

FIRST-ORDER METHODS FOR SPARSE COVARIANCE SELECTION*

ALEXANDRE D'ASPREMONT[†], ONUREENA BANERJEE[‡], AND LAURENT EL GHAOUI[‡]

Abstract. Given a sample covariance matrix, we solve a maximum likelihood problem penalized by the number of nonzero coefficients in the inverse covariance matrix. Our objective is to find a sparse representation of the sample data and to highlight conditional independence relationships between the sample variables. We first formulate a convex relaxation of this combinatorial problem, we then detail two efficient first-order algorithms with low memory requirements to solve large-scale, dense problem instances.

Key words. covariance selection, semidefinite programming, coordinate descent

AMS subject classifications. 90C22, 62H20, 90C27

DOI. 10.1137/060670985

1. Introduction. We discuss a problem of model selection.¹ 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 [7], setting a certain number of coefficients in the inverse covariance matrix Σ^{-1} to zero, a procedure known as *covariance selection*, improves the stability of this estimation procedure by reducing the number of parameters to estimate and highlight structure in the underlying model.

Here, we focus on the problem of *discovering* this pattern of zeroes in the inverse covariance matrix. We seek to trade-off the log-likelihood of the solution with the number of zeroes in its inverse, and solve the following estimation problem:

$$(1.1) \quad \begin{aligned} & \text{maximize} && \log \det X - \langle \Sigma, X \rangle - \rho \mathbf{Card}(X) \\ & \text{subject to} && \alpha \mathbf{I}_n \preceq X \preceq \beta \mathbf{I}_n \end{aligned}$$

in the variable $X \in \mathbf{S}_n$, where $\Sigma \in \mathbf{S}_n^+$ is the sample covariance matrix, $\mathbf{Card}(X)$ is the cardinality of X , i.e., the number of nonzero components in X , $\rho > 0$ is a parameter controlling the trade-off between log-likelihood and cardinality, and finally, $\alpha, \beta > 0$ fix bounds on the eigenvalues of the solution.

Zeroes in the inverse covariance matrix correspond to conditionally independent variables in the model and this approach can be used to simultaneously determine a robust estimate of the covariance matrix and, perhaps more importantly, discover *structure* in the underlying graphical model. In particular, we can view (1.1) as a model selection problem using Akaike (AIC, see [1]) or Bayes (BIC, see [5]) information criterions. Both these problems can be written as in (1.1) with $\rho = 2/N$ for the AIC problem and $\rho = 2 \log(N/2)/N$ for the BIC problem, where N is the sample size.

*Received by the editors September 28, 2006; accepted for publication (in revised form) by L. Vandenberghe May 31, 2007; published electronically January 25, 2008. This research was supported by NSF grant DMS-0625352, EUROCONTROL grant C20083E/BM/05, and a gift from Google, Inc.

<http://www.siam.org/journals/simax/30-1/67098.html>

[†]ORFE Department, Princeton University, Princeton, NJ 08544 (aspremon@princeton.edu).

[‡]EECS Department, UC Berkeley, Berkeley, CA 94720 (onureena@eecs.berkeley.edu, elghaoui@eecs.berkeley.edu).

¹A subset of the results discussed here appeared in the Proceedings of the International Conference on Machine Learning, Pittsburgh, PA, 2006.

This has applications in speech recognition (see [2, 3]) or gene networks analysis (see [9, 8], for example).

The $\mathbf{Card}(X)$ penalty term makes the estimation problem (1.1) combinatorial (NP-hard in fact), and our first objective here is to derive a convex relaxation to this problem which can be solved efficiently. We then derive two first-order algorithms geared towards memory efficiency and large-scale, *dense* problem instances.

In [3], Bilmes proposed a method for covariance selection based on choosing statistical dependencies according to conditional mutual information computed using training data. Other recent work involves identifying those Gaussian graphical models that are best supported by the data and any available prior information on the covariance matrix. This approach is used by [13, 9] on gene expression data. Recently, [6, 12] also considered penalized maximum likelihood estimation for covariance selection. In contrast to our results here, [12] works on the Cholesky decomposition of X using an iterative (heuristic) algorithm to minimize a nonconvex penalized likelihood problem, while [6] proposes a set of large scale interior point algorithms to solve sparse problems, i.e., problems for which the conditional independence structure is already known.

The paper is organized as follows, in section 2, we detail our convex relaxation of problem (1.1) and study the dual. In section 3, we derive two efficient algorithms to solve it. Finally, in section 4 we describe some numerical results.

