

Sparse Approximation via Penalty Decomposition Methods ^{*}

Zhaosong Lu[†]

Yong Zhang[‡]

September 8, 2010 (Revised: February 19, 2012; March 13, 2013; August 27, 2013)

Abstract

In this paper we consider sparse approximation problems, that is, general l_0 minimization problems with the l_0 -“norm” of a vector being a part of constraints or objective function. In particular, we first study the first-order optimality conditions for these problems. We then propose penalty decomposition (PD) methods for solving them in which a sequence of penalty subproblems are solved by a block coordinate descent (BCD) method. Under some suitable assumptions, we establish that any accumulation point of the sequence generated by the PD methods satisfies the first-order optimality conditions of the problems. Furthermore, for the problems in which the l_0 part is the only nonconvex part, we show that such an accumulation point is a local minimizer of the problems. In addition, we show that any accumulation point of the sequence generated by the BCD method is a block coordinate minimizer of the penalty subproblem. Moreover, for the problems in which the l_0 part is the only nonconvex part, we establish that such an accumulation point is a local minimizer of the penalty subproblem. Finally, we test the performance of our PD methods by applying them to sparse logistic regression, sparse inverse covariance selection, and compressed sensing problems. The computational results demonstrate that when solutions of same cardinality are sought, our approach applied to the l_0 -based models generally has better solution quality and/or speed than the existing approaches that are applied to the corresponding l_1 -based models.

Key words: l_0 minimization, penalty decomposition methods, block coordinate descent method, compressed sensing, sparse logistic regression, sparse inverse covariance selection

1 Introduction

Nowadays, there are numerous applications in which sparse solutions are concerned. For example, in compressed sensing, a large sparse signal is decoded by using a relatively small number of linear measurements, which can be formulated as finding a sparse solution to a system of linear equalities and/or inequalities. The similar ideas have also been widely used in linear regression. Recently, sparse inverse covariance selection becomes an important tool in discovering the conditional independence in graphical models. One popular approach for sparse inverse covariance selection is to find an approximate sparse inverse covariance while maximizing the log-likelihood (see, for example, [16]). Similarly, sparse logistic regression has been proposed as a promising method for feature selection in classification problems in which a sparse solution is sought to minimize the average logistic loss (see, for example, [41]).

^{*}This work was supported in part by NSERC Discovery Grant.

[†]Department of Mathematics, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada. (email: zhaosong@sfu.ca).

[‡]Department of Mathematics, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada. (email: yza30@sfu.ca).

Mathematically, all these applications can be formulated into the following l_0 minimization problems:

$$\min_{x \in \mathcal{X}} \{f(x) : g(x) \leq 0, h(x) = 0, \|x_J\|_0 \leq r\}, \quad (1)$$

$$\min_{x \in \mathcal{X}} \{f(x) + \nu \|x_J\|_0 : g(x) \leq 0, h(x) = 0\} \quad (2)$$

for some integer $r \geq 0$ and $\nu \geq 0$ controlling the sparsity (or cardinality) of the solution, where \mathcal{X} is a closed convex set in the n -dimensional Euclidean space \mathbb{R}^n , $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $h : \mathbb{R}^n \rightarrow \mathbb{R}^p$ are continuously differentiable functions, and $\|x_J\|_0$ denotes the cardinality of the subvector formed by the entries of x indexed by J . Some algorithms are proposed for solving special cases of these problems. For example, the iterative hard thresholding algorithms [26, 5, 6] and matching pursuit algorithms [38, 52] are developed for solving the l_0 -regularized least squares problems arising in compressed sensing, but they cannot be applied to the general l_0 minimization problems (1) and (2). In the literature, one popular approach for dealing with (1) and (2) is to replace $\|\cdot\|_0$ by the l_1 -norm $\|\cdot\|_1$ and solve the resulting relaxation problems instead (see, for example, [14, 41, 10, 51]). For some applications such as compressed sensing, it has been shown in [8] that under some suitable assumptions this approach is capable of solving (1) and (2). Recently, another relaxation approach has been proposed to solve problems (1) and (2) in which $\|\cdot\|_0$ is replaced by l_p -“norm” $\|\cdot\|_p$ for some $p \in (0, 1)$ (see, for example, [9, 11, 12]). In general, it is not clear about the solution quality of these approaches. Indeed, for the example given in the Appendix, the l_p relaxation approach for $p \in (0, 1]$ fails to recover the sparse solution.

In this paper we propose penalty decomposition (PD) methods for solving problems (1) and (2) in which a sequence of penalty subproblems are solved by a block coordinate descent (BCD) method. Under some suitable assumptions, we establish that any accumulation point of the sequence generated by the PD method satisfies the first-order optimality conditions of (1) and (2). Furthermore, when h 's are affine, and f and g 's are convex, we show that such an accumulation point is a local minimizer of the problems. In addition, we show that any accumulation point of the sequence generated by the BCD method is a block coordinate minimizer of the penalty subproblem. Moreover, when h 's are affine, and f and g 's are convex, we establish that such an accumulation point is a local minimizer of the penalty subproblem. Finally, we test the performance of our PD methods by applying them to sparse logistic regression, sparse inverse covariance selection, and compressed sensing problems. The computational results demonstrate that when solutions of same cardinality are sought, our approach applied to the l_0 -based models generally has better solution quality and/or speed than the existing approaches that are applied to the corresponding l_1 -based models.

The rest of this paper is organized as follows. In Subsection 1.1, we introduce the notation that is used throughout the paper. In Section 2, we establish the first-order optimality conditions for general l_0 minimization problems. In Section 3, we study a class of special l_0 minimization problems. We develop the PD methods for general l_0 minimization problems in Section 4 and establish some convergence results for them. In Section 5, we conduct numerical experiments to test the performance of our PD methods for solving sparse logistic regression, sparse inverse covariance selection, and compressed sensing problems. Finally, we present some concluding remarks in section 6.

1.1 Notation

In this paper, the symbols \mathbb{R}^n and \mathbb{R}_+^n denote the n -dimensional Euclidean space and the nonnegative orthant of \mathbb{R}^n , respectively. Given a vector $v \in \mathbb{R}^n$, the nonnegative part of v is denoted by $v^+ = \max(v, 0)$, where the maximization operates entry-wise. For any real vector, $\|\cdot\|_0$ and $\|\cdot\|$ denote the

cardinality (i.e., the number of nonzero entries) and the Euclidean norm of the vector, respectively. Given an index set $L \subseteq \{1, \dots, n\}$, $|L|$ denotes the size of L , and the elements of L are denoted by $L(1), \dots, L(|L|)$, which are always arranged in ascending order. x_L denotes the subvector formed by the entries of x indexed by L . Likewise, X_L denotes the submatrix formed by the columns of X indexed by L . In addition, For any two sets A and B , the subtraction of A and B is given by $A \setminus B = \{x \in A : x \notin B\}$. Given a closed set $C \subseteq \mathbb{R}^n$, let $\mathcal{N}_C(x)$ and $\mathcal{T}_C(x)$ denote the normal and tangent cones of C at any $x \in C$, respectively. The space of all $m \times n$ matrices with real entries is denoted by $\mathbb{R}^{m \times n}$, and the space of symmetric $n \times n$ matrices is denoted by \mathcal{S}^n . We denote by I the identity matrix, whose dimension should be clear from the context. If $X \in \mathcal{S}^n$ is positive semidefinite (resp., definite), we write $X \succeq 0$ (resp., $X \succ 0$). The cone of positive semidefinite (resp., definite) matrices is denoted by \mathcal{S}_+^n (resp., \mathcal{S}_{++}^n). \mathcal{D} is an operator which maps a vector to a diagonal matrix whose diagonal consists of the vector. Given an $n \times n$ matrix X , $\tilde{\mathcal{D}}(X)$ denotes a diagonal matrix whose i th diagonal element is X_{ii} for $i = 1, \dots, n$.

2 First-order optimality conditions

In this section we study the first-order optimality conditions for problems (1) and (2). In particular, we first discuss the first-order necessary conditions for them. Then we study the first-order sufficient conditions for them when the l_0 part is the only nonconvex part.

We now establish the first-order necessary optimality conditions for problems (1) and (2).

Theorem 2.1 *Assume that x^* is a local minimizer of problem (1). Let $J^* \subseteq J$ be an index set with $|J^*| = r$ such that $x_j^* = 0$ for all $j \in \bar{J}^*$, where $\bar{J}^* = J \setminus J^*$. Suppose that the following Robinson condition*

$$\left\{ \left[\begin{array}{c} g'(x^*)d - v \\ h'(x^*)d \\ (I_{\bar{J}^*})^T d \end{array} \right] : d \in \mathcal{T}_{\mathcal{X}}(x^*), v \in \mathbb{R}^m, v_i \leq 0, i \in \mathcal{A}(x^*) \right\} = \mathbb{R}^m \times \mathbb{R}^p \times \mathbb{R}^{|J|-r} \quad (3)$$

holds, where $g'(x^*)$ and $h'(x^*)$ denote the Jacobian of the functions $g = (g_1, \dots, g_m)$ and $h = (h_1, \dots, h_p)$ at x^* , respectively, and

$$\mathcal{A}(x^*) = \{1 \leq i \leq m : g_i(x^*) = 0\}. \quad (4)$$

Then, there exists $(\lambda^*, \mu^*, z^*) \in \mathbb{R}^m \times \mathbb{R}^p \times \mathbb{R}^n$ together with x^* satisfying

$$\begin{aligned} -\nabla f(x^*) - \nabla g(x^*)\lambda^* - \nabla h(x^*)\mu^* - z^* &\in \mathcal{N}_{\mathcal{X}}(x^*), \\ \lambda_i^* &\geq 0, \lambda_i^* g_i(x^*) = 0, i = 1, \dots, m; \quad z_j^* = 0, j \in \bar{J} \cup J^*. \end{aligned} \quad (5)$$

where \bar{J} is the complement of J in $\{1, \dots, n\}$.

Proof. By the assumption that x^* is a local minimizer of problem (1), one can observe that x^* is also a local minimizer of the following problem:

$$\min_{x \in \mathcal{X}} \{f(x) : g(x) \leq 0, h(x) = 0, x_{\bar{J}^*} = 0\}. \quad (6)$$

Using this observation, (3) and Theorem 3.25 of [47], we see that the conclusion holds. \blacksquare

Theorem 2.2 Assume that x^* is a local minimizer of problem (2). Let $J^* = \{j \in J : x_j^* \neq 0\}$ and $\bar{J}^* = J \setminus J^*$. suppose that the following Robinson condition

$$\left\{ \left[\begin{array}{c} g'(x^*)d - v \\ h'(x^*)d \\ (I_{\bar{J}^*})^T d \end{array} \right] : d \in \mathcal{T}_{\mathcal{X}}(x^*), v \in \mathfrak{R}^m, v_i \leq 0, i \in \mathcal{A}(x^*) \right\} = \mathfrak{R}^m \times \mathfrak{R}^p \times \mathfrak{R}^{|\bar{J}^*|} \quad (7)$$

holds, where $\mathcal{A}(x^*)$ is defined in (4). Then, there exists $(\lambda^*, \mu^*, z^*) \in \mathfrak{R}^m \times \mathfrak{R}^p \times \mathfrak{R}^n$ together with x^* satisfying (5).

Proof. It is not hard to observe that x^* is a local minimizer of problem (2) if and only if x^* is a local minimizer of problem (6). Using this observation, (7) and Theorem 3.25 of [47], we see that the conclusion holds. ■

We next establish the first-order sufficient optimality conditions for problems (1) and (2) when the l_0 part is the only nonconvex part.

Theorem 2.3 Assume that h 's are affine functions, and f and g 's are convex functions. Let x^* be a feasible point of problem (1), and let $\mathcal{J}^* = \{J^* \subseteq J : |J^*| = r, x_j^* = 0, \forall j \in J \setminus J^*\}$. Suppose that for any $J^* \in \mathcal{J}^*$, there exists some $(\lambda^*, \mu^*, z^*) \in \mathfrak{R}^m \times \mathfrak{R}^p \times \mathfrak{R}^n$ such that (5) holds. Then, x^* is a local minimizer of problem (1).

Proof. It follows from the above assumptions and Theorem 3.34 of [47] that x^* is a minimizer of problem (6) for all $\bar{J}^* \in \{J \setminus J^* : J^* \in \mathcal{J}^*\}$. Hence, there exists $\epsilon > 0$ such that $f(x) \geq f(x^*)$ for all $x \in \cup_{J^* \in \mathcal{J}^*} \mathcal{O}_{J^*}(x^*; \epsilon)$, where

$$\mathcal{O}_{J^*}(x^*; \epsilon) = \{x \in \mathcal{X} : g(x) \leq 0, h(x) = 0, x_{\bar{J}^*} = 0, \|x - x^*\| < \epsilon\}$$

with $\bar{J}^* = J \setminus J^*$. One can observe from (1) that for any $x \in \mathcal{O}(x^*; \epsilon)$, where

$$\mathcal{O}(x^*; \epsilon) = \{x \in \mathcal{X} : g(x) \leq 0, h(x) = 0, \|x_J\|_0 \leq r, \|x - x^*\| < \epsilon\},$$

there exists $J^* \in \mathcal{J}^*$ such that $x \in \mathcal{O}_{J^*}(x^*; \epsilon)$ and hence $f(x) \geq f(x^*)$. It implies that the conclusion holds. ■

Theorem 2.4 Assume that h 's are affine functions, and f and g 's are convex functions. Let x^* be a feasible point of problem (2), and let $J^* = \{j \in J : x_j^* \neq 0\}$. Suppose that for such J^* , there exists some $(\lambda^*, \mu^*, z^*) \in \mathfrak{R}^m \times \mathfrak{R}^p \times \mathfrak{R}^n$ such that (5) holds. Then, x^* is a local minimizer of problem (2).

Proof. By virtue of the above assumptions and Theorem 3.34 of [47], we know that x^* is a minimizer of problem (6) with $\bar{J}^* = J \setminus J^*$. Also, we observe that any point is a local minimizer of problem (2) if and only if it is a local minimizer of problem (6). It then implies that x^* is a local minimizer of (2). ■

Remark. The second-order necessary or sufficient optimality conditions for problems (1) and (2) can be similarly established as above. ■

3 A class of special l_0 minimization

In this section we show that a class of special l_0 minimization problems have closed-form solutions, which can be used to develop penalty decomposition methods for solving general l_0 minimization problems.

