
A proximal Newton framework for composite minimization: Graph learning without Cholesky decompositions and matrix inversions

Quoc Tran Dinh
Anastasios Kyrillidis
Volkan Cevher

QUOC.TRANDINH@EPFL.CH
ANASTASIOS.KYRILLIDIS@EPFL.CH
VOLKAN.CEVHER@EPFL.CH

LIONS, École Polytechnique Fédérale de Lausanne, Switzerland

Abstract

We propose an algorithmic framework for convex minimization problems of composite functions with two terms: a self-concordant part and a possibly nonsmooth regularization part. Our method is a new proximal Newton algorithm with local quadratic convergence rate. As a specific problem instance, we consider sparse precision matrix estimation problems in graph learning. Via a careful dual formulation and a novel analytic step-size selection, we instantiate an algorithm within our framework for graph learning that avoids *Cholesky decompositions and matrix inversions*, making it attractive for parallel and distributed implementations.

1. Introduction

Sparse inverse covariance matrix estimation is a key step in graph learning. To understand the setup, let us consider learning a Gaussian Markov random field (GMRF) of p nodes/variables from a dataset $\mathcal{D} := \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$, where $\mathbf{x}_j \in \mathcal{D}$ is a p -dimensional random vector, drawn from the Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Let $\boldsymbol{\Theta} = \boldsymbol{\Sigma}^{-1}$ be the inverse covariance (or the precision) matrix for the model. To satisfy the conditional independencies with respect to the GMRF, $\boldsymbol{\Theta}$ must have zero in $\boldsymbol{\Theta}_{ij}$ corresponding to the absence of an edge between nodes i and j (Dempster, 1972).

To learn the underlying graph structure from $\boldsymbol{\Sigma}^{-1}$, one can use the empirical covariance matrix $\hat{\boldsymbol{\Sigma}}$. Unfortunately, this approach is fundamentally ill-posed since the empirical estimates converge to the true covariance at a $(1/\sqrt{m})$ -rate (Dempster, 1972). Hence, inferring

the true graph structure accurately requires an overwhelming number of samples. Unsurprisingly, we usually have fewer samples than the ambient dimension, compounding the difficulty of estimation.

While the possible GMRF structures are exponentially large, the most interesting graphs are rather simple with a sparse set of edges. Provable learning of such graphs can be achieved by ℓ_1 -norm regularization in the maximum log-likelihood estimation:

$$\boldsymbol{\Theta}^* \in \arg \min_{\boldsymbol{\Theta} > 0} \left\{ \underbrace{-\log \det(\boldsymbol{\Theta}) + \text{tr}(\hat{\boldsymbol{\Sigma}}\boldsymbol{\Theta})}_{=: f(\boldsymbol{\Theta})} + \rho \underbrace{\|\text{vec}(\boldsymbol{\Theta})\|_1}_{=: g(\boldsymbol{\Theta})} \right\}, \quad (1)$$

where $\rho > 0$ is a parameter to balance the fidelity error and the sparsity of the solution and vec is the vectorization operator. Here, $f(\boldsymbol{\Theta})$ corresponds to the empirical log-likelihood and $g(\boldsymbol{\Theta})$ is the sparsity-promoting term. Under this setting, the authors in (Ravikumar et al., 2011) prove that $m = \mathcal{O}(d^2 \log p)$ is sufficient for correctly estimating the GMRF, where d is the graph node-degree. Moreover, the above formulation still makes sense for learning other graph models, such as the Ising model, due to the connection of $f(\boldsymbol{\Theta})$ to the Bregman distance (Banerjee et al., 2008).

Numerical solution methods for solving problem (1) have been widely studied in the literature. For instance, in (Banerjee et al., 2008; Scheinberg & Rish, 2009; Scheinberg et al., 2010; Hsieh et al., 2011; Rolfs et al., 2012; Olsen et al., 2012) the authors proposed first order primal and dual approaches to (1) and used state-of-the-art structural convex optimization techniques such as coordinate descent methods and Lasso-based procedures. Alternatively, the authors in (Hsieh et al., 2011; Olsen et al., 2012) focused on the second order methods and, practically, achieved fast methods with a high accuracy. In (Scheinberg et al., 2010; Yuan, 2012), the authors studied alternating direction methods to solve (1), while the work in (Li & Toh, 2010) is based on interior point-type methods. Algo-

rithmic approaches where more structure is known a priori can be found in (Lu, 2010).

The complexity of the state-of-the-art approaches mentioned above is dominated by the Cholesky decomposition ($\mathcal{O}(p^3)$ in general), which currently creates an important scalability bottleneck. This decomposition appears mandatory since all these approaches employ a guess-and-check step-size selection procedures to ensure the iterates remain in the positive definite (PD) cone and the inversion of a $p \times p$ matrix, whose theoretical cost normally scales with the cost of $p \times p$ matrix multiplications ($\mathcal{O}(p^3)$ direct, $\mathcal{O}(p^{2.807})$ Strassen, and $\mathcal{O}(p^{2.376})$ Coppersmith-Winograd). The inversion operation is seemingly mandatory in the optimization of (1) since the calculation of the descent direction $\nabla f(\Theta_i) := -\Theta_i^{-1} + \widehat{\Sigma}$ requires it, and quadratic cost approximations to $f(\Theta)$ also need it. Via Cholesky decompositions, one can first check if the current solution satisfies the PD cone constraint and then recycle the decomposition for inversion for the next iteration.

Contributions: We propose a new proximal-Newton framework for solving the general problem of (1) by only assuming that $f(\cdot)$ is self-concordant. Our algorithm consists of two phases. In Phase 1, we apply a damped proximal Newton scheme with a new, analytic step-size selection procedure, and prove that our objective function always decreases at least a certain fixed amount. As a result, we avoid globalization strategies such as backtracking line-search or trust-region procedures in the existing methods. Moreover, our step-size selection is optimal in the sense that it cannot be improved without additional assumptions on the problem structure. In Phase 2, we simply apply the full step proximal-Newton iteration as we get into its provable quadratic convergence region which we can compute explicitly. Moreover, we do not require any additional assumption such as the uniform boundedness of the Hessian as in (Lee et al., 2012).

In the context of graph learning, we discuss a specific instance of our framework, which avoids Cholesky decompositions and matrix inversions *altogether*. Hence, the per iteration complexity of our approach is dominated by the cost of $p \times p$ matrix multiplications. This is because (i) our analytic step-size selection procedure ensures the positive definiteness of the iterates doing away with *global strategies* such as line-search which demands the objective evaluations (via Cholesky), and (ii) we avoid calculating the gradient explicitly, and hence matrix inversion by a careful dual formulation. As a result, our approach is attractive for distributed and parallel implementations.