2. Problem setup.

2.1. Convex relaxation. Given a sample covariance matrix $\Sigma \in \mathbf{S}_n^+$, we can write the following convex relaxation to the estimation problem (1.1):

$$(2.1) \quad \begin{aligned} & \text{maximize} && \log \det X - \langle \Sigma, X \rangle - \rho \mathbf{1}^T |X| \mathbf{1} \\ & \text{subject to} && \alpha \mathbf{I}_n \preceq X \preceq \beta \mathbf{I}_n, \end{aligned}$$

with variable $X \in \mathbf{S}^n$, where $\mathbf{1}$ is the n -vector of ones, so that $\mathbf{1}^T |X| \mathbf{1} = \sum_{i,j=1}^n |X_{ij}|$. The penalty term involving the sum of absolute values of the entries of X is a proxy for the number of its nonzero elements: the function $\mathbf{1}^T |X| \mathbf{1}$ can be seen as the largest convex lower bound on $\mathbf{Card}(X)$ on the hypercube, an argument used by [11] for rank minimization. It is also often used in regression techniques, such as the LASSO studied in [19], when sparsity of the solution is a concern. This relaxation is provably tight in certain cases (see [10]). In our model, the bounds (α, β) on the eigenvalues of X are fixed and user-chosen. Although we allow $\alpha = 0$ and $\beta = +\infty$, such bounds are useful in practice to control the condition number of the solution.

When $\alpha = 0$ and $\beta = +\infty$, for $\rho = 0$, provided $\Sigma \succ 0$, problems (1.1) and (2.1) have a unique solution $X^* = \Sigma^{-1}$, and the corresponding maximum-likelihood estimate is Σ . Due to noise in the data, in practice, the sample estimate Σ may not have a sparse inverse, even if the underlying graphical model exhibits conditional independence properties. By striking a trade-off between the maximality of the likelihood and the number of nonzero elements in the inverse covariance matrix, our approach is potentially useful at *discovering structure*, precisely conditional independence properties in the data. This means that we have to focus on the case where the matrix X is *dense*. At the same time, it serves as a regularization technique: when Σ is rank-deficient, there is no well-defined maximum-likelihood estimate, whereas the solution to problem (2.1) is always unique and well defined for $\rho > 0$, as seen later.

2.2. Dual problem, robustness. We can rewrite the relaxation (2.1) as the following min-max problem:

$$(2.2) \quad \max_{\{X: \alpha \mathbf{I}_n \preceq X \preceq \beta \mathbf{I}_n\}} \min_{\{U: |U_{ij}| \leq \rho\}} \log \det X - \langle \Sigma + U, X \rangle,$$

which gives a natural interpretation of problem (2.1) as a worst-case *robust maximum likelihood* problem with componentwise bounded, additive noise on the sample covariance matrix Σ . The corresponding Lagrangian is given by

$$L(X, U, P, Q) = \log \det X - \mathbf{Tr}((\Sigma + U + Q - P)X) - \alpha \mathbf{Tr} P + \beta \mathbf{Tr} Q,$$

and we get the following dual to (2.1):

$$(2.3) \quad \begin{aligned} & \text{minimize} && -\log \det(\Sigma + U + Q - P) - n + \alpha \mathbf{Tr} P - \beta \mathbf{Tr} Q \\ & \text{subject to} && P, Q \succeq 0, \quad |U_{ij}| \leq \rho, \quad i, j = 1, \dots, n, \end{aligned}$$

in the variables $U, P, Q \in \mathbf{S}_n$. In what follows, we denote by $\|X\|_2$ the largest singular value of the matrix X and by $\|X\|_F$ its Frobenius norm. When $\alpha = 0$ and $\beta = +\infty$, the first-order optimality conditions impose $X(\Sigma + U) = \mathbf{I}_n$, hence we always have

$$X \succeq \alpha(n) \mathbf{I}_n \quad \text{with} \quad \alpha(n) := \frac{1}{\|\Sigma\|_2 + n\rho};$$

zero duality gap also means $\mathbf{Tr}(\Sigma X) = n - \rho \mathbf{1}^T X \mathbf{1}$. Because X and Σ are both positive semidefinite, we get

$$\|X\|_2 \leq \|X\|_F \leq \mathbf{1}^T X \mathbf{1} \leq n/\rho,$$

which, together with $\mathbf{Tr}(\Sigma X) \geq \lambda_{\min}(\Sigma) \|X\|_2$, means $\|X\|_2 \leq n/\lambda_{\min}(\Sigma)$. Finally then, we must always have

$$X \preceq \beta(n) \mathbf{I}_n \quad \text{with} \quad \beta(n) := n \min \left(\frac{1}{\rho}, \|\Sigma^{-1}\|_2 \right)$$

and $0 < \alpha(n) \leq \lambda(X) \leq \beta(n) < +\infty$ at the optimum. Setting $\alpha = 0$ and $\beta = +\infty$ in problem (2.1) is then equivalent to setting $\alpha = \alpha(n)$ and $\beta = \beta(n)$. Since the objective function of problem (2.1) is strictly convex when $0 < \alpha(n) \leq \lambda(X) \leq \beta(n) < +\infty$, this shows that (2.1) always has a unique solution.

