

SOLVING LOG-DETERMINANT OPTIMIZATION PROBLEMS BY A NEWTON-CG PRIMAL PROXIMAL POINT ALGORITHM*

CHENGJING WANG[†], DEFENG SUN[‡], AND KIM-CHUAN TOH[§]

Abstract. We propose a Newton-CG primal proximal point algorithm (PPA) for solving large scale log-determinant optimization problems. Our algorithm employs the essential ideas of PPA, the Newton method, and the preconditioned CG solver. When applying the Newton method to solve the inner subproblem, we find that the log-determinant term plays the role of a smoothing term as in the traditional smoothing Newton technique. Focusing on the problem of maximum likelihood sparse estimation of a Gaussian graphical model, we demonstrate that our algorithm performs favorably compared to existing state-of-the-art algorithms and is much preferred when a high quality solution is required for problems with many equality constraints.

Key words. log-determinant optimization problem, sparse inverse covariance selection, proximal point algorithm, Newton's method

AMS subject classifications. 90C06, 90C22, 90C25, 65F10

DOI. 10.1137/090772514

1. Introduction. In this paper, by defining $\log 0 := -\infty$, we consider the following standard primal and dual log-determinant (log-det) problems:

$$(P) \quad \min_X \{ \langle C, X \rangle - \mu \log \det X : \mathcal{A}(X) = b, X \succeq 0 \},$$

$$(D) \quad \max_{y, Z} \{ b^T y + \mu \log \det Z + n\mu(1 - \log \mu) : Z + \mathcal{A}^T y = C, Z \succeq 0 \},$$

where $C \in \mathcal{S}^n$, $b \in \mathcal{R}^m$, $\mu \geq 0$ is a given parameter, $\mathcal{A} : \mathcal{S}^n \rightarrow \mathcal{R}^m$ is a linear map, and $\mathcal{A}^T : \mathcal{R}^m \rightarrow \mathcal{S}^n$ is the adjoint of \mathcal{A} . We assume that \mathcal{A} is surjective, and hence $\mathcal{A}\mathcal{A}^T$ is nonsingular. Note that the linear maps \mathcal{A} and \mathcal{A}^T can be expressed, respectively, as

$$(1) \quad \mathcal{A}(X) = [\langle A_1, X \rangle, \dots, \langle A_m, X \rangle]^T, \quad \mathcal{A}^T(y) = \sum_{k=1}^m y_k A_k,$$

where A_k , $k = 1, \dots, m$, are given matrices in \mathcal{S}^n . For an explanation of all other main notation, see subsection 1.1.

It is clear that the log-det problem (P) is a convex optimization problem; i.e., the objective function $\langle C, X \rangle - \mu \log \det X$ is convex (on \mathcal{S}_+^n), and the feasible region is convex. The log-det problems (P) and (D) can be considered as a generalization of linear semidefinite programming (SDP) problems. One can see that, in the limiting case where $\mu = 0$, they reduce, respectively, to the standard primal and dual linear SDP problems. Log-det problems arise in many practical applications such

*Received by the editors September 30, 2009; accepted for publication (in revised form) July 30, 2010; published electronically October 14, 2010.

<http://www.siam.org/journals/siopt/20-6/77251.html>

[†]College of Mathematics, Southwest Jiaotong University, Chengdu 610031, People's Republic of China (renascencewang@hotmail.com).

[‡]Department of Mathematics and NUS Risk Management Institute, National University of Singapore, 10 Lower Kent Ridge Road, Singapore 119076, Singapore (matsundf@nus.edu.sg). This author's research was partially supported by the Academic Research Fund under grant R-146-000-104-112.

[§]Department of Mathematics, National University of Singapore, 10 Lower Kent Ridge Road, Singapore 119076, Singapore (mattokc@nus.edu.sg) and Singapore-MIT Alliance, 4 Engineering Drive 3, Singapore 117576, Singapore.

as computational geometry, statistics, system identification, experiment design, and information and communication theory. Thus the algorithms we develop here can potentially find wide applications. One may refer to [5, 31, 28] for an extensive account of applications of the log-det problem.

For small- and medium-sized log-det problems, including linear SDP problems, it is widely accepted that interior-point methods (IPMs) with direct solvers are generally very efficient and robust; see, for example, [28, 30]. For log-det problems with m large and n moderate (say no more than 2000), the limitations faced by IPMs with direct solvers become very severe due to the need to compute, store, and factorize the $m \times m$ Schur matrices that are typically dense.

Recently, Zhao, Sun, and Toh [36] proposed a Newton-CG augmented Lagrangian (NAL) method for solving linear SDP problems. This method can be very efficient when the problems are primal and dual nondegenerate. The NAL method is essentially a proximal point method applied to the primal problem where the inner subproblems are solved by an inexact semismooth Newton method using a preconditioned CG (PCG) solver. Recent studies conducted by Sun, Sun, and Zhang [26] and Chan and Sun [6] revealed that, under the constraint nondegenerate conditions for (P) and (D) (i.e., the primal and dual nondegeneracy conditions in the IPM literature, e.g., [1]), the NAL method can be regarded locally as an approximate generalized Newton method applied to a semismooth equation. The latter result may explain to a large extent why the NAL method can be very efficient.

As the log-det problem (P) is an extension of the primal linear SDP problem, it is natural for us to further use the NAL method developed for linear SDP problems to solve log-det problems. Following what has been done in linear SDP problems, our approach is to apply a Newton-CG primal proximal point algorithm (PPA) to (P) and then use an inexact Newton method to solve the inner subproblems by using a PCG solver to compute inexact Newton directions. We note that when solving the inner subproblems in the NAL method for linear SDP problems [36], a semismooth Newton method has to be used since the objective functions are differentiable but not twice continuously differentiable. For log-det problems, however, the objective functions in the inner subproblems are twice continuously differentiable (actually, analytic) due to the fact that the term $-\mu \log \det X$ acts as a smoothing term. This interesting phenomenon implies that the standard Newton method can be used to solve the inner subproblem. It also reveals a close connection between adding the log-barrier term $-\mu \log \det X$ to a linear SDP problem and the technique of smoothing the Karush-Kuhn-Tucker (KKT) conditions [6].

In [22, 23], Rockafellar established a general theory on the global convergence and local linear rate of convergence of the sequence generated by the proximal point and augmented Lagrangian methods for solving convex optimization problems including (P) and (D). Borrowing Rockafellar's results, we can establish global convergence and a local convergence rate for our Newton-CG PPA method for (P) and (D) without much difficulty.

In problem (P), we deal only with a matrix variable, but the PPA method we develop in this paper can easily be extended to more general log-det problems to include vector variables. Although this kind of extension seems trivial, the numerical experiments in section 6 include these problems; thus we still list them as follows:

$$(2) \quad \min \{ \langle C, X \rangle - \mu \log \det X + c^T x - \nu \log x : \mathcal{A}(X) + Bx = b, X \succeq 0, x \geq 0 \},$$

$$\max \{ b^T y + \mu \log \det Z + \nu \log z + \kappa : Z + \mathcal{A}^T y = C, z + B^T y = c, Z \succeq 0, z \geq 0 \},$$

where $\nu > 0$ is a given parameter, $c \in \mathcal{R}^l$ and $B \in \mathcal{R}^{m \times l}$ are given data, and $\kappa = n\mu(1 - \log \mu) + l\nu(1 - \log \nu)$.

In the implementation of our Newton-CG PPA method, we focus on the maximum likelihood sparse estimation of a Gaussian graphical model (GGM). Given n variables drawn from a Gaussian distribution $\mathcal{N}(0, C)$ where the true covariance matrix C is unknown, we estimate C from a sample covariance matrix Σ by maximizing its log-likelihood. Following [9], setting a certain number of coefficients in the inverse covariance matrix Σ^{-1} to zeros, a procedure known as *covariance selection*, improves the stability of this estimation procedure by reducing the number of parameters to be estimated and highlighting structure. This class of problems includes two subclasses. The first is that the coefficients to be set to zero are *completely* known, and it can be formulated as follows:

$$(3) \quad \min_X \left\{ \langle S, X \rangle - \log \det X : X_{ij} = 0 \ \forall (i, j) \in \Omega^c, X \succeq 0 \right\},$$

where Ω^c is the set of pairs of nodes (i, j) in a graph that are connected by an edge, and $S \in \mathcal{S}^n$ is a given sample covariance matrix. Problem (3) is also known as a sparse covariance selection problem. In [7], Dahl, Vandenberghe, and Roychowdhury showed that when the underlying dependency graph is nearly chordal, an inexact Newton method combined with an appropriate PCG solver can be quite efficient in solving (3) with n up to 2000 but on very sparse data (for instance, when $n = 2000$, the number of upper nonzeros is only about $4000 \sim 6000$). But for general large scale problems of the form (3), little research has been done in finding efficient algorithms to solve the problems.