Paper outline: In Section 2, we first recall some

fundamental concepts of convex optimization and self-concordant functions. Then, we describe the basic optimization set up and show the unique solvability of the problem. In Section 3 we outline our algorithmic framework and describe its analytical complexity. We also deal with the solution of the subproblems by applying the new dual approach in this section. Section 4 presents an application of our theory to graph selection problems. Experimental results on real graph learning problems can be found in Section 5.

2. Preliminaries

Basic definitions: We reserve lower-case and bold lower-case letters for scalar and vector representation, respectively. Upper-case bold letters denote matrices. Let $\mathbf{vec}: \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p^2}$ be the vectorization operator which maps a matrix to a single column, and $\mathbf{mat}: \mathbb{R}^{p^2} \rightarrow \mathbb{R}^{p \times p}$ is the inverse mapping of \mathbf{vec} which transforms a vector to a matrix. For a closed convex function f , we denote its domain by $\text{dom}(f)$, $\text{dom}(f) := \{x \in \mathbb{R}^n \mid f(x) < +\infty\}$.

Definition 2.1 (Self-concordant functions (Definition 2.1.1, pp. 12, (Nesterov & Nemirovski, 1994))). *A convex function $h: \mathbb{R} \rightarrow \mathbb{R}$ is (standard) self-concordant if $|h'''(x)| \leq 2h''(x)^{3/2}$, $\forall x \in \mathbb{R}$. Furthermore, a function $h: \mathbb{R}^n \rightarrow \mathbb{R}$ is self-concordant if, for any $t \in \mathbb{R}$, the function $\phi(t) := h(\mathbf{x} + t\mathbf{v})$ is self-concordant for all $\mathbf{x} \in \text{dom}(f)$ and $\mathbf{v} \in \mathbb{R}^n$.*

Let $h \in \mathcal{C}^3(\text{dom}(f))$ be a strictly convex and self-concordant function. For a given vector $\mathbf{v} \in \mathbb{R}^n$, the local norm around $\mathbf{x} \in \text{dom}(f)$ with respect to $h(\cdot)$ is defined as $\|\mathbf{v}\|_{\mathbf{x}} := (\mathbf{v}^T \nabla^2 h(\mathbf{x}) \mathbf{v})^{1/2}$ while the corresponding dual norm is given as $\|\mathbf{v}\|_{\mathbf{x}}^* := (\mathbf{v}^T \nabla^2 h(\mathbf{x})^{-1} \mathbf{v})^{1/2}$. Let $\omega: \mathbb{R} \rightarrow \mathbb{R}_+$ be a function defined as $\omega(t) := t - \ln(1+t)$ and $\omega_*: [0, 1] \rightarrow \mathbb{R}_+$ be a function defined as $\omega_*(t) := -t - \ln(1-t)$. The functions ω and ω_* are both nonnegative, strictly convex and increasing. Based on (Nesterov, 2004)[Theorems 4.1.7 & 4.1.8], we recall the following estimates:

$$\omega(\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}) + \nabla h(\mathbf{x})^T (\mathbf{y} - \mathbf{x}) + h(\mathbf{x}) \leq h(\mathbf{y}), \quad (2)$$

$$h(\mathbf{y}) \leq h(\mathbf{x}) + \nabla h(\mathbf{x})^T (\mathbf{y} - \mathbf{x}) + \omega_*(\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}}), \quad (3)$$

where (2) holds for all $\mathbf{x}, \mathbf{y} \in \text{dom}(f)$, and (3) holds for all $\mathbf{x}, \mathbf{y} \in \text{dom}(f)$ such that $\|\mathbf{y} - \mathbf{x}\|_{\mathbf{x}} < 1$.

Problem statement: In this paper, we consider the following structural convex optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \left\{ F(\mathbf{x}) \mid F(\mathbf{x}) := f(\mathbf{x}) + g(\mathbf{x}) \right\}, \quad (4)$$

where $f(\mathbf{x})$ is a *convex, self-concordant* function and $g(\mathbf{x})$ is a *proper, lower semicontinuous* and possibly

nonsmooth convex regularization term. It is easy to certify that problem (1) can be transformed into (4) by using the transformation $\mathbf{x} := \text{vec}(\Theta)$:

$$f(\mathbf{x}) := \begin{cases} -\log \det(\text{mat}(\mathbf{x})) + \text{tr}(\widehat{\Sigma} \text{mat}(\mathbf{x})), & \text{mat}(\mathbf{x}) \succ 0, \\ +\infty & \text{otherwise,} \end{cases}$$

$$g(\mathbf{x}) := \rho \|\mathbf{x}\|_1 \text{ and } n := p^2.$$

Proximity operator: A basic tool to handle nonsmooth convex functions is the proximity operator: let g be a proper lower semicontinuous, possibly nonsmooth and convex in \mathbb{R}^n . We denote by $\partial g(\mathbf{x})$ the subdifferential of g at \mathbf{x} . Let f be a self-concordant function and $\mathbf{x} \in \text{dom}(f)$ be fixed. We define $P_g^{\bar{\mathbf{x}}}(\mathbf{u}) := (\nabla^2 f(\bar{\mathbf{x}}) + \partial g)^{-1}(\mathbf{u})$ for $\mathbf{u} \in \mathbb{R}^n$. This operator is a nonexpansive mapping, i.e.,

$$\|P_g^{\bar{\mathbf{x}}}(\mathbf{u}) - P_g^{\bar{\mathbf{x}}}(\mathbf{v})\|_{\mathbf{x}} \leq \|\mathbf{u} - \mathbf{v}\|_{\mathbf{x}}^*, \quad \forall \mathbf{u}, \mathbf{v}. \quad (5)$$

Unique solvability of the problem: We generalize the result in (Hsieh et al., 2011) to show that problem (4) is uniquely solvable.

Lemma 2.2. *For some $\mathbf{x} \in \text{dom}(F)$, let $\lambda(\mathbf{x}) := \|\nabla f(\mathbf{x}) + \mathbf{v}\|_{\mathbf{x}}^* < 1$ for $\mathbf{v} \in \partial g(\mathbf{x})$. Then the solution \mathbf{x}^* of (4) exists and is unique.*

The proof of this lemma can be done similarly as Theorem 4.1.11, pp. 187 in (Nesterov, 2004). For completeness, we provide it in the supplementary document.

