Finding Dantzig selectors with a proximity operator based fixed-point algorithm

Ashley Prater

Air Force Research Laboratory, Information Directorate, Rome, NY 13441, USA

Lixin Shen*

Department of Mathematics, Syracuse University, Syracuse, NY 13244, USA.

Bruce W. Suter

Air Force Research Laboratory, Information Directorate, Rome, NY 13441, USA

Abstract

A simple iterative method for finding the Dantzig selector, designed for linear regression problems, is introduced. The method consists of two stages. The first stage approximates the Dantzig selector through a fixed-point formulation of solutions to the Dantzig selector problem; the second stage constructs a new estimator by regressing data onto the support of the approximated Dantzig selector. The proposed method is compared to an alternating direction method. The results of numerical simulations using both the proposed method and the alternating direction method on synthetic and real-world data sets are presented. The numerical simulations demonstrate that the two methods produce results of similar quality; however the proposed method tends to be significantly faster.

Keywords: Dantzig selector, proximity operator, fixed-point algorithm, alternating direction method.

1. Introduction

This paper considers the problem of estimating a vector of parameter $\beta \in \mathbb{R}^p$ from the linear problem

$$y = X\beta + z, \quad (1)$$

where $y \in \mathbb{R}^n$ is a vector of observations, X an $n \times p$ predictor matrix, and z a vector of independent normal random variables. The goal is to find a relevant parametric vector $\beta^* \in \mathbb{R}^p$ among many potential candidates and to obtain high prediction accuracy.

The ℓ_1 penalized least squares estimator for problem (1) has been the focus of a great deal of attention for variable selection and estimation in high-dimensional linear regression when the number of variables is much larger than the sample size [11, 21, 24, 27, 28, 31].

*Corresponding author.

Email addresses: ashley.prater.3@us.af.mil (Ashley Prater), lshen03@syr.edu (Lixin Shen), bruce.suter@us.af.mil (Bruce W. Suter)

Recently the Dantzig selector was proposed in [7] for problem (1). The Dantzig selector $\hat{\beta} \in \mathbb{R}^p$ is a solution to the optimization problem

$$\hat{\beta} \in \operatorname{argmin}\{\|\beta\|_1 : \|D^{-1}X^\top(X\beta - y)\|_\infty \leq \delta\}, \quad (2)$$

with a fixed parameter $\delta > 0$ and a diagonal matrix D where the diagonal entries are equal to the ℓ_2 norm of the columns of X . Here, we write $\|x\|_q$ for the ℓ_q norm of $x \in \mathbb{R}^p$, $1 \leq q \leq \infty$. Optimal ℓ_2 rate properties for $\|\hat{\beta} - \beta^*\|_2$ were established under a sparsity scenario and impressive empirical performance on real world problems involving large values of p was shown in [7]. Since then the Dantzig selector has received a considerable amount of attention. Discussions on the Dantzig selector can be found in [4, 6, 8, 12, 14, 22, 25]. In [16], an algorithm was proposed for fitting the entire coefficient path of the Dantzig selector with a similar computational cost to the least angle algorithm that is used to compute the ℓ_1 minimization via the LASSO technique. The Dantzig selector is a convex, but not strictly convex, optimization problem. Unique solutions are in general not guaranteed. Conditions ensuring the uniqueness of the Dantzig selector were presented in [10]. In [18] a new class of Dantzig selectors for linear regression problems for right-censored outcomes was proposed.

The importance of the Dantzig selector in linear regressions has been demonstrated in the aforementioned work. Efficient methods for solving problem (2), which however were not emphasized in the current literature, are highly needed. In [7], the problem is cast as a linear program which is solved by using a primal-dual interior point algorithm [5]. As it is well known, interior point methods are not efficient for large-scale problems. In [3], the problem is cast as linear cone programming problem for which a smooth approximation to its dual problem is solved by an optimal first-order method [2, 23]. Recently, an alternating direction method (ADM) for finding the Dantzig selector was studied in [19]. Numerical experiments showed that this method usually outperforms the method in [3] in terms of CPU time while producing solutions of comparable quality. The problem was rewritten in [19] in a form to which ADM can be easily applied. ADM itself is an iterative algorithm. In each iterate, two subproblems are needed to be solved successively. One of the subproblems has a closed form solution, while the other does not and is approximated by a nonmonotone gradient method proposed in [20]. To alleviate the difficulty caused by the subproblem without a closed form solution, a linearized ADM was proposed for the Dantzig selector and was shown to be efficient for solving both synthetic and real world data sets in [30].