The second subclass of the GGM is that the coefficients to be set to zeros are *partially* known, and it is formulated as follows:

$$(4) \quad \min_X \left\{ \langle S, X \rangle - \log \det X + \sum_{(i,j) \notin \Omega^c} \rho_{ij} |X_{ij}| : X_{ij} = 0 \ \forall (i, j) \in \Omega^c, X \succeq 0 \right\}.$$

In [8], d'Aspremont, Banerjee, and El Ghaoui, among the earliest, proposed to apply Nesterov's smooth approximation (NSA) scheme to solve (4) for the case where $\Omega^c = \emptyset$. Subsequently, Lu [13, 14] suggested an adaptive Nesterov's smooth (ANS) method to solve (4). The ANS method is currently one of the most effective methods for solving large scale problems (e.g., $n \geq 1000$, $m \geq 500000$) of the form (4). In the ANS method, the equality constraints in (4) are removed and included in the objective function via the penalty approach. The main idea in the ANS method is basically to apply a variant of Nesterov's smooth method [20] to solve the penalized problem subject to the single constraint $X \succeq 0$. In fact, both the ANS and NSA methods have the same principal idea, but the latter runs much slowly than the former. In contrast to IPMs, the greatest merit of the ANS method is that it needs much lower storage and computational cost per iteration. In [13], the ANS method was demonstrated to be rather efficient in solving randomly generated problems of form (4), while obtaining solutions with low/moderate accuracy. However, as the ANS method is a first-order method, it may require a huge computing cost to obtain high accuracy solutions. In addition, since the penalty approach is used in the ANS method to solve (4), the number of iterations may increase drastically if the penalty parameter is updated frequently. Another limitation of the ANS method introduced in [13] is that it can deal only with the special equality constraints in (4). It appears to be difficult to extend the ANS method to deal with more general equality constraints of the form $\mathcal{A}(X) = b$.

After this paper's first round of review, one referee brought [25] and [35] to our attention. These two papers also dealt with (4) but without equality constraints. According to our numerical experiments, for examples without equality constraints given in section 6 (the real data), only the alternating direction method (ADM) proposed by Yuan [35] seems competitive with the PPA and the ANS method (the ADM is slightly slower than the PPA but faster than the ANS method). It is worth further investigation to determine whether the ADM approach of Yuan [35] can be used to solve the general model (4) efficiently.

Our numerical results show that for both problems (3) and (4), our Newton-CG PPA method can be very efficient and robust in solving large scale problems generated as in [8] and [13]. Indeed, we are able to solve sparse covariance selection problems with n up to 2000 and m up to 1.7×10^6 in about 26 minutes. For both problems, our method consistently outperforms the ANS method by a substantial margin, especially when the problems are large and the required accuracy tolerances are relatively high.

The remaining part of this paper is organized as follows. In section 2, we give some preliminaries including a brief introduction on concepts related to the PPA. In sections 3 and 4, we present the details of the PPA method and the Newton-CG algorithm. In section 5, we give the convergence analysis of our PPA method. The numerical performance of our algorithm is presented in section 6. Finally, we give some concluding remarks in section 7.

1.1. Notation. In this paper, all vectors are assumed to be finite dimensional. The symbol \mathcal{R}^n denotes the n -dimensional Euclidean space. The set of all $m \times n$ matrices with real entries is denoted by $\mathcal{R}^{m \times n}$. The space of all symmetric $n \times n$ matrices is denoted by \mathcal{S}^n , $\langle \cdot, \cdot \rangle$ stands for the standard trace inner product in \mathcal{S}^n , and $\|\cdot\|$ denotes the Frobenius norm. Moreover, \mathcal{S}_+^n (resp., \mathcal{S}_{++}^n) is the cone of $n \times n$ symmetric positive semidefinite (resp., definite) matrices. If $X \in \mathcal{S}_+^n$, we also write $X \succeq 0$. Given matrices X and Y in $\mathcal{R}^{m \times n}$, $X \circ Y$ denotes the Hadamard product of X and Y .

2. Preliminaries. For the sake of subsequent discussions, we first introduce some concepts related to the proximal point method based on the classic papers by Rockafellar [22, 23].

Let H be a real Hilbert space with an inner product $\langle \cdot, \cdot \rangle$. A multifunction $T : H \rightrightarrows H$ is said to be a *monotone operator* if

$$(5) \quad \langle z - z', w - w' \rangle \geq 0 \quad \text{whenever } w \in T(z), w' \in T(z').$$

It is said to be *maximal monotone* if, in addition, the graph

$$G(T) = \{(z, w) \in H \times H \mid w \in T(z)\}$$

is not properly contained in the graph of any other monotone operator $T' : H \rightrightarrows H$. For example, if T is the subdifferential ∂f of a lower semicontinuous convex function $f : H \rightarrow (-\infty, +\infty]$, $f \not\equiv +\infty$, then T is maximal monotone (see Minty [17] or Moreau [18]) and the relation $0 \in T(z)$ means that $f(z) = \min f$.

Rockafellar [22] studied a fundamental algorithm for solving

$$(6) \quad 0 \in T(z)$$

in the case of an arbitrary maximal monotone operator T . The operator $P = (I + \lambda T)^{-1}$ is known to be single-valued from all of H into H , where $\lambda > 0$ is a given

parameter. It is also *nonexpansive*, i.e.,

$$\|P(z) - P(z')\| \leq \|z - z'\|,$$

and one has $P(z) = z$ if and only if $0 \in T(z)$. The operator P is called the *proximal mapping* associated with λT , following the terminology of Moreau [18] for the case of $T = \partial f$.

The PPA generates, for any starting point z^0 , a sequence $\{z^k\}$ in H by the approximate rule:

$$z^{k+1} \approx (I + \lambda_k T)^{-1}(z^k).$$

Here $\{\lambda_k\}$ is a sequence of positive real numbers. In the case of $T = \partial f$, this procedure reduces to

$$(7) \quad z^{k+1} \approx \arg \min_z \left\{ f(z) + \frac{1}{2\lambda_k} \|z - z^k\|^2 \right\},$$

which was first introduced by Martinet in [15]. In the literature of the PPA, the quadratic regularization term $\frac{1}{2\lambda_k} \|z - z^k\|^2$ is not the only choice and can be done in a more generic form. For example, it may be constructed in terms of the φ -divergence functional and Bregman's measure of distance [27]. It is interesting to note that the term $-\log \det X$ can also be regarded as a generation function for one kind of Bregman divergence [10]. In this paper, this term is fixed and is not treated as a regularization term. What we need is exactly the quadratic regularization term. So we will state only convergence analysis of the PPA under this framework. For more general results, one may consult [27] and the references therein.

DEFINITION 2.1 (cf. [22]). *For a maximal monotone operator T , we say that its inverse T^{-1} is Lipschitz continuous at the origin (with modulus $a \geq 0$) if there is a unique solution \bar{z} to $z = T^{-1}(0)$ and for some $\tau > 0$ we have*

$$\|z - \bar{z}\| \leq a\|w\|, \quad \text{where } z \in T^{-1}(w) \text{ and } \|w\| \leq \tau.$$

We state the following lemma, which will be needed later in the derivation of the PPA method for solving (P).

LEMMA 2.1. *Let Y be an $n \times n$ symmetric matrix with eigenvalue decomposition $Y = PDP^T$ with $D = \text{diag}(d)$. We assume that $d_1 \geq \dots \geq d_r > 0 \geq d_{r+1} \dots \geq d_n$. Let $\gamma > 0$ be given. For the two scalar functions $\phi_\gamma^+(x) := (\sqrt{x^2 + 4\gamma} + x)/2$ and $\phi_\gamma^-(x) := (\sqrt{x^2 + 4\gamma} - x)/2$ for all $x \in \mathcal{R}$, we define their matrix counterparts:*

$$(8) \quad Y_1 = \phi_\gamma^+(Y) := P \text{diag}(\phi_\gamma^+(d)) P^T \quad \text{and} \quad Y_2 = \phi_\gamma^-(Y) := P \text{diag}(\phi_\gamma^-(d)) P^T.$$

Then we have the following:

(a) *The following decomposition holds: $Y = Y_1 - Y_2$, where $Y_1, Y_2 \succ 0$ and $Y_1 Y_2 = \gamma I$.*

(b) *ϕ_γ^+ is continuously differentiable everywhere in \mathcal{S}^n , and its derivative $(\phi_\gamma^+)'(Y)[H]$ at Y for any $H \in \mathcal{S}^n$ is given by*

$$(\phi_\gamma^+)'(Y)[H] = P(\Omega \circ (P^T H P))P^T,$$

where $\Omega \in \mathcal{S}^n$ is defined by

$$\Omega_{ij} = \frac{\phi_\gamma^+(d_i) + \phi_\gamma^+(d_j)}{\sqrt{d_i^2 + 4\gamma} + \sqrt{d_j^2 + 4\gamma}}, \quad i, j = 1, \dots, n.$$

(c) $(\phi_\gamma^+)'(Y)[Y_1 + Y_2] = \phi_\gamma^+(Y)$.

Proof. (a) It is easy to verify that the decomposition holds. (b) The result follows from [3, Chap. V.3.3] and the fact that

$$\frac{\phi_\gamma^+(d_i) - \phi_\gamma^+(d_j)}{d_i - d_j} = \frac{\phi_\gamma^+(d_i) + \phi_\gamma^+(d_j)}{\sqrt{d_i^2 + 4\gamma} + \sqrt{d_j^2 + 4\gamma}}, \quad d_i \neq d_j.$$

(c) We have $(\phi_\gamma^+)'(Y)[Y_1 + Y_2] = P(\Omega \circ \text{diag}(\phi_\gamma^+(d) + \phi_\gamma^-(d)))P^T = P\text{diag}(\phi_\gamma^+(d))P^T$, and the required result follows. \square

3. The primal proximal point algorithm. Define the feasible sets of (P) and (D), respectively, by

$$\mathcal{F}_P = \{X \in \mathcal{S}^n : \mathcal{A}(X) = b, X \succeq 0\}, \quad \mathcal{F}_D = \{(y, Z) \in \mathcal{R}^m \times \mathcal{S}^n : Z + \mathcal{A}^T y = C, Z \succeq 0\}.$$

Throughout this paper, we assume that the following conditions for (P) and (D) hold.