3. Algorithms. In this section, we present two algorithms for solving problem (2.1), one based on an optimal first-order method developed in [18], the other based on a block-coordinate gradient method. Of course, problem (2.1) is convex and can readily be solved using interior point methods (see [4], for example). However, such second-order methods become quickly impractical for solving (1.1), since the corresponding complexity to compute an ϵ -suboptimal solution is $O(n^6 \log(1/\epsilon))$. Note, however, that we cannot expect to do better than $O(n^3)$, which is the cost of solving the nonpenalized problem for dense covariance matrices Σ .

3.1. Smooth optimization. The recently-developed first-order algorithms due to [18] trade-off a better dependence on problem size against a worst dependence on accuracy, usually $1/\epsilon$ instead of its logarithm and the method we describe next has a complexity of $O(n^{4.5}/\epsilon)$. In addition, the memory requirement of these first-order methods is much lower than that of interior-point methods, which involve forming a dense Hessian, and hence, become quickly prohibitive with a problem having $O(n^2)$ variables.

Nesterov's model. The algorithm in [18] supposes that the function to minimize conforms to a certain representation. This is the case for our problem here, so we first write (2.2) in the saddle-function format described in [18]:

$$\min_{X \in \mathcal{Q}_1} -\log \det X + \langle \Sigma, X \rangle + \rho \mathbf{1}^T |X| \mathbf{1} \equiv \min_{X \in \mathcal{Q}_1} \max_{U \in \mathcal{Q}_2} \hat{f}(X) + \langle A(X), U \rangle,$$

where we define $\hat{f}(X) = -\log \det X + \langle \Sigma, X \rangle$, $A = \rho I_{n^2}$, and

$$\mathcal{Q}_1 := \{X \in \mathcal{S}^n : \alpha \mathbf{I}_n \preceq X \preceq \beta \mathbf{I}_n\}, \quad \mathcal{Q}_2 := \{U \in \mathcal{S}^n : \|U\|_\infty \leq 1\}.$$

The adjoint of this problem, corresponding to the dual problem (2.3), is then written:

$$(3.1) \quad \max_{U \in \mathcal{Q}_2} \phi(U) \quad \text{where} \quad \phi(U) := \min_{X \in \mathcal{Q}_1} -\log \det X + \langle \Sigma + U, X \rangle.$$

When a function can be represented in this saddle function format, the method described in [18] combines two steps.

Regularization: By adding a strongly convex penalty to the saddle function representation of f , the algorithm first computes a smooth ϵ -approximation of f with Lipschitz continuous gradient. This can be seen as a generalized Moreau–Yosida regularization step (see [14], for example).

Optimal first order minimization: The algorithm then applies the optimal first-order scheme for functions with Lipschitz continuous gradients detailed in [16] to the regularized function. Each iteration requires efficiently computing the regularized function value and its gradient. In all the semidefinite programming applications detailed here, this can be done extremely efficiently, with a complexity of $O(n^3)$ and memory requirements in $O(n^2)$. The method is only efficient if all these steps can be performed explicitly or at least very efficiently. As we will see below, this is the case here.

Prox-functions and related parameters. To \mathcal{Q}_1 and \mathcal{Q}_2 we now associate norms and so-called prox-functions. For \mathcal{Q}_1 , we use the Frobenius norm and a prox-function:

$$d_1(X) = -\log \det X + n \log \beta.$$

The function d_1 is strongly convex on \mathcal{Q}_1 , with a convexity parameter of $\sigma_1 = 1/\beta^2$, in the sense that $\nabla^2 d_1(X)[H, H] = \text{Tr}(X^{-1} H X^{-1} H) \geq \beta^{-2} \|H\|_F^2$ for every H . Furthermore, the center of the set, $X_0 := \arg \min_{X \in \mathcal{Q}_1} d_1(X)$ is $X_0 = \beta \mathbf{I}_n$ and satisfies $d_1(X_0) = 0$. With our choice, we have $D_1 := \max_{X \in \mathcal{Q}_1} d_1(X) = n \log(\beta/\alpha)$.

