac{1}{2\gamma} \langle \alpha^2 \Omega_\alpha^2 - \alpha \Omega^2 \odot \Delta, \Delta \rangle \right) \leq \frac{\alpha}{2\gamma} \langle \max\{\Omega_\alpha^2, \Omega^2\} \odot \Delta, \Delta \rangle \leq \frac{\alpha}{2\gamma} \langle \Delta, \Delta \rangle, \quad (57)$$

where Ω_α and Ω are defined in (49)-(50). Going back to the bound on the decrease in F in (53) we get:

$$F(\mathbf{L}^{(t)} + \alpha\Delta^{(t)}) - F(\mathbf{L}^{(t)}) \leq -\frac{\alpha(1+\alpha)}{2} \langle \Delta, \Delta \rangle_{\mathbf{M}^{(t)}} + \frac{1}{2} \langle \alpha\Delta^{(t)}, \alpha\Delta^{(t)} \rangle_{\nabla^2 f(\mathbf{L}^{(t)})} + \frac{\alpha}{2\gamma} \langle \Delta, \Delta \rangle + O(\|\Delta\|^3) \quad (58)$$

$$= -\frac{\alpha}{2} \langle \Delta, \Delta \rangle_{(\mathbf{M}^{(t)} - \gamma^{-1}\mathbf{I})} - \frac{\alpha^2}{2} \langle \Delta, \Delta \rangle_{(\mathbf{M}^{(t)} - \nabla^2 f(\mathbf{L}^{(t)}))} + O(\|\Delta\|^3). \quad (59)$$

Two conditions are needed to keep the above bound negative and guarantee a decrease in the objective. One is by having $\mathbf{M}^{(t)} \succ \gamma^{-1}\mathbf{I}$, which is easy to guarantee, e.g., via the choice of ε in our case. The other condition, assuming that we have $\mathbf{M}^{(t)} \succ \gamma^{-1}\mathbf{I}$, can be written as

$$\alpha(\mathbf{M} - \gamma^{-1}) + \alpha^2 \nabla^2 f(\mathbf{L}^{(t)}) \succ 0,$$

and then we have

$$F(\mathbf{L}^{(t)} + \alpha\Delta^{(t)}) - F(\mathbf{L}^{(t)}) \leq -\alpha(\lambda_{\min}^{\mathbf{M}} - \frac{1}{\gamma} - \alpha\theta) \|\Delta^{(t)}\|_F^2. \quad (60)$$

where θ is the upper bound on the Hessian, and $\lambda_{\min}^{\mathbf{M}} > \gamma^{-1}$ is the smallest eigenvalue of \mathbf{M} . This value is negative for any $0 \leq \alpha_{\min} < \frac{\lambda_{\min}^{\mathbf{M}} - \gamma^{-1}}{\theta}$. \square

The result above is relevant to Algorithm 1, where we have $\mathbf{M}^{(t)} = \nabla^2 f(\mathbf{L}^{(t)}) + \text{diag}(\epsilon^2)$. Hence, the second term in (59) almost vanishes if ϵ is small. On the other hand, ϵ allows us to make sure that the first term in (59) remains negative. We note that one can also monitor which entries are in the concave region of the MCP (smaller than $\gamma\lambda$ in magnitude) and set some value of $\varepsilon_{i,j}$ for them only, following (57).

The analysis above essentially is also suitable for Algorithm 2 if we set $\beta_t = 0$, i.e., use a projected quasi-Newton method with \mathbf{D} as a diagonal approximation to the Hessian. This is because in Algorithm 2, the Newton problem is also a smooth objective with an MCP prior and positivity constraints, and the change of variables to δ is only for convenience of symmetry and does not change the problem.