Assumption 3.1. Problem (P) satisfies the condition

(9) $\exists X_0 \in \mathcal{S}_{++}^n$ such that $\mathcal{A}(X_0) = b$.

Assumption 3.2. Problem (D) satisfies the condition

(10) $\exists (y_0, Z_0) \in \mathcal{R}^m \times \mathcal{S}_{++}^n$ such that $Z_0 + \mathcal{A}^T y_0 = C$.

Under the above assumptions, the problem (P) has a unique optimal solution, denoted by \bar{X} , and the problem (D) has a unique optimal solution, denoted by (\bar{y}, \bar{Z}) . In addition, the following KKT conditions are necessary and sufficient for the optimality of (P) and (D):

$$\begin{aligned} \mathcal{A}(X) - b &= 0, \\ Z + \mathcal{A}^T y - C &= 0, \\ XZ &= \mu I, \quad X \succeq 0, \quad Z \succeq 0. \end{aligned} \tag{11}$$

The last condition in (11) can easily be seen to be equivalent to the condition

(12) $\phi_\gamma^+(X - \lambda Z) = X$ with $\gamma := \lambda\mu$

for any given $\lambda > 0$, where ϕ_γ^+ is defined by (8) in Lemma 2.1. Recall that in [36] for the linear SDP case (where $\mu = 0$), the complementarity condition $XZ = 0$, with $X, Z \succeq 0$, is equivalent to $\Pi_+(X - \lambda Z) = X$, for any $\lambda > 0$, where $\Pi_+(\cdot)$ is the metric projector onto \mathcal{S}_+^n . One can see from (12) that when $\mu > 0$, the log-barrier term $-\mu \log \det X$ in (P) contributes to a smoothing term in the projector Π_+ .

LEMMA 3.1. *Given any $Y \in \mathcal{S}^n$ and $\lambda > 0$, we have*

(13) $\min_{Z \succ 0} \left\{ \frac{1}{2\lambda} \|Y - Z\|^2 - \mu \log \det Z \right\} = \frac{1}{2\lambda} \|\phi_\gamma^-(Y)\|^2 - \mu \log \det(\phi_\gamma^+(Y)),$

where $\gamma = \lambda\mu$.

Proof. Note that the minimization problem in (13) is an unconstrained problem and the objective function is strictly convex and continuously differentiable. Thus any stationary point would be the unique minimizer of the problem. The stationary point, if it exists, is the solution of the following equation:

(14) $Y = Z - \gamma Z^{-1}.$

By Lemma 2.1(a), we see that $Z_* := \phi_\gamma^+(Y)$ satisfies (14). Thus the optimization problem in (13) has a unique minimizer and the minimum objective function value is given by

$$\frac{1}{2\lambda} \|Y - Z_*\|^2 - \mu \log \det Z_* = \frac{1}{2\lambda} \|\phi_\gamma^-(Y)\|^2 - \mu \log \det(\phi_\gamma^+(Y)).$$

This completes the proof. \square

Let $l(X; y) : \mathcal{S}^n \times \mathcal{R}^m \rightarrow \mathcal{R}$ be the ordinary Lagrangian function for (P) in extended form:

$$(15) \quad l(X; y) = \begin{cases} \langle C, X \rangle - \mu \log \det X + \langle y, b - \mathcal{A}(X) \rangle & \text{if } X \in \mathcal{S}_+^n, \\ \infty & \text{otherwise.} \end{cases}$$

The essential objective function in (P) is given by

$$(16) \quad f(X) = \max_{y \in \mathcal{R}^m} l(X; y) = \begin{cases} \langle C, X \rangle - \mu \log \det X & \text{if } X \in \mathcal{F}_P, \\ \infty & \text{otherwise.} \end{cases}$$

For later developments, we define the following maximal monotone operator associated with $l(X, y)$:

$$T_l(X, y) := \{(U, v) \in \mathcal{S}^n \times \mathcal{R}^m \mid (U, -v) \in \partial l(X, y), (X, y) \in \mathcal{S}^n \times \mathcal{R}^m\}.$$

Let F_λ be the Moreau–Yosida regularization (see [18, 34]) of f in (16) associated with $\lambda > 0$, i.e.,

$$(17) \quad F_\lambda(X) = \min_{Y \in \mathcal{S}^n} \left\{ f(Y) + \frac{1}{2\lambda} \|Y - X\|^2 \right\} = \min_{Y \in \mathcal{S}_{++}^n} \left\{ f(Y) + \frac{1}{2\lambda} \|Y - X\|^2 \right\}.$$

From (16), we have

$$(18) \quad \begin{aligned} F_\lambda(X) &= \min_{Y \in \mathcal{S}_{++}^n} \sup_{y \in \mathcal{R}^m} \left\{ l(Y; y) + \frac{1}{2\lambda} \|Y - X\|^2 \right\} \\ &= \sup_{y \in \mathcal{R}^m} \min_{Y \in \mathcal{S}_{++}^n} \left\{ l(Y; y) + \frac{1}{2\lambda} \|Y - X\|^2 \right\} = \sup_{y \in \mathcal{R}^m} \Theta_\lambda(X, y), \end{aligned}$$

where

$$\begin{aligned} \Theta_\lambda(X, y) &= \min_{Y \in \mathcal{S}_{++}^n} \left\{ l(Y; y) + \frac{1}{2\lambda} \|Y - X\|^2 \right\} \\ &= b^T y + \min_{Y \in \mathcal{S}_{++}^n} \left\{ \langle C - \mathcal{A}^T y, Y \rangle - \mu \log \det Y + \frac{1}{2\lambda} \|Y - X\|^2 \right\} \\ &= b^T y - \frac{1}{2\lambda} \|W_\lambda(X, y)\|^2 + \frac{1}{2\lambda} \|X\|^2 + \min_{Y \in \mathcal{S}_{++}^n} \left\{ \frac{1}{2\lambda} \|Y - W_\lambda(X, y)\|^2 - \mu \log \det Y \right\}. \end{aligned}$$

Here $W_\lambda(X, y) := X - \lambda(C - \mathcal{A}^T y)$. Note that the interchange of min and sup in (18) follows from [21, Thm. 37.3]. By Lemma 3.1, the minimum objective value in the above minimization problem is attained at $Y_* = \phi_\gamma^+(W_\lambda(X, y))$. Thus we have

$$(19) \quad \begin{aligned} \Theta_\lambda(X, y) &= b^T y + \frac{1}{2\lambda} \|X\|^2 - \frac{1}{2\lambda} \|\phi_\gamma^+(W_\lambda(X, y))\|^2 \\ &\quad - \mu \log \det \phi_\gamma^+(W_\lambda(X, y)) + n\mu. \end{aligned}$$

Note that for a given X , the function $\Theta_\lambda(X, \cdot)$ is analytic; cf. [29]. Its first- and second-order derivatives with respect to y can be computed as in the following lemma.

LEMMA 3.2. *For any $y \in \mathcal{R}^m$ and $X \succ 0$, we have*

$$(20) \quad \nabla_y \Theta_\lambda(X, y) = b - \mathcal{A} \phi_\gamma^+(W_\lambda(X, y)),$$

$$(21) \quad \nabla_{yy}^2 \Theta_\lambda(X, y) = -\lambda \mathcal{A} (\phi_\gamma^+)'(W_\lambda(X, y)) \mathcal{A}^T.$$

Proof. To simplify notation, we use W to denote $W_\lambda(X, y)$ in this proof. To prove (20), note that

$$\begin{aligned} \nabla_y \Theta_\lambda(X, y) &= b - \mathcal{A} (\phi_\gamma^+)'(W) [\phi_\gamma^+(W)] - \lambda \mu \mathcal{A} (\phi_\gamma^+)'(W) [(\phi_\gamma^+(W))^{-1}] \\ &= b - \mathcal{A} (\phi_\gamma^+)'(W) [\phi_\gamma^+(W) + \phi_\gamma^-(W)]. \end{aligned}$$

By Lemma 2.1(c), the required result follows. From (20), the result in (21) follows readily. \square

Let $y_\lambda(X)$ be such that

$$y_\lambda(X) \in \arg \sup_{y \in \mathcal{R}^m} \Theta_\lambda(X, y).$$

Then we know that $\phi_\gamma^+(W_\lambda(X, y_\lambda(X)))$ is the unique optimal solution to (17). Consequently, we have that $F_\lambda(X) = \Theta_\lambda(X, y_\lambda(X))$ and

$$(22) \quad \nabla F_\lambda(X) = \frac{1}{\lambda} (X - \phi_\gamma^+(W_\lambda(X, y_\lambda(X)))) = C - \mathcal{A}^T y - \frac{1}{\lambda} \phi_\gamma^-(W_\lambda(X, y_\lambda(X))).$$

Given $X^0 \in \mathcal{S}_{++}^n$, the exact PPA for solving (P) is given by

$$(23) \quad X^{k+1} = (I + \lambda_k T_f)^{-1}(X^k) = \arg \min_{X \in \mathcal{S}_{++}^n} \left\{ f(X) + \frac{1}{2\lambda_k} \|X - X^k\|^2 \right\},$$

where $T_f = \partial f$. It can be shown [24, Thm. 2.26] that

$$(24) \quad X^{k+1} = X^k - \lambda_k \nabla F_{\lambda_k}(X^k) = \phi_{\gamma_k}^+(W_\lambda(X^k, y_{\lambda_k}(X^k))),$$

where $\gamma_k = \lambda_k \mu$.