To \mathcal{Q}_2 , we also associate the Frobenius norm and the prox-function $d_2(U) = \|U\|_F^2/2$. With this choice, the center U_0 of \mathcal{Q}_2 is $U_0 = 0$. Furthermore, the function d_2 is strongly convex on its domain with a convexity parameter with respect to the 2-norm $\sigma_2 = 1$, and we have $D_2 := \max_{U \in \mathcal{Q}_2} d_2(U) = n^2/2$.

The function \hat{f} has a gradient that is Lipschitz-continuous with respect to the Frobenius norm on the set \mathcal{Q}_1 with Lipschitz constant $M = 1/\alpha^2$. Finally, the norm (induced by the Frobenius norm) of the operator $A = \rho I_{n^2}$ is ρ .

Smooth minimization. The method is based on replacing the objective of the original problem, $f(X)$, with $f_\epsilon(X)$, where $\epsilon > 0$ is the desired accuracy, and f_ϵ is a penalized function involving the prox-function d_2 , defined as

$$(3.2) \quad f_\epsilon(X) := \hat{f}(X) + \max_{U \in \mathcal{Q}_2} \{\langle X, U \rangle - (\epsilon/2 D_2) d_2(U)\}.$$

The above function turns out to be a smooth uniform approximation to f everywhere, with maximal error $\epsilon/2$. Furthermore, the function f_ϵ has a Lipschitz-continuous gradient with Lipschitz constant given by $L(\epsilon) := M + D_2\|A\|^2/(2\sigma_2\epsilon)$. A specific first-order algorithm detailed in [16] for smooth, constrained convex minimization is then applied to the function f_ϵ to get a convergence rate in $O(1/\epsilon)$.

Nesterov's algorithm. Choose $\epsilon > 0$ and set $X_0 = \beta \mathbf{I}_n$, the algorithm then updates primal and dual iterates Y_k and \hat{U}_k using the following steps:

1. Compute $\nabla f_\epsilon(X_k) = -X^{-1} + \Sigma + U^*(X_k)$, where $U^*(X)$ solves (3.2).
2. Find $Y_k = \operatorname{argmin}_{Y \in \mathcal{Q}_1} \{\langle \nabla f_\epsilon(X_k), Y - X_k \rangle + \frac{1}{2}L(\epsilon)\|Y - X_k\|_F^2\}$.
3. Find $Z_k = \operatorname{argmin}_{Z \in \mathcal{Q}_1} \left\{ \frac{L(\epsilon)d_1(Z)}{\sigma_1} + \sum_{i=0}^k \frac{i+1}{2} (f_\epsilon(X_i) + \langle \nabla f_\epsilon(X_i), Z - X_i \rangle) \right\}$.
4. Update $X_k = \frac{2}{k+3}Z_k + \frac{k+1}{k+3}Y_k$ and $\hat{U}_k = \frac{k\hat{U}_{k-1} + 2U^*(X_k)}{(k+2)}$.
5. Repeat until the duality gap is less than the target precision:

$$-\log \det Y_k + \langle \Sigma, Y_k \rangle + \rho \mathbf{1}^T |Y_k| \mathbf{1} - \phi(\hat{U}_k) \leq \epsilon.$$

The key to the method's success is that steps 1–3 and 5 can be performed explicitly and only involve an eigenvalue decomposition. Step one above computes the (smooth) function value and gradient. The second step computes the *gradient mapping*, which matches the gradient step for unconstrained problems (see [17, p. 86]). Steps three and four update an *estimate sequence* [17, p. 72] of f_μ whose minimum can be computed explicitly and gives an increasingly tight upper bound on the minimum of f_μ . We now present these steps in detail for our problem.

Step 1. The first step requires computing the gradient of the function

$$f_\epsilon(X) = \hat{f}(X) + \max_{u \in \mathcal{Q}_2} \langle X, U \rangle - (\epsilon/2D_2)d_2(U).$$

This function can be expressed in closed form as $f_\epsilon(X) = \hat{f}(X) + \sum_{i,j} \psi_\mu(X_{ij})$, where

$$\psi_\epsilon(x) := \begin{cases} |x| - (\epsilon/4D_2) & \text{if } |x| \geq (\epsilon/2\rho D_2), \\ D_2x^2/\epsilon & \text{otherwise,} \end{cases}$$

which is simply the Moreau–Yosida regularization of the absolute value and the gradient of the function at X is

$$\nabla f_\mu(X) = -X^{-1} + \Sigma + U^*(X),$$

with

$$U^*(X) := \max(\min(X/\mu, \rho), -\rho),$$

with \min and \max understood componentwise. The cost of this step is dominated by that of computing the inverse of X , which is $O(n^3)$.