3. Two-phase proximal Newton method

Our algorithmic framework is simply a proximal-Newton method which generates an iterative sequence $\{\mathbf{x}^k\}_{k \geq 0}$ starting from $\mathbf{x}^0 \in \text{dom}(F)$. The new point \mathbf{x}^{k+1} is computed as $\mathbf{x}^{k+1} = \mathbf{x}^k + \alpha_k \mathbf{d}^k$, where $\alpha_k \in (0, 1]$ is a step size and \mathbf{d}^k is the proximal-Newton-type direction as the solution to the subproblem:

$$\min_{\mathbf{d}} \{Q(\mathbf{d}; \mathbf{x}^k) + g(\mathbf{x}^k + \mathbf{d})\}. \quad (Q(\mathbf{x}^k))$$

Here, $Q(\mathbf{d}; \mathbf{x}^k)$ is the following quadratic surrogate of the function f around \mathbf{x}^k :

$$Q(\mathbf{d}; \mathbf{x}^k) := f(\mathbf{x}^k) + \nabla f(\mathbf{x}^k)^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}^k) \mathbf{d}. \quad (6)$$

We denote \mathbf{d}^k the unique solution of $Q(\mathbf{x}^k)$. The optimality condition for $Q(\mathbf{x}^k)$ is written as follows:

$$\mathbf{0} \in \partial g(\mathbf{x}^k + \mathbf{d}^k) + \nabla f(\mathbf{x}^k) + \nabla^2 f(\mathbf{x}^k) \mathbf{d}^k. \quad (7)$$

Fixed-point characterization. For given $\mathbf{x} \in \text{dom}(F)$, if we define $S(\mathbf{x}) := \nabla^2 f(\mathbf{x}) \mathbf{x} - \nabla f(\mathbf{x})$ then the unique solution \mathbf{d}^k of $Q(\mathbf{x}^k)$ can be computed as

$$\mathbf{d}^k := \left(P_g^{\mathbf{x}^k} \circ S \right) (\mathbf{x}^k) - \mathbf{x}^k = -(\mathbb{I} - R_g)(\mathbf{x}^k). \quad (8)$$

Here, $R_g(\cdot) := (P_g^{\mathbf{x}} \circ S)(\cdot) \equiv P_g^{\mathbf{x}}(S(\cdot))$. The next lemma shows that the fixed point of R_g is the unique solution of (4). The proof is straightforward, and is omitted.

Lemma 3.1. *Let R_g be a mapping defined by (8). Then \mathbf{x}^* is the unique solution of (4) if and only if \mathbf{x}^* is the fixed-point of R_g , i.e., $\mathbf{x}^* = R_g(\mathbf{x}^*)$.*

Lemma 3.1 suggests that we can generate an iterative sequence based on the fixed-point principle. Under certain assumptions, one can ensure that R_g is contractive and the sequence generated by this scheme is convergent. Hence, we characterize this below.

3.1. Full-step proximal-Newton scheme

Here, we show that if we start sufficiently close to the solution \mathbf{x}^* , then we can compute the next iteration \mathbf{x}^{k+1} with full-step $\alpha_{k+1} = 1$, i.e.,

$$\mathbf{x}^{k+1} := \mathbf{x}^k + \mathbf{d}^k, \quad (9)$$

where \mathbf{d}^k is the unique solution to $Q(\mathbf{x}^k)$. We call this scheme the **full-step proximal Newton (FPN)** scheme. For any $k \geq 0$, let us define

$$\lambda_k := \|\mathbf{x}^{k+1} - \mathbf{x}^k\|_{\mathbf{x}^k}. \quad (10)$$

We refer to this quantity as the *proximal Newton decrement*. The following theorem establishes the local quadratic convergence of the FPN scheme (9).

Theorem 3.2. *For a given \mathbf{x}^k , let \mathbf{x}^{k+1} be the point generated by the full-step proximal Newton scheme (9) and λ_k be defined by (10). Then, if $\lambda_k < 1 - \frac{1}{\sqrt{2}} \approx 0.292893$, it holds that*

$$\lambda_{k+1} \leq (1 - 4\lambda_k + 2\lambda_k^2)^{-1} \lambda_k^2. \quad (11)$$

Consequently, the sequence $\{\mathbf{x}^k\}_{k \geq 0}$ generated by the FPN scheme (9) starting from $\mathbf{x}^0 \in \text{dom}(F)$ such that $\lambda_0 \leq \sigma \leq \bar{\sigma} := \frac{5 - \sqrt{17}}{4} \approx 0.219224$, locally converges to the unique solution of (4) at a quadratic rate.

The proof of Theorem 3.2 can be found in the supplementary document.

3.2. Damped proximal Newton scheme

We now establish that, with an appropriate choice of the step-size $\alpha \in (0, 1]$, the iterative sequence $\{\mathbf{x}^k\}_{k \geq 0}$ generated by the *damped proximal Newton* scheme

$$\mathbf{x}^{k+1} := \mathbf{x}^k + \alpha_k \mathbf{d}^k \quad (12)$$

is a decreasing sequence, i.e., $F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) - \omega(\sigma)$ whenever $\lambda_k \geq \sigma$, where $\sigma > 0$ is fixed. First, we show the following property for the new iteration \mathbf{x}^{k+1} .

Lemma 3.3. *Suppose that \mathbf{x}^{k+1} is a point generated by (12). Then, we have*

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) - [\alpha_k \lambda_k^2 - \omega^*(\alpha_k \lambda_k)], \quad (13)$$

provided that $\alpha_k \lambda_k < 1$.