In this paper, the Dantzig selectors for problem (2) are found by an algorithm based upon proximity operators. We first rewrite the problem as an unconstrained structural optimization problem via an indicator function. The resulting problem is then solved by a primal-dual algorithm. In comparison with the one given in [19], our proposed algorithm is easy to implement. Ours achieves comparable quality results while consuming much less CPU time.

The outline of the paper is organized as follows. In Section 2 we present our fixed-point theory based proximity operator algorithm for solving problem (2). In Section 3, we present numerical experiments comparing the accuracy and efficiency of the proposed algorithm with ADM proposed in [19]. The first set of experiments uses simulated sparse signals and the second set uses samples of biomarker data to predict the diagnosis of leukemia patients. Section 4 concludes the paper.

The following notation will be used in the rest of the paper. For any vector $u \in \mathbb{R}^d$, let u_i and $u(i)$ both denote the i -th component of u . Also for any vector $u \in \mathbb{R}^d$, $|u|$ is the component-wise absolute values of u , that is the i -th component of $|u|$ is $|u_i|$, while $\text{sign}(u)$ is the vector whose i -th component is 1 if $u_i > 0$ and -1 otherwise. Given two vectors u and v in $\mathbb{R}^d$, $x \circ y$ denotes the Hadamard (component-wise) product of u and v , $\max\{u, v\}$ denotes the vector whose i -th entry is $\max\{u_i, v_i\}$, and $\min\{u, v\}$ denotes the vector whose i -th entry is $\min\{u_i, v_i\}$. Let $\mathbb{1}$ denote the vector of all ones whose dimension should be clear from the context.

The set of all natural numbers is denoted by $\mathbb{N}$. For the usual d -dimensional Euclidean space denoted by $\mathbb{R}^d$ we define $\langle x, y \rangle := \sum_{i=1}^d x_i y_i$, for $x, y \in \mathbb{R}^d$, the standard inner product in $\mathbb{R}^d$. We denote by $\|\cdot\|_1$, $\|\cdot\|_2$, and $\|\cdot\|_\infty$ the ℓ_1 norm, ℓ_2 norm, and the ℓ_∞ norm of a vector, respectively. The class of all lower semicontinuous convex functions $f : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ such that $\text{dom} f := \{x \in \mathbb{R}^d : f(x) < +\infty\} \neq \emptyset$ is denoted by $\Gamma_0(\mathbb{R}^d)$. For a closed convex set $\mathcal{C}$ of $\mathbb{R}^d$, its indicator function $\iota_{\mathcal{C}}$ is in $\Gamma_0(\mathbb{R}^d)$ and is defined as

$$\iota_{\mathcal{C}}(u) := \begin{cases} 0, & \text{if } u \in \mathcal{C}, \\ +\infty, & \text{otherwise.} \end{cases}$$

For a function $f \in \Gamma_0(\mathbb{R}^d)$, $\text{argmin}_{x \in C} f(x)$ is the set of points of the given argument in C for which f attains its minimum value, i.e., $\text{argmin}_{x \in C} f(x) = \{x \in C : f(y) \geq f(x) \text{ for all } y \in C\}$.

2. The Dantzig Selector with Proximity Algorithms

In this section, we develop a proximity algorithm for solving the optimization problem (2). We begin with reviewing two existing works on this problem, namely the alternating direction method (ADM) proposed in [19] and the linearized alternating direction method of multipliers (LADM) proposed in [30].