Step 2. This step involves a problem of the form

$$T_{\mathcal{Q}_1}(X) = \arg \min_{Y \in \mathcal{Q}_1} \langle \nabla f_\epsilon(X), Y - X \rangle + \frac{1}{2}L\|Y - X\|_F^2,$$

where $X \in \mathcal{Q}_1$ is given. This problem can be reduced to one of projection on \mathcal{Q}_1 , namely

$$\min_{Y \in \mathcal{Q}_1} \|Y - G\|_F^2,$$

where $G := X - L^{-1}\nabla f_\epsilon(X)$. Using the rotational invariance of this problem, we reduce it to a vector problem:

$$\min_\lambda \sum_i (\lambda_i - \gamma_i)^2 : \alpha \leq \lambda_i \leq \beta, \quad i = 1, \dots, n,$$

where γ is the vector of eigenvalues of G . This problem admits a simple explicit solution:

$$\lambda_i = \min(\max(\gamma_i, \alpha), \beta), \quad i = 1, \dots, n.$$

The corresponding solution is then $Y = V^T \mathbf{diag}(\lambda)V$, where $G = V^T \mathbf{diag}(\gamma)V$ is the eigenvalue decomposition of G . The cost of this step is dominated by the cost of forming the eigenvalue decomposition of G , which is $O(n^3)$.

Step 3. The third step involves solving a problem of the form

$$(3.3) \quad Z := \arg \max_{X \in \mathcal{Q}_1} d_1(X) + \langle S, X \rangle,$$

where S is given. Again, due to the rotational invariance of the objective and feasible set, we can reduce the problem to a one-dimensional problem:

$$\min_\lambda \sum_i \sigma_i \lambda_i - \log \lambda_i : \alpha \leq \lambda_i \leq \beta,$$

where σ contains the eigenvalues of S . This problem has a simple, explicit solution:

$$\lambda_i = \min(\max(1/\sigma_i, \alpha), \beta), \quad i = 1, \dots, n.$$

The corresponding solution is then $Y = V^T \mathbf{diag}(\lambda)V$, where $S = V^T \mathbf{diag}(\sigma)V$ is the eigenvalue decomposition of S . Again, the cost of this step is dominated by the cost of forming the eigenvalue decomposition of S , which is $O(n^3)$.

Computing $\phi(\hat{U}_k)$. For a given matrix \hat{U}_k the function ϕ in (3.1) is computed as

$$\phi(\hat{U}_k) = \min_{X \in \mathcal{Q}_1} -\log \det X + \langle \Sigma + \hat{U}_k, X \rangle.$$

This means projecting $(\Sigma + \hat{U}_k)^{-1}$ on \mathcal{Q}_1 only involves an eigenvalue decomposition.

Complexity estimate. To summarize, for step 1, the gradient of f_ϵ is readily computed in closed form, via the computation of the inverse of X . Step 2 essentially amounts to projecting on \mathcal{Q}_1 and requires an eigenvalue problem to be solved; likewise for step 3. In fact, each iteration costs $O(n^3)$. The number of iterations necessary to achieve an objective with absolute accuracy less than ϵ is then given by

$$(3.4) \quad N(\epsilon) := 4\|A\| \frac{1}{\epsilon} \sqrt{\frac{D_1 D_2}{\sigma_1 \sigma_2}} + \sqrt{\frac{M D_1}{\sigma_1 \epsilon}} = \frac{\kappa \sqrt{n(\log \kappa)}}{\epsilon} (4n\alpha\rho + \sqrt{\epsilon}),$$

where $\kappa = \beta/\alpha$ bounds the solution's condition number. Thus, the overall complexity when $\rho > 0$ is in $O(n^{4.5}/\epsilon)$, as claimed.