Proposition 3.1 *Let $\mathcal{X}_i \subseteq \Re$ and $\phi_i : \Re \rightarrow \Re$ for $i = 1, \dots, n$ be given. Suppose that r is a positive integer and $0 \in \mathcal{X}_i$ for all i . Consider the following l_0 minimization problem:*

$$\min \left\{ \phi(x) = \sum_{i=1}^n \phi_i(x_i) : \|x\|_0 \leq r, x \in \mathcal{X}_1 \times \dots \times \mathcal{X}_n \right\}. \quad (8)$$

Let $\tilde{x}_i^* \in \text{Arg min}\{\phi_i(x_i) : x_i \in \mathcal{X}_i\}$ and $I^* \subseteq \{1, \dots, n\}$ be the index set corresponding to r largest values of $\{v_i^*\}_{i=1}^n$, where $v_i^* = \phi_i(0) - \phi_i(\tilde{x}_i^*)$ for $i = 1, \dots, n$. Then, x^* is an optimal solution of problem (8), where x^* is defined as follows:

$$x_i^* = \begin{cases} \tilde{x}_i^* & \text{if } i \in I^*; \\ 0 & \text{otherwise,} \end{cases} \quad i = 1, \dots, n.$$

Proof. By the assumption that $0 \in \mathcal{X}_i$ for all i , and the definitions of x^* , \tilde{x}^* and I^* , we see that $x^* \in \mathcal{X}_1 \times \dots \times \mathcal{X}_n$ and $\|x^*\|_0 \leq r$. Hence, x^* is a feasible solution of (8). It remains to show that $\phi(x) \geq \phi(x^*)$ for any feasible point x of (8). Indeed, let x be arbitrarily chosen such that $\|x\|_0 \leq r$ and $x \in \mathcal{X}_1 \times \dots \times \mathcal{X}_n$, and let $L = \{i : x_i \neq 0\}$. Clearly, $|L| \leq r = |I^*|$. Let \bar{I}^* and \bar{L} denote the complement of I^* and L in $\{1, \dots, n\}$, respectively. It then follows that

$$|\bar{L} \cap I^*| = |I^*| - |I^* \cap L| \geq |L| - |I^* \cap L| = |L \cap \bar{I}^*|.$$

In view of the definitions of x^* , \tilde{x}^* , I^* , \bar{I}^* , L and \bar{L} , we further have

$$\begin{aligned} \phi(x) - \phi(x^*) &= \sum_{i \in L \cap I^*} (\phi_i(x_i) - \phi_i(x_i^*)) + \sum_{i \in \bar{L} \cap \bar{I}^*} (\phi_i(x_i) - \phi_i(x_i^*)) \\ &\quad + \sum_{i \in \bar{L} \cap I^*} (\phi_i(x_i) - \phi_i(x_i^*)) + \sum_{i \in L \cap \bar{I}^*} (\phi_i(x_i) - \phi_i(x_i^*)), \\ &= \sum_{i \in L \cap I^*} (\phi_i(x_i) - \phi_i(\tilde{x}_i^*)) + \sum_{i \in \bar{L} \cap \bar{I}^*} (\phi_i(0) - \phi_i(0)) \\ &\quad + \sum_{i \in \bar{L} \cap I^*} (\phi_i(0) - \phi_i(\tilde{x}_i^*)) + \sum_{i \in L \cap \bar{I}^*} (\phi_i(x_i) - \phi_i(0)), \\ &\geq \sum_{i \in \bar{L} \cap I^*} (\phi_i(0) - \phi_i(\tilde{x}_i^*)) + \sum_{i \in L \cap \bar{I}^*} (\phi_i(\tilde{x}_i^*) - \phi_i(0)), \\ &= \sum_{i \in \bar{L} \cap I^*} (\phi_i(0) - \phi_i(\tilde{x}_i^*)) - \sum_{i \in L \cap \bar{I}^*} (\phi_i(0) - \phi_i(\tilde{x}_i^*)) \geq 0, \end{aligned}$$

where the last inequality follows from the definition of I^* and the relation $|\bar{L} \cap I^*| \geq |L \cap \bar{I}^*|$. Thus, we see that $\phi(x) \geq \phi(x^*)$ for any feasible point x of (8), which implies that the conclusion holds. \blacksquare

It is straightforward to establish the following result.

Proposition 3.2 *Let $\mathcal{X}_i \subseteq \Re$ and $\phi_i : \Re \rightarrow \Re$ for $i = 1, \dots, n$ be given. Suppose that $\nu \geq 0$ and $0 \in \mathcal{X}_i$ for all i . Consider the following l_0 minimization problem:*

$$\min \left\{ \nu \|x\|_0 + \sum_{i=1}^n \phi_i(x_i) : x \in \mathcal{X}_1 \times \dots \times \mathcal{X}_n \right\}. \quad (9)$$

Let $\tilde{x}_i^* \in \text{Arg min}\{\phi_i(x_i) : x_i \in \mathcal{X}_i\}$ and $v_i^* = \phi_i(0) - \nu - \phi_i(\tilde{x}_i^*)$ for $i = 1, \dots, n$. Then, x^* is an optimal solution of problem (9), where x^* is defined as follows:

$$x_i^* = \begin{cases} \tilde{x}_i^* & \text{if } v_i^* \geq 0; \\ 0 & \text{otherwise,} \end{cases} \quad i = 1, \dots, n.$$

4 Penalty decomposition methods for general l_0 minimization

In this section we propose penalty decomposition (PD) methods for solving general l_0 minimization problems (1) and (2) and establish their convergence. Throughout this section, we make the following assumption for problems (1) and (2).

Assumption 1 *Problems (1) and (2) are feasible, and moreover, at least a feasible solution, denoted by x^{feas} , is known.*

This assumption will be used to design the PD methods with nice convergence properties. It can be dropped, but the theoretical convergence of the corresponding PD methods may become weaker. We shall also mention that, for numerous real applications, x^{feas} is readily available or can be observed from the physical background of problems. For example, all application problems discussed in Section 5 have a trivial feasible solution. On the other hand, for some problems which do not have a trivial feasible solution, one can always approximate them by the problems which have a trivial feasible solution. For instance, problem (1) can be approximately solved as the following problem:

$$\min_{x \in \mathcal{X}} \{f(x) + \rho(\|u^+\|^2 + \|v\|^2) : g(x) - u \leq 0, h(x) - v = 0, \|x_J\|_0 \leq r\}$$

for some large ρ . The latter problem has a trivial feasible solution when \mathcal{X} is sufficiently simple.

4.1 Penalty decomposition method for problem (1)

In this subsection we propose a PD method for solving problem (1) and establish its convergence.

We observe that (1) can be equivalently reformulated as

$$\min_{x \in \mathcal{X}, y \in \mathcal{Y}} \{f(x) : g(x) \leq 0, h(x) = 0, x_J - y = 0\}, \quad (10)$$

where

$$\mathcal{Y} = \{y \in \mathbb{R}^{|J|} : \|y\|_0 \leq r\}.$$

The associated quadratic penalty function is defined as follows:

$$q_\varrho(x, y) = f(x) + \frac{\varrho}{2}(\|[g(x)]^+\|^2 + \|h(x)\|^2 + \|x_J - y\|^2) \quad \forall x \in \mathcal{X}, y \in \mathcal{Y} \quad (11)$$

for some penalty parameter $\varrho > 0$.

We are now ready to propose a PD method for solving problem (10) (or equivalently, (1)) in which each penalty subproblem is approximately solved by a block coordinate descent (BCD) method.

Penalty decomposition method for (1):

Let $\{\epsilon_k\}$ be a positive decreasing sequence. Let $\varrho_0 > 0$, $\sigma > 1$ be given. Choose an arbitrary $y_0^0 \in \mathcal{Y}$ and a constant $\Upsilon \geq \max\{f(x^{\text{feas}}), \min_{x \in \mathcal{X}} q_{\varrho_0}(x, y_0^0)\}$. Set $k = 0$.

- 1) Set $l = 0$ and apply the BCD method to find an approximate solution $(x^k, y^k) \in \mathcal{X} \times \mathcal{Y}$ for the penalty subproblem

$$\min\{q_{\varrho_k}(x, y) : x \in \mathcal{X}, y \in \mathcal{Y}\} \quad (12)$$

by performing steps 1a)-1d):

- 1a) Solve $x_{l+1}^k \in \text{Arg min}_{x \in \mathcal{X}} q_{\rho_k}(x, y_l^k)$.
- 1b) Solve $y_{l+1}^k \in \text{Arg min}_{y \in \mathcal{Y}} q_{\rho_k}(x_{l+1}^k, y)$.
- 1c) Set $(x^k, y^k) := (x_{l+1}^k, y_{l+1}^k)$. If (x^k, y^k) satisfies

$$\|\mathcal{P}_{\mathcal{X}}(x^k - \nabla_x q_{\rho_k}(x^k, y^k)) - x^k\| \leq \epsilon_k, \quad (13)$$

then go to step 2).

- 1d) Set $l \leftarrow l + 1$ and go to step 1a).
- 2) Set $\rho_{k+1} := \sigma \rho_k$.
- 3) If $\min_{x \in \mathcal{X}} q_{\rho_{k+1}}(x, y^k) > \Upsilon$, set $y_0^{k+1} := x^{\text{feas}}$. Otherwise, set $y_0^{k+1} := y^k$.
- 4) Set $k \leftarrow k + 1$ and go to step 1).

end

Remark. The condition (13) will be used to establish the global convergence of the above method. It may not be easily verifiable unless \mathcal{X} is simple. On the other hand, we observe that the sequence $\{q_{\rho_k}(x_l^k, y_l^k)\}$ is non-increasing for any fixed k . In practice, it is thus reasonable to terminate the BCD method based on the progress of $\{q_{\rho_k}(x_l^k, y_l^k)\}$. Another practical termination criterion for the BCD method is based on the relative change of the sequence $\{(x_l^k, y_l^k)\}$, that is,

$$\max \left\{ \frac{\|x_l^k - x_{l-1}^k\|_{\infty}}{\max(\|x_l^k\|_{\infty}, 1)}, \frac{\|y_l^k - y_{l-1}^k\|_{\infty}}{\max(\|y_l^k\|_{\infty}, 1)} \right\} \leq \epsilon_I \quad (14)$$

for some $\epsilon_I > 0$. In addition, we can terminate the outer iterations of the PD method once

$$\|x^k - y^k\|_{\infty} \leq \epsilon_O \quad (15)$$

for some $\epsilon_O > 0$. Given that problem (12) is nonconvex, the BCD method may converge to a stationary point. To enhance the performance of the BCD method, one may execute it multiple times by restarting from a suitable perturbation of the current best approximate solution. For example, at the k th outer iteration, let (x^k, y^k) be the current best approximate solution of (12) found by the BCD method, and let $r_k = \|y^k\|_0$. Assume that $r_k > 1$. Before starting the $(k+1)$ th outer iteration, one can re-apply the BCD method starting from $y_0^k \in \text{Arg min}\{\|y - y^k\| : \|y\|_0 \leq r_k - 1\}$ and obtain a new approximate solution $(\tilde{x}^k, \tilde{y}^k)$ of (12). If $q_{\rho_k}(\tilde{x}^k, \tilde{y}^k)$ is “sufficiently” smaller than $q_{\rho_k}(x^k, y^k)$, one can set $(x^k, y^k) := (\tilde{x}^k, \tilde{y}^k)$ and repeat the above process. Otherwise, one can terminate the k th outer iteration and start the next outer iteration. Finally, it follows from Proposition 3.1 that the subproblem in step 1b) has a closed-form solution. \blacksquare

We next establish a convergence result regarding the inner iterations of the above PD method. In particular, we will show that an approximate solution (x^k, y^k) of problem (12) satisfying (13) can be found by the BCD method described in steps 1a)-1d). For notational convenience, we omit the index k from (12) and consider the BCD method for solving the problem

$$\min\{q_{\rho}(x, y) : x \in \mathcal{X}, y \in \mathcal{Y}\} \quad (16)$$

instead. Accordingly, we rename the iterates of the above BCD method and present it as follows.

Block coordinate descent method for (16):

Choose an arbitrary initial point $y^0 \in \mathcal{Y}$. Set $l = 0$.

- 1) Solve $x^{l+1} \in \text{Arg min}_{x \in \mathcal{X}} q_\rho(x, y^l)$.
- 2) Solve $y^{l+1} \in \text{Arg min}_{y \in \mathcal{Y}} q_\rho(x^{l+1}, y)$.
- 3) Set $l \leftarrow l + 1$ and go to step 1).

end

Lemma 4.1 *Suppose that $(x^*, y^*) \in \mathfrak{R}^n \times \mathfrak{R}^{|J|}$ is a block coordinate minimizer of problem (16), that is,*

$$x^* \in \text{Arg min}_{x \in \mathcal{X}} q_\rho(x, y^*), \quad y^* \in \text{Arg min}_{y \in \mathcal{Y}} q_\rho(x^*, y). \quad (17)$$

Furthermore, assume that h 's are affine functions, and f and g 's are convex functions. Then, (x^, y^*) is a local minimizer of problem (16).*

Proof. Let $K = \{i : y_i^* \neq 0\}$, and let h_x, h_y be any two vectors such that $x^* + h_x \in \mathcal{X}$, $y^* + h_y \in \mathcal{Y}$ and $|(h_y)_i| < |y_i^*|$ for all $i \in K$. Claim that

$$(y^* - x_J^*)^T h_y = 0. \quad (18)$$

If $\|x_J^*\|_0 > r$, we observe from the second relation of (17) and Proposition 3.1 that $\|y^*\|_0 = r$ and $y_i^* = x_{J(i)}^*$ for all $i \in K$, which, together with $y^* + h_y \in \mathcal{Y}$ and $|(h_y)_i| < |y_i^*|$ for all $i \in K$, implies that $(h_y)_i = 0$ for all $i \notin K$ and hence (18) holds. On the other hand, if $\|x_J^*\|_0 \leq r$, one can observe that $y^* = x_J^*$ and thus (18) also holds. In addition, by the assumption that h 's are affine functions, and f and g 's are convex functions, we know that q_ρ is convex. It then follows from the first relation of (17) and the first-order optimality condition that $[\nabla_x q_\rho(x^*, y^*)]^T h_x \geq 0$. Using this inequality along with (18) and the convexity of q_ρ , we have

$$\begin{aligned} q_\rho(x^* + h_x, y^* + h_y) &\geq q_\rho(x^*, y^*) + [\nabla_x q_\rho(x^*, y^*)]^T h_x + [\nabla_y q_\rho(x^*, y^*)]^T h_y \\ &= q_\rho(x^*, y^*) + [\nabla_x q_\rho(x^*, y^*)]^T h_x + \varrho(y^* - x_J^*)^T h_y \geq q_\rho(x^*, y^*), \end{aligned}$$

which together with the above choice of h_x and h_y implies that (x^*, y^*) is a local minimizer of (16). ■

Theorem 4.2 *Let $\{(x^l, y^l)\}$ be the sequence generated by the above BCD method, and let $\epsilon > 0$ be given. Suppose that (x^*, y^*) is an accumulation point of $\{(x^l, y^l)\}$. Then the following statements hold:*

- (a) (x^*, y^*) is a block coordinate minimizer of problem (16).
- (b) There exists some $l > 0$ such that

$$\|\mathcal{P}_{\mathcal{X}}(x^l - \nabla_x q_\rho(x^l, y^l)) - x^l\| < \epsilon.$$

(c) Furthermore, if h 's are affine functions, and f and g 's are convex functions, then (x^*, y^*) is a local minimizer of problem (16).