Problem (2) can be reformulated as follows:

$$\min_{\beta \in \mathbb{R}^p, \tau \in \{\tau : \|D^{-1}\tau\|_\infty \leq \delta\}} \{\|\beta\|_1 : X^\top(X\beta - y) = \tau\}, \quad (3)$$

where $\tau \in \mathbb{R}^p$ is an auxiliary variable. The augmented Lagrangian function for problem (3) is

$$L_c(\beta, \tau, \gamma) := \|\beta\|_1 + \langle \gamma, X^\top(X\beta - y) - \tau \rangle + \frac{c}{2} \|X^\top(X\beta - y) - \tau\|_2^2,$$

where $\gamma \in \mathbb{R}^p$ is the Lagrange multiplier and $c > 0$ is a penalty parameter.

The iterative scheme of ADM for problem (3) is

$$\begin{cases} \tau^{k+1} \leftarrow \text{argmin}_{\tau \in \{\tau : \|D^{-1}\tau\|_\infty \leq \delta\}} L_c(\beta^k, \tau, \gamma^k), \\ \beta^{k+1} \leftarrow \text{argmin}_{\beta \in \mathbb{R}^p} L_c(\beta, \tau^{k+1}, \gamma^k), \\ \gamma^{k+1} = \gamma^k + c(X^\top(X\beta^{k+1} - y) - \tau^{k+1}), \end{cases}$$

which, with some elementary manipulations, can equivalently be written as

$$\begin{cases} \tau^{k+1} \leftarrow \text{argmin}_{\tau \in \{\tau : \|D^{-1}\tau\|_\infty \leq \delta\}} \|\tau - (X^\top(X\beta^k - y) + \frac{\gamma^k}{c})\|_2^2, \\ \beta^{k+1} \leftarrow \text{argmin}_{\beta \in \mathbb{R}^p} \{\|\beta\|_1 + \frac{c}{2} \|X^\top(X\beta - y) - \tau^{k+1} + \frac{\gamma^k}{c}\|_2^2\}, \\ \gamma^{k+1} = \gamma^k + c(X^\top(X\beta^{k+1} - y) - \tau^{k+1}). \end{cases} \quad (4)$$

The τ -related subproblem in (4) has a closed form solution, but the β -related subproblem does not and is solved approximately by using the nonmonotone gradient method in [19].

The iterative scheme of LADM for problem (3) with $D = I$ is

$$\begin{cases} \beta^{k+1} \leftarrow \operatorname{argmin}_{\beta \in \mathbb{R}^p} \{ \|\beta\|_1 + c \langle v^k, \beta - \beta^k \rangle + \frac{\ell}{2} \|\beta - \beta^k\|_2^2 \}, \\ \tau^{k+1} \leftarrow \operatorname{argmin}_{\tau \in \{ \tau : \|\tau\|_\infty \leq \delta \}} \|\tau - (X^\top (X\beta^{k+1} - y) + \frac{\gamma^k}{c})\|_2^2, \\ \gamma^{k+1} = \gamma^k + c(X^\top (X\beta^{k+1} - y) - \tau^{k+1}), \end{cases} \quad (5)$$

where $\ell > 0$ is a proximal parameter and $v^k := X^\top X(X^\top (X\beta^k - y) - \tau^k + \frac{\gamma^k}{c})$. Note that the order of updating τ^{k+1} and β^{k+1} in ADM is reversed in LADM. For the β -subproblem in LADM, the last two terms in the objective function can be viewed as the linearization of the quadratic term $\frac{c}{2} \|X^\top (X\beta - y) - \tau^k + \frac{\gamma^k}{c}\|_2^2$ with respect to β at β^k after dropping a constant. Furthermore, the β -subproblem, after completing the square of these two terms and ignoring the resulting constant term, is the same as

$$\beta^{k+1} \leftarrow \operatorname{argmin}_{\beta \in \mathbb{R}^p} \{ \|\beta\|_1 + \frac{\ell}{2} \|\beta - (\beta^k - \frac{c}{\ell} v^k)\|_2^2 \},$$

which has a closed form solution. The τ -subproblem has a closed form solution as in ADM. Therefore, LADM can be easily and efficiently implemented. It was shown in [30] that for any $c > 0$ and $\ell > c\|X^\top X\|_2^2$, the sequence $\{(\beta^k, \tau^k, \gamma^k) : k \in \mathbb{N}\}$ with any initial iterate $(\beta^0, \tau^0, \gamma^0)$ converges. Furthermore, the limit of the sequence $\{(\beta^k, \tau^k) : k \in \mathbb{N}\}$ is a solution of the Dantzig selector problem (3).