The exact PPA outlined in (23) is impractical for computational purposes. Hence we consider an inexact PPA for solving (P), which has the following template.

Algorithm 1: The primal PPA. Given a tolerance $\varepsilon > 0$. Input $X^0 \in \mathcal{S}_{++}^n$ and $\lambda_0 > 0$. Set $k := 0$. Iterate:

Step 1. Find an approximate maximizer

$$(25) \quad y^{k+1} \approx \arg \sup_{y \in \mathcal{R}^m} \left\{ \theta_k(y) := \Theta_{\lambda_k}(X^k, y) \right\},$$

where $\Theta_{\lambda_k}(X^k, y)$ is defined as in (20).

Step 2. Compute

$$(26) \quad X^{k+1} = \phi_{\gamma_k}^+(W_{\lambda_k}(X^k, y^{k+1})), \quad Z^{k+1} = \frac{1}{\lambda_k} \phi_{\gamma_k}^-(W_{\lambda_k}(X^k, y^{k+1})).$$

Step 3. If $\|(X^k - X^{k+1})/\lambda_k\| \leq \varepsilon$, stop; else, $\lambda_k = 2\lambda_k$; end.

Remark 3.1. Note that $b - \mathcal{A}(X^{k+1}) = b - \mathcal{A}\phi_{\gamma_k}^+(W_{\lambda_k}(X^k, y^{k+1})) = \nabla_y \Theta_{\lambda_k}(X^k, y^{k+1}) \approx 0$.

Remark 3.2. Observe that the function $\Theta_\lambda(X, y)$ is twice continuously differentiable (actually, analytic) in y . In contrast, its counterpart $L_\sigma(y, X)$ for a linear SDP in [36] fails to be twice continuously differentiable in y , and only the Clarke's generalized Jacobian of $\nabla_y L_\sigma(y, X)$ (i.e., $\partial \nabla_y L_\sigma(y, X)$) can be obtained. This difference can be attributed to the term $-\mu \log \det X$ in problem (P). In other words, $-\mu \log \det X$ works as a smoothing term that turns $L_\sigma(y, X)$ (which is not twice continuously differentiable) into an analytic function in y . This idea is different from the traditional smoothing technique of using a smoothing function on the KKT conditions since the latter are not motivated by adding a smoothing term to an objective function. Our derivation of $\Theta_\lambda(y, X)$ shows that the smoothing technique of using a squared smoothing function $\phi_\gamma^+(x) = (\sqrt{x^2 + 4\gamma} + x)/2$ can indeed be derived by adding the log-barrier term to the objective function.

The advantage of viewing the smoothing technique from the perspective of adding a log-barrier term is that the error between the minimum objective function values of the perturbed problem and the original problem can be estimated. In the traditional smoothing technique for the KKT conditions, there is no obvious means of estimating the error between the objective function value of the solution computed from the smoothed KKT conditions and the true minimum objective function value. We believe the connection we discovered here could be useful for the error analysis of the smoothing technique applied to the KKT conditions.

For the sake of subsequent convergence analysis, we present the following proposition.

PROPOSITION 3.1. *Suppose that (P) satisfies (9). Let $\bar{X} \in S_{++}^n$ be the unique optimal solution to (P), i.e., $\bar{X} = T_f^{-1}(0)$. Then T_f^{-1} is Lipschitz continuous at the origin.*

Proof. From [23, Prop. 3], it suffices to show that the following quadratic growth condition holds at \bar{X} for some positive constant α :

$$(27) \quad f(X) \geq f(\bar{X}) + \alpha \|X - \bar{X}\|^2 \quad \forall X \in \mathcal{N} \text{ such that } X \in \mathcal{F}_P,$$

where \mathcal{N} is a neighborhood of \bar{X} in S_{++}^n . From [4, Thm. 3.137], to prove (27), it suffices to show that the second-order sufficient condition for (P) holds.

Now for $X \in S_{++}^n$, we have

$$\langle \Delta X, \nabla_{XX}^2 l(X; y)(\Delta X) \rangle = \mu \langle X^{-1} \Delta X X^{-1}, \Delta X \rangle \geq \mu \lambda_{\max}^{-2}(X) \|\Delta X\|^2 \quad \forall \Delta X \in \mathcal{S}^n,$$

where $\lambda_{\max}(X)$ is the maximal eigenvalue of X , which is equivalent to

$$(28) \quad \langle \Delta X, \nabla_{XX}^2 l(X; y)(\Delta X) \rangle > 0 \quad \forall \Delta X \in \mathcal{S}^n \setminus \{0\}.$$

Certainly, (28) implies the second-order sufficient condition for problem (P). \square

We can also prove in parallel that the maximal monotone operator T_l is Lipschitz continuous at the origin.

4. The Newton-CG method for inner problems. In the algorithm framework proposed in section 3, we have to compute $y^{k+1} \approx \arg \sup \{\theta_k(y) : y \in \mathcal{R}^m\}$. In this paper, we will introduce the Newton-CG method to achieve this goal.

Algorithm 2: The Newton-CG method.

Step 1. Given $\mu \in (0, \frac{1}{2})$, $\tau_1, \tau_2 \in (0, 1)$, and $\delta \in (0, 1)$, choose $y^0 \in \mathcal{R}^m$.

Step 2. For $j = 0, 1, 2, \dots$,

Step 1.1. Apply the PCG method to find an approximate solution d^j of

$$(29) \quad (\nabla_{yy}^2 \theta_k(y^j) - \epsilon_j I)d = -\nabla_y \theta_k(y^j),$$

where $\epsilon_j := \tau_1 \min\{\tau_2, \|\nabla_y \theta_k(y^j)\|\}$.

Step 1.2. Set $\alpha_j = \delta^{m_j}$, where m_j is the first nonnegative integer m for which

$$\theta_k(y^j + \delta^m d^j) \geq \theta_k(y^j) + \mu \delta^m \langle \nabla_y \theta_k(y^j), d^j \rangle.$$

Step 1.3. Set $y^{j+1} = y^j + \alpha_j d^j$.

From (21) and the positive definiteness property of $\phi'_+(W(y; X))$ (for some properties of the projection operator one may refer to [16]), we have that $-\nabla_{yy}^2 \theta_k(y^j)$ is always positive definite; then $-\nabla_{yy}^2 \theta_k(y^j) + \epsilon_j I$ is positive definite as long as $\nabla_y \theta_k(y^j) \neq 0$. So we can always apply the PCG method to (29). Of course, the direction d^j generated from (29) is always an ascent direction. With respect to the analysis of the global convergence and local quadratic convergence rate of the above algorithm, we will not present the details, and one may refer to section 3.3 of [36] since it is very similar to the semismooth Newton-CG algorithm used in that paper. The difference lies in that d^j obtained from (29) in this paper is an approximate Newton direction; in contrast, d^j obtained from (61) in [36] is a semismooth Newton direction.

5. Convergence analysis. Global convergence and the local convergence rate of our Newton-CG PPA method to problems (P) and (D) can be directly derived from Rockafellar’s papers [22, 23] without much difficulty. For the sake of completeness, we state only the results below.

Since we cannot solve the inner problems exactly, we will use the following stopping criteria considered by Rockafellar [22, 23] for terminating Algorithm 2:

- (A) $\sup \theta_k(y) - \theta_k(y^{k+1}) \leq \epsilon_k^2 / 2\lambda_k, \quad \epsilon_k \geq 0, \sum_{k=0}^{\infty} \epsilon_k < \infty;$
- (B) $\sup \theta_k(y) - \theta_k(y^{k+1}) \leq \delta_k^2 / 2\lambda_k \|X^{k+1} - X^k\|^2, \quad \delta_k \geq 0, \sum_{k=0}^{\infty} \delta_k < \infty;$
- (B') $\|\nabla_y \theta_k(y^{k+1})\| \leq \delta'_k / \lambda_k \|X^{k+1} - X^k\|, \quad 0 \leq \delta'_k \rightarrow 0.$

In view of Proposition 3.1, we can directly obtain from [22, 23] the following convergence results.

THEOREM 5.1. *Let Algorithm 1 be executed with stopping criterion (A). If (D) satisfies condition (10), then the generated sequence $\{X^k\} \subset \mathcal{S}_{++}^n$ is bounded and $\{X^k\}$ converges to \bar{X} , where \bar{X} is the unique optimal solution to (P), and $\{y^k\}$ is asymptotically maximizing for (D) with $\min(P) = \sup(D)$.*

If $\{X^k\}$ is bounded and (P) satisfies condition (9), then the sequence $\{y^k\}$ is also bounded, and the accumulation point of the sequence $\{y^k\}$ is the unique optimal solution to (D).

THEOREM 5.2. *Let Algorithm 1 be executed with stopping criteria (A) and (B). Assume that (D) satisfies condition (10) and (P) satisfies condition (9). Then the generated sequence $\{X^k\} \subset \mathcal{S}_{++}^n$ is bounded and $\{X^k\}$ converges to the unique solution \bar{X} to (P) with $\min(\text{P}) = \sup(\text{D})$, and*

$$\|X^{k+1} - \bar{X}\| \leq \theta_k \|X^k - \bar{X}\| \quad \forall k \text{ sufficiently large,}$$

where

$$\theta_k = [a_f(a_f^2 + \sigma_k^2)^{-1/2} + \delta_k](1 - \delta_k)^{-1} \rightarrow \theta_\infty = a_f(a_f^2 + \sigma_\infty^2)^{-1/2} < 1, \quad \sigma_k \rightarrow \sigma_\infty,$$

and a_f is a Lipschitz constant of T_f^{-1} at the origin. The conclusions of Theorem 5.1 about $\{y^k\}$ are valid.