3.2. Block-coordinate gradient methods. In this section, we focus on the particular case where $\alpha = 0$ and $\beta = +\infty$ (hence, implicitly $\alpha = \alpha(n)$ and $\beta = \beta(n)$) and derive gradient minimization algorithms that take advantage of the problem structure. We consider the following problem:

$$(3.5) \quad \max_X \log \det X - \langle \Sigma, X \rangle - \rho \mathbf{1}^T |X| \mathbf{1}$$

in the variable $X \in \mathbf{S}_n$, where $\rho > 0$ again controls the trade-off between log-likelihood and sparsity of the inverse covariance matrix. Its dual is given by

$$(3.6) \quad \begin{aligned} & \text{minimize} && -\log \det(\Sigma + U) - n \\ & \text{subject to} && |U_{ij}| \leq \rho, \quad i, j = 1, \dots, n, \end{aligned}$$

in the variable $U \in \mathbf{S}_n$. We partition the matrices X and U in block format:

$$X = \begin{pmatrix} Z & x \\ x^T & y \end{pmatrix} \quad \text{and} \quad U = \begin{pmatrix} V & u \\ u^T & w \end{pmatrix},$$

where $Z \succ 0$ and U are fixed and $x, u \in \mathbf{R}^{(n-1)}$, $y, w \in \mathbf{R}$ are the variables (row and column) we are updating. We also partition the sample matrix according to the same block structure:

$$\Sigma = \begin{pmatrix} A & b \\ b^T & c \end{pmatrix},$$

where $A \in \mathbf{S}_{(n-1)}$, $b \in \mathbf{R}^{(n-1)}$, $c \in \mathbf{R}$. In the methods that follow, we will update only one column (and corresponding row) at a time and without loss of generality we can always assume that we are updating the last one.

Block-coordinate descent. The dual problem (3.6):

$$\begin{aligned} & \text{minimize} && -\log \det(\Sigma + U) - n \\ & \text{subject to} && |U_{ij}| \leq \rho, \quad i, j = 1, \dots, n, \end{aligned}$$

in the variable $U \in \mathbf{S}_n$, can be written in block format as

$$\begin{aligned} & \text{minimize} && -\log \det(A + V) - \log((w + c) - (b + u)^T(A + V)^{-1}(b + u)) - n \\ & \text{subject to} && |w| \leq \rho, |u_i| \leq \rho, \quad i = 1, \dots, n, \end{aligned}$$

in the variables $u \in \mathbf{R}^{(n-1)}$ and $w \in \mathbf{R}$ (V is fixed at each iteration). We directly get $w = \rho$ so the diagonal of the optimal solution must be $\rho \mathbf{1}$. The main step at each iteration is then a box constrained quadratic program (QP):

$$(3.7) \quad \begin{aligned} & \text{minimize} && (b + u)^T(A + V)^{-1}(b + u) \\ & \text{subject to} && |u_i| \leq \rho, \quad i = 1, \dots, n, \end{aligned}$$

in the variable $u \in \mathbf{R}^{(n-1)}$. To summarize, the block coordinate descent algorithm proceeds as follows:

1. Pick the row and column to update.
2. Compute $(A + V)^{-1}$.
3. Solve the box constrained QP in (3.7).
4. Repeat until duality gap is less than precision: $\langle \Sigma, X \rangle - n + \rho \mathbf{1}^T |X| \mathbf{1} \leq \epsilon$.

At each iteration, we need to compute the inverse of the submatrix $(A + V) \in \mathbf{S}_{(n-1)}$, but we can update this inverse using the Sherman–Woodbury–Morrison formula on two rank-two updates; hence, it is only necessary to compute a full inverse at the first iteration.

Block-coordinate ascent. For a fixed Z , problem (3.5) is equivalent to

$$\begin{aligned} & \text{maximize} && \log(y - x^T Z^{-1} x) - 2b^T x - y(c + \rho) - 2\rho\|x\|_1 \\ & \text{subject to} && y - x^T Z^{-1} x > 0, \quad y > 0, \end{aligned}$$

in the variables $x \in \mathbf{R}^{(n-1)}$, $y \in \mathbf{R}$, where $Z \succ 0$ (given) and the Schur complement constraints imply $X \succ 0$. We can solve for the optimal y explicitly and the problem in x becomes

$$\max_x -x^T Q x - 2b^T x - 2\rho\|x\|_1,$$

where $Q := (c + \rho)Z^{-1}$. Its dual is also box-constrained QP:

$$\begin{aligned} & \text{minimize} && (b + u)^T Z (b + u) \\ & \text{subject to} && \|u\|_\infty \leq \rho, \end{aligned}$$

in the variable $u \in \mathbf{R}^{(n-1)}$. At the optimum for this QP, we must have

$$x = -\frac{1}{(c + \rho)} Z(b + u), \quad \text{and} \quad y = \frac{1}{(c + \rho)} + \frac{1}{(c + \rho)^2} (b + u)^T Z (b + u),$$

and we iterate as above.