In the following, we present our fixed-point theory based proximity operator algorithm for solving the optimization problem (2). For simplicity of exposition, with the matrices X and D , the vector y , and the constant δ appearing in problem (2), we set

$$A := D^{-1}X^\top X, \quad b := D^{-1}X^\top y, \quad \mathcal{C} := \{\beta \in \mathbb{R}^p : \|\beta - b\|_\infty \leq \delta\}. \quad (6)$$

Then the optimization problem (2) can be rewritten as

$$\hat{\beta} \in \operatorname{argmin} \{ \|\beta\|_1 + \iota_{\mathcal{C}}(A\beta) : \beta \in \mathbb{R}^p \}. \quad (7)$$

The objective function of this problem is convex and coercive thanks to the ℓ_1 -norm being coercive. Hence a solution to problem (7) exists and can be characterized in terms of proximity operator. To this end, we review the definition of proximity operator.

For a function $f \in \Gamma_0(\mathbb{R}^d)$, the proximity operator of f with parameter λ , denoted by $\operatorname{prox}_{\lambda f}$, is a mapping from $\mathbb{R}^d$ to itself, defined for a given point $x \in \mathbb{R}^d$ by

$$\operatorname{prox}_{\lambda f}(x) := \operatorname{argmin} \left\{ \frac{1}{2\lambda} \|u - x\|_2^2 + f(u) : u \in \mathbb{R}^d \right\}.$$

Now, we can present a characterization of solutions of problem (7) that is simply derived from Fermat's rule.

Theorem 2.1. *Let the $p \times p$ matrix A and the vector $b \in \mathbb{R}^p$ be given as in (6). If $\beta \in \mathbb{R}^p$ is a solution to problem (7), then for any $\alpha > 0$ and $\lambda > 0$ there exists a vector $\tau \in \mathbb{R}^p$ such that*

$$\beta = \text{prox}_{\frac{1}{\alpha}\|\cdot\|_1} \left(\beta - \frac{\lambda}{\alpha} A^\top \tau \right), \quad (8)$$

$$\tau = (I - \text{prox}_{\iota_c})(A\beta + \tau). \quad (9)$$

Conversely, if there exist $\alpha > 0$ and $\lambda > 0$ such that $\beta, \tau \in \mathbb{R}^p$ satisfy equations (8) and (9), then β is a solution of problem (7).

Proof: The proof of the result follows straightforwardly a general result in [17, Proposition 1]. For completeness, we present its proof here. First, we assume that β is a solution to problem (7). By Fermat's rule and the chain rule of subdifferentiation, $0 \in \partial\|\cdot\|_1(\beta) + A^\top \partial\iota_c(A\beta)$. Then for any $\alpha > 0$ and $\lambda > 0$ there exists $\tau \in \frac{1}{\lambda} \partial\iota_c(A\beta)$ such that $-\frac{\lambda}{\alpha} A^\top \tau \in \partial(\frac{1}{\alpha}\|\cdot\|_1)(\beta)$, that is, in terms of proximity operator, equation (8). Since the set $\partial\iota_c(A\beta)$ is a cone, then $\tau \in \frac{1}{\lambda} \partial\iota_c(A\beta)$ implies $\tau \in \partial\iota_c(A\beta)$ which is essentially equivalent to equation (9).