Moreover, if the stopping criterion (B') is also used, then in addition to the above conclusions the sequence $\{y^k\} \rightarrow \bar{y}$, where \bar{y} is the unique optimal solution to (D) and one has

$$\|y^{k+1} - \bar{y}\| \leq \theta'_k \|X^{k+1} - X^k\| \quad \forall k \text{ sufficiently large,}$$

where

$$\theta'_k = a_l(1 + \delta'_k)/\sigma_k \rightarrow \delta_\infty = a_l/\sigma_\infty$$

and a_l is a Lipschitz constant of T_l^{-1} at the origin.

Remark 5.1. In Algorithm 1 we can also add the term $-\frac{1}{2\lambda_k}\|y - y^k\|^2$ to $\theta_k(y)$. Actually, in our MATLAB code, one can optionally add this term. This actually corresponds to the PPA of multipliers considered in [23, sect. 5]. Convergence analysis for this improvement can be conducted in a way parallel to that of Algorithm 1.

Note that in the stopping criteria (A) and (B), $\sup \theta_k(y)$ is an unknown value. Since $\hat{\theta}_k(y) := \theta_k(y) - \frac{1}{2\lambda_k}\|y - y^k\|^2$ is a strongly concave function with modulus $\frac{1}{\lambda_k}$, one has the estimation

$$\sup \hat{\theta}_k(y) - \hat{\theta}_k(y^{k+1}) \leq \frac{1}{2}\lambda_k \|\nabla_y \hat{\theta}_k(y^{k+1})\|^2;$$

thus criteria (A) and (B) can be practically modified as follows:

$$\begin{aligned} \|\nabla_y \hat{\theta}_k(y^{k+1})\| &\leq \epsilon_k, \quad \epsilon_k \geq 0, \quad \sum_{k=0}^{\infty} \epsilon_k < \infty; \\ \|\nabla_y \hat{\theta}_k(y^{k+1})\| &\leq \delta_k \|X^{k+1} - X^k\|, \quad \delta_k \geq 0, \quad \sum_{k=0}^{\infty} \delta_k < \infty. \end{aligned}$$

6. Numerical experiments. In this section, we present some numerical results to demonstrate the performance of our PPA on (3) and (4), for Gaussian graphical models with both synthetic and real data. We implemented the PPA in MATLAB. All runs are performed on an Intel Xeon 3.20GHz PC with 4GB memory, running Linux and MATLAB (Version 7.6).

We measure the infeasibilities and optimality for the primal and dual problems (P) and (D) as follows:

$$(30) \quad R_D = \frac{\|C - \mathcal{A}^T y - Z\|}{1 + \|C\|}, \quad R_P = \frac{\|b - \mathcal{A}(X)\|}{1 + \|b\|}, \quad R_G = \frac{|\text{pobj} - \text{dobj}|}{1 + |\text{pobj}| + |\text{dobj}|},$$

where $\text{pobj} = \langle C, X \rangle - \mu \log \det X$ and $\text{dobj} = b^T y + \mu \log \det Z + n\mu(1 - \log \mu)$. The above measures are similar to those adopted in [36]. In our numerical experiments, we stop the PPA when

$$(31) \quad \max\{R_D, R_P\} \leq \text{To1},$$

where To1 is a pre-specified accuracy tolerance. Note that the third equation $XZ = \mu I$ in (11) holds up to machine precision because of the way we define Z in (26) in the PPA. Unless otherwise specified, we set $\text{To1} = 10^{-6}$ as the default. We choose the initial iterate $X^0 = I$, and $\lambda_0 = 1$.

We should note that in the PPA, computing the full eigenvalue decomposition of the matrix $W_{\lambda_k}(X^k, y)$ to evaluate the function $\Theta_{\lambda_k}(X^k, y)$ in (25) may constitute a major part of the overall computation. Thus it is essential for us to use an eigenvalue decomposition routine that is as efficient as possible. In our implementation, we use the LAPACK routine `dsyevd.f` (based on a divide-and-conquer strategy) to compute the full eigenvalue decomposition of a symmetric matrix. On our machine, it is about 7 to 10 times faster than the MATLAB `eig` routine when n is larger than 500. In the ANS method of [14, 13], having an efficient eigenvalue decomposition routine is even more crucial. Thus in our experiments, we also use the faster eigenvalue routine for the ANS method.

We focus our numerical experiments on the problems (3) and (4). The problem (4) is not expressed in the standard form given in (2), but it can easily be expressed as such by introducing additional constraints and variables. To be precise, the standard form reformulation of (4) is given as follows:

$$(32) \quad \begin{aligned} & \text{minimize} && \langle C, X \rangle - \log \det X + \rho^T x^+ + \rho^T x^- - \nu \log x^+ - \nu \log x^- \\ & \text{subject to} && X_{ij} = 0 \quad \forall (i, j) \in \Omega^c, \\ & && X_{ij} - x_{ij}^+ + x_{ij}^- = 0 \quad \forall (i, j) \notin \Omega^c, \\ & && X \succeq 0, \quad x^+, x^- \geq 0, \end{aligned}$$

where we set $\nu = 10^{-16}$ (originally, $\nu = 0$; however, for the convenience of the theory, we set $\nu = 10^{-16}$ in practical computation), $\rho, x^+, x^- \in \mathcal{R}^{m_2}$, $m_2 = m - m_1$, $m = \frac{n(n+1)}{2}$, and $m_1 = \frac{1}{2}|\Omega^c|$ ($|\Omega^c|$ denotes the number of the elements in the index set Ω^c).

We should emphasize that our algorithm is sensitive to the scaling of the data, especially for problem (4). Thus in our implementation, we first scale the data by setting $A_k \leftarrow A_k/\|A_k\|$, $C \leftarrow C/\|C\|$, and $b \leftarrow b/\|b\|$.

In this paper, we mainly compare the performance of our PPA with that of the ANS method in [13, 14], whose MATLAB codes are available at <http://www.math.sfu.ca/~zhaosong/>. The reason for comparing only with the ANS method is because it is currently the most advanced first-order method developed for solving the covariance selection problems (3) and (4). In the ANS method, the convergence criterion is controlled by two parameters, ϵ_o, ϵ_c , which stand for the errors in the objective value and primal infeasibility, respectively.

As mentioned in [33, 32], we may evaluate the performance of an estimator $\widehat{\Sigma}$ of the true covariance matrix Σ by a normalized L_2 -loss function which is defined as follows:

$$L_2^{\text{loss}} = \|\Sigma^{-1}\widehat{\Sigma} - I\|_F/n.$$

Thus in our numerical experiments, we also report the above value when it is possible to do so.

6.1. Some acceleration techniques.

6.1.1. A diagonal preconditioner for (29). To achieve faster convergence for the CG method to solve (29), one may apply a proper preconditioner to the linear system. But a suitable balance between having an effective preconditioner and additional computational cost must be determined. In our implementation, we devise an easy-to-compute diagonal preconditioner by using an idea first developed in [11].

Let M denote the coefficient matrix in the left-hand side of (29), which has the form

$$M := -\lambda \mathbf{A} \mathbf{T} \mathbf{A}^T - \epsilon I,$$

where \mathbf{A} and \mathbf{T} denote the matrix representations of the linear maps \mathcal{A} and \mathcal{T} with respect to the standard bases in \mathcal{S}^n and \mathcal{R}^m , respectively. Here $\mathcal{T} : \mathcal{S}^n \rightarrow \mathcal{S}^n$ is the linear operator defined by $\mathcal{T}(X) = P(\Omega \circ (P^T X P))P^T$.

Recall that the standard basis in \mathcal{S}^n is given by $\{E_{ij} := \alpha_{ij}(e_i e_j^T + e_j e_i^T) : 1 \leq i \leq j \leq n\}$, where e_i is the i th unit vector in \mathcal{R}^n and $\alpha_{ij} = 1/\sqrt{2}$ if $i \neq j$ and $\alpha_{ij} = 1/2$ otherwise. Then the diagonal element of \mathbf{T} with respect to the basis element E_{ij} is given by

$$(33) \quad \mathbf{T}_{(ij),(ij)} = \langle P^T E_{ij} P, \Omega \circ (P^T E_{ij} P) \rangle = \begin{cases} ((P \circ P)\Omega(P \circ P)^T)_{ij} + \langle v^{(ij)}, \Omega v^{(ij)} \rangle & \text{if } i \neq j, \\ ((P \circ P)\Omega(P \circ P)^T)_{ij} & \text{otherwise,} \end{cases}$$

where $v^{(ij)} = P_i \circ P_j$ and P_i, P_j are the i th and j th rows of P , respectively. From (33) we can see that to compute all the diagonal entries of \mathbf{T} , the computing cost of $O(n^4)$ flops is needed. Fortunately, the first term

$$\mathbf{d}_{(ij)} := ((P \circ P)\Omega(P \circ P)^T)_{ij}$$

in the right-hand side of (33) is typically a very good approximation of $\mathbf{T}_{(ij),(ij)}$. More importantly, computing all the elements $\mathbf{d}_{(ij)}$, for $1 \leq i \leq j \leq n$, needs only $O(n^3)$ flops since only the matrix product $(P \circ P)\Omega(P \circ P)^T$ is involved. We propose the following diagonal preconditioner for M :

$$M_D := -\lambda \text{diag}(\mathbf{A} \text{diag}(\mathbf{d}) \mathbf{A}^T) - \epsilon I.$$

6.1.2. An alternating Newton PPA. In this subsection, we introduce an alternating Newton PPA to accelerate the convergence of the primal PPA for solving problem (P).