Proof. Let $\mathbf{y}^k = \mathbf{x}^k + \mathbf{d}^k$, where \mathbf{d}^k is the unique solution of $\mathcal{Q}(\mathbf{x}^k)$. It follows from the optimality condition of (7) that there exists $\mathbf{v}_k \in \partial g(\mathbf{y}^k)$ such that

$$\mathbf{v}_k = -\nabla f(\mathbf{x}^k) - \nabla^2 f(\mathbf{x}^k)(\mathbf{y}^k - \mathbf{x}^k). \quad (14)$$

Since f is self-concordant, by (3), for any \mathbf{x}^{k+1} such that $\lambda_k = \|\mathbf{x}^{k+1} - \mathbf{x}^k\|_{\mathbf{x}^k} < 1$ we have

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) + \nabla f(\mathbf{x}^k)^T (\mathbf{x}^{k+1} - \mathbf{x}^k) + \omega^*(\|\mathbf{x}^{k+1} - \mathbf{x}^k\|_{\mathbf{x}^k}) + g(\mathbf{x}^{k+1}) - g(\mathbf{x}^k). \quad (15)$$

Since g is convex, $\alpha \in [0, 1]$, by using (14) we have

$$\begin{aligned} g(\mathbf{x}^{k+1}) - g(\mathbf{x}^k) &= g((1 - \alpha_k)\mathbf{x}^k + \alpha_k \mathbf{y}^k) - g(\mathbf{x}^k) \\ &\leq \alpha_k [g(\mathbf{y}^k) - g(\mathbf{x}^k)] \\ &\leq \alpha \mathbf{v}_k^T (\mathbf{y}^k - \mathbf{x}^k) \\ &= \alpha_k \mathbf{v}_k^T \mathbf{d}^k \\ &\stackrel{(14)}{=} -\alpha_k \nabla f(\mathbf{x}^k)^T \mathbf{d}^k - \alpha_k \|\mathbf{d}^k\|_{\bar{\mathbf{x}}^k}^2. \end{aligned} \quad (16)$$

Now, substituting (16) into (15) and noting that $\mathbf{x}^{k+1} - \mathbf{x}^k = \alpha_k \mathbf{d}^k$ we obtain the following result

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) + \omega^*(\alpha \|\mathbf{d}^k\|_{\mathbf{x}^k}) - \alpha_k (\mathbf{d}^k)^T \nabla^2 f(\mathbf{x}^k) \mathbf{d}^k = F(\mathbf{x}^k) - [\alpha_k \lambda_k^2 - \omega^*(\alpha_k \lambda_k)], \quad (17)$$

which is indeed (13), provided that $\alpha_k \lambda_k < 1$. \square

The following theorem provides an explicit formula for the step size α_k .

Theorem 3.4. *Let \mathbf{x}^{k+1} be a new point generated by the scheme (12) and λ_k be defined by (10). Then, if we choose $\alpha_k := (1 + \lambda_k)^{-1} \in (0, 1]$ then*

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) - \omega(\lambda_k). \quad (18)$$

Moreover, the step $\alpha_k = (1 + \lambda_k)^{-1}$ is optimal.

Proof. By the choice of α_k , we have $\alpha_k \lambda_k = (1 + \lambda_k)^{-1} \lambda_k < 1$. By using the estimate (13) we have

$$F(\mathbf{x}^{k+1}) \leq F(\mathbf{x}^k) - (1 + \lambda_k)^{-1} \lambda_k^2 + \omega^*((1 + \lambda_k)^{-1} \lambda_k).$$

Since $\frac{t^2}{1+t} - \omega^*(\frac{t}{1+t}) = \omega(t)$ for any $t > 0$, the last inequality implies (18). Finally, we note that the function $\varphi(\alpha) := \alpha \lambda (1 + \lambda) + \ln(1 - \alpha \lambda)$ is maximized at $\alpha_k = (1 + \lambda_k)^{-1}$, showing that α_k is optimal. \square

Theorem 3.4 shows that the damped proximal Newton scheme generates a new point \mathbf{x}^{k+1} that decreases F of (4) at least $\omega(\sigma)$ at each iteration, whenever $\lambda_k \geq \sigma$.

Quadratic convergence: Similar to the full-step proximal-Newton scheme (9), we can also show the quadratic convergence of the damped proximal-Newton scheme (12). This statement is summarized in the following theorem.

Theorem 3.5. *For a given $\mathbf{x}^k \in \text{dom}(F)$, let \mathbf{x}^{k+1} be a new point generated by the scheme (12) with $\alpha_k := (1 + \lambda_k)^{-1}$. Then, if $\lambda_k < 1 - \frac{1}{\sqrt{2}}$, it holds that*

$$\lambda_{k+1} \leq (1 - 4\lambda_k + 2\lambda_k^2)^{-1} (1 - \lambda_k^2) \lambda_k^2. \quad (19)$$

Hence, the sequence $\{\mathbf{x}^k\}_{k \geq 0}$ generated by (12) with $\alpha_k = (1 + \lambda_k)^{-1}$ starting from $\mathbf{x}^0 \in \text{dom}(F)$ such that $\lambda_0 \leq \sigma \leq \bar{\sigma} := 0.221876$ locally converges to \mathbf{x}^* , the unique solution of (4) at a quadratic rate.

The proof of Theorem 3.5 can be found in the supplementary document. Note that the value $\bar{\sigma}$ in Theorem 3.5 is larger than in Theorem 3.2. However, both values are not tight.

3.3. The algorithm pseudocode

As proved by Theorems 3.4 and 3.5, we can use the damped proximal-Newton scheme to build the algorithm. Now, we present a two-phase proximal-Newton algorithm. We first select a constant $\sigma \in (0, \bar{\sigma}]$. At each iteration, we compute the new point \mathbf{x}^{k+1} by using the damped proximal Newton scheme (12) until we get $\lambda_k \leq \sigma$. Then, we switch to the full-step Newton scheme and perform it until the convergence is achieved. These steps are described in Algorithm 1.

Note that the radius σ of the quadratic convergence region in Algorithm 1 can be fixed at its upper bound $\bar{\sigma}$. The maximum number of iterations j_{\max} and k_{\max} can also be specified, if necessary.

3.4. Iteration-complexity analysis

We analyze the complexity of Algorithm 1 by separating Phase 1 and Phase 2. This analysis is summarized in the following theorem.

Theorem 3.6. *The maximum number of iterations required in Phase 1 does not exceed $j_{\max} := \left\lfloor \frac{F(\mathbf{x}^0) - F(\mathbf{x}^*)}{\omega(\sigma)} \right\rfloor + 1$, where \mathbf{x}^* is the unique solution of (4). The maximum number of iterations required in Phase 2 to obtain $\lambda_k \leq \varepsilon$ does not exceed $k_{\max} := O(\ln \ln(\frac{c}{\varepsilon}))$, where $c := (1 - 4\sigma + 2\sigma^2)^{-1} > 0$.*

Proof. Since $\lambda_j \geq \sigma$ for all $j \geq 0$ in Phase 1, it follows from Theorem 3.4 that $F(\mathbf{x}^{j+1}) \leq F(\mathbf{x}^j) - \omega(\sigma)$. By induction we have $F(\mathbf{x}^*) \leq F(\mathbf{x}^{j_{\max}}) \leq F(\mathbf{x}^0) - j_{\max} \omega(\sigma)$. This implies that $j_{\max} \leq \left\lfloor \frac{F(\mathbf{x}^0) - F(\mathbf{x}^*)}{\omega(\sigma)} \right\rfloor + 1$.