Conversely, if equations (8) and (9) are satisfied, we then have $-\frac{\lambda}{\alpha} A^\top \tau \in \partial(\frac{1}{\alpha}\|\cdot\|_1)(\beta)$ and $\tau \in \partial\iota_c(A\beta)$ accordingly. Using the fact that the set $\partial\iota_c(A\beta)$ is a cone again, the second inclusion $\tau \in \partial\iota_c(A\beta)$ implies $\frac{\lambda}{\alpha} \tau \in \frac{1}{\alpha} \partial\iota_c(A\beta)$. Multiplying $A^\top$ to both sides of the previous inclusion and using the chain rule $\partial(\iota_c \circ A)(\beta) = A^\top \partial\iota_c(A\beta)$, we have that $\frac{\lambda}{\alpha} A^\top \tau \in \frac{1}{\alpha} \partial(\iota_c \circ A)(\beta)$. Since $-\frac{\lambda}{\alpha} A^\top \tau \in \partial(\frac{1}{\alpha}\|\cdot\|_1)(\beta)$, we obtain $0 \in \partial\|\cdot\|_1(\beta) + \partial(\iota_c \circ A)(\beta)$. This shows that β is a solution to problem (7). $\square$

We comment on the computation of the proximity operators $\text{prox}_{\frac{1}{\alpha}\|\cdot\|_1}$ and prox_{ι_c} appearing in equations (8) and (9). The proximity operator $\text{prox}_{\frac{1}{\alpha}\|\cdot\|_1}$ at any $u \in \mathbb{R}^p$ is the well-known soft-thresholding operator given as follow:

$$\text{prox}_{\frac{1}{\alpha}\|\cdot\|_1}(u) = \text{sign}(u) \circ \max \left\{ |u| - \frac{1}{\alpha} \mathbb{1}, 0 \right\}. \quad (10)$$

To compute the proximity operator prox_{ι_c} , we state the following result.

Lemma 2.2. *Let δ be a constant, let b be a vector in $\mathbb{R}^p$, and let the set $\mathcal{C}$ be given in (6). Then for any vector $v \in \mathbb{R}^p$,*

$$\text{prox}_{\iota_c}(v) = b + \min\{\max\{v - b, -\delta \mathbb{1}\}, \delta \mathbb{1}\} \quad (11)$$

and

$$(I - \text{prox}_{\iota_c})(v) = \text{prox}_{\delta\|\cdot\|_1}(v - b). \quad (12)$$

Proof: It is well-known that the proximity operator prox_{ι_c} is the projection operator onto the set $\mathcal{C}$. Since the set $\mathcal{C}$ is the cube with b as its center and 2δ as the length of its side in $\mathbb{R}^p$, $\text{prox}_{\iota_c}(v)$, given by (11), is essentially the projection of the vector $v \in \mathbb{R}^p$ onto $\mathcal{C}$. Further, it holds that $(I - \text{prox}_{\iota_c})(v) = (v - b) - \min\{\max\{(v - b), -\delta \mathbb{1}\}, \delta \mathbb{1}\}$. From this identity, we can directly check that for each i from 1 to n

$$(v - b)_i - \min\{\max\{(v - b)_i, -\delta\}, \delta\} = \text{sign}((v - b)_i) \cdot \max\{|(v - b)_i| - \delta, 0\},$$

which, by using equation (10), is $\text{prox}_{\delta\|\cdot\|_1}((v-b)_i)$. This completes the proof. $\square$

We remark that equation (12) is simply the Moreau identity applied to the indicator function $\iota_{\mathcal{C}}$ (see, e.g., [1, Theorem 14.3] or [26, Theorem 31.5]).

As a result of Lemma 2.2, equation (9) can be rewritten as follows:

$$\tau = \text{prox}_{\delta\|\cdot\|_1}(A\beta + \tau - b). \quad (13)$$

Therefore, by Theorem 2.1, finding a solution β to problem (7) amounts to solving the coupled fixed-point equations (8) and (13).