Proof. We first show that statement (a) holds. Indeed, one can observe that

$$q_\varrho(x^{l+1}, y^l) \leq q_\varrho(x, y^l) \quad \forall x \in \mathcal{X}, \quad (19)$$

$$q_\varrho(x^l, y^l) \leq q_\varrho(x^l, y) \quad \forall y \in \mathcal{Y}. \quad (20)$$

It follows that

$$q_\varrho(x^{l+1}, y^{l+1}) \leq q_\varrho(x^{l+1}, y^l) \leq q_\varrho(x^l, y^l) \quad \forall l \geq 1. \quad (21)$$

Hence, the sequence $\{q_\varrho(x^l, y^l)\}$ is non-increasing. Since (x^*, y^*) is an accumulation point of $\{(x^l, y^l)\}$, there exists a subsequence L such that $\lim_{l \in L \rightarrow \infty} (x^l, y^l) = (x^*, y^*)$. We then observe that $\{q_\varrho(x^l, y^l)\}_{l \in L}$ is bounded, which together with the monotonicity of $\{q_\varrho(x^l, y^l)\}$ implies that $\{q_\varrho(x^l, y^l)\}$ is bounded below and hence $\lim_{l \rightarrow \infty} q_\varrho(x^l, y^l)$ exists. This observation, (21) and the continuity of $q_\varrho(\cdot, \cdot)$ yield

$$\lim_{l \rightarrow \infty} q_\varrho(x^{l+1}, y^l) = \lim_{l \rightarrow \infty} q_\varrho(x^l, y^l) = \lim_{l \in L \rightarrow \infty} q_\varrho(x^l, y^l) = q_\varrho(x^*, y^*).$$

Using these relations, the continuity of $q_\varrho(\cdot, \cdot)$, and taking limits on both sides of (19) and (20) as $l \in L \rightarrow \infty$, we have

$$q_\varrho(x^*, y^*) \leq q_\varrho(x, y^*) \quad \forall x \in \mathcal{X}, \quad (22)$$

$$q_\varrho(x^*, y^*) \leq q_\varrho(x^*, y) \quad \forall y \in \mathcal{Y}. \quad (23)$$

In addition, from the definition of \mathcal{Y} , we know that $\|y^l\|_0 \leq r$, which immediately implies $\|y^*\|_0 \leq r$. Also, $x^* \in \mathcal{X}$ due to the closedness of \mathcal{X} . This together with (22) and (23) implies that (x^*, y^*) is a block coordinate minimizer of (16) and hence statement (a) holds. Using (22) and the first-order optimality condition, we have

$$\|\mathcal{P}_{\mathcal{X}}(x^* - \nabla_x q_\varrho(x^*, y^*)) - x^*\| = 0.$$

By the continuity of $\mathcal{P}_{\mathcal{X}}(\cdot)$ and $\nabla_x q_\varrho(\cdot, \cdot)$, and the relation $\lim_{l \in L \rightarrow \infty} (x^l, y^l) = (x^*, y^*)$, one can see that

$$\lim_{l \in L \rightarrow \infty} \|\mathcal{P}_{\mathcal{X}}(x^l - \nabla_x q_\varrho(x^l, y^l)) - x^l\| = 0,$$

and hence, statement (b) immediately follows. In addition, statement (c) holds due to statement (a) and Lemma 4.1. \blacksquare

The following theorem establishes the convergence of the outer iterations of the PD method for solving problem (1). In particular, we show that under some suitable assumption, any accumulation point of the sequence generated by the PD method satisfies the first-order optimality conditions of (1). Moreover, when the l_0 part is the only nonconvex part, we show that under some assumption, the accumulation point is a local minimizer of (1).

Theorem 4.3 *Assume that $\epsilon_k \rightarrow 0$. Let $\{(x^k, y^k)\}$ be the sequence generated by the above PD method, $I_k = \{i_1^k, \dots, i_r^k\}$ be a set of r distinct indices in $\{1, \dots, |J|\}$ such that $(y^k)_i = 0$ for any $i \notin I_k$, and let $J_k = \{J(i) : i \in I_k\}$. Suppose that the level set $\mathcal{X}_\Upsilon := \{x \in \mathcal{X} : f(x) \leq \Upsilon\}$ is compact. Then, the following statements hold:*

(a) *The sequence $\{(x^k, y^k)\}$ is bounded.*

(b) Suppose (x^*, y^*) is an accumulation point of $\{(x^k, y^k)\}$. Then, $x^* = y^*$ and x^* is a feasible point of problem (1). Moreover, there exists a subsequence K such that $\{(x^k, y^k)\}_{k \in K} \rightarrow (x^*, y^*)$, $I_k = I^*$ and $J_k = J^* := \{J(i) : i \in I^*\}$ for some index set $I^* \subseteq \{1, \dots, |J|\}$ when $k \in K$ is sufficiently large.

(c) Let x^* , K and J^* be defined above, and let $\bar{J}^* = J \setminus J^*$. Suppose that the Robinson condition (3) holds at x^* for such \bar{J}^* . Then, $\{(\lambda^k, \mu^k, \varpi^k)\}_{k \in K}$ is bounded, where

$$\lambda^k = \varrho_k [g(x^k)]^+, \quad \mu^k = \varrho_k h(x^k), \quad \varpi^k = \varrho_k (x_J^k - y^k). \quad (24)$$

Moreover, each accumulation point $(\lambda^*, \mu^*, \varpi^*)$ of $\{(\lambda^k, \mu^k, \varpi^k)\}_{k \in K}$ together with x^* satisfies the first-order optimality conditions (5) with $z_j^* = \varpi_j^*$ for all $j = J(i) \in \bar{J}^*$. Further, if $\|x_J^*\|_0 = r$, h 's are affine functions, and f and g 's are convex functions, then x^* is a local minimizer of problem (1).

Proof. In view of (11) and our choice of y_0^k that is specified in step 3), one can observe that

$$f(x^k) + \frac{\varrho_k}{2} (\| [g(x^k)]^+ \|^2 + \|h(x^k)\|^2 + \|x_J^k - y^k\|^2) = q_{\varrho_k}(x^k, y^k) \leq \min_{x \in \mathcal{X}} q_{\varrho_k}(x, y_0^k) \leq \Upsilon \quad \forall k. \quad (25)$$

It immediately implies that $\{x^k\} \subseteq \mathcal{X}_\Upsilon$, and hence, $\{x^k\}$ is bounded. Moreover, we can obtain from (25) that

$$\|x_J^k - y^k\|^2 \leq 2[\Upsilon - f(x^k)]/\varrho_k \leq 2[\Upsilon - \min_{x \in \mathcal{X}_\Upsilon} f(x)]/\varrho_0,$$

which together with the boundedness of $\{x^k\}$ yields that $\{y^k\}$ is bounded. Therefore, statement (a) follows. We next show that statement (b) also holds. Since (x^*, y^*) is an accumulation point of $\{(x^k, y^k)\}$, there exists a subsequence $\{(x^k, y^k)\}_{k \in \bar{K}} \rightarrow (x^*, y^*)$. Recall that I_k is an index set. It follows that $\{(i_1^k, \dots, i_r^k)\}_{k \in \bar{K}}$ is bounded for all k . Thus there exists a subsequence $K \subseteq \bar{K}$ such that $\{(i_1^k, \dots, i_r^k)\}_{k \in K} \rightarrow (i_1^*, \dots, i_r^*)$ for some r distinct indices i_1^*, \dots, i_r^* . Since i_1^k, \dots, i_r^k are r distinct integers, one can easily conclude that $(i_1^k, \dots, i_r^k) = (i_1^*, \dots, i_r^*)$ for sufficiently large $k \in K$. Let $I^* = \{i_1^*, \dots, i_r^*\}$. It then follows that $I_k = I^*$ and $J_k = J^*$ when $k \in K$ is sufficiently large, and moreover, $\{(x^k, y^k)\}_{k \in K} \rightarrow (x^*, y^*)$. Therefore, statement (b) holds. Finally, we show that statement (c) holds. Indeed, let s^k be the vector such that

$$\mathcal{P}_{\mathcal{X}}(x^k - \nabla_x q_{\varrho_k}(x^k, y^k)) = x^k + s^k.$$

It then follows from (13) that $\|s^k\| \leq \epsilon_k$ for all k , which together with $\lim_{k \rightarrow \infty} \epsilon_k = 0$ implies $\lim_{k \rightarrow \infty} s^k = 0$. By a well-known property of the projection map $\mathcal{P}_{\mathcal{X}}$, we have

$$(x - x^k - s^k)^T [x^k - \nabla_x q_{\varrho_k}(x^k, y^k) - x^k - s^k] \leq 0, \quad \forall x \in \mathcal{X}.$$

Hence, we obtain that

$$-\nabla_x q_{\varrho_k}(x^k, y^k) - s^k \in \mathcal{N}_{\mathcal{X}}(x^k + s^k). \quad (26)$$

Using this relation, (26), (24) and the definition of q_{ϱ} , we have

$$-\nabla f(x^k) - \nabla g(x^k) \lambda^k - \nabla h(x^k) \mu^k - I_J \varpi^k - s^k \in \mathcal{N}_{\mathcal{X}}(x^k + s^k). \quad (27)$$

We now claim that $\{(\lambda^k, \mu^k, \varpi^k)\}_{k \in K}$ is bounded. Suppose for contradiction that it is unbounded. By passing to a subsequence if necessary, we can assume that $\{\|(\lambda^k, \mu^k, \varpi^k)\|\}_{k \in K} \rightarrow \infty$. Let $(\bar{\lambda}^k, \bar{\mu}^k, \bar{\varpi}^k) =$

$(\lambda^k, \mu^k, \varpi^k)/\|(\lambda^k, \mu^k, \varpi^k)\|$. Without loss of generality, we assume that $\{(\bar{\lambda}^k, \bar{\mu}^k, \bar{\varpi}^k)\}_{k \in K} \rightarrow (\bar{\lambda}, \bar{\mu}, \bar{\varpi})$ (otherwise, one can consider its convergent subsequence). Clearly, $\|(\bar{\lambda}, \bar{\mu}, \bar{\varpi})\| = 1$. Dividing both sides of (27) by $\|(\lambda^k, \mu^k, \varpi^k)\|$, taking limits as $k \in K \rightarrow \infty$, and using the relation $\lim_{k \in K \rightarrow \infty} s^k = 0$ and the semicontinuity of $\mathcal{N}_{\mathcal{X}}(\cdot)$, we obtain that

$$-\nabla g(x^*)\bar{\lambda} - \nabla h(x^*)\bar{\mu} - I_J\bar{\varpi} \in \mathcal{N}_{\mathcal{X}}(x^*). \quad (28)$$

We can see from (4) and (24) that $\bar{\lambda} \in \mathfrak{R}_+^m$, and $\bar{\lambda}_i = 0$ for $i \notin \mathcal{A}(x^*)$. Also, from Proposition 3.1 and the definitions of y_k, I_k and J_k , one can observe that $x_{J_k}^k = y_{I_k}^k$ and hence $\varpi_{I_k}^k = 0$. In addition, we know from statement (b) that $I_k = I^*$ when $k \in K$ is sufficiently large. Hence, $\bar{\varpi}_{I^*} = 0$. Since Robinson's condition (3) is satisfied at x^* , there exist $d \in \mathcal{T}_{\mathcal{X}}(x^*)$ and $v \in \mathfrak{R}^m$ such that $v_i \leq 0$ for $i \in \mathcal{A}(x^*)$, and

$$g'(x^*)d - v = -\bar{\lambda}, \quad h'(x^*)d = -\bar{\mu}, \quad (I_{\bar{J}^*})^T d = -\bar{\varpi}_{\bar{J}^*},$$

where \bar{J}^* is the complement of I^* in $\{1, \dots, |J|\}$. Recall that $\bar{\lambda} \in \mathfrak{R}_+^m$, $\bar{\lambda}_i = 0$ for $i \notin \mathcal{A}(x^*)$, and $v_i \leq 0$ for $i \in \mathcal{A}(x^*)$. Hence, $v^T \bar{\lambda} \leq 0$. In addition, since $\bar{\varpi}_{I^*} = 0$, one has $I_J \bar{\varpi} = I_{\bar{J}^*} \bar{\varpi}_{\bar{J}^*}$. Using these relations, (28), and the facts that $d \in \mathcal{T}_{\mathcal{X}}(x^*)$ and $\bar{\varpi}_{I^*} = 0$, we have

$$\begin{aligned} \|\bar{\lambda}\|^2 + \|\bar{\mu}\|^2 + \|\bar{\varpi}\|^2 &= -[(-\bar{\lambda})^T \bar{\lambda} + (-\bar{\mu})^T \bar{\mu} + (-\bar{\varpi}_{\bar{J}^*})^T \bar{\varpi}_{\bar{J}^*}] \\ &= -[(g'(x^*)d - v)^T \bar{\lambda} + (h'(x^*)d)^T \bar{\mu} + ((I_{\bar{J}^*})^T d)^T \bar{\varpi}_{\bar{J}^*}] \\ &= d^T (-\nabla g(x^*)\bar{\lambda} - \nabla h(x^*)\bar{\mu} - I_J \bar{\varpi}) + v^T \bar{\lambda} \leq 0. \end{aligned}$$

It yields $(\bar{\lambda}, \bar{\mu}, \bar{\varpi}) = (0, 0, 0)$, which contradicts the identity $\|(\bar{\lambda}, \bar{\mu}, \bar{\varpi})\| = 1$. Therefore, the subsequence $\{(\lambda^k, \mu^k, \varpi^k)\}_{k \in K}$ is bounded. Let $(\lambda^*, \mu^*, \varpi^*)$ be an accumulation point of $\{(\lambda^k, \mu^k, \varpi^k)\}_{k \in K}$. By passing to a subsequence if necessary, we can assume that $(\lambda^k, \mu^k, \varpi^k) \rightarrow (\lambda^*, \mu^*, \varpi^*)$ as $k \in K \rightarrow \infty$. Taking limits on both sides of (27) as $k \in K \rightarrow \infty$, and using the relations $\lim_{k \in K \rightarrow \infty} s^k = 0$ and the semicontinuity of $\mathcal{N}_{\mathcal{X}}(\cdot)$, we see that the first relation of (5) holds with $z^* = I_J \varpi^*$. By a similar argument as above, one can show that $\varpi_{I^*}^* = 0$. This together with the definitions of J^* and \bar{J}^* implies that z^* satisfies

$$z_j^* = \begin{cases} 0 & \text{if } j \in \bar{J} \cup J^*, \\ \varpi_i^* & \text{if } j = J(i) \in \bar{J}^*, \end{cases}$$

where \bar{J} is the complement of J in $\{1, \dots, n\}$. In addition, we see from (24) that $\lambda_i^k \geq 0$ and $\lambda_i^k g_i(x^k) = 0$ for all i , which immediately lead to the second relation of (5). Hence, $(\lambda^*, \mu^*, \varpi^*)$ together with x^* satisfies (5). Suppose now that $\|x_j^*\|_0 = r$. Then, $\mathcal{J}^* = \{\tilde{J}^* \subseteq J : |\tilde{J}^*| = r, x_j^* = 0, \forall j \notin \tilde{J}^*\} = \{J^*\}$. Therefore, the assumptions of Theorem 2.3 hold. It then follows from Theorem 2.3 that x^* is a local minimizer of (1). \blacksquare

4.2 Penalty decomposition method for problem (2)

In this subsection we propose a PD method for solving problem (2) and establish some convergence results for it.