Algorithm 1 (*Proximal Newton algorithm*)

Initialization:

Require a starting point $\mathbf{x}^0 \in \text{dom}(F)$ and a constant $\sigma \in (0, \bar{\sigma}]$, where $\bar{\sigma} := \frac{(5-\sqrt{17})}{4} \approx 0.219224$.

Phase 1: (*Damped proximal Newton iterations*).

for $j = 0$ **to** j_{\max} **do**

1. Compute the proximal-Newton search direction \mathbf{d}^j as the unique solution of $\mathcal{Q}(\mathbf{x}^j)$.
2. Compute $\lambda_j := \|\mathbf{d}^j\|_{\mathbf{x}^j}$.
3. **if** $\lambda_j \leq \sigma$ **then** terminate Phase 1.
4. Otherwise, update the next iteration $\mathbf{x}^{j+1} := \mathbf{x}^j + \alpha_j \mathbf{d}^j$, where $\alpha_j := (1 + \lambda_j)^{-1}$.

end for

Phase 2: (*Full-step proximal Newton iterations*).

Set $\mathbf{x}^0 := \mathbf{x}^j$ from Phase 1 and choose a desired accuracy $\varepsilon > 0$.

for $k = 0$ **to** k_{\max} **do**

1. Compute the proximal-Newton direction \mathbf{d}^k as the unique solution of $\mathcal{Q}(\mathbf{x}^k)$.
2. Compute $\lambda_k := \|\mathbf{d}^k\|_{\mathbf{x}^k}$.
3. **if** $\lambda_k \leq \varepsilon$ **then** terminate Phase 2.
4. Otherwise, update $\mathbf{x}^{k+1} := \mathbf{x}^k + \mathbf{d}^k$.

end for

$F(\mathbf{x}^*)/\omega(\sigma)$. Hence, we can fix

$$j_{\max} := \left\lceil \frac{F(\mathbf{x}^0) - F(\mathbf{x}^*)}{\omega(\sigma)} \right\rceil + 1.$$

Let $c := (1 - 4\sigma + 2\sigma^2)^{-1} > 0$. By induction, it follows from Theorem 3.2 that we have $\lambda_k \leq (c)^{2^k-1} \lambda_0^{2^k} \leq (c)^{2^k-1} \sigma^{2^k}$. In order to ensure $\lambda_k \leq \varepsilon$, we require $(c)^{2^k-1} \sigma^{2^k} \leq \varepsilon$, which leads to $k \leq O(\ln \ln(c/\varepsilon))$. Hence, we can show that $k_{\max} := O(\ln \ln(c/\varepsilon))$. \square

We note that we do not use j_{\max} as a stopping criterion of Phase 1 of Algorithm 1. In practice, we only need an upper bound of this quantity. If we fix σ at $\bar{\sigma}$ then $c \approx 4.561553$ and the complexity of Phase 2 becomes $O(\ln \ln(\frac{4.5}{\varepsilon}))$.

3.5. Dual solution approach of the subproblem

In this subsection we consider a specific instance of g : $g(\mathbf{x}) := \rho \|\mathbf{x}\|_1$. First, we derive a dual formulation of the convex subproblem $\mathcal{Q}(\mathbf{x}^k)$. For notational convenience, we let $\mathbf{q}_k := \nabla f(\mathbf{x}^k)$, $\mathbf{H}_k := \nabla^2 f(\mathbf{x}^k)$. Then, the convex subproblem $\mathcal{Q}(\mathbf{x}^k)$ can be written equivalently as

$$\min_{\mathbf{y} \in \mathbb{R}^n} \left\{ \frac{1}{2} \mathbf{y}^T \mathbf{H}_k \mathbf{y} + (\mathbf{q}_k - \mathbf{H}_k \mathbf{x}^k)^T \mathbf{y} + \rho \|\mathbf{y}\|_1 \right\}. \quad (20)$$

By using the min-max principle, we can write (20) as

$$\max_{\|\mathbf{u}\|_{\infty} \leq 1} \min_{\mathbf{y} \in \mathbb{R}^n} \left\{ \frac{1}{2} \mathbf{y}^T \mathbf{H}_k \mathbf{y} + (\mathbf{q}_k - \mathbf{H}_k \mathbf{x}^k)^T \mathbf{y} + \rho \mathbf{u}^T \mathbf{y} \right\}. \quad (21)$$

Solving the inner minimization in (21) we obtain:

$$\min_{\|\mathbf{u}\|_{\infty} \leq 1} \left\{ \frac{1}{2} \mathbf{u}^T \mathbf{H}_k^{-1} \mathbf{u} + \tilde{\mathbf{q}}_k^T \mathbf{u} \right\}, \quad (22)$$

where $\tilde{\mathbf{q}}_k := \frac{1}{\rho} (\mathbf{H}_k^{-1} \mathbf{q}_k - \mathbf{x}^k)$. Note that the objective function $\varphi_k(\mathbf{u}) := \frac{1}{2} \mathbf{u}^T \mathbf{H}_k^{-1} \mathbf{u} + \tilde{\mathbf{q}}_k^T \mathbf{u}$ of (22) is strongly convex. One can apply the fast projected gradient methods with linear convergence rate in (Nesterov, 2007; Beck & Teboulle, 2009) for solving this problem.

In order to recover the solution of the primal subproblem $\mathcal{Q}(\mathbf{x}^k)$, we note that the solution of the parametric minimization problem in (21) is given by $\mathbf{y}_k^*(\mathbf{u}) := \mathbf{x}^k - \mathbf{H}_k^{-1} (\mathbf{q}_k + \rho \mathbf{u})$. Let \mathbf{u}_k^* be the optimal solution of (22). We can recover the primal proximal-Newton search direction \mathbf{d}^k of the subproblem $\mathcal{Q}(\mathbf{x}^k)$ as

$$\mathbf{d}^k = -\nabla^2 f(\mathbf{x}^k)^{-1} [\nabla f(\mathbf{x}^k) + \rho \mathbf{u}_k^*]. \quad (23)$$

To compute the quantity $\lambda_k := \|\mathbf{d}^k\|_{\mathbf{x}^k}$ in Algorithm 1, we use (23) such that

$$\lambda_k = \|\nabla f(\mathbf{x}^k) + \rho \mathbf{u}_k^*\|_{\mathbf{x}^k}^*. \quad (24)$$

Note that computing λ_k by (24) requires the inverse of the Hessian matrix $\nabla^2 f(\mathbf{x}^k)$.