Two iterative schemes can be derived from equations (8) and (13). Let us write equation (8) as $\beta = \text{prox}_{\frac{1}{\alpha}\|\cdot\|_1}(\beta - \frac{\lambda}{\alpha}A^\top(2\tau - \tau))$. With any initial estimates $\tau^{-1} = \tau^0$ and β^0 , the first iterative scheme based upon equations (8) and (13) is as follows:

$$\begin{cases} \beta^{k+1} = \text{prox}_{\frac{1}{\alpha}\|\cdot\|_1}(\beta^k - \frac{\lambda}{\alpha}A^\top(2\tau^k - \tau^{k-1})), \\ \tau^{k+1} = \text{prox}_{\delta\|\cdot\|_1}(A\beta^{k+1} + \tau^k - b). \end{cases} \quad (14)$$

We would like to comment the connection of this scheme with some existing ones. The dual formulation of (7), as derived in [19], is

$$\max_{\tau \in \mathbb{R}^p} \{-\langle b, \tau \rangle - \delta\|\tau\|_1 : \|A^\top \tau\|_\infty \leq 1\}.$$

Applying the primal-dual hybrid gradient method (see [13, Equation 2.18]) to the above dual formulation yields exactly the iterative scheme (14). It was further pointed out in [9] that the iterative scheme (14) is essentially the same as the linearized ADM applying to problem (7). In other words, the iterative scheme (14) is the same as (5) in the case of $D = I$.

Now, let us introduce the second iterative scheme for problem (7). Let us write equation (13) as $\tau = \text{prox}_{\delta\|\cdot\|_1}(A(2\beta - \beta) + \tau - b)$. With any initial estimates $\beta^{-1} = \beta^0$ and τ^0 , the second iterative scheme based upon equations (8) and (13) is as follows:

$$\begin{cases} \tau^{k+1} = \text{prox}_{\delta\|\cdot\|_1}(A(2\beta^k - \beta^{k-1}) + \tau^k - b), \\ \beta^{k+1} = \text{prox}_{\frac{1}{\alpha}\|\cdot\|_1}(\beta^k - \frac{\lambda}{\alpha}A^\top \tau^{k+1}). \end{cases} \quad (15)$$

We would like to point it out that the scheme (15) is the PDHGMu algorithm in [13] and the Chambolle-Pock primal-dual algorithm in [9] applied to the original Dantzig selector problem.

The sequence $\{(\beta^k, \tau^k) : k \in \mathbb{N}\}$ generated by the iterative schemes (14) and (15) will converge for any initial seed when $\lambda/\alpha < 1/\|A\|_2^2$. The proof of this convergence result can be found in [9, 17]. Hence, the limit of the sequence $\{(\beta^k, \tau^k) : k \in \mathbb{N}\}$ is a fixed-point of equations (8) and (9). In particular, the limit of the sequence $\{\beta^k : k \in \mathbb{N}\}$ is a solution to problem (7) by Theorem 2.1.

As noted in [7], the Dantzig selector often slightly underestimates the true values of the nonzero parameters. To correct this bias and increase performance in practical settings, a postprocessing procedure was proposed in [7]. Assume that β^∞ is the limit of the sequence $\{\beta^k : k \in \mathbb{N}\}$ that is generated through the iterative scheme (15). This postprocessing consists of two steps. The first step is to estimate $\Lambda := \{i : \beta_i^\infty \neq 0\}$, the support of

the vector β^∞ . Let X_Λ be the $n \times |\Lambda|$ submatrix obtained by extracting the columns of X corresponding to the indices in Λ , and let $\hat{\beta}_\Lambda$ be the $|\Lambda|$ -dimensional vector obtained by extracting the coordinates of $\hat{\beta} \in \mathbb{R}^p$ corresponding to the indices in Λ . The second step of the postprocessing is to construct the estimator $\hat{\beta} \in \mathbb{R}^p$ such that

$$\hat{\beta}_\Lambda = \operatorname{argmin}\{\|X_\Lambda \beta - y\|_2 : \beta \in \mathbb{R}^p\}$$

and set the other coordinates to zero. If the matrix $X_\Lambda^\top X_\Lambda$ is invertible then $\hat{\beta}_\Lambda = (X_\Lambda^\top X_\Lambda)^{-1} X_\Lambda^\top y$.