For the PPA, we know that the outer iteration is actually a gradient method, i.e.,

$$(34) \quad X^{k+1} = X^k - \lambda_k \nabla F_{\lambda_k}(X^k) = \phi_{\gamma_k}^+(W_{\lambda_k}(X^k, y_{\lambda_k}(X^k))).$$

To improve the convergence rate of the outer iteration, we may use Newton's method instead. Since

$$\nabla F_{\lambda}(X) = \frac{1}{\lambda} \left(X - \phi_{\gamma}^+(W_{\lambda}(X, y_{\lambda}(X))) \right),$$

we have that

$$(35) \quad \nabla^2 F_\lambda(X)[H] = \frac{1}{\lambda} [H - (\phi_\gamma^+)'(W)(H + \lambda \mathcal{A}^T y'_\gamma(X; H))] \quad \forall H \in \mathcal{S}^n.$$

Here we used W to denote $W_\lambda(X, y_\lambda(X))$ to simplify the notation. Thus, the Newton direction H^k is the solution to the following linear system of equations:

$$(36) \quad \nabla^2 F_\lambda(X^k)[H^k] = -\nabla F_\lambda(X^k).$$

Once H^k has been computed, the new iteration is updated by (if the iteration X^k is sufficiently close to the optimal solution)

$$(37) \quad X^{k+1} = X^k + H^k.$$

In practical computation, we adopt the update based on Newton's method in (37) (without line search) when the relative primal infeasibility and dual feasibility are both less than 10^{-2} ; otherwise, we use the update in (34). We call the PPA with its outer iteration updated possibly by Newton's update (37) the alternating Newton PPA (ANPPA).

Next we discuss how the Newton system (36) can be solved by the CG method. To apply the CG method, all we need is to be able to evaluate $\nabla^2 F_\lambda(X^k)[H]$ given any $H \in \mathcal{S}^n$. Now we discuss in detail how this can be done. Observe that since

$$0 = \nabla_y \Theta_\lambda(X, y_\lambda(X)) = b - \mathcal{A} \phi_\gamma^+(W_\lambda(X, y_\lambda(X))),$$

it follows that

$$0 = -\mathcal{A}(\phi_\gamma^+)'(W)[H + \lambda \mathcal{A}^T y'_\gamma(X; H)]$$

and

$$(38) \quad \lambda \mathcal{A}(\phi_\gamma^+)'(W) \mathcal{A}^T y'_\gamma(X; H) = -\mathcal{A}(\phi_\gamma^+)'(W)[H].$$

Thus from the linear system (38), we can compute $y'_\gamma(X; H)$ given any X and H . By substituting it into (35), we can evaluate the expression $\nabla^2 F_\lambda(X)[H]$ for any given X and H .

To summarize, given X and H , we can evaluate $\nabla^2 F_\lambda(X)[H]$ by solving the linear system (38) for $y'_\gamma(X; H)$. Observe that the linear system (38) has exactly the same form as the linear system (29); thus whatever techniques we have developed to solve (29) can be used to solve (38). Note that we terminate the CG method for solving (36) when the relative residual norm is less than 5×10^{-2} .

6.2. Synthetic experiments I. All instances used in this section were randomly generated in a manner similar to that described in d'Aspremont, Banerjee, and El Ghaoui [8]. Indeed, we generate a random sparse positive definite matrix $\Sigma^{-1} \in \mathcal{S}_{++}^n$ with a density of about 10% nonzero entries as follows. First we generate an $n \times n$ random sparse matrix U with nonzero entries set to ± 1 . Then we set

$$A = U' * U; \quad d = \text{diag}(A); \quad A = \max(\min(A - \text{diag}(d), 1), -1);$$

$$B = A + \text{diag}(1 + d);$$

$$\Sigma^{-1} = B + \max(-1.2 * \min(\text{eig}(B)), 0.001) * I;$$

TABLE 1
Performance of the ANPPA on (3) with synthetic data (I).

Problem	m n	$it/itsub/pcg$	pobj	dobj	$R_P/R_D/R_G/L_2^{\text{loss}}$	Time
rand-500	112172 500	12 23 9.9	-1.85356498 2	-1.85356176 2	4.1-7 2.9-7 8.7-7 1.8-2	43.1
rand-1000	441294 1000	11 22 10.0	-1.22155120 2	-1.22155034 2	8.0-8 4.7-8 3.5-7 6.7-2	247.0
rand-1500	979620 1500	11 22 8.9	-4.32799516 2	-4.32799143 2	2.2-7 1.4-7 4.3-7 6.8-2	698.0
rand-2000	1719589 2000	12 23 8.7	-1.34583154 3	-1.34583004 3	8.1-7 3.5-7 5.6-7 4.3-2	1581.4

The sample covariance matrix S for (3) is generated in a manner similar to that in [8, 13] via the following script:

$$\mathbf{E} = 2 * \text{rand}(n) - 1; \quad \mathbf{E} = 0.5 * (\mathbf{E} + \mathbf{E}');$$

$$\mathbf{S} = \Sigma + 0.15 * (\|\Sigma\|_{\mathbb{F}} / \|\mathbf{E}\|_{\mathbb{F}}) * \mathbf{E};$$

$$\mathbf{S} = \mathbf{S} + \max(-\min(\text{eig}(\mathbf{S})), 0.001) * \mathbf{I};$$

The set Ω^c is generated as in the MATLAB codes developed by Lu for the paper [13], specifically,

$$\Omega^c = \{(i, j) : (\Sigma^{-1})_{ij} = 0, |i - j| \geq 5\}.$$

In this synthetic experiment, we apply the ANPPA to problem (3). The performance of the ANPPA is presented in Table 1. For each instance in the table, we report the matrix dimension (n); the number of linear constraints (m); the number of outer iterations (it), the total number of subproblems ($itsub$) solved by the ANPPA, and the average number of PCG steps (pcg) taken to solve each linear system in (29); the primal ($pobj$) and dual ($dobj$) objective values; the primal (R_P) and dual (R_D) infeasibilities, the relative gap (R_G), and L_2^{loss} ; and the time (in seconds) taken. We may observe from the table that the ANPPA can very efficiently solve problem (3) with synthetic data.

In Table 2, we compare the performance of our ANPPA and the ANS method on the first three instances reported in Table 1. For the ANPPA, To1 in (31) is set to 10^{-6} , 10^{-7} , and 10^{-8} ; for ANS, ϵ_o is set to 10^{-1} , 10^{-2} , and 10^{-3} , and ϵ_c is set to 10^{-4} , so that the gap ($= |pobj - dobj|$) can fall below 10^{-1} , 10^{-2} , and 10^{-3} , respectively. For each instance in the table, we give the matrix dimension (n) and the number of linear constraints (m), the gaps achieved, and the times taken (in seconds). From Table 2, we can see that both methods are able to solve all instances within a reasonable amount of time. However, the ANPPA consistently outperforms the ANS method by a factor which ranges from 3 to 15.

6.3. Synthetic experiments II. We note that the procedure used in [8] to generate the data matrix S for (3) is not in line with the standard practice in statistics. But since the covariance selection problem is a problem in statistics, we prefer to generate the data matrix S according to the standard practice; see, for example, [33, 32]. Thus in this subsection, the true covariance matrices Σ and the index sets Ω^c are generated exactly the same way as in the previous subsection. But the sample covariance matrices S are generated differently. For each test problem, we sample $2n$ instances from the multivariate Gaussian distribution $N(0, \Sigma)$ to generate a sample covariance matrix S .

In the first synthetic experiment, we apply the ANPPA to problem (3). The performance of the ANPPA is presented in Table 3. Again, the ANPPA can very

TABLE 2
Comparison of the ANPPA and the ANS method on (3) with synthetic data (I).

Problem	$m n$	Tolerance		Iteration		Gap		Time	
		PPA (To1)	ANS (ϵ_o, ϵ_c)	PPA	ANS	PPA	ANS	PPA	ANS
rand-500	112172 500	3×10^{-6}	$(10^{-1}, 10^{-4})$	23	534	3.22-4	9.70-2	43.1	133.8
		3×10^{-7}	$(10^{-2}, 10^{-4})$	27	1255	1.18-6	9.94-3	55.2	285.9
		3×10^{-8}	$(10^{-3}, 10^{-4})$	27	2945	1.18-6	9.94-4	55.2	646.5
rand-1000	441294 1000	3×10^{-6}	$(10^{-1}, 10^{-4})$	22	826	8.52-5	9.66-2	247.0	1137.5
		3×10^{-7}	$(10^{-2}, 10^{-4})$	22	1916	8.52-5	9.94-3	247.0	2407.9
		3×10^{-8}	$(10^{-3}, 10^{-4})$	25	3924	2.39-6	9.94-4	304.8	4757.5
rand-1500	979620 1500	3×10^{-6}	$(10^{-1}, 10^{-4})$	22	778	3.73-4	9.94-2	698.0	3303.7
		3×10^{-7}	$(10^{-2}, 10^{-4})$	22	1741	3.73-4	9.94-3	698.0	6706.3
		3×10^{-8}	$(10^{-3}, 10^{-4})$	25	3452	1.17-5	9.94-4	857.0	12766.2

TABLE 3
Performance of the ANPPA on (3) with synthetic data (II).