4. Application to graph selection

In this section, we customize the theory framework of Algorithm 1 by using only Phase 1 to solve the graph selection problem (1).

Quantification: For clarity, we retain the matrix variable form as presented in (1). We note that $f(\Theta)$ is a self-concordant convex function, while $g(\Theta)$ is a proper, lower semicontinuous and nonsmooth convex function. Thus, our theory presented above can be applied to (1). Given the current estimate $\Theta_i \succ 0$, we have $\nabla f(\Theta_i) = \widehat{\Sigma} - \Theta_i^{-1}$ and $\nabla^2 f(\Theta_i) = \Theta_i^{-1} \otimes \Theta_i^{-1}$. Under this setting, the dual subproblem (22) becomes:

$$\mathbf{U}_i^* = \arg \min_{\|\text{vec}(\mathbf{U})\|_{\infty} \leq 1} \left\{ \frac{1}{2} \text{tr}((\Theta_i \mathbf{U})^2) + \text{tr}(\tilde{\mathbf{Q}}_i \mathbf{U}) \right\}, \quad (25)$$

where $\tilde{\mathbf{Q}}_i := \rho^{-1} [\Theta_i \widehat{\Sigma} \Theta_i - 2\Theta_i]$. Given the dual solution \mathbf{U}_i^* of (25), the primal proximal-Newton search direction (i.e. the solution of $\mathcal{Q}(\mathbf{x}^k)$) is computed as

$$\Delta_i := - \left((\Theta_i \widehat{\Sigma} - \mathbb{I}) \Theta_i + \rho \Theta_i \mathbf{U}_i^* \Theta_i \right). \quad (26)$$

The quantity λ_i defined in (24) can be computed by

$$\lambda_i := (p - 2 \cdot \text{tr}(\mathbf{W}_i) + \text{tr}(\mathbf{W}_i^2))^{1/2}. \quad (27)$$

where $\mathbf{W}_i := \Theta_i(\widehat{\Sigma} + \rho\mathbf{U}_i^*)$.

The graph learning algorithm: Algorithm 2 summarizes the proposed scheme for graph selection.

Algorithm 2 (*Dual PN for graph selection (DPNGS)*)

Input: Matrix $\Sigma \succ 0$ and a given tolerance $\varepsilon > 0$.

Output: An approximate solution Θ_i of (1).

Initialization: Find a starting point $\Theta_0 \succ 0$.

for $i = 0$ **to** i_{\max} **do**

1. Set $\tilde{\mathbf{Q}}_i := \rho^{-1} (\Theta_i \widehat{\Sigma} \Theta_i - 2\Theta_i)$.
2. Compute \mathbf{U}_i^* in (25).
3. Compute λ_i by (27), where $\mathbf{W}_i := \Theta_i(\widehat{\Sigma} + \rho\mathbf{U}_i^*)$.
4. If $\lambda_i \leq \varepsilon$ terminate.
5. Compute $\Delta_i := -((\Theta_i \widehat{\Sigma} - \mathbb{I})\Theta_i + \rho\Theta_i\mathbf{U}_i^*\Theta_i)$.
6. Set $\alpha_i := (1 + \lambda_i)^{-1}$.
7. Update $\Theta_{i+1} := \Theta_i + \alpha_i\Delta_i$.

end for

Overall, Algorithm 2 does not require any matrix inversion operation. It only needs matrix-vector and matrix-matrix calculations, making the parallelization of the code easier. We note that due to the predefined step-size selection α_i in Algorithm 1 we do not need to do any backtracking line-search step. This advantage can avoid some overhead computation regarding the evaluation of the objective function which is usually expensive in this application.

Arithmetical complexity analysis: Since the analytical complexity is provided in Theorem 3.6, we only analyze the arithmetical complexity of Algorithm 2 here. As we work through the dual problem, the primal solution is dense even if majority of the entries are rather small (e.g., smaller than 10^{-6}).¹ Hence, the arithmetical complexity of Algorithm 2 is dominated by the complexity of $p \times p$ matrix multiplications.

For instance, the computation of $\tilde{\mathbf{Q}}_i$ and Δ_i require basic matrix multiplications. For the computation of λ_i , we require two trace operations: $\text{tr}(\mathbf{W}_i)$ in $\mathcal{O}(p)$ time-complexity and $\text{tr}(\mathbf{W}_i^2)$ in $\mathcal{O}(p^2)$ time-complexity. We note here that, while \mathbf{W}_i is a *dense* matrix, the trace operation requires only the computation of the diagonal elements of \mathbf{W}_i^2 . Given Θ_i , α_i and Δ_i , Θ_{i+1} requires $\mathcal{O}(p^2)$ time-complexity.

To compute (25), we can use the fast projected gradient method (FPGM) (Nesterov, 2007; Beck &

¹In our MATLAB code, we made no attempts for sparsification of the primal solution. The overall complexity of the algorithm can be improved via thresholding tricks.

Teboulle, 2009) with step size $1/L$ where L is the Lipschitz constant of the gradient of the objective function in (25). It is easy to observe that $L_i := \gamma_{\max}^2(\Theta_i)$ where $\gamma_{\max}(\Theta_i)$ is the largest eigenvalue of Θ_i . For sparse Θ_i , we can approximately compute $\gamma_{\max}(\Theta_i)$ is $\mathcal{O}(p^2)$ by using *iterative power methods* (typically, 10 iterations suffice). The projection onto $\|\text{vec}(\mathbf{U})\|_{\infty} \leq 1$ clips the elements by unity in $\mathcal{O}(p^2)$ time. Thus, the time overhead due to acceleration is within $\mathcal{O}(p^2)$.

Given the above, FPGM requires a constant number of iterations k_{\max} , which is independent of the dimension p , to achieve an ε_{in} solution accuracy. Overall, the time-complexity for the solution in (25) is $\mathcal{O}(k_{\max}M)$, where M is the cost of matrix multiplication.

Remark 4.1 (Parallel and distributed implementation ability). *In Algorithm 2, the outer loop does not require any Cholesky decomposition or matrix inversion. Suppose that the fast projected gradient method is applied to solve the dual subproblem (25). The main operation needed in the whole algorithm is matrix-matrix multiplication of the form $\Theta_i\mathbf{U}\Theta_i$, where Θ_i and \mathbf{U} are symmetric positive definite. This operation can naturally be computed in a parallel or distributed manner. For more details, we refer the reader to Chapter 1 in (Bertsekas & Tsitsiklis, 1989).*

5. Numerical experiments

In this section we test DPNGS (Algorithm 2 in Section 4) and compare it with the state-of-the-art graph selection algorithm QUadratic Inverse Covariance (QUIC) algorithm (Hsieh et al., 2011) on a real world data set. The QUIC algorithm is also a Newton-based method, which in addition exploits the sparsity in solving its primal subproblems. We note that QUIC was implemented in C while our codes in this work are implemented in MATLAB.