Putting all above discussion together, a complete two-stage procedure for finding a solution of problem (7) is described in Algorithm 1.

Algorithm 1 (Two-stage scheme for problem (7))

Input: Set the fixed parameters

$$y \in \mathbb{R}^n, \quad A \in \mathbb{R}^{p \times p}, \quad b \in \mathbb{R}^p, \quad \delta, \alpha, \text{tol} \in \mathbb{R}_+, \quad \text{and } \lambda = 0.999\alpha/\|A\|_2^2.$$

Initialization: Set the initial parameters

$$\tau^0 = 0, \quad \beta^{-1} = \beta^0 = 0, \quad \text{and } k = 0.$$

Stage-I: Generate the sequence $\{(\tau^k, \beta^k) : k \in \mathbb{N}\}$ using Equations (10) and (15).

while (stopping criterion not met) **do**

$$\begin{aligned} \tau^{k+1} &\leftarrow \operatorname{prox}_{\delta\|\cdot\|_1}(A(2\beta^k - \beta^{k-1}) + \tau^k - b), \\ \beta^{k+1} &\leftarrow \operatorname{prox}_{\frac{1}{\alpha}\|\cdot\|_1}(\beta^k - \frac{\lambda}{\alpha}A^\top \tau^{k+1}), \\ k &\leftarrow k + 1. \end{aligned}$$

end while

Stage-II: Let $(\tau^\infty, \beta^\infty)$ be the last set of parameters computed in Stage-I.

- Approximate $\operatorname{supp}(\beta^\infty)$ by $\Lambda = \{j : |\beta^\infty(j)| < \text{tol}\}$.
- Compute $\hat{v} = \operatorname{argmin}_{v \in \mathbb{R}^{|\Lambda|}} \{\|X_\Lambda v - y\|_2\}$.
- Extend $\hat{v}$ to form the Dantzig selector $\hat{\beta}$ on Λ :

$$\begin{cases} \hat{\beta}(\Lambda(i)) = \hat{v}(i), & \text{for } i = 1 : |\Lambda|, \\ \hat{\beta}(j) = 0, & \text{for } j \notin \Lambda. \end{cases}$$

Stage-I of Algorithm 1 terminates once the sequence $\{(\tau^k, \beta^k) : k \in \mathbb{N}\}$ reaches a stationary point. To estimate when this occurs, terminate the iterations when either of the following stopping criteria are met:

1. The relative change between successive terms in the sequence $\{\beta^k\}$ falls below a spec-

ified tolerance;

$$\frac{\|\beta^{k+1} - \beta^k\|_2}{\|\beta^k\|_2} < \varepsilon,$$

for some $\varepsilon > 0$, or

2. The support of the sequence $\{\beta^k\}$ is stationary for a specified number of successive iterations;

$$\text{supp}(\beta^k) = \text{supp}(\beta^{k+1}) = \dots = \text{supp}(\beta^{k+\eta}),$$

for a fixed $\eta \in \mathbb{N}$ and some positive integer k .

Stage-I is the largest contributor to the computational complexity of Algorithm 1, with each iteration requiring a few matrix-vector multiplications, each of cost $\mathcal{O}(np)$. In comparison, each outer loop of ADM computing the τ and γ -related subproblems has complexity $\mathcal{O}(np)$, while each inner loop of ADM approximating the β -related subproblem has complexity $\mathcal{O}(np)$ as well. In general, Algorithm 1 and ADM will use a different number of iterations to terminate their iterative stages, so their overall complexities cannot be directly compared. However, the numerical experiments in the next section indicate that Algorithm 1 tends to have less overall complexity than ADM since Algorithm 1 has a shorter runtime even in situations where it requires more iterations.