We observe that problem (2) can be equivalently reformulated as

$$\min_{x \in \mathcal{X}, y \in \mathfrak{R}^{|J|}} \{f(x) + \nu \|y\|_0 : g(x) \leq 0, h(x) = 0, x_J - y = 0\}. \quad (29)$$

The associated quadratic penalty function for (29) is defined as

$$p_\varrho(x, y) := f(x) + \nu \|y\|_0 + \frac{\varrho}{2} (\| [g(x)]^+ \|^2 + \|h(x)\|^2 + \|x_J - y\|^2) \quad \forall x \in \mathcal{X}, y \in \mathfrak{R}^{|J|} \quad (30)$$

for some penalty parameter $\varrho > 0$.

We are now ready to present the PD method for solving (29) (or, equivalently, (2)) in which each penalty subproblem is approximately solved by a BCD method.

Penalty decomposition method for (2):

Let $\{\epsilon_k\}$ be a positive decreasing sequence. Let $\varrho_0 > 0$, $\sigma > 1$ be given, and let q_ϱ be defined in (11). Choose an arbitrary $y_0^0 \in \mathfrak{R}^{|J|}$ and a constant Υ such that $\Upsilon \geq \max\{f(x^{\text{feas}}) + \nu \|x^{\text{feas}}\|_0, \min_{x \in \mathcal{X}} p_{\varrho_0}(x, y_0^0)\}$. Set $k = 0$.

- 1) Set $l = 0$ and apply the BCD method to find an approximate solution $(x^k, y^k) \in \mathcal{X} \times \mathfrak{R}^{|J|}$ for the penalty subproblem

$$\min\{p_{\varrho_k}(x, y) : x \in \mathcal{X}, y \in \mathfrak{R}^{|J|}\} \quad (31)$$

by performing steps 1a)-1d):

- 1a) Solve $x_{l+1}^k \in \text{Arg} \min_{x \in \mathcal{X}} p_{\varrho_k}(x, y_l^k)$.

- 1b) Solve $y_{l+1}^k \in \text{Arg} \min_{y \in \mathfrak{R}^{|J|}} p_{\varrho_k}(x_{l+1}^k, y)$.

- 1c) Set $(x^k, y^k) := (x_{l+1}^k, y_{l+1}^k)$. If (x^k, y^k) satisfies

$$\|\mathcal{P}_{\mathcal{X}}(x^k - \nabla_x q_{\varrho_k}(x^k, y^k)) - x^k\| \leq \epsilon_k, \quad (32)$$

then go to step 2).

- 1d) Set $l \leftarrow l + 1$ and go to step 1a).

- 2) Set $\varrho_{k+1} := \sigma \varrho_k$.

- 3) If $\min_{x \in \mathcal{X}} p_{\varrho_{k+1}}(x, y^k) > \Upsilon$, set $y_0^{k+1} := x^{\text{feas}}$. Otherwise, set $y_0^{k+1} := y^k$.

- 4) Set $k \leftarrow k + 1$ and go to step 1).

end

Remark. The practical termination criteria proposed in Subsection 4.1 can also be applied to this PD method. In addition, one can apply a similar strategy as mentioned in Subsection 4.1 to enhance the performance of the BCD method for solving (31). Finally, in view of Proposition 3.2, the BCD subproblem in step 1b) has a closed-form solution. ■

We next establish a convergence result regarding the inner iterations of the above PD method. In particular, we will show that an approximate solution (x^k, y^k) of problem (31) satisfying (32) can be found by the BCD method described in steps 1a)-1d). For convenience of presentation, we omit the index k from (31) and consider the BCD method for solving the following problem:

$$\min\{p_\varrho(x, y) : x \in \mathcal{X}, y \in \mathfrak{R}^{|J|}\} \quad (33)$$

instead. Accordingly, we rename the iterates of the above BCD method. We can observe that the resulting BCD method is the same as the one presented in Subsection 4.1 except that p_ϱ and $\mathfrak{R}^{|J|}$ replace q_ϱ and \mathcal{Y} , respectively. For the sake of brevity, we omit the presentation of this BCD method.

Lemma 4.4 Suppose that $(x^*, y^*) \in \mathfrak{R}^n \times \mathfrak{R}^{|J|}$ is a block coordinate minimizer of problem (33), that is,

$$x^* \in \operatorname{Arg} \min_{x \in \mathcal{X}} p_\varrho(x, y^*), \quad y^* \in \operatorname{Arg} \min_{y \in \mathfrak{R}^{|J|}} p_\varrho(x^*, y). \quad (34)$$

Furthermore, assume that h 's are affine functions, and f and g 's are convex functions. Then, (x^*, y^*) is a local minimizer of problem (33).

Proof. Let $K = \{i : y_i^* \neq 0\}$, and let h_x, h_y be any two vectors such that $x^* + h_x \in \mathcal{X}$, $|(h_y)_i| < \nu/(\rho|x_{J(i)}^*| + 1)$ for any $i \notin K$ and $|(h_y)_i| < |y_i^*|$ for all $i \in K$. We observe from the second relation of (34) and Proposition 3.2 that $y_i^* = x_{J(i)}^*$ for all $i \in K$. Also, for the above choice of h_y , one has $y_i^* + (h_y)_i \neq 0$ for all $i \in K$. Hence, $\|y_i^* + (h_y)_i\|_0 = \|y_i^*\|_0$ for every $i \in K$. Using these relations and the definition of h_y , we can see that

$$\rho(y^* - x_J^*)^T h_y + \nu\|y^* + h_y\|_0 - \nu\|y^*\|_0 = -\rho \sum_{i \notin K} x_{J(i)}^* (h_y)_i + \nu \sum_{i \notin K} \|(h_y)_i\|_0 \geq 0. \quad (35)$$

Let q_ϱ be defined in (11). By the first relation of (34) and a similar argument as in Lemma 4.1, we have $[\nabla_x q_\varrho(x^*, y^*)]^T h_x \geq 0$. Using this inequality along with (35) and the convexity of q_ϱ , we have

$$\begin{aligned} p_\varrho(x^* + h_x, y^* + h_y) &= q_\varrho(x^* + h_x, y^* + h_y) + \nu\|y^* + h_y\|_0 \\ &\geq q_\varrho(x^*, y^*) + [\nabla_x q_\varrho(x^*, y^*)]^T h_x + [\nabla_y q_\varrho(x^*, y^*)]^T h_y + \nu\|y^* + h_y\|_0 \\ &\geq p_\varrho(x^*, y^*) + \varrho(y^* - x_J^*)^T h_y + \nu\|y^* + h_y\|_0 - \nu\|y^*\|_0 \geq p_\varrho(x^*, y^*), \end{aligned}$$

which together with our choice of h_x and h_y implies that (x^*, y^*) is a local minimizer of (33). \blacksquare

Theorem 4.5 Let $\{(x^l, y^l)\}$ be the sequence generated by the above BCD method, and let $\epsilon > 0$ be given. Suppose that (x^*, y^*) is an accumulation point of $\{(x^l, y^l)\}$. Then the following statements hold:

- (a) (x^*, y^*) is a block coordinate minimizer of problem (33).
- (b) There exists some $l > 0$ such that

$$\|\mathcal{P}_{\mathcal{X}}(x^l - \nabla_x q_\varrho(x^l, y^l)) - x^l\| < \epsilon$$

where the function q_ϱ is defined in (11).

- (c) Furthermore, if h 's are affine functions, and f and g 's are convex functions, then (x^*, y^*) is a local minimizer of problem (33).

Proof. We first show that statement (a) holds. Indeed, one can observe that

$$p_\varrho(x^{l+1}, y^l) \leq p_\varrho(x, y^l) \quad \forall x \in \mathcal{X}, \quad (36)$$

$$p_\varrho(x^l, y^l) \leq p_\varrho(x^l, y) \quad \forall y \in \mathfrak{R}^{|J|}. \quad (37)$$

Using these relations and a similar argument as in the proof of Theorem 4.2, one can show that $\lim_{l \rightarrow \infty} p_\varrho(x^l, y^l)$ exists, and moreover,

$$\lim_{l \rightarrow \infty} p_\varrho(x^l, y^l) = \lim_{l \rightarrow \infty} p_\varrho(x^{l+1}, y^l). \quad (38)$$

Since (x^*, y^*) is an accumulation point of $\{(x^l, y^l)\}$, there exists a subsequence L such that $\lim_{l \in L \rightarrow \infty} (x^l, y^l) = (x^*, y^*)$ and moreover, $x^* \in \mathcal{X}$ due to the closedness of \mathcal{X} . For notational convenience, let

$$F(x) := f(x) + \frac{\rho}{2}(\|[g(x)]^+\|^2 + \|h(x)\|^2).$$

It then follows from (30) that

$$p_\rho(x, y) = F(x) + \nu \|y\|_0 + \frac{\rho}{2} \|x_J - y\|^2, \quad \forall x \in \mathcal{X}, y \in \mathfrak{R}^{|J|}. \quad (39)$$

Since $\lim_{l \in L} y^l = y^*$, one has $\|y^l\|_0 \geq \|y^*\|_0$ for sufficiently large $l \in L$. Using this relation, (37) and (39), we obtain that, when $l \in L$ is sufficiently large,

$$p_\rho(x^l, y) \geq p_\rho(x^l, y^l) = F(x^l) + \nu \|y^l\|_0 + \frac{\rho}{2} \|x_J^l - y^l\|^2 \geq F(x^l) + \nu \|y^*\|_0 + \frac{\rho}{2} \|x_J^l - y^l\|^2.$$

Upon taking limits on both sides of the above inequality as $l \in L \rightarrow \infty$ and using the continuity of F , one has

$$p_\rho(x^*, y) \geq F(x^*) + \nu \|y^*\|_0 + \frac{\rho}{2} \|x_J^* - y^*\|^2 = p_\rho(x^*, y^*), \quad \forall y \in \mathfrak{R}^{|J|}. \quad (40)$$

In addition, it follows from (36) and (39) that

$$\begin{aligned} F(x) + \frac{1}{2} \|x_J - y^l\|^2 &= p_\rho(x, y^l) - \nu \|y^l\|_0 \geq p_\rho(x^{l+1}, y^l) - \nu \|y^l\|_0 \\ &= F(x^{l+1}) + \frac{1}{2} \|x_J^{l+1} - y^l\|^2, \quad \forall x \in \mathcal{X}. \end{aligned} \quad (41)$$

Since $\{\|y^l\|_0\}_{l \in L}$ is bounded, there exists a subsequence $\bar{L} \subseteq L$ such that $\lim_{l \in \bar{L} \rightarrow \infty} \|y^l\|_0$ exists. Then we have

$$\begin{aligned} \lim_{l \in \bar{L} \rightarrow \infty} F(x^{l+1}) + \frac{1}{2} \|x_J^{l+1} - y^l\|^2 &= \lim_{l \in \bar{L} \rightarrow \infty} p_\rho(x^{l+1}, y^l) - \nu \|y^l\|_0 \\ &= \lim_{l \in \bar{L} \rightarrow \infty} p_\rho(x^{l+1}, y^l) - \nu \lim_{l \in \bar{L} \rightarrow \infty} \|y^l\|_0 = \lim_{l \in \bar{L} \rightarrow \infty} p_\rho(x^l, y^l) - \nu \lim_{l \in \bar{L} \rightarrow \infty} \|y^l\|_0 \\ &= \lim_{l \in \bar{L} \rightarrow \infty} p_\rho(x^l, y^l) - \nu \|y^*\|_0 = \lim_{l \in \bar{L} \rightarrow \infty} F(x^l) + \frac{1}{2} \|x_J^l - y^l\|^2 = F(x^*) + \frac{1}{2} \|x_J^* - y^*\|^2, \end{aligned}$$

where the third equality is due to (38). Using this relation and taking limits on both sides of (41) as $l \in \bar{L} \rightarrow \infty$, we further have

$$F(x) + \frac{1}{2} \|x_J - y^*\|^2 \geq F(x^*) + \frac{1}{2} \|x_J^* - y^*\|^2, \quad \forall x \in \mathcal{X},$$

which together with (30) yields

$$p_\rho(x, y^*) \geq p_\rho(x^*, y^*), \quad \forall x \in \mathcal{X}.$$

This relation along with (40) implies that (x^*, y^*) is a block coordinate minimizer of (33) and hence statement (a) holds. Statement (b) can be similarly proved as that of Theorem 4.2. In addition, statement (c) holds due to statement (a) and Lemma 4.4. \blacksquare

Remark. A similar result as in statement (c) is recently established in [59] for the BCD method when applied to solve the *unconstrained* problem:

$$\min_{x, y} \frac{1}{2} \|Ax - b\|^2 + \frac{\rho}{2} \|Wx - y\|^2 + \sum_i \nu_i \|y_i\|_0 \quad (42)$$

under the assumption that $A^T A \succ 0$, $W^T W = I$, $\varrho > 0$, and $\nu_i \geq 0$ for all i . The proof of [59] strongly relies on this assumption and the fact that the BCD subproblems have closed-form solutions. We believe that it cannot be extended to problem (33). In addition, it is not hard to observe that problem (42) can be equivalently reformulated into a problem in the form of (33) and thus the convergence of the BCD method for (42) directly follows from Theorem 4.5 above. ■

We next establish the convergence of the outer iterations of the PD method for solving problem (2). In particular, we show that under some suitable assumption, any accumulation point of the sequence generated by the PD method satisfies the first-order optimality conditions of (2). Moreover, when the l_0 part is the only nonconvex part, we show that the accumulation point is a local minimizer of (2).

Theorem 4.6 *Assume that $\epsilon_k \rightarrow 0$. Let $\{(x^k, y^k)\}$ be the sequence generated by the above PD method. Suppose that the level set $\mathcal{X}_\Upsilon := \{x \in \mathcal{X} : f(x) \leq \Upsilon\}$ is compact. Then, the following statements hold:*

- (a) *The sequence $\{(x^k, y^k)\}$ is bounded;*
- (b) *Suppose (x^*, y^*) is an accumulation point of $\{(x^k, y^k)\}$. Then, $x^* = y^*$ and x^* is a feasible point of problem (2).*
- (c) *Let (x^*, y^*) be defined above. Suppose that $\{(x^k, y^k)\}_{k \in K} \rightarrow (x^*, y^*)$ for some subsequence K . Let $J^* = \{j \in J : x_j^* \neq 0\}$, $\bar{J}^* = J \setminus J^*$. Assume that the Robinson condition (7) holds at x^* for such \bar{J}^* . Then, $\{(\lambda^k, \mu^k, \varpi^k)\}_{k \in K}$ is bounded, where*

$$\lambda^k = \varrho_k [g(x^k)]^+, \quad \mu^k = \varrho_k h(x^k), \quad \varpi^k = \varrho_k (x_J^k - y^k).$$

Moreover, each accumulation point $(\lambda^, \mu^*, \varpi^*)$ of $\{(\lambda^k, \mu^k, \varpi^k)\}_{k \in K}$ together with x^* satisfies the first-order optimality condition (5) with $z_j^* = \varpi_i^*$ for all $j = J(i) \in \bar{J}^*$. Further, if h 's are affine functions, and f and g 's are convex functions, then x^* is a local minimizer of problem (2).*

Proof. Statement (a) and (b) can be similarly proved as those of Theorem 4.3. We now show that statement (c) holds. Let $I^* = \{i : J(i) \in J^*\}$. From Proposition 3.2 and the definitions of y^k and J^* , we can observe that $y_{J^*}^k = x_{J^*}^k$ when $k \in K$ is sufficiently large. Hence, $\varpi_{J^*}^k = 0$ for sufficiently large $k \in K$. The rest of the proof for the first two conclusions of this statement is similar to that of statement (c) of Theorem 4.3. The last conclusion of this statement holds due to its second conclusion and Theorem 2.4. ■

5 Numerical results

In this section, we conduct numerical experiments to test the performance of our PD methods proposed in Section 4 by applying them to sparse logistic regression, sparse inverse covariance selection, and compressed sensing problems. The codes of all the methods implemented in this section are written in Matlab, which are available online at www.math.sfu.ca/~zhaosong. All experiments are performed in Matlab 7.11.0 (2010b) on a workstation with an Intel Xeon E5410 CPU (2.33 GHz) and 8GB RAM running Red Hat Enterprise Linux (kernel 2.6.18).