Implementation details: We test DPNGS on MATLAB 2011b running on a PC Intel Xeon X5690 at 3.47GHz per core with 94Gb RAM. To solve (25), we use the FPGM scheme as detailed in the supplementary material. We terminate FPGM if either $\|\mathbf{U}_{k+1} - \mathbf{U}_k\|_{\text{F}} \leq \varepsilon_{\text{in}} \max\{\|\mathbf{U}_k\|_{\text{F}}, 1\}$ or the number of iterations reaches k_{\max} where $\varepsilon_{\text{in}} > 0$ and k_{\max} will be specified later. The stopping criterion of the outer loop is $\lambda_i \leq 10^{-6}$ and the maximum number of outer iterations is chosen as $i_{\max} := 200$. We test the following three variants of DPNGS: DPNGS [$\varepsilon_{\text{in}} = 10^{-6}$ and $k_{\max} = 1000$], DPNGS(5) [$\varepsilon_{\text{in}} = 10^{-4}$ and $k_{\max} = 5$], and DPNGS(10) [$\varepsilon_{\text{in}} = 10^{-5}$ and $k_{\max} = 10$]. The DPNGS(5) and DPNGS(10) variants can be considered as inexact variants of DPNGS.

Real-world data: In our experiments, we use the

real biology data preprocessed by (Li & Toh, 2010) to compare the performance of the DPNGS variants above and QUIC (Hsieh et al., 2011) for 5 problems: *Lymph* ($p = 587$), *Estrogen* ($p = 692$), *Arabidopsis* ($p = 834$), *Leukemia* ($p = 1225$) and *Hereditary* ($p = 1869$). This dataset can be found at http://ima.umn.edu/~maxxa007/send_SICS/.

Convergence behaviour analysis: First, we verify the convergence behaviour of Algorithm 2 by analyzing the quadratic convergence of the quantity λ_i , where λ_i is defined by (27). Our analysis is based on the *Lymph* problem with $p = 587$ variables. We note that λ_i reveals the weighted norm of the proximal-gradient mapping of the problem. The convergence behaviour is plotted in Figure 1 for three different values of ρ , namely $\rho = 0.25$, $\rho = 0.1$, $\rho = 0.05$ and $\rho = 0.01$. Figure 1 shows that whenever the values of λ_i gets




Figure 1. Quadratic convergence of DPNGS

into the quadratic region, it converges with only a few iterations. As ρ becomes smaller, we need more iterations to get into the quadratic convergence region.

Next, we illustrate the step-size α_i of DPNGS. Figure 2 shows the increasing behaviour of the step size on the same dataset. Since $\alpha_i = (1 + \lambda_i)^{-1}$, it converges quickly at the last iterations. We also com-




Figure 2. The step size of DPNGS

pare the objective value decrement of both algorithms in y -log-scale in Figure 3. Using the same tolerance

level, we reach the objective value -4.141662×10^2 after 69 iterations while QUIC needs 159 iterations. Moreover, Figure 3 shows the quadratic convergence of our approach in contrast to QUIC; the latter requires many more iterations to slightly improve the objective. Figure 4 is the histogram of the solution in




Figure 3. The difference of the objective values of DPNGS and QUIC in y -log-scale

log scale reported by DPNGS and QUIC. Due to the dual solution approach, DPNGS reports an approximate solution with similar sparsity pattern as the one of QUIC. However, our solution has many small numbers instead of zero as in QUIC as revealed in Figure 4. *This seems to be the main weakness of the dual approach:* it obviates matrix inversions by avoiding the primal problem, which can return solutions with exact zeros thanks to its soft-thresholding prox-operator.

As a result, DPNGS carries around extremely small coefficients (almost of them smaller than 5×10^{-5}) often preventing it from achieving the same objective level as the numerical experiments on the full data set shows. At the same time, since the approach does not rely on coordinate descent on active sets, it appears much less sensitive to the choice of ρ . This could be an advantage of DPNGS in applications requiring smaller ρ values. If exact sparsity is needed, then a single primal iteration suffices to remove the small coefficients.




Figure 4. The histogram of the coefficient absolute values of the solution in log-scale of DPNGS and QUIC (right).

Numerical experiments on the full dataset: We now report the numerical experiments on the biology

Table 1. Summary of comparison results on real world datasets.

Algorithm	$\rho = 0.25$			$\rho = 0.1$			$\rho = 0.05$			$\rho = 0.01$		
	#iter	time[s]	$F(\Theta_i)$	#iter	time[s]	$F(\Theta_i)$	#iter	time[s]	$F(\Theta_i)$	#iter	time[s]	$F(\Theta_i)$
Lymph Problem ($p = 587$)												
DPNGS	19	49.028	613.26	29	61.548	341.89	40	66.635	133.60	69	104.259	-414.17
DPNGS(10)	39	7.470	613.42	34	8.257	342.12	43	8.678	133.87	78	35.543	-413.82
DPNGS(5)	61	7.067	615.87	30	4.323	344.72	41	5.862	136.37	69	123.552	-414.17
QUIC (C code)	22	8.392	613.25	44	33.202	341.88	82	176.135	133.60	201	2103.788	-414.17
Estrogen Problem ($p = 692$)												
DPNGS	24	141.027	627.87	39	171.721	251.20	52	167.460	-11.59	83	205.262	-643.21
DPNGS(10)	56	15.500	628.10	49	14.092	251.52	59	19.262	-11.25	90	28.930	-642.85
DPNGS(5)	39	9.310	631.53	46	8.388	254.61	51	9.332	-7.69	81	42.955	-639.54
QUIC (C code)	19	7.060	627.85	43	49.235	251.19	81	244.242	-11.60	-	-	-
Arabidopsis Problem ($p = 834$)												
DPNGS	26	174.947	728.57	43	220.365	228.16	61	253.180	-146.10	100	430.505	-1086.57
DPNGS(10)	48	22.268	728.96	45	22.404	228.57	60	26.007	-145.72	200	101.428	-1038.60
DPNGS(5)	38	9.826	733.67	44	11.113	233.04	57	18.378	-141.84	95	73.948	-1083.53
QUIC (C code)	21	19.684	728.52	49	116.016	228.14	95	562.532	-146.13	-	-	-
Leukemia Problem ($p = 1255$)												
DPNGS	28	669.548	1143.79	48	624.145	386.37	71	726.688	-279.93	130	1398.133	-2071.33
DPNGS(10)	65	82.497	1144.66	48	60.108	387.26	68	84.017	-279.12	126	166.567	-2070.02
DPNGS(5)	49	38.317	1154.13	48	37.273	395.08	70	50.886	-271.01	124	258.090	-2060.25
QUIC (C code)	18	69.826	1143.76	41	344.199	386.33	76	1385.577	-280.07	-	-	-
Hereditary Problem ($p = 1869$)												
DPNGS	41	2645.875	1258.31	82	3805.608	-348.49	113	5445.974	-1609.59	183	9020.237	-4569.85
DPNGS(10)	63	242.528	1261.15	80	297.131	-345.47	126	435.159	-1606.67	190	732.802	-4566.66
DPNGS(5)	58	129.821	1290.34	79	169.817	-313.87	126	439.386	-1606.67	179	1140.932	-4537.95
QUIC (C code)	21	437.252	1258.00	45	1197.895	-348.80	84	3182.211	-1609.92	-	-	-