Smooth optimization for box-constrained QPs. The two block-coordinate methods detailed in this section both amount to solving a sequence of box-constrained quadratic program of the form

$$(3.8) \quad \begin{aligned} & \text{minimize} && x^T A x + b^T x \\ & \text{subject to} && \|x\|_\infty \leq \rho, \end{aligned}$$

in the variable $x \in \mathbf{R}^n$. The objective function has a Lipschitz continuous gradient with constant $L = 2\lambda^{\max}(A)$ on the box $\mathcal{B} = \{x \in \mathbf{R}^n : \|x\|_\infty \leq \rho\}$, where we can define a prox function $(1/2)\|x\|^2$ which is strongly convex with constant one and bounded above by $(1/2)n\rho^2$ on \mathcal{B} . From [16] or [18], we know that solving (3.8) up to a precision ϵ will require at most $2\rho\sqrt{n\lambda^{\max}(A)}/\sqrt{\epsilon}$ iterations of the first-order method detailed in [16], with each iteration equivalent to a matrix-vector product and a projection on the box \mathcal{B} . This means that the total complexity of solving (3.8) is given by

$$O\left(\rho n^{2.5} \sqrt{\frac{\lambda^{\max}(A)}{\epsilon}}\right).$$

Complexity estimate. Following [15], with block coordinate descent corresponding to coordinate descent with the almost cyclic rule (defined in [15], it simply means here that we go through each index at least once per outer iteration) and using the fact that $\log \det(X)$ satisfies the strict concavity assumptions in [15, assumption A2], we can show that the convergence rate of the block coordinate descent method is at least linear. Each iteration requires solving a box-constrained QP and takes $O(n^3 \log(1/\epsilon))$ operations using an interior point solver or $O(n^{2.5}/\sqrt{\epsilon})$ using the optimal first-order scheme in [16]. We cannot use the same argument to show convergence of block coordinate ascent but empirical performance is comparable. In practice we have found that a small number of sweeps through all columns, independent of problem size n , is sufficient for convergence.

Implementation. The block coordinate descent methods implemented here correspond to coordinate descent using the almost cyclic rule, alternative row/column selection rules could improve the convergence speed. Also, each iteration of the block coordinate descent method corresponds to two rank-two updates of the inverse matrix; hence, the cost of maintaining the inverse submatrix using the Sherman–Woodbury–Morrison formula is only $O(n^2)$.

4. Numerical results. In this section we test the performance of the methods detailed above on some randomly generated examples. We first form a sparse matrix A with a diagonal equal to one and a few randomly chosen, nonzero off-diagonal terms equal to $+1$ or -1 . We then form the matrix

$$B = A^{-1} + \sigma V$$

where $V \in \mathbf{S}_n$ is a symmetric, i.i.d. uniform random matrix. Finally, we make B positive definite by shifting its eigenvalues, and use this noisy, random matrix to test our covariance selection methods.




FIG. 4.1. *Recovering the sparsity pattern.* We plot the original inverse covariance matrix A , the solution to problem (2.1), and the noisy inverse B^{-1} .

In Figure 4.1, we plot the sparsity patterns of the original inverse covariance matrix A , the solution to problem (2.1), and the noisy inverse B^{-1} in a randomly generated example with $n = 30$, $\sigma = 0.15$, and $\rho = 0.5$. In Figure 4.2 we represent the dependence structure of interest rates (sampled over a year) inferred from the inverse covariance matrix. Each node represents a particular interest rate maturity and the nodes are linked if the corresponding coefficient in the inverse covariance matrix is nonzero (i.e., they are conditionally dependent). We compare the solution to problem (2.1) on this matrix for $\rho = 0$ and $\rho = 0.1$ and notice that in the sparse solution the rates appear clearly clustered by maturity.

In Figure 4.3, we study computing times for various choices of algorithms and problem sizes. On the *left*, we plot CPU time to reduce the duality gap by a factor 10^{-2} versus problem size n , on randomly generated problems, using the coordinate descent code and the optimal first-order for solving box QPs. On the *right*, we plot duality gap versus CPU time for both smooth minimization and block-coordinate algorithms for a randomly generated problem of size $n = 250$. For the smooth minimization code, we set $\alpha = 1/\lambda^{\max}(B)$ and we plot computing time for both $\beta = 1/(2\lambda^{\min}(B))$ (smooth. opt. 1/2) and $\beta = 2/\lambda^{\min}(B)$ (smooth. opt. 2). In the examples of Figure 4.3, we notice that the numerical cost of our methods grows experimentally as $O(n^3)$.