5.1 Sparse logistic regression problem

In this subsection, we apply the PD method studied in Subsection 4.1 to solve sparse logistic regression problem, which has numerous applications in machine learning, computer vision, data mining, bioinformatics, and neural signal processing (see, for example, [3, 54, 33, 43, 22, 44]).

Given n samples $\{z^1, \dots, z^n\}$ with p features, and n binary outcomes b_1, \dots, b_n , let $a^i = b_i z^i$ for $i = 1, \dots, n$. The *average logistic loss* function is defined as

$$l_{\text{avg}}(v, w) := \sum_{i=1}^n \theta(w^T a^i + v b_i) / n$$

for some model variables $v \in \Re$ and $w \in \Re^p$, where θ is the *logistic loss* function

$$\theta(t) := \log(1 + \exp(-t)).$$

Then the *sparse logistic regression* problem can be formulated as

$$\min_{v, w} \{l_{\text{avg}}(v, w) : \|w\|_0 \leq r\}, \quad (43)$$

where $r \in [1, p]$ is some integer for controlling the sparsity of the solution. In the literature, one common approach for finding an approximate solution to (43) is by solving the following l_1 regularization problem:

$$\min_{v, w} l_{\text{avg}}(v, w) + \lambda \|w\|_1 \quad (44)$$

for some regularization parameter $\lambda \geq 0$ (see, for example, [29, 19, 42, 31, 34, 49]). Our aim below is to apply the PD method studied in Subsection 4.1 to solve (43) directly.

Letting $x = (v, w)$, $J = \{2, \dots, p+1\}$ and $f(x) = l_{\text{avg}}(x_1, x_J)$, we can see that problem (43) is in the form of (1). Therefore, the PD method proposed in Subsection 4.1 can be suitably applied to solve (43). Also, we observe that the main computation effort of the PD method when applied to (43) lies in solving the subproblem arising in step 1a), which is in the form of

$$\min_x \left\{ l_{\text{avg}}(x_1, x_J) + \frac{\varrho}{2} \|x - c\|^2 : x \in \Re^{p+1} \right\} \quad (45)$$

for some $\varrho > 0$ and $c \in \Re^{p+1}$. To efficiently solve (45), we apply the nonmonotone projected gradient method proposed in [4, Algorithm 2.2]; in particular, we set its parameter $M = 2$ and terminate the method when $\|\nabla F(x)\| / \max\{|F(x)|, 1\} \leq 10^{-4}$, where $F(x)$ denotes the objective function of (45).

We now address the initialization and the termination criteria for our PD method when applied to (43). In particular, we randomly generate $z \in \Re^{p+1}$ such that $\|z_J\|_0 \leq r$ and set the initial point $y_0^0 = z$. We choose the initial penalty parameter ϱ_0 to be 0.1, and set the parameter $\sigma = \sqrt{10}$. In addition, we use (14) and (15) as the inner and outer termination criteria for the PD method and set their accuracy parameters ϵ_I and ϵ_O to be 5×10^{-4} and 10^{-3} , respectively.

We next conduct numerical experiments to test the performance of our PD method for solving (43) on some real and random data. We also compare the quality of the approximate solutions of (43) obtained by our method with that of (44) found by a first-order solver SLEP [34]. For the latter method, we set `opts.mFlag=1`, `opts.lFlag=1` and `opts.tFlag=2`. And the rest of its parameters are set by default.

In the first experiment, we compare the solution quality of our PD method with SLEP on three small- or medium-sized benchmark data sets which are from the UCI machine learning bench market repository [40] and other sources [23]. The first data set is the colon tumor gene expression data [23]

Table 1: Computational results on three real data sets

Data	Features p	Samples n	λ/λ_{\max}	r	SLEP			PD		
					l_{avg}	Error (%)	Time	l_{avg}	Error (%)	Time
Colon	2000	62	0.5	7	0.4398	17.74	0.2	0.4126	12.9	9.1
			0.1	22	0.1326	1.61	0.5	0.0150	0	6.0
			0.05	25	0.0664	0	0.6	0.0108	0	5.0
			0.01	28	0.0134	0	1.3	0.0057	0	5.4
Ionosphere	34	351	0.5	3	0.4804	17.38	0.1	0.3466	13.39	0.7
			0.1	11	0.3062	11.40	0.1	0.2490	9.12	1.0
			0.05	14	0.2505	9.12	0.1	0.2002	8.26	1.1
			0.01	24	0.1846	6.55	0.4	0.1710	5.98	1.7
Advertisements	1430	2359	0.5	3	0.2915	12.04	2.3	0.2578	7.21	31.9
			0.1	36	0.1399	4.11	14.2	0.1110	4.11	56.0
			0.05	67	0.1042	2.92	21.6	0.0681	2.92	74.1
			0.01	197	0.0475	1.10	153.0	0.0249	1.10	77.4

with more features than samples; the second one is the ionosphere data [40] with less features than samples; and the third one is the Internet advertisements data [40] with roughly same magnitude of features as samples. We discard the samples with missing data and standardize each data set so that the sample mean is zero and the sample variance is one. For each data set, we first apply SLEP to solve problem (44) with four different values of λ , which are the same ones as used in [29], namely, $0.5\lambda_{\max}$, $0.1\lambda_{\max}$, $0.05\lambda_{\max}$, and $0.01\lambda_{\max}$, where λ_{\max} is the upper bound on the useful range of λ that is defined in [29]. For each such λ , let w_{λ}^* be the approximate optimal w obtained by SLEP. We then apply our PD method to solve problem (43) with $r = \|w_{\lambda}^*\|_0$ so that the resulting approximate optimal w is at least as sparse as w_{λ}^* .

To compare the solution quality of the above two methods, we introduce a criterion, that is, *error rate*. Given any model variables (v, w) and a sample vector $z \in \mathbb{R}^p$, the outcome predicted by (v, w) for z is given by

$$\phi(z) = \text{sgn}(w^T z + v),$$

where

$$\text{sgn}(t) = \begin{cases} +1 & \text{if } t > 0, \\ -1 & \text{otherwise.} \end{cases}$$

Recall that z^i and b_i are the given samples and outcomes for $i = 1, \dots, n$. The *error rate* of (v, w) for predicting the outcomes b_1, \dots, b_n is defined as

$$\text{Error} := \left\{ \sum_{i=1}^n \|\phi(z^i) - b_i\|_0 / n \right\} \times 100\%. \quad (46)$$

The computational results are presented in Table 1. In detail, the name and dimensions of each data set are given in the first three columns. The fourth column gives the ratio between λ and its upper bound λ_{\max} . The fifth column lists the value of r , that is, the cardinality of w_{λ}^* which is defined above. In addition, the average logistic loss, the error rate and the CPU time (in seconds) for both SLEP and PD are reported in columns six to eleven. We can observe that, although SLEP is faster than the PD method in most cases, the PD method substantially outperforms SLEP in terms of the solution quality since it generally achieves lower average logistic loss and error rate while the cardinality of both solutions is the same.

The out-of-sample error rate is often used to evaluate the quality of a model vector, which is a slight modification of (46) by taking sum over the testing samples rather than the training samples. It




Figure 1: Sparse recovery.

usually depends on the quality and amount of training samples. For example, when the ratio between number of training samples and features is small, the out-of-sample error rate is usually high for most of models. Due to this reason, the above data sets may not be appropriate for evaluating out-of-sample error rate. Instead, we download a real data called “German” from the UCI machine learning benchmark repository [40]. This data set contains 1,000 samples and 24 features, which has a reasonably high sample-to-feature ratio. It is thus a suitable data set to evaluate out-of-sample error rate. We randomly partition those samples into two parts: one consisting of 900 samples used as training data and another consisting of 100 samples used as testing data. Similarly as above, we first apply SLEP to (44) with a sequence of suitably chosen λ to obtain solutions with cardinalities from 1 to 24. For PD, we simply set r to be 1 to 24. In this way, the solutions of PD and SLEP are of the same cardinality. The results of this experiment are presented in Figure 1. We can see that PD generally outperforms SLEP in terms of solution quality since it achieves smaller average logistic loss and lower error rate for both training and testing data.

In next experiment, we test our PD method on the random data sets of three different sizes. For each size, we randomly generate the data set consisting of 100 instances. In particular, the first data set has more features than samples; the second data set has more samples than features; and the last data set has equal number of features as samples. The samples $\{z^1, \dots, z^n\}$ and the corresponding outcomes b_1, \dots, b_n are generated in the same manner as described in [29]. In detail, for each instance we choose equal number of positive and negative samples, that is, $m_+ = m_- = m/2$, where m_+ (resp., m_-) is the number of samples with outcome $+1$ (resp., -1). The features of positive (resp., negative) samples are independent and identically distributed, drawn from a normal distribution $N(\mu, 1)$, where μ is in turn drawn from a uniform distribution on $[0, 1]$ (resp., $[-1, 0]$). For each such instance, similar to the previous experiment, we first apply SLEP to solve problem (44) with five different values of λ , which are $0.9\lambda_{\max}$, $0.7\lambda_{\max}$, $0.5\lambda_{\max}$, $0.3\lambda_{\max}$ and $0.1\lambda_{\max}$. For each such λ , let w_λ^* be the approximate optimal w obtained by SLEP. We then apply our PD method to solve problem (43) with $r = \|w_\lambda^*\|_0$ so that the resulting approximate optimal w is at least as sparse as w_λ^* . The average results of each data set over 100 instances are reported in Table 2. We also observe that the PD method is slower than SLEP, but it has better solution quality than SLEP in terms of average logistic loss and error rate.

In summary, the above experiments demonstrate that the quality of the approximate solution of (43) obtained by our PD method is generally better than that of (44) found by SLEP when the same cardinality is considered. This observation is actually not surprising as (44) is a relaxation of (43).

5.2 Sparse inverse covariance selection problem

In this subsection, we apply the PD method proposed in Subsection 4.1 to solve the sparse inverse covariance selection problem, which has numerous real-world applications such as speech recognition

Table 2: Computational results on random data sets

Size $n \times p$	λ/λ_{\max}	r	SLEP			PD		
			l_{avg}	Error (%)	Time	l_{avg}	Error (%)	Time
2000×5000	0.9	27.7	0.6359	8.01	4.3	0.1802	6.98	89.6
	0.7	91.4	0.5046	3.43	11.2	0.0550	2.37	159.6
	0.5	161.1	0.3827	1.98	17.7	0.0075	0.07	295.3
	0.3	238.9	0.2639	1.23	21.8	0.0022	0	216.1
	0.1	330.1	0.1289	0.46	19.6	0.0015	0	130.5
5000×2000	0.9	17.7	0.6380	8.49	3.5	0.2254	8.09	154.3
	0.7	65.7	0.5036	3.16	10.1	0.0372	1.55	296.6
	0.5	121.0	0.3764	1.48	16.4	0.0042	0	299.8
	0.3	180.8	0.2517	0.57	22.3	0.0018	0	190.0
	0.1	255.2	0.1114	0.04	23.7	0.0013	0	124.2
5000×5000	0.9	30.7	0.6341	7.02	4.8	0.1761	6.55	125.8
	0.7	105.7	0.5022	2.95	12.9	0.0355	1.47	255.3
	0.5	192.0	0.3793	1.63	20.3	0.0042	0	325.4
	0.3	278.3	0.2592	0.88	25.1	0.0020	0	187.4
	0.1	397.0	0.1231	0.19	24.4	0.0015	0	113.4

and gene network analysis (see, for example, [2, 18]).

Given a sample covariance matrix $\Sigma \in \mathcal{S}_{++}^p$ and a set Ω consisting of pairs of known conditionally independent nodes, the sparse inverse covariance selection problem can be formulated as

$$\begin{aligned}
 \max_{X \succeq 0} \quad & \log \det X - \langle \Sigma, X \rangle \\
 \text{s.t.} \quad & \sum_{(i,j) \in \bar{\Omega}} \|X_{ij}\|_0 \leq r, \\
 & X_{ij} = 0 \quad \forall (i,j) \in \Omega,
 \end{aligned} \tag{47}$$

where $\bar{\Omega} = \{(i,j) : (i,j) \notin \Omega, i \neq j\}$, and $r \in [1, |\bar{\Omega}|]$ is some integer for controlling the sparsity of the solution. In the literature, one common approach for finding an approximate solution to (47) is by solving the following l_1 regularization problem:

$$\begin{aligned}
 \max_{X \succeq 0} \quad & \log \det X - \langle \Sigma, X \rangle - \sum_{(i,j) \in \bar{\Omega}} \rho_{ij} |X_{ij}| \\
 \text{s.t.} \quad & X_{ij} = 0 \quad \forall (i,j) \in \Omega,
 \end{aligned} \tag{48}$$

where $\{\rho_{ij}\}_{(i,j) \in \bar{\Omega}}$ is a set of regularization parameters (see, for example, [14, 15, 1, 36, 37, 21, 57, 35]). Our goal below is to apply the PD method studied in Subsection 4.1 to solve (47) directly.