dataset and compare the methods. We test both algorithms with four different values of ρ , namely $\rho = 0.25$, $\rho = 0.1$, $\rho = 0.05$ and $\rho = 0.01$. The numerical results are reported in Table 1. Since QUIC exceeds the maximum number of iterations $i_{\max} = 200$ and takes a long time, we do not report the results corresponding to $\rho = 0.01$. We note again that at the time of this ICML submission, our implementation is done in MATLAB, while QUIC code is (carefully!) implemented in C. Hence, our timings may improve.

We highlight several interesting results from Table 1. First, QUIC obtains the highest accuracy results in most cases, which we attribute to the “lack of soft thresholding” in our algorithm. As the DPNGS algorithm carries around a score of extremely small numbers (effectively making the solution dense in the numerical sense), its solutions are close to QUIC’s solutions within numerical precision. Moreover, QUIC is extremely efficient when the ρ value is large, since it exploits the sparsity of the putative solutions via coordinate descent. Unsurprisingly, QUIC slows down significantly as ρ is decreased.

DPNGS(5) and DPNGS(10) can obtain near optimal solutions quite rapidly. In particular, DPNGS(10) seems to be the most competitive across the board, often taking a fraction of QUIC’s time to provide a very close solution. Hence, one can expect these schemes

to be used for initializing other algorithms. For instance, QUIC can be a good candidate. We observed in all cases that, in the first few iterations, QUIC performs several Cholesky decompositions to stay within the positive definite cone. As the complexity of such operation is large, our step-size selection within QUIC or a DPNGS(10) initialization can be helpful.

6. Conclusions

In this paper, we present the new composite self-concordant optimization framework. As a concrete application, we demonstrate that graph learning is possible without any Cholesky decompositions via analytic step-size selection as well as without matrix inversions via a careful dual formulation within our framework. By exploiting the self-concordance in the composite graph learning objective, we provide an optimal step-size for this class of composite minimization with proximal Newton methods. We show that within the dual formulation of the Newton subproblem, we do not need to explicitly calculate the gradient as it appears in a multiplication form with the Hessian. Thanks to the special structure of this multiplication, we avoid matrix inversions in graph learning. Overall, we expect our optimization framework to have more applications in signal processing/machine learning and be amenable to various parallelization techniques, beyond the ones considered in the graph learning problem.

References

- Banerjee, O., El Ghaoui, L., and d’Aspremont, A. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *The Journal of Machine Learning Research*, 9:485–516, 2008.
- Beck, A. and Teboulle, M. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM J. Imaging Sciences*, 2(1):183–202, 2009.
- Bertsekas, D.P. and Tsitsiklis, J. N. *Parallel and distributed computation: Numerical methods*. Prentice Hall, 1989.
- Boyd, S. and Vandenberghe, L. *Convex Optimization*. University Press, Cambridge, 2004.
- Dempster, A. P. Covariance selection. *Biometrics*, 28: 157–175, 1972.
- Hsieh, C. J., Sustik, M.A., Dhillon, I.S., and Ravikumar, P. Sparse inverse covariance matrix estimation using quadratic approximation. *Advances in Neural Information Processing Systems (NIPS)*, 24:1–18, 2011.
- Lee, J.D., Sun, Y., and Saunders, M.A. Proximal newton-type methods for convex optimization. *Tech. Report.*, pp. 1–25, 2012.
- Li, L. and Toh, K.C. An inexact interior point method for l_1 -regularized sparse covariance selection. *Mathematical Programming Computation*, 2(3):291–315, 2010.
- Lu, Z. Adaptive first-order methods for general sparse inverse covariance selection. *SIAM Journal on Matrix Analysis and Applications*, 31(4):2000–2016, 2010.
- Nesterov, Y. *Introductory lectures on convex optimization: a basic course*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, 2004.
- Nesterov, Y. Gradient methods for minimizing composite objective function. *CORE Discussion paper*, 76, 2007.
- Nesterov, Y. and Nemirovski, A. *Interior-point Polynomial Algorithms in Convex Programming*. Society for Industrial Mathematics, 1994.
- Olsen, P.A., Oztoprak, F., Nocedal, J., and Rennie, S.J. Newton-like methods for sparse inverse covariance estimation. *Optimization Online*, 2012.
- Ravikumar, P., Wainwright, M. J., Raskutti, G., and Yu, B. High-dimensional covariance estimation by minimizing l_1 -penalized log-determinant divergence. *Electron. J. Statist.*, 5:935–988, 2011.
- Rolfs, B., Rajaratnam, B., Guillot, D., Wong, I., and Maleki, A. Iterative thresholding algorithm for sparse inverse covariance estimation. In *Advances in Neural Information Processing Systems 25*, pp. 1583–1591, 2012.
- Scheinberg, K. and Rish, I. Sinco-a greedy coordinate ascent method for sparse inverse covariance selection problem. *preprint*, 2009.
- Scheinberg, K., Ma, S., and Goldfarb, D. Sparse inverse covariance selection via alternating linearization methods. *arXiv preprint arXiv:1011.0097*, 2010.
- Yuan, X. Alternating direction method for covariance selection models. *Journal of Scientific Computing*, 51(2):261–273, 2012.