Problem	$m n$	$it/it_{sub}/pcg$	pobj	dobj	$R_P/R_D/R_G/L_2^{loss}$	Time
rand-500	112172 500	13 27 13.2	-3.13591727 2	-3.13591672 2	2.6-8 3.8-8 8.8-8 1.7-2	61.4
rand-1000	441294 1000	13 29 18.9	-9.74364421 2	-9.74360450 2	3.1-7 6.0-7 2.0-6 2.0-2	468.8
rand-1500	979620 1500	13 29 15.8	-1.91034252 3	-1.91033842 3	7.1-7 3.5-7 1.1-6 1.8-2	1384.5
rand-2000	1719589 2000	15 34 15.9	-3.00395927 3	-3.00395919 3	1.7-8 5.0-9 1.4-8 1.5-2	3696.2




FIG. 1. Accuracy versus time for the random instance rand-500 reported in Table 3.

efficiently solve problem (3) with S generated from $2n$ samples of the Gaussian distribution $N(0, \Sigma)$. Comparing with Table 1, it appears that the log-det problems in Table 3 are harder to solve when n is large.

In Figure 1, we show that the ANPPA can also obtain a very accurate solution for the instance rand-500 reported in Table 3 without incurring a substantial amount of additional computing time. As can be seen from the figure, the time taken only grows almost linearly when the required accuracy is geometrically reduced.

In Table 4, we compare the performance of our ANPPA and the ANS method on

TABLE 4
Comparison of the ANPPA and the ANS method on (3) with synthetic data (II).

Problem	$m n$	Tolerance		Iteration		Gap		Time	
		PPA (To1)	ANS (ϵ_o, ϵ_c)	PPA	ANS	PPA	ANS	PPA	ANS
rand-500	112172 500	3×10^{-6}	$(10^{-1}, 10^{-4})$	27	2358	5.53-5	9.79-2	61.4	518.3
		3×10^{-7}	$(10^{-2}, 10^{-4})$	27	5779	5.53-5	9.94-3	61.4	1233.2
		3×10^{-8}	$(10^{-3}, 10^{-4})$	30	12796	9.04-8	9.94-4	74.1	2712.3
rand-1000	441294 1000	3×10^{-6}	$(10^{-1}, 10^{-4})$	29	3776	3.97-3	9.91-2	468.8	4499.8
		3×10^{-7}	$(10^{-2}, 10^{-4})$	33	10055	3.86-6	9.94-3	583.5	11715.4
		3×10^{-8}	$(10^{-3}, 10^{-4})$	33	22519	3.86-6	9.94-4	583.5	26173.1
rand-1500	979620 1500	3×10^{-6}	$(10^{-1}, 10^{-4})$	29	3691	4.10-3	9.94-2	1384.5	13601.7
		3×10^{-7}	$(10^{-2}, 10^{-4})$	33	9027	1.05-4	9.94-3	1699.4	32440.2
		3×10^{-8}	$(10^{-3}, 10^{-4})$	33	18408	1.05-4	9.94-4	1699.4	65773.7

TABLE 5
Performance of the ANPPA on (4) with synthetic data (II).

Problem	$m n$	$it/itsub/pcg$			pobj	dobj	$R_P/R_D/R_G/L_2^{\text{loss}}$				Time
		it	itsub	pcg			R_P	R_D	R_G	L_2^{loss}	
rand-500	112172 500	15	50	12.5	-3.11255742 2	-3.11256007 2	1.9-8	1.7-7	4.2-7	1.7-2	105.7
rand-1000	441294 1000	17	60	18.3	-9.70441465 2	-9.70441034 2	2.1-8	1.4-7	2.2-7	2.0-2	942.2
rand-1500	979620 1500	18	55	16.3	-1.90500086 3	-1.90499588 3	8.4-8	4.7-7	1.3-6	1.8-2	2498.9
rand-2000	1719589 2000	19	53	16.8	-2.99725089 3	-2.99724734 3	1.9-8	4.4-7	5.9-7	1.5-2	5429.7

the first three instances reported in Table 3. For the ANPPA, To1 in (31) is set to 3×10^{-6} , 3×10^{-7} , and 3×10^{-8} ; for ANS, ϵ_o is set to 10^{-1} , 10^{-2} , and 10^{-3} , and ϵ_c is set to 10^{-4} , so that the gap ($= |\text{pobj} - \text{dobj}|$) can fall below 10^{-1} , 10^{-2} , and 10^{-3} , respectively. From Table 4, we can see that the ANPPA consistently outperforms the ANS method by a substantial margin, which ranges from a factor of 8 to 44. It is interesting to note that while the computing time of the ANPPA grows only modestly when the required accuracy tolerance is reduced by a factor of 10, the computing time for the ANS method grows by at least a factor of 2.

In the second synthetic experiment, we consider problem (4). We set $\rho_{ij} = 1/n^{1.5}$ for all $(i, j) \notin \Omega^c$. We note that the parameters ρ_{ij} are chosen empirically so as to give a reasonably good value for $\|\Sigma - \hat{\Sigma}\|_F$.

In Tables 5 and 6, we report the results in a format similar to those in Tables 3 and 4, respectively. Again, we may observe from the tables that the ANPPA outperformed the ANS method by a substantial margin.

6.4. Real data experiments. In this section, we compare the ANPPA and the ANS method on two gene expression data sets. Since the authors of [2] have already considered these data sets, we can refer the reader to [2] for the choice of the parameters ρ_{ij} .

6.4.1. Rosetta Inpharmatics Compendium. We applied our ANPPA and the ANS method to the Rosetta Inpharmatics Compendium of gene expression profiles described by Hughes et al. [12]. The data set contains 253 samples with $n = 6136$ variables. We aim to estimate the sparse covariance matrix of a Gaussian graphic model whose conditional independence is unknown. Naturally, we formulate it as problem (4), with $\Omega^c = \emptyset$. As for the parameters, we set $\rho_{ij} = 0.0313$ as in [2].

TABLE 6
Comparison of the ANPPA and the ANS method on (4) with synthetic data (II).

Problem	$m \mid n$	Tolerance		Iteration		Gap		Time	
		PPA (To1)	ANS (ϵ_o, ϵ_c)	PPA	ANS	PPA	ANS	PPA	ANS
rand-500	112172 500	3×10^{-6}	($10^{-1}, 10^{-4}$)	50	2327	2.65-4	9.90-2	105.7	510.3
		3×10^{-7}	($10^{-2}, 10^{-4}$)	50	5818	2.65-4	9.94-3	105.7	1236.9
		3×10^{-8}	($10^{-3}, 10^{-4}$)	55	12962	6.65-7	9.94-4	121.4	2747.0
rand-1000	441294 1000	3×10^{-6}	($10^{-1}, 10^{-4}$)	60	3741	4.31-4	9.88-2	942.2	4460.8
		3×10^{-7}	($10^{-2}, 10^{-4}$)	60	9931	4.31-4	9.94-3	942.2	11562.8
		3×10^{-8}	($10^{-3}, 10^{-4}$)	65	22620	1.88-5	9.94-4	1090.4	26278.6
randd-1500	979620 1500	3×10^{-6}	($10^{-1}, 10^{-4}$)	55	3745	4.98-3	9.90-2	2498.9	13830.3
		3×10^{-7}	($10^{-2}, 10^{-4}$)	59	9572	2.25-4	9.94-3	2855.2	34208.0
		3×10^{-8}	($10^{-3}, 10^{-4}$)	59	20895	2.25-4	9.94-4	2855.2	73997.1

TABLE 7
Comparison of the ANPPA and ANS method on (4) with $\Omega^c = \emptyset$ for the Rosetta Inpharmatics Compendium data.

Problem	$m \mid n$	Tolerance		Iteration		Primal objective value		Time	
		PPA (To1)	ANS (ϵ_o)	PPA	ANS	PPA	ANS	PPA	ANS
Rosetta	500	10^{-6}	10^{-3}	40	636	-7.42642518 2	-7.42642052 2	57.0	127.6
Rosetta	1000	10^{-6}	10^{-3}	40	763	-1.66546366 3	-1.66546478 3	298.8	881.6
Rosetta	1500	10^{-6}	10^{-3}	42	972	-2.64937351 3	-2.64937721 3	914.4	3424.7

As our ANPPA can handle only problems with matrix dimensions up to about 3000, we test only on a subset of the data. We create 3 subsets by taking 500, 1000, and 1500 variables with the highest variances, respectively. Note that as the variances vary widely, we normalized the sample covariance matrices to have unit variances on the diagonal.

In the experiments, we set To1 = 10^{-6} for the ANPPA, and $(\epsilon_o, \epsilon_c) = (10^{-2}, 10^{-6})$ for the ANS method.

The performances of the ANPPA and ANS methods for the Rosetta Inpharmatics Compendium of gene expression profiles are presented in Table 7. From Table 7, we can see that although both methods can solve the problem, the ANPPA is about 3.7 times faster than the ANS method when $n = 1500$.