10.5 Convergence

Before stating the main convergence theorem, we prove the following auxiliary lemma:

Lemma 4. *Let $\{\mathbf{L}^{(t)}\}$ be series of points produced by (35) followed by the linesearch (13), with $\lambda_{\max}^{\mathbf{M}}\mathbf{I} \preceq \mathbf{M}^{(t)} \preceq \lambda_{\min}^{\mathbf{M}}\mathbf{I} \preceq 0$. Let $\{\mathbf{L}^{(t_j)}\}$ be any converging subseries of $\{\mathbf{L}^{(t)}\}$ and let $\bar{\mathbf{L}}$ be its limit. Then $\bar{\mathbf{L}}$ is a stationary point of F in (9).*

Proof. Since the sub-series $\{\mathbf{L}^{(t_j)}\}$ converges to $\bar{\mathbf{L}}$, then $\{F(\mathbf{L}^{(t_j)})\}$ converges to $F(\bar{\mathbf{L}})$. Therefore, $\{F(\mathbf{L}^{(t_j+1)}) - F(\mathbf{L}^{(t_j)})\} \rightarrow 0$, which implies following (60) that $\|\mathbf{L}^{(t_j)} - \mathbf{L}^{(t_j+1)}\|_F^2 \rightarrow 0$ and $\lim_{j \rightarrow \infty} \mathbf{L}^{(t_j+1)} = \bar{\mathbf{L}}$. According to (35) and the linesearch (13), $\mathbf{L}^{(t+1)} = \mathbf{L}^{(t)} + \alpha^{(t)}\Delta^{(t)}$ and because, $\|\mathbf{L}^{(t_j+1)} - \mathbf{L}^{(t_j)}\|_F^2 \rightarrow 0$ and $\alpha \geq \alpha_{\min} > 0$, this implies $\|\Delta^{(t)}\|_F^2 \rightarrow 0$. We know that $\Delta^{(t)}$ satisfies (43), and by taking the limit as $j \rightarrow \infty$ and using Lemma 2 we get that $\bar{\mathbf{L}}$ is a stationary point of (9). \square

Now, after providing the auxiliary lemmas above, we are ready to state the convergence theorem.




Figure 7: Convergence plots for the Lymph, Arabidopsis, Leukemia, and Hereditary BC datasets (left to right). The relative decrease is measured by $\frac{|F(\mathbf{L}^{(t)}) - F(\mathbf{L}^{(0)})|}{|F(\mathbf{L}^*) - F(\mathbf{L}^{(0)})|}$

Theorem 2. Let $\{\mathbf{L}^{(t)}\}$ be a series of points produced by a proximal quadratic minimization, constrained by \mathcal{F} with some SPD matrices $0 \prec \lambda_{\min}^{\mathbf{M}} \mathbf{I} \preceq \mathbf{M}^{(t)} \preceq \lambda_{\max}^{\mathbf{M}} \mathbf{I}$ as the Hessian (as in (35)), followed by the linesearch (13), starting from $\mathbf{L}^{(0)} \in \mathcal{F}$. Then any limit point $\bar{\mathbf{L}}$ of $\{\mathbf{L}^{(t)}\}$ is a stationary point of (9).

Proof. By Lemma 3, the series $F(\mathbf{L}^{(t)})$ is monotonically decreasing, hence it is also convergent as F is bounded from below. Since the level set \mathcal{R} is compact by Lemma 1, and $\mathbf{L}^{(t)}$ is bounded in \mathcal{R} , there is a converging subseries $\{\mathbf{L}^{(t_j)}\}$, which let $\bar{\mathbf{L}}$ be its limit point. By Lemma 4, $\bar{\mathbf{L}}$ is a stationary point of F in (9). Since F is continuous then $F(\mathbf{L}^{(t_j)}) \rightarrow F(\bar{\mathbf{L}})$. Since the limit of $\{\mathbf{L}^{(t)}\}$ and $F(\mathbf{L}^{(t)})$ equal to that of any of their corresponding subseries, then $\{\mathbf{L}^{(t)}\} \rightarrow \bar{\mathbf{L}}$ and $F(\mathbf{L}^{(t)}) \rightarrow F(\bar{\mathbf{L}})$. \square

11 RESULTS USING REAL DATA

In this subsection, we examine several gene expression datasets that are commonly used in research related to model selection, classification, and graphical models. These datasets are genetic regulatory networks, in which individual genes are represented as nodes in a graph, with the edges denoting the conditional dependencies between their expression profiles. It is important to note that these datasets can consist of thousands of genes while the number of samples is limited. This requires sparsity assumptions to enable the topology identification. Detailed information regarding the datasets, namely Lymph, Arabidopsis, Leukemia, and Hereditary BC datasets, can be found in (Li and Toh, 2010). As there is no ground truth in this case, we show convergence plots of the different methods in Fig. 7. It can be seen that for these real data simulations, the proposed method converges significantly faster than the other methods.