FIG. 4.2. We plot the network formed using the solution to problem (2.1) on an interest rate covariance matrix for $\rho = 0$ (left) and $\rho = 0.1$ (right). In the sparse solution the rates appear clearly clustered by maturity.




FIG. 4.3. Computing time. Left: We plot CPU time to reduce the duality gap by a factor 10^{-2} versus problem size n , on randomly generated problems, using the coordinate descent code and the optimal first-order algorithm for solving box QPs (dashed line, circles), smooth minimization (dotted line, squares), and a simple conjugate gradient method with a Polak–Ribiere update, without preconditioning (solid line, stars). Right: We plot the duality gap versus CPU time for both smooth minimization and block-coordinate algorithms for a problem of size $n = 250$.

Acknowledgments. The authors would like to thank Francis Bach, Peter Bartlett, and Martin Wainwright for enlightening discussions on the topic. We are grateful to the (anonymous) referees for important comments and for finding an error in an earlier version of the paper.

REFERENCES

- [1] J. AKAIKE, *Information theory and an extension of the maximum likelihood principle*, in Second International Symposium on Information Theory, B. N. Petrov and F. Csaki, eds., Akademiai Kiado, Budapest, Hungary, 1973, pp. 267–281.
- [2] J. A. BILMES, *Natural Statistic Models for Automatic Speech Recognition*, Ph.D. thesis, UC Berkeley, Dept. of EECS, CS Division, Berkeley, CA, 1999.
- [3] J. A. BILMES, *Factored sparse inverse covariance matrices*, IEEE International Conference on Acoustics, Speech, and Signal Processing, Istanbul, Turkey, 2000, pp. 1009–1012.
- [4] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, 2004.

- [5] K. P. BURNHAM AND D. R. ANDERSON, *Multimodel inference: Understanding AIC or BIC in model selection*, Sociol. Methods Res., 33 (2004), pp. 261–304.
- [6] J. DAHL, V. ROYCHOWDHURY, AND L. VANDENBERGHE, *Maximum Likelihood Estimation of Gaussian Graphical Models: Numerical Implementation and Topology Selection*, UCLA preprint, 2005.
- [7] A. DEMPSTER, *Covariance selection*, Biometrics, 28 (1972), pp. 157–175.
- [8] A. DOBRA, C. HANS, B. JONES, J. R. NEVINS, G. YAO, AND M. WEST, *Sparse graphical models for exploring gene expression data*, J. Multivariate Anal., 90 (2004), pp. 196–212.
- [9] A. DOBRA AND M. WEST, *Bayesian covariance selection*, ISDS Working paper, 2004.
- [10] D. L. DONOHO AND J. TANNER, *Sparse nonnegative solutions of underdetermined linear equations by linear programming*, Proc. Natl. Acad. Sci. USA, 102 (2005), pp. 9446–9451.
- [11] M. FAZEL, H. HINDI, AND S. BOYD, *A rank minimization heuristic with application to minimum order system approximation*, in Proceedings of the American Control Conference, Arlington, VA, 6 (2001), pp. 4734–4739.
- [12] J. Z. HUANG, N. LIU, AND M. POURAHMADI, *Covariance selection and estimation via penalized normal likelihood*, Biometrika, 93 (2007), pp. 85–98.
- [13] B. JONES, C. CARVALHO, A. DOBRA, C. HANS, C. CARTER, AND M. WEST, *Experiments in stochastic computation for high-dimensional graphical models*, Statist. Sci., 20 (2005), pp. 388–400.
- [14] C. LEMARÉCHAL AND C. SAGASTIZÁBAL, *Practical aspects of the Moreau-Yosida regularization: Theoretical preliminaries*, SIAM J. Optim., 7 (1997), pp. 367–385.
- [15] Z. Q. LUO AND P. TSENG, *On the convergence of the coordinate descent method for convex differentiable minimization*, J. Optim. Theory Appl., 72 (1992), pp. 7–35.
- [16] Y. NESTEROV, *A method of solving a convex programming problem with convergence rate $O(1/k^2)$* , Soviet Math. Dokl., 27 (1983), pp. 372–376.
- [17] Y. NESTEROV, *Introductory Lectures on Convex Optimization*, Springer, 2003.
- [18] Y. NESTEROV, *Smooth minimization of nonsmooth functions*, Math. Program., Ser. A, 103 (2005), pp. 127–152.
- [19] R. TIBSHIRANI, *Regression shrinkage and selection via the LASSO*, J. Roy. Statist. Soc., Ser. B, 58 (1996), pp. 267–288.