6.4.2. Iconix microarray data. Next we analyzed the performances of the ANPPA and ANS methods on a subset of a 10000 gene microarray data set obtained from 255 drug-treated rat livers; see Natsoulis et al. [19] for details. In our first test problem, we took 200 variables with the highest variances from the large set to form the sample covariance matrix S . The other 2 test problems were created by considering 500 and 1000 variables with the highest variances in the large data set. As in the last data set, we normalized the sample covariance matrices to have unit variances on the diagonal.

For the same reason as mentioned earlier, we set $\Omega^c = \emptyset$ in problem (4). We set $\rho_{ij} = 0.0853$ as in [2].

TABLE 8

Comparison of the ANPPA and ANS method on (4) with $\Omega^c = \emptyset$ for the Iconix microarray data.

Problem	m n	Tolerance		Iteration		Primal objective value		Time	
		PPA (Tol)	ANS (ϵ_o)	PPA	ANS	PPA	ANS	PPA	ANS
Iconix	200	10^{-6}	10^{-3}	43	1805	-6.13127781 0	-6.13036186 0	17.7	50.7
Iconix	500	10^{-6}	10^{-3}	54	3809	5.31683802 1	5.31688551 1	222.9	795.2
Iconix	1000	10^{-6}	10^{-3}	65	6646	1.78893452 2	1.78892330 2	1585.2	7847.3

The performance of the ANPPA and ANS methods for the Iconix microarray data is presented in Table 8. From the table, we see that the ANPPA is about 5 times faster than the ANS method when $n = 1000$.

7. Concluding remarks. We designed a primal PPA to solve log-det optimization problems. Rigorous convergence results for the PPA are obtained from the classical results for a generic PPA. We also considered accelerating the outer iteration of PPA by Newton's method and designed an alternating Newton primal PPA. Extensive numerical experiments conducted on log-det problems arising from sparse estimation of inverse covariance matrices in Gaussian graphical models using synthetic data and real data demonstrated that our ANPPA is very efficient.

In contrast to the case for the linear SDP problems, the log-det term used in this paper plays a key role of a smoothing term such that the standard smooth Newton method can be used to solve the inner problem. The key discovery of this paper is the connection of the log-det smoothing term to the technique of using the squared smoothing function. It opens up a new door for dealing with nonsmooth equations and understanding the smoothing technique more deeply.

Acknowledgments. We thank Onureena Banerjee for providing us with part of the test data and helpful suggestions, and Zhaosong Lu and Xiaoming Yuan for sharing with us their MATLAB codes and fruitful discussions. We also thank the two anonymous referees and the associate editor for their helpful comments and suggestions, which improved the quality of this paper.

REFERENCES

- [1] F. ALIZADEH, J.-P. A. HAEBERLY, AND M. L. OVERTON, *Complementarity and nondegeneracy in semidefinite programming*, Math. Programming Ser. B, 77 (1997), pp. 111–128.
- [2] O. BANERJEE, L. EL GHAOU, AND A. D'ASPREMONT, *Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data*, J. Mach. Learn. Res., 9 (2008), pp. 485–516.
- [3] R. BHATIA, *Matrix Analysis*, Springer-Verlag, New York, 1997.
- [4] J. F. BONNANS AND A. SHAPIRO, *Perturbation Analysis of Optimization Problems*, Springer-Verlag, New York, 2000.
- [5] S. BOYD, L. EL GHAOU, E. FERON, AND V. BALAKRISHNAN, *Linear Matrix Inequalities in System and Control Theory*, Stud. Appl. Math. 15, SIAM, Philadelphia, 1994.
- [6] Z. X. CHAN AND D. SUN, *Constraint nondegeneracy, strong regularity, and nonsingularity in semidefinite programming*, SIAM J. Optim., 19 (2008), pp. 370–396.
- [7] J. DAHL, L. VANDENBERGHE, AND V. ROYCHOWDHURY, *Covariance selection for non-chordal graphs via chordal embedding*, Optim. Methods Softw., 23 (2008), pp. 501–520.
- [8] A. D'ASPREMONT, O. BANERJEE, AND L. EL GHAOU, *First-order methods for sparse covariance selection*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 56–66.
- [9] A. DEMPSTER, *Covariance selection*, Biometrics, 28 (1972), pp. 157–175.

- [10] I. S. DHILLON AND J. A. TROPP, *Matrix nearness problems with Bregman divergences*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 1120–1146.
- [11] Y. GAO AND D. SUN, *Calibrating least squares semidefinite programming with equality and inequality constraints*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 1432–1457.
- [12] T. R. HUGHES, M. J. MARTON, A. R. JONES, C. J. ROBERTS, R. STOUGHTON, C. D. ARMOUR, H. A. BENNETT, E. COFFEY, H. DAI, Y. D. HE, M. J. KIDD, A. M. KING, M. R. MEYER, D. SLADE, P. Y. LUM, S. B. STEPANIANTS, D. D. SHOEMAKER, D. GACHOTTE, K. CHAKRABURTTY, J. SIMON, M. BARD, AND S. H. FRIEND, *Functional discovery via a compendium of expression profiles*, Cell, 102 (2000), pp. 109–126.
- [13] Z. LU, *Smooth optimization approach for sparse covariance selection*, SIAM J. Optim., 19 (2009), pp. 1807–1827.
- [14] Z. LU, *Adaptive first-order methods for general sparse inverse covariance selection*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 2000–2016.
- [15] B. MARTINET, *Régularisation d'inéquations variationnelles par approximations successives*, Rev. Française Informat. Recherche Opérationnelle, 4 (1970), pp. 154–158.
- [16] F. MENG, D. SUN, AND G. ZHAO, *Semismoothness of solutions to generalized equations and the Moreau-Yosida regularization*, Math. Program. Ser. B, 104 (2005), pp. 561–581.
- [17] G. J. MINTY, *On the monotonicity of the gradient of a convex function*, Pacific J. Math., 14 (1964), pp. 243–247.
- [18] J. J. MOREAU, *Proximité et dualité dans un espace Hilbertien*, Bull. Soc. Math. France, 93 (1965), pp. 273–299.
- [19] G. NATSOULIS, C. I. PEARSON, J. GOLLUB, B. P. EYNON, J. FERNG, R. NAIR, R. IDURY, M. D. LEE, M. R. FIELDEN, R. J. BRENNAN, A. H. ROTER, AND K. JARNAGIN, *The liver pharmacological and xenobiotic gene response repertoire*, Mol. Syst. Biol., 175 (2008), pp. 1–12.
- [20] YU. NESTEROV, *Smooth minimization of non-smooth functions*, Math. Program. Ser. A, 103 (2005), pp. 127–152.
- [21] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.
- [22] R. T. ROCKAFELLAR, *Monotone operators and the proximal point algorithm*, SIAM J. Control Optim., 14 (1976), pp. 877–898.
- [23] R. T. ROCKAFELLAR, *Augmented Lagrangians and applications of the proximal point algorithm in convex programming*, Math. Oper. Res., 1 (1976), pp. 97–116.
- [24] R. T. ROCKAFELLAR AND R. J. B. WETS, *Variational Analysis*, Springer-Verlag, Berlin, 1998.
- [25] K. SCHEINBERG AND I. RISH, *Learning sparse Gaussian Markov networks using a greedy coordinate ascent approach*, in Machine Learning and Knowledge Discovery in Databases, Lecture Notes in Comput. Sci. 6323, J. L. Balcázar, F. Bonchi, A. Gionis, and M. Sebag, eds., Springer-Verlag, Berlin, 2010, pp. 196–212.
- [26] D. SUN, J. SUN, AND L. ZHANG, *The rate of convergence of the augmented Lagrangian method for nonlinear semidefinite programming*, Math. Program. Ser. A, 114 (2008), pp. 349–391.
- [27] M. TEBoulLE, *Entropic proximal mappings with applications to nonlinear programming*, Math. Oper. Res., 17 (1992), pp. 670–690.
- [28] K. C. TOH, *Primal-dual path-following algorithms for determinant maximization problems with linear matrix inequalities*, Comput. Optim. Appl., 14 (1999), pp. 309–330.
- [29] N.-K. TSING, M. K. H. FAN, AND E. I. VERRIEST, *On analyticity of functions involving eigenvalues*, Linear Algebra Appl., 207 (1994), pp. 159–180.
- [30] R. H. TÛTÜNCÜ, K. C. TOH, AND M. J. TODD, *Solving semidefinite-quadratic-linear programs using SDPT3*, Math. Program. Ser. B, 95 (2003), pp. 189–217.
- [31] L. VANDENBERGHE, S. BOYD, AND S.-P. WU, *Determinant maximization with linear matrix inequality constraints*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 499–533.
- [32] F. WONG, C. K. CARTER, AND R. KOHN, *Efficient estimation of covariance selection models*, Biometrika, 90 (2003), pp. 809–830.
- [33] W. B. WU AND M. POURAHMADI, *Nonparametric estimation of large covariance matrices of longitudinal data*, Biometrika, 90 (2003), pp. 831–844.
- [34] K. YOSIDA, *Functional Analysis*, Springer-Verlag, Berlin, 1964.
- [35] X. YUAN, *Alternating Direction Methods for Sparse Covariance Selection*, http://www.optimization-online.org/DB_HTML/2009/09/2390.html (4 September 2009).
- [36] X.-Y. ZHAO, D. SUN, AND K.-C. TOH, *A Newton-CG augmented Lagrangian method for semidefinite programming*, SIAM J. Optim., 20 (2010), pp. 1737–1765.