Letting $\mathcal{X} = \{X \in \mathcal{S}_+^p : X_{ij} = 0, (i,j) \in \Omega\}$ and $J = \bar{\Omega}$, we clearly see that problem (47) is in the form of (1) and thus it can be suitably solved by the PD method proposed in Subsection 4.1 with

$$\mathcal{Y} = \left\{ Y \in \mathcal{S}^p : \sum_{(i,j) \in \bar{\Omega}} \|Y_{ij}\|_0 \leq r \right\}.$$

Notice that the main computation effort of the PD method when applied to (47) lies in solving the subproblem arising in step 1a), which is in the form of

$$\min_{X \succeq 0} \left\{ -\log \det X + \frac{\rho}{2} \|X - C\|_F^2 : X_{ij} = 0 \quad \forall (i,j) \in \Omega \right\} \tag{49}$$

for some $\varrho > 0$ and $C \in \mathcal{S}^p$. Given that problem (49) generally does not have a closed-form solution, we now slightly modify the above sets \mathcal{X} and \mathcal{Y} by replacing them by

$$\mathcal{X} = \mathcal{S}_+^p, \quad \mathcal{Y} = \left\{ Y \in \mathcal{S}^p : \sum_{(i,j) \in \tilde{\Omega}} \|Y_{ij}\|_0 \leq r, Y_{ij} = 0, (i,j) \in \Omega \right\},$$

respectively, and then apply the PD method presented in Subsection 4.1 to solve (47). For this PD method, the subproblem arising in step 1a) is now in the form of

$$\min_X \left\{ -\log \det X + \frac{\varrho}{2} \|X - C\|_F^2 : X \succeq 0 \right\} \quad (50)$$

for some $\varrho > 0$ and $C \in \mathcal{S}^p$. It can be shown that problem (50) has a closed-form solution, which is given by $V \mathcal{D}(x^*) V^T$, where $x_i^* = (\lambda_i + \sqrt{\lambda_i^2 + 4/\varrho})/2$ for all i and $V \mathcal{D}(\lambda) V^T$ is the eigenvalue decomposition of C for some $\lambda \in \mathbb{R}^p$ (see, for example, [57]). Also, it follows from Proposition 3.1 that the subproblem arising in step 1b) for the above \mathcal{Y} has a closed-form solution.

We now address the initialization and the termination criteria for the above PD method. In particular, we set the initial point $Y_0^0 = (\tilde{\mathcal{D}}(\Sigma))^{-1}$, the initial penalty parameter $\varrho_0 = 1$, and the parameter $\sigma = \sqrt{10}$. In addition, we use (15) and

$$\frac{|q_{\varrho_k}(x_{l+1}^k, y_{l+1}^k) - q_{\varrho_k}(x_l^k, y_l^k)|}{\max\{|q_{\varrho_k}(x_l^k, y_l^k)|, 1\}} \leq \epsilon_I$$

as the outer and inner termination criteria for the PD method, and set the associated accuracy parameters $\epsilon_O = 10^{-4}$ and $\epsilon_I = 10^{-4}, 10^{-3}$ for the random and real data below, respectively.

We next conduct numerical experiments to test the performance of our PD method for solving (47) on some random and real data. We also compare the quality of the approximate solutions of (47) obtained by our method with that of (48) found by the proximal point algorithm (PPA) [57] and the Newton based method called QUIC¹ [27]. Our PD method and PPA call the LAPACK routine `dsyevd.f` [30] for computing the full eigenvalue decomposition of a symmetric matrix, which is usually faster than the Matlab's `eig` routine when p is larger than 500. For PPA, we set $\text{Tol} = 10^{-6}$ and use the default values for all other parameters. In addition, for QUIC, we use the default values for all parameters.

In the first experiment, we compare the solution quality of our PD method with PPA and QUIC on a set of random instances which are generated in a similar manner as described in [14, 36, 37, 57, 35]. In particular, we first generate a true covariance matrix $\Sigma^t \in \mathcal{S}_{++}^p$ such that its inverse $(\Sigma^t)^{-1}$ is with the prescribed density δ , and set

$$\Omega = \left\{ (i, j) : (\Sigma^t)^{-1}_{ij} = 0, |i - j| \geq \lfloor p/2 \rfloor \right\}.$$

We then generate a matrix $B \in \mathcal{S}^p$ by letting

$$B = \Sigma^t + \tau V,$$

where $V \in \mathcal{S}^p$ contains pseudo-random values drawn from a uniform distribution on the interval $[-1, 1]$, and τ is a small positive number. Finally, we obtain the following sample covariance matrix:

$$\Sigma = B - \min\{\lambda_{\min}(B) - \vartheta, 0\}I,$$

¹QUIC can only solve the problem (48) with $\Omega = \emptyset$.

Table 3: Computational results for $\delta = 10\%$

Problem		PPA					PD		
p	$ \Omega $	$\rho_{\bar{\Omega}}$	r	Likelihood	Loss	Time	Likelihood	Loss	Time
500	56724	0.01	183876	-950.88	2.4594	34.1	-936.45	2.3920	2.5
		0.10	45018	-999.89	2.5749	44.8	-978.61	2.4498	5.3
		1.00	5540	-1046.44	2.9190	66.2	-1032.79	2.6380	24.8
		10.0	2608	-1471.67	4.2442	75.1	-1129.50	2.8845	55.5
1000	226702	0.01	745470	-2247.14	3.1240	150.2	-2220.47	3.0486	13.1
		0.10	186602	-2344.03	3.2291	158.7	-2301.12	3.1224	19.8
		1.00	29110	-2405.88	3.5034	349.8	-2371.68	3.2743	59.1
		10.0	9604	-3094.57	4.6834	395.9	-2515.80	3.4243	129.5
1500	509978	0.01	1686128	-3647.71	3.4894	373.7	-3607.23	3.4083	35.7
		0.10	438146	-3799.02	3.5933	303.6	-3731.17	3.5059	44.9
		1.00	61222	-3873.93	3.8319	907.4	-3832.88	3.6226	155.3
		10.0	17360	-4780.33	4.9264	698.8	-3924.94	3.7146	328.0
2000	905240	0.01	3012206	-5177.80	3.7803	780.0	-5126.09	3.7046	65.5
		0.10	822714	-5375.21	3.8797	657.5	-5282.37	3.7901	94.3
		1.00	126604	-5457.90	4.0919	907.4	-5424.66	3.9713	200.2
		10.0	29954	-6535.54	5.1130	1397.4	-5532.03	4.0019	588.0

where ϑ is a small positive number. Specifically, we choose $\tau = 0.15$, $\vartheta = 1.0e - 4$, $\delta = 10\%$, 50% and 100% , respectively. Clearly, $\delta = 100\%$ means that $\Omega = \emptyset$, that is, none of zero entries of the actual sparse inverse covariance matrix is known beforehand. In addition, for all $(i, j) \in \bar{\Omega}$, we set $\rho_{ij} = \rho_{\bar{\Omega}}$ for some $\rho_{\bar{\Omega}} > 0$. For each instance, we first apply PPA or QUIC ² (48) for four values of $\rho_{\bar{\Omega}}$, which are 0.01, 0.1, 1, and 10. For each $\rho_{\bar{\Omega}}$, let \tilde{X}^* be the solution obtained by PPA or QUIC. We then apply our PD method to solve problem (47) with $r = \sum_{(i,j) \in \bar{\Omega}} \|\tilde{X}_{ij}^*\|_0$ so that the resulting solution is at least as sparse as \tilde{X}^* .

As mentioned in [35], to evaluate how well the true inverse covariance matrix $(\Sigma^t)^{-1}$ is recovered by a matrix $X \in \mathcal{S}_{++}^p$, one can compute the *normalized entropy loss* which is defined as follows:

$$\text{Loss} := \frac{1}{p} (\langle \Sigma^t, X \rangle - \log \det(\Sigma^t X) - p).$$

The results of PPA, QUIC and the PD method on these instances are presented in Tables 3-6, respectively. In each table, the order p of Σ is given in column one. The size of Ω is given in column two. The values of $\rho_{\bar{\Omega}}$ and r are given in columns three and four. The log-likelihood (i.e., the objective value of (47)), the normalized entropy loss and the CPU time (in seconds) of PPA or QUIC and the PD method are given in the last six columns, respectively. We observe that our PD method is substantially faster than PPA for these instances. Moreover, it outperforms PPA in terms of solution quality since it achieves larger log-likelihood and smaller normalized entropy loss. **We also observe that PD outperforms QUIC in terms of speed, log-likelihood and normalized entropy loss.**

In the second experiment, we aim to compare the performance of our PD method with PPA and QUIC on two gene expression data sets that have been widely used in the literature (see, for example, [24, 45, 58, 17, 35]). We first pre-process the data by the same procedure as described in [35] to obtain a sample covariance matrix Σ , and set $\Omega = \emptyset$ and $\rho_{ij} = \rho_{\bar{\Omega}}$ for some $\rho_{\bar{\Omega}} > 0$. We now apply PPA or QUIC to solve problem (48) with $\rho_{\bar{\Omega}} = 0.01, 0.05, 0.1, 0.5, 0.7$ and 0.9 , respectively. For each $\rho_{\bar{\Omega}}$, we choose r to be the number of nonzero off-diagonal entries of the solution of PPA or QUIC, which

²We only compare PD with QUIC on the instances with $\delta = 100\%$, namely, $\Omega = \emptyset$ since the latter method only solves the problem (48) with $\Omega = \emptyset$.

Table 4: Computational results for $\delta = 50\%$

Problem		$\rho_{\bar{\Omega}}$	r	PPA			PD		
p	$ \Omega $			Likelihood	Loss	Time	Likelihood	Loss	Time
500	37738	0.01	202226	-947.33	3.1774	37.2	-935.11	3.1134	2.2
		0.10	50118	-1001.23	3.3040	41.8	-978.03	3.1662	4.7
		1.00	11810	-1052.09	3.6779	81.1	-101.80	3.2889	14.5
		10.0	5032	-1500.00	5.0486	71.1	-1041.64	3.3966	28.1
1000	152512	0.01	816070	-2225.875	3.8864	149.7	-2201.98	3.8126	12.1
		0.10	203686	-2335.81	4.0029	131.0	-2288.11	3.8913	17.2
		1.00	46928	-2400.81	4.2945	372.7	-2349.02	4.0085	44.1
		10.0	17370	-3128.63	5.5159	265.2	-2390.09	4.1138	84.3
1500	340656	0.01	1851266	-3649.78	4.2553	361.2	-3616.72	4.1787	32.0
		0.10	475146	-3815.09	4.3668	303.4	-3743.19	4.2725	42.3
		1.00	42902	-3895.09	4.6025	1341.0	-3874.68	4.4823	155.8
		10.0	7430	-4759.67	5.6739	881.2	-4253.34	4.6876	468.6
2000	605990	0.01	3301648	-5149.12	4.5763	801.3	-5104.27	4.5006	61.7
		0.10	893410	-5371.26	4.6851	620.0	-5269.06	4.5969	82.4
		1.00	153984	-5456.54	4.9033	1426.0	-5406.89	4.7614	175.9
		10.0	33456	-6560.54	5.9405	1552.3	-5512.48	4.7982	565.5

Table 5: Computational results for $\delta = 100\%$

Problem		$\rho_{\bar{\Omega}}$	r	PPA			PD		
p	$ \Omega $			Likelihood	Loss	Time	Likelihood	Loss	Time
500	0	0.01	238232	-930.00	3.5345	36.0	-918.52	3.4838	1.3
		0.10	57064	-1000.78	3.6826	43.6	-973.06	3.5313	4.0
		1.00	15474	-1053.04	4.0675	76.1	-1006.95	3.6500	10.6
		10.0	7448	-1511.88	5.4613	51.4	-1023.82	3.7319	18.1
1000	0	0.01	963400	-2188.06	4.1983	156.3	-2161.58	4.1383	5.3
		0.10	231424	-2335.09	4.3387	122.4	-2277.90	4.2045	16.8
		1.00	47528	-2401.69	4.6304	329.6	-2349.74	4.3449	42.6
		10.0	18156	-3127.94	5.8521	244.1	-2388.22	4.4466	79.0
1500	0	0.01	2181060	-3585.21	4.5878	364.1	-3545.43	4.5260	12.3
		0.10	551150	-3806.07	4.7234	288.2	-3717.25	4.6059	41.3
		1.00	102512	-3883.94	4.9709	912.8	-3826.26	4.7537	93.5
		10.0	31526	-4821.26	6.0886	848.7	-3898.50	4.8824	185.4
2000	0	0.01	3892592	-5075.44	4.8867	734.1	-5021.95	4.8222	23.8
		0.10	1027584	-5367.86	5.0183	590.6	-5246.45	4.9138	76.1
		1.00	122394	-5456.64	5.2330	1705.8	-5422.48	5.1168	197.8
		10.0	25298	-6531.08	6.2571	1803.4	-5636.74	5.3492	417.1

Table 6: Computational results for $\delta = 100\%$

Problem		$\rho_{\bar{\Omega}}$	r	QUIC			PD		
p	$ \Omega $			Likelihood	Loss	Time	Likelihood	Loss	Time
500	0	0.01	222292	-929.98	3.5345	76477.7	-918.79	3.4839	2.1
		0.10	40936	-1000.78	3.6826	9697.7	-983.39	3.5570	8.0
		1.00	14614	-1053.01	4.0675	1682.6	-1008.00	3.6562	113.1
		10.00	7330	-1511.88	5.4613	584.0	-1024.48	3.7337	190.1
1000	0	0.01	887936	-2176.69	4.1871	490839.9	-2150.02	4.1265	7.4
		0.10	140166	-2326.70	4.3293	44398.7	-2297.48	4.2447	25.3
		1.00	24810	-2394.04	4.6113	12000.9	-2362.15	4.3909	78.9
		10.00	8564	-3088.29	5.8051	13486.7	-2505.80	4.5468	169.3

Table 7: Computational results on two real data sets

Data	Genes p	Samples n	$\rho_{\bar{\Omega}}$	r	PPA			PD		
					Likelihood	Loss	Time	Likelihood	Loss	Time
Lymph	587	148	0.01	144294	790.12	23.24	101.5	1035.24	22.79	38.0
			0.05	67474	174.86	24.35	85.2	716.97	23.27	31.5
			0.10	38504	-47.03	24.73	66.7	389.65	23.85	26.1
			0.50	4440	-561.38	25.52	33.2	-260.32	24.91	24.8
			0.70	940	-642.05	25.63	26.9	-511.70	25.30	22.0
			0.90	146	-684.59	25.70	22.0	-598.05	25.51	14.9
Leukemia	1255	72	0.01	249216	3229.75	28.25	705.7	3555.38	28.12	177.1
			0.05	169144	1308.38	29.85	491.1	2996.95	28.45	189.2
			0.10	107180	505.02	30.53	501.4	2531.62	28.82	202.8
			0.50	37914	-931.59	31.65	345.9	797.23	30.16	256.6
			0.70	4764	-1367.22	31.84	125.7	-1012.48	31.48	271.6
			0.90	24	-1465.70	31.90	110.6	-1301.99	31.68	187.8

Table 8: Computational results on two real data sets

Data	Genes p	Samples n	$\rho_{\bar{\Omega}}$	r	QUIC			PD		
					Likelihood	Loss	Time	Likelihood	Loss	Time
Lymph	587	148	0.01	141472	790.09	23.24	3110.2	1030.93	22.79	40.3
			0.05	65372	174.85	24.35	125.3	704.10	23.30	32.5
			0.10	37026	-47.04	24.73	25.9	379.93	23.86	28.5
			0.50	4030	-561.38	25.52	0.6	-283.53	24.91	26.1
			0.70	858	-642.05	25.63	0.4	-539.43	25.38	22.3
			0.90	130	-684.59	25.70	0.2	-608.54	25.59	17.1
Leukemia	1255	72	0.01	238770	3229.62	28.26	18964.5	3445.32	28.33	187.3
			0.05	161204	1308.37	29.86	1348.4	2948.86	28.48	192.8
			0.10	101070	505.01	30.54	453.9	2456.69	28.87	219.5
			0.50	33526	-931.59	31.66	22.0	691.12	30.24	266.0
			0.70	3830	-1367.22	31.85	1.8	-1089.35	31.49	286.4
			0.90	24	-1465.70	31.90	0.9	-1301.99	31.68	197.4

implies that the solution of the PD method when applied to (47) is at least as sparse as that of PPA or QUIC. As the true covariance matrix Σ^t is unknown for these data sets, we now modify the normalized entropy loss defined above by replacing Σ^t by Σ . The results of PPA, QUIC and our PD method on these two data sets are presented in Tables 7 and 8. In detail, the name and dimension of each data set are given in the first three columns. The values of $\rho_{\bar{\Omega}}$ and r are listed in the fourth and fifth columns. The log-likelihood, the normalized entropy loss and the CPU time (in seconds) of PPA or QUIC and the PD method are given in the last six columns, respectively. We can observe that our PD method is generally faster than PPA. Moreover, our PD method outperforms PPA in terms of log-likelihood and normalized entropy loss. Also, we observe that PD substantially outperforms QUIC in terms of log-likelihood and normalized entropy loss. In addition, QUIC is faster than PD when the solution is sufficiently sparse.

Our third experiment is similar to the one conducted in [14, 37]. We intend to compare sparse recoverability of our PD method with PPA. To this aim, we specialize $p = 30$ and $(\Sigma^t)^{-1} \in S_{++}^p$ to be the matrix with diagonal entries around one and a few randomly chosen, nonzero off-diagonal entries equal to +1 or -1. And the sample covariance matrix Σ is then similarly generated as above. In addition, we set $\Omega = \{(i, j) : (\Sigma^t)_{ij}^{-1} = 0, |i - j| \geq 15\}$ and $\rho_{ij} = \rho_{\bar{\Omega}}$ for all $(i, j) \in \bar{\Omega}$, where $\rho_{\bar{\Omega}}$ is the smallest number such that the approximate solution obtained by PPA shares the same number of nonzero off-diagonal entries as $(\Sigma^t)^{-1}$. For problem (47), we choose $r = \sum_{(i,j) \in \bar{\Omega}} \|(\Sigma^t)_{ij}^{-1}\|_0$ (i.e., the

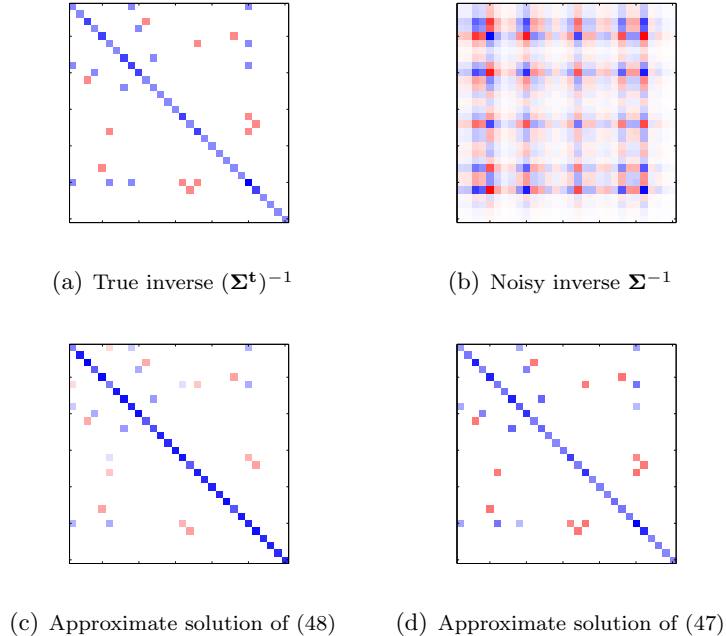


Figure 2: Sparse recovery.

Table 9: Numerical results for sparse recovery

	nnz	Likelihood	Loss
PPA	24	-35.45	0.178
PD	24	-29.56	0.008

number of nonzero off-diagonal entries of $(\Sigma^t)^{-1}$. PPA and the PD method are then applied to solve (48) and (47) with the aforementioned ρ_{ij} and r , respectively. In Figure 2, we plot the sparsity patterns of the original inverse covariance matrix $(\Sigma^t)^{-1}$, the noisy inverse sample covariance matrix Σ^{-1} , and the approximate solutions to (48) and (47) obtained by PPA and our PD method, respectively. We first observe that the sparsity of both solutions is the same as $(\Sigma^t)^{-1}$. Moreover, the solution of our PD method completely recovers the sparsity patterns of $(\Sigma^t)^{-1}$, but the solution of PPA misrecovers a few patterns. In addition, we present the log-likelihood and the normalized entropy loss of these solutions in Table 9. One can see that the solution of our PD method achieves much larger log-likelihood and smaller normalized entropy loss.

As a summary, the above experiments show that the quality of the approximate solution of (47) obtained by our PD method is generally better than that of (48) found by PPA and QUIC when the same cardinality is considered.

5.3 Compressed sensing

In this subsection, we apply the PD methods proposed in Section 4 to solve the compressed sensing (CS) problem, which has important applications in signal processing (see, for example, [13, 50, 32, 48, 10, 39, 53]).

When the observation is noise free, the CS problem can be formulated as

$$\min_{x \in \mathfrak{R}^p} \{\|x\|_0 : Ax = b\}, \quad (51)$$

where $A \in \mathfrak{R}^{n \times p}$ is a data matrix and $b \in \mathfrak{R}^n$ is an observation vector. One popular approach for finding an approximate solution to (51) is to solve the following l_1 regularization problem:

$$\min_{x \in \mathfrak{R}^p} \{\|x\|_1 : Ax = b\}, \quad (52)$$

(see, for example, [55, 10]). Our aim below is to apply the PD method studied in Subsection 4.2 to solve problem (51) directly.

Clearly, problem (51) is in the form of (2) and thus the PD method proposed in Subsection 4.2 can be suitably applied to solve (51). Also, one can observe that the main computation effort of the PD method when applied to (51) lies in solving the subproblem arising in step 1a), which is in the form of

$$\min_x \{\|x - c\|^2 : Ax = b\} \quad (53)$$

for some $c \in \mathfrak{R}^p$. It is well known that problem (53) has a closed-form solution given by

$$x^* = c - A^T(AA^T)^{-1}(Ac - b).$$

We now address the initialization and the termination criteria for the PD method. In particular, we choose y_0^0 to be a feasible point of (51) obtained by executing the Matlab command $A \setminus b$. Also, we set the initial penalty parameter $\varrho_0 = 0.1$ and the parameter $\sigma = 10$. In addition, we use (14) and

$$\frac{\|x^k - y^k\|_\infty}{\max\{|p_{\varrho^k}(x^k, y^k)|, 1\}} \leq \epsilon_O$$

as the inner and outer termination criteria, and set the associated accuracy parameters $\epsilon_I = 10^{-5}$ and $\epsilon_O = 10^{-6}$, respectively.

We next conduct experiment to test the performance of our PD method for finding a sparse approximate solution to problem (51) on the data sets from Sparco [56].³ We also compare the quality of such sparse approximate solution with that of the one found by a first-order solver SPGL1 [55] applied to (52). For the latter method, we use the default value for all parameters. To evaluate the quality of these sparse approximate solutions, we adopt a similar criterion as described in [46, 7]. Indeed, suppose that B be a basis matrix for a given signal f for which we wish to find a sparse recovery. Given a sparse vector x , the corresponding sparse approximate signal is $f_x = Bx$. The associated mean squared error is defined as

$$\text{MSE} := \|f_x - f\|/p,$$

We only report the computational results for 10 data sets in Table 10 since the performance difference between PD and SPGL1 on the other data sets is similar to that on these data sets. In detail, the data name and the size of data are given in the first three columns. The MSE, the solution cardinality and the CPU time for both methods are reported in columns four to nine, respectively. It can be observed that SPGL1 is generally faster than PD, but PD generally provides much more sparse solution while with lower MSE. Thus, the resulting signal by PD is less noisy. For example, we plot in Figure 3 the results for data **blknheavi** whose actual sparse representation is pre-known. It can be seen that the signal recovered by SPGL1 has more noise than the one by PD.

Table 10: Computational results on data from Sparco

Data	Size		PD			SPGL1		
	p	n	MSE	nnz	Time	MSE	nnz	Time
blkheavi	128	128	$1.28e-07$	12	0.1	$4.53e-03$	128	1.5
jitter	1000	200	$4.50e-08$	3	0.1	$1.38e-07$	28	0.1
gausspike	1024	256	$2.63e-07$	32	1.5	$1.09e-08$	143	0.2
sgnspike	2560	600	$3.84e-08$	20	0.2	$8.08e-08$	101	0.1
blkheavi	1024	1024	$3.19e-09$	12	0.4	$4.22e-03$	1024	14.3
cosspike	2048	1024	$8.97e-07$	121	0.3	$8.40e-08$	413	0.2
angiogram	10000	10000	$2.74e-06$	575	2.4	$5.77e-07$	1094	0.5
blurspike	16384	16384	$2.97e-03$	7906	10.1	$3.15e-03$	16384	20.6
srcsep1	57344	29166	$3.41e-05$	9736	743.6	$1.39e-08$	33887	102.5
srcsep2	86016	29166	$6.93e-04$	12485	1005.6	$2.80e-04$	52539	136.2

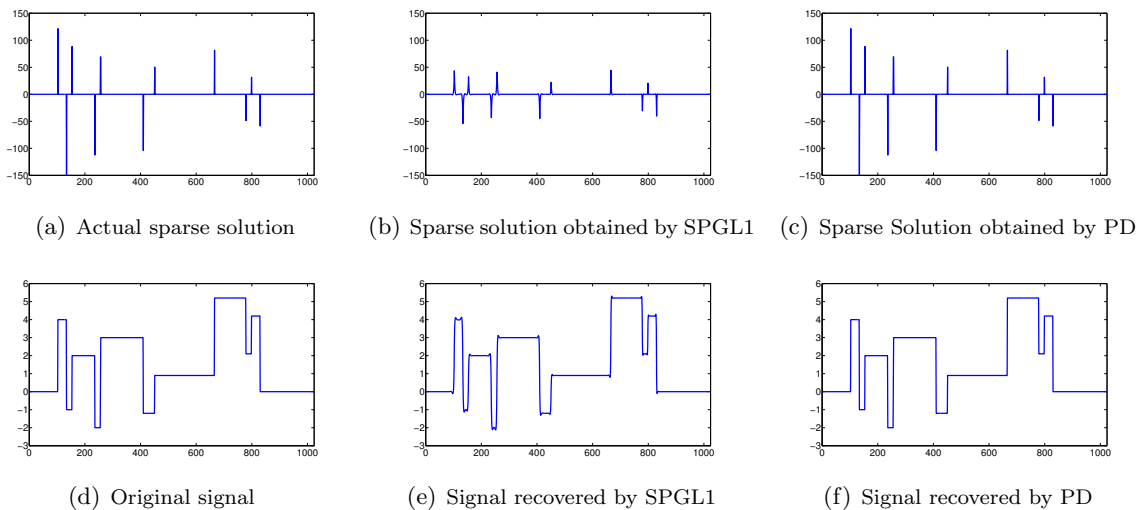


Figure 3: Sparse recovery.

In the remainder of this subsection we consider the CS problem with noisy observation. In this case, the CS problem can be formulated as

$$\min_{x \in \mathbb{R}^p} \left\{ \frac{1}{2} \|Ax - b\|^2 : \|x\|_0 \leq r \right\}, \quad (54)$$

where $A \in \mathbb{R}^{n \times p}$ is a data matrix, $b \in \mathbb{R}^n$ is an observation vector, and $r \in [1, p]$ is some integer for controlling the sparsity of the solution. One popular approach for finding an approximate solution to (54) is to solve the following l_1 regularization problem:

$$\min_{x \in \mathbb{R}^p} \frac{1}{2} \|Ax - b\|^2 + \lambda \|x\|_1, \quad (55)$$

where $\lambda \geq 0$ is a regularization parameter (see, for example, [20, 25, 28]). Our goal below is to apply the PD method studied in Subsection 4.1 to solve (54) directly.

Clearly, problem (54) is in the form of (1) and thus the PD method proposed in Subsection 4 can be suitably applied to solve (54). The main computation effort of the PD method when applied to (54)

³Roughly speaking, a sparse approximate solution x to (51) means that x is sparse and $Ax \approx b$.

lies in solving the subproblem arising in step 1a), which is an unconstrained quadratic programming problem that can be solved by the conjugate gradient method. We now address the initialization and the termination criteria for the PD method. In particular, we randomly choose an initial point $y_0^0 \in \mathbb{R}^p$ such that $\|y_0^0\|_0 \leq r$. Also, we set the initial penalty parameter $\varrho_0 = 1$ and the parameter $\sigma = \sqrt{10}$. In addition, we use

$$\frac{|q_{\varrho_k}(x_{l+1}^k, y_{l+1}^k) - q_{\varrho_k}(x_l^k, y_l^k)|}{\max\{|q_{\varrho_k}(x_l^k, y_l^k)|, 1\}} \leq \epsilon_I$$

and

$$\frac{\|x^k - y^k\|_\infty}{\max\{|q_{\varrho_k}(x^k, y^k)|, 1\}} \leq \epsilon_O$$

as the inner and outer termination criteria for the PD method, and set their associated accuracy parameters $\epsilon_I = 10^{-2}$ and $\epsilon_O = 10^{-3}$.

We next conduct numerical experiments to test the performance of our PD method for solving problem (54) on random data. We also compare the quality of the approximate solutions of (54) obtained by our PD method and the iterative hard-thresholding algorithm (IHT) [5, 6] with that of (55) found by a first-order solver GPSR [20]. For IHT, we set $stopTol = 10^{-6}$ and use the default values for all other parameters. And for GPSR, all the parameters are set as their default values.

We first randomly generate a data matrix $A \in \mathbb{R}^{n \times p}$ and an observation vector $b \in \mathbb{R}^n$ according to a standard Gaussian distribution. Then we apply GPSR to problem (55) with a set of p distinct λ 's so that the cardinality of the resulting approximate solution gradually increases from 1 to p . Accordingly, we apply our PD method and IHT to problem (54) with $r = 1, \dots, p$. It shall be mentioned that a warm-start strategy is applied to all three methods. That is, an approximate solution of problem (54) (resp., (52)) for current r (resp., λ) is used as the initial point for the PD method and IHT (resp., GPSR) when applied to the problem for next r (resp., λ). The average computational results of both methods over 100 random instances with $(n, p) = (1024, 4096)$ are plotted in Figure 4. In detail, we plot the average residual $\|Ax - b\|$ against the cardinality in the left graph and the average accumulated CPU time⁴ (in seconds) against the cardinality in the right graph. We observe that the residuals of the approximate solutions of (55) obtained by our PD method and IHT are almost equal and substantially smaller than that of (54) found by GPSR when the same cardinality is considered. In addition, we can see that GPSR is faster than the other two methods.

We also conduct a similar experiment as above except that A is randomly generated with orthonormal rows. The results are plotted in Figure 5. We observe that the PD method and IHT are generally slower than GPSR, but they have better solution quality than GPSR in terms of residuals.

6 Concluding remarks

In this paper we propose penalty decomposition methods for general l_0 minimization problems in which each subproblem is solved by a block coordinate descend method. Under some suitable assumptions, we establish that any accumulation point of the sequence generated by the PD methods satisfies the first-order optimality conditions of the problems. Furthermore, for the problems in which the l_0 part is the only nonconvex part, we show that such an accumulation point is a local minimizer of the problems. The computational results on compressed sensing, sparse logistic regression and sparse inverse covariance selection problems demonstrate that when solutions of same cardinality are sought, our approach applied

⁴For a cardinality r , the corresponding accumulated CPU time is the total CPU time used to compute approximate solutions of problem (54) or (52) with cardinality from 1 to r .

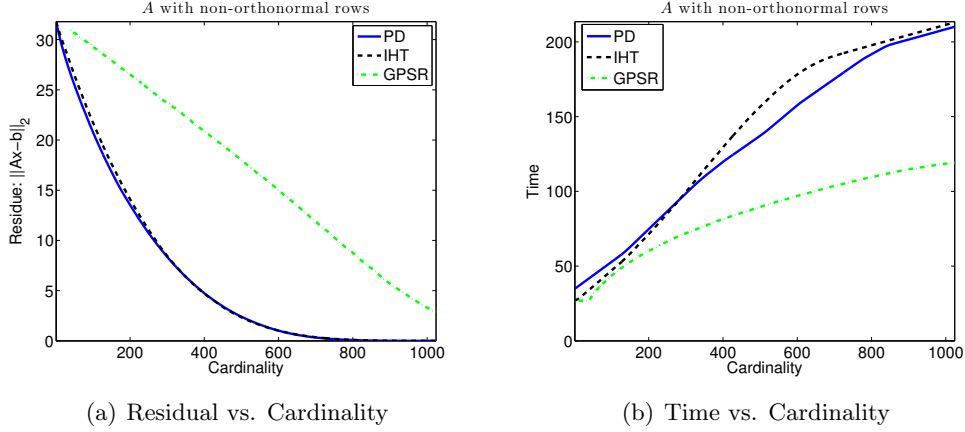


Figure 4: Trade-off curves.

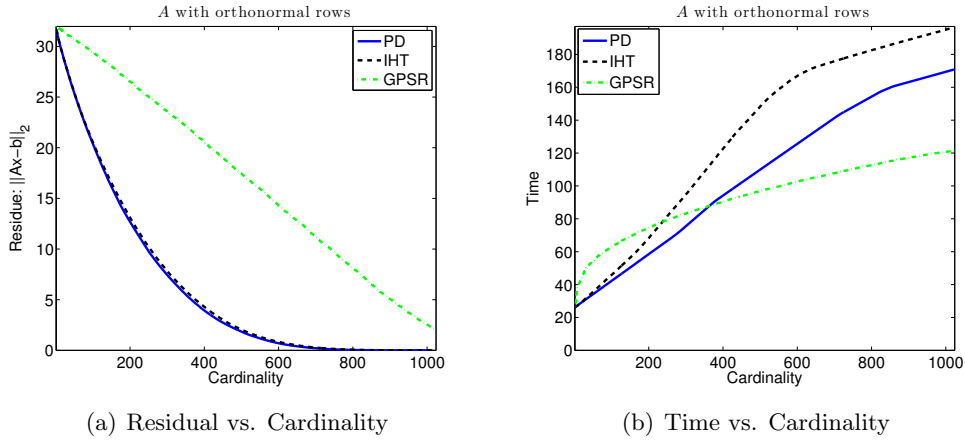


Figure 5: Trade-off curves.

to the l_0 -based models generally has better solution quality and/or speed than the existing approaches that are applied to the corresponding l_1 -based models.

We shall remark that the augmented Lagrangian decomposition methods can be developed for solving l_0 minimization problems (1) and (2) simply by replacing the quadratic penalty functions in the PD methods by augmented Lagrangian functions. Nevertheless, as observed in our experiments, their practical performance is generally worse than the PD methods.

Acknowledgement

The authors would like to thank the Associate Editor and two anonymous reviewers for numerous insightful comments and suggestions, which have greatly improved the paper.

Appendix

In this appendix we provide an example to demonstrate that the l_p -norm relaxation approaches for $p \in (0, 1]$ may fail to recover the sparse solution.

Let $p \in (0, 1]$ be arbitrarily chosen. Given any $b^1, b^2 \in \mathfrak{R}^n$, let $b = b^1 + b^2$, $\alpha = \|(b^1; b^2)\|_p$ and $A = [b^1, b^2, \alpha I_n, \alpha I_n]$, where I_n denotes the $n \times n$ identity matrix and $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$ for all $x \in \mathfrak{R}^n$. Consider the linear system $Ax = b$. It is easy to observe that this system has the sparse solution $x^s = (1, 1, 0, \dots, 0)^T$. However, x^s cannot be recovered by solving the l_p -“norm” regularization problem:

$$f^* = \min_x \left\{ f(x) := \frac{1}{2} \|Ax - b\|^2 + \nu \|x\|_p \right\}$$

for any $\nu > 0$. Indeed, let $\bar{x} = (0, 0, b^1/\alpha, b^2/\alpha)^T$. Then, we have $f(x^s) = 2^{1/p}\nu$ and $f(\bar{x}) = \nu$, which implies that $f(x^s) > f(\bar{x}) \geq f^*$. Thus, x^s cannot be an optimal solution of the above problem for any $\nu > 0$. Moreover, the relative error between $f(x^s)$ and f^* is fairly large since

$$(f(x^s) - f^*)/f^* \geq (f(x^s) - f(\bar{x}))/f(\bar{x}) = 2^{1/p} - 1 \geq 1.$$

Therefore, the true sparse solution x^s may not even be a “good” approximate solution to the l_p -“norm” regularization problem.

References

- [1] O. Banerjee, L. E. Ghaoui, and A. D’Aspremont. Model selection through sparse maximum likelihood estimation. *J. Mach. Learn. Res.*, 9:485-516, 2008.
- [2] J. A. Bilmes. Factored sparse inverse covariance matrices. *International Conference on Acoustics, Speech and Signal processing*, Washington, D.C., 1009-1012, 2000.
- [3] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
- [4] E. G. Birgin, J. M. Martínez, and M. Raydan. Nonmonotone spectral projected gradient methods on convex sets. *SIAM J. Optimiz.*, 4:1196–1211, 2000.
- [5] T. Blumensath and M. E. Davies. Iterative thresholding for sparse approximations. *J. FOURIER ANAL. APPL.*, 14:629–654, 2008.
- [6] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comput. Harmon. Anal.*, 27(3):265–274, 2009.
- [7] E. J. Candés and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 2009.
- [8] E. J. Candés, J. Romberg and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE T. Inform. Theory*, 52:489-509, 2006.
- [9] R. Chartrand. Exact reconstruction of sparse signals via nonconvex minimization. *IEEE Signal Proc. Lett.*, 14:707-710, 2007.
- [10] S. Chen, D. Donoho and M. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20:33-61, 1998.

- [11] X. Chen, F. Xu and Y. Ye. Lower bound theory of nonzero entries in solutions of l_2 - l_p Minimization. *SIAM J. Sci. Comput.*, 32:2832-2852, 2010.
- [12] X. Chen and W. Zhou. Convergence of reweighted l_1 minimization algorithms and unique solution of truncated l_p minimization. Technical report, 2010.
- [13] J. Claerbout and F. Muir. Robust modelling of erratic data. *Geophysics*, 38:826-844, 1973.
- [14] A. D'Aspremont, O. Banerjee and L. E. Ghaoui. First-order methods for sparse covariance selection. *SIAM J. Matrix Anal. A.*, 30(1):56-66, 2008.
- [15] J. Dahl, L. Vandenberghe and V. Roychowdhury. Covariance selection for nonchordal graphs via chordal embedding. *Optim. Method. Softw.*, 23(4):501-520, 2008.
- [16] A. Dempster. Covariance selection. *Biometrics*, 28:157-175, 1978.
- [17] A. Dobra. Dependency networks for genome-wide data. *Biostatistics*, 8(1):1-28, 2007.
- [18] 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.*, 90:196-212, 2004.
- [19] B. Efron, T. Hastie, I. Johnstone and R. Tibshirani. Least angle regression. *Ann. Stat.*, 32(2):407-499, 2004.
- [20] M. A. T. Figueiredo, R. D. Nowak and S. J. Wright. Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems. *IEEE J. Sel. Top. Signal.: Special Issue on Convex Optimization Methods for Signal Processing*, 1(4):586-598, 2007.
- [21] J. Friedman, T. Hastie and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostat.*, 9(3):432-441, 2008.
- [22] A. D. Gerson, L. C. Parra and P. Sajda. Cortical origins of response time variability during rapid discrimination of visual objects. *Neuroimage*, 28(2):342-353, 2005.
- [23] G. Golub and C. Van Loan. *Matrix Computations*, volume 13 of *Studies in Applied Mathematics*. John Hopkins University Press, third edition, 1996.
- [24] T. R. Golub, D. K. Slonim, P. Tamayo, C. Huard, M. Gaasenbeek, J. P. Mesirov, H. Coller, M. Loh, J. R. Downing, M. A. Caligiuri, C. D. Bloomfield, and E. S. Lander. Molecular classification of cancer: class discovery and class prediction by expression monitoring. *Science*, 286:531-537, 1999.
- [25] E. T. Hale, W. Yin and Y. Zhang. Fixed-point continuation applied to compressed sensing: Implementation and numerical experiments. *J. Comput. Math*, 28(2):170-194, 2010.
- [26] K. K. Herrity, A. C. Gilbert and J. A. Tropp. Sparse approximation via iterative thresholding. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2006.
- [27] C. J. Hsien, M. A. Sustik, I. S. Dhillon and P. Ravikumar. Sparse Inverse Covariance Matrix Estimation Using Quadratic Approximation. *Advances in Neural Information Processing Systems*, 2011.
- [28] S. J. Kim, K. Koh, M. Lustig, S. Boyd and D. Gorinevsky. An interior-point method for large-scale l_1 -regularized least squares. *IEEE J. Sel. Top. Signal.*, 1(4):606-617, December 2007.

- [29] K. Koh, S. J. Kim and S. Boyd. An interior-point method for large-scale l_1 -regularized logistic regression. *J. Mach. Learn. Res.*, 8:1519-1555, 2007.
- [30] Linear Algebra PACKage. Available at <http://www.netlib.org/lapack/index.html>.
- [31] S. Lee, H. Lee, P. Abbeel and A. Ng. Efficient l_1 -regularized logistic regression. In *21th National Conference on Artificial Intelligence (AAAI)*, 2006.
- [32] S. Levy and P. Fullagar. Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution. *Geophysics*, 46:1235-1243, 1981.
- [33] J. G. Liao and K. V. Chin. Logistic regression for disease classification using microarray data: model selection in a large p and small n case. *Bioinformatics*, 23(15):1945-1951, 2007.
- [34] J. Liu, S. Ji and J. Ye. *SLEP: Sparse learning with efficient projections*. Arizona State University, 2009. Available at <http://www.public.asu.edu/~jye02/Software/SLEP>.
- [35] L. Li and K. C. Toh. An inexact interior point method for l_1 -regularized sparse covariance selection. *Math. Program. Comput.*, 2:291-315, 2010.
- [36] Z. Lu. Smooth optimization approach for sparse covariance selection. *SIAM J. Optimiz.*, 19(4):1807-1827, 2009.
- [37] Z. Lu. Adaptive first-order methods for general sparse inverse covariance selection. *SIAM J. Matrix Anal. A.*, 31(4):2000-2016, 2010.
- [38] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE T. Image Process.*, 41(12):3397-3415, 1993.
- [39] A. Miller. Subset selection in regression. Chapman and Hall, London, 2002.
- [40] D. Newman, S. Hettich, C. Blake and C. Merz. UCI repository of machine learning databases, 1998. Available at www.ics.uci.edu/~mlearn/MLRepository.html.
- [41] A. Y. Ng. Feature selection, l_1 vs. l_2 regularization, and rotational invariance. In *Proceedings of the Twenty-First International Conference on Machine learning (ICML)*, 72-85, 2004.
- [42] M. Y. Park and T. Hastie. *Regularization path algorithms for detecting gene interactions*. Department of Statistics, Stanford University, 2006.
- [43] L. C. Parra, C. D. Spence, A. D. Gerson and P. Sajda. Recipes for the linear analysis of EEG. *Neuroimage*, 28(2):326-341, 2005.
- [44] M. G. Philiastides and P. Sajda. Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cereb. Cortex*, 16(4):509-518, 2006.
- [45] J. Pittman, E. Huang, H. Dressman, C. F. Horng, S. H. Cheng, M. H. Tsou, C. M. Chen, A. Bild, E. S. Iversen, A. T. Huang, J. R. Nevins and M. West. Integrated modeling of clinical and gene expression information for personalized prediction of disease outcomes. *P. Natl. Acad. Sci. USA*, 101(22):8431-8436, 2004.
- [46] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 2007.

- [47] A. Ruszczyński. *Nonlinear Optimization*. Princeton University Press, 2006.
- [48] F. Santosa and W. Symes. Linear inversion of band-limited reflection histograms. *SIAM J. Sci. Stat. Comp.*, 7:1307-1330, 1986.
- [49] J. Shi, W. Yin, S. Osher and P. Sajda. A fast hybrid algorithm for large-scale l_1 -regularized logistic regression. *J. Mach. Learn. Res.*, 11:713-741, 2010.
- [50] H. Taylor, S. Bank and J. McCoy. Deconvolution with the l_1 -norm. *Geophysics*, 44:39-52, 1979.
- [51] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Stat. Soc. B*, 58(1):267-288, 1996.
- [52] J. A. Tropp. Greed is good: algorithmic results for sparse approximation. *IEEE T. Inform. Theory*, 50(10):2231-2242, 2004.
- [53] J. Tropp. Just relax: Convex programming methods for identifying sparse signals. *IEEE T. Inform. Theory*, 51:1030-1051, 2006.
- [54] Y. Tsuruoka, J. McNaught, J. Tsujii and S. Ananiadou. Learning string similarity measures for gene/protein name dictionary look-up using logistic regression. *Bioinformatics*, 23(20):2768-2774, 2007.
- [55] E. Van Den Berg and M. P. Friedlander. Probing the Pareto frontier for basis pursuit solutions. *SIAM J. Sci. Comp.*, 31(2):890-912, 2008.
- [56] E. Van Den Berg, M. P. Friedlander, G. Hennenfent, F. Herrmann, R. Saab and O. Yilmaz. Algorithm 890: Sparco: a testing framework for sparse reconstruction. *ACM T. Math. Software*, 35(4):1-16, 2009.
- [57] C. Wang, D. Sun and K. C. Toh. Solving log-determinant optimization problems by a Newton-CG proximal point algorithm. *SIAM J. Optimiz.*, 20(6):2994-3013, 2010.
- [58] K. Y. Yeung, R. E. Bumgarner and A. E. Raftery. Bayesian model averaging: development of an improved multi-class, gene selection and classification tool for microarray data. *Bioinformatics*, 21(10):2394-2402, 2005.
- [59] Y. Zhang, B. Dong and Z. Lu. l_0 minimization for wavelet frame based image restoration. *Math. Comput.*, 82: 995-1015, 2013.