3. Numerical Experiments

In the following experiments, we apply the proposed proximity operator based approach presented in Algorithm 1 and the alternating direction method (ADM) presented in [19] to solve the Dantzig selector problem (2) using both synthetic and real data sets. The experiments using synthetic data are performed in MATLAB R2013a on single nodes of the Condor Supercomputer, hosted at AFRL/RIT Affiliated Resource Center. The full capabilities of Condor were not taken advantage of; we ran the algorithms in serial using single nodes to emulate a typical high end consumer workstation. Each utilized node is equipped with an Intel Xeon X5650 6 core CPU, with 2.67 GHz and 6×8 GB RAM. The experiments using the real data set are performed in MATLAB R2014a on a PC with an Intel Core i7-3630QM 2.40 GHz processor and 16 GB RAM running Windows 7 Enterprise.

Example 3.1. *Synthetic Data Set*

In this series of simulations, sparse coefficient vectors are generated and then recovered from noisy random linear observations using both Algorithm 1 and ADM. The parameters used are $n = 720m$, $p = 2560m$ and $s = 80m$ for $m \in \{2, 3, \dots, 10\}$, and $\sigma \in \{0.01, 0.05, 0.10, 0.15\}$ corresponding to 1%, 5%, 10% and 15% noise levels. For each combination of m and σ , 100 simulations each of Algorithm 1 and ADM are performed. All other parameters for ADM are selected following the guidelines in [19, Section 3] and the parameters selected for the initialization stage of Algorithm 1 are $tol = 2\sigma$, $\alpha = 0.2\|A\|_2^2$ and $\delta = \sigma\sqrt{2\log p}$. The parameters for the stopping criteria are $\varepsilon = 10^{-4}$ and $\eta = \max\{\lceil 4\log(\alpha)\log(\sigma) + 2\alpha \rceil, 5\}$. The parameters tol, δ and η depend on the noise level σ , which in practice may not be known a priori. However, the noise level may be

well-approximated using existing methods. In the event that the noise level is not accurately approximated, the speed of convergence of Algorithm 1 will be affected, but the accuracy should not suffer much.

The $n \times p$ sensing matrices X are generated for each simulation with independent Gaussian entries normalized so each column has unit ℓ_2 norm. To generate the coefficient vector, for each simulation a support set S of size $|S| = s$ is selected uniformly at random. Then the vector β with indices in S is defined according to $\beta_{S(i)} = \epsilon_i(1 + |a_i|)$, where $\{a_i\}$ is a collection of independently and identically distributed random variables sampled from the standard normal distribution and $\{\epsilon_i\}$ is a collection of independently and identically distributed random variables sampled from the uniform distribution on $\{-1, 1\}$. For $i \notin S$, set $\beta_i = 0$. Then Algorithm 1 and ADM are used to approximate the Dantzig selector $\hat{\beta}$ from the observations $y = X\beta + z$, where z is a collection of independent and identically distributed random variables sampled from the normal distribution with mean zero and standard deviation σ .

Figure 3.1: A demonstration of the accuracy of the Dantzig selector recovered using Algorithm 1 and ADM with and without postprocessing for a single simulation of Experiment 3.1 with parameters $\sigma = 0.05$ with $(n, p, s) = (720, 2560, 80)$.

The accuracy of the Dantzig selector recovered in the simulations is measured by

$$\rho := \left(\frac{\|\beta - \hat{\beta}\|_2^2}{\sum_{j=1}^p \min\{\beta_j^2, \sigma^2\}} \right)^{1/2}, \quad (16)$$

where β denotes the true parameter and $\hat{\beta}$ denotes the parameter recovered using either Algorithm 1 or ADM. The denominator term of Equation (16) is the expected mean squared-error of the ideal estimator [7]. Therefore, $\rho \geq 0$, and a smaller ρ implies a more accurate estimator.

Figure 3.2: A comparison of ρ , computed as in Equation (16) which measures the accuracy of the approximated Dantzig selectors, for Algorithm 1 and ADM for noise levels $\sigma = 0.01, 0.05, 0.10$ and 0.15 in Example 3.1. In each plot, the points along the curve represent the mean number of iterations required for each parameter m over 100 simulations, and the points on the vertical lines represent one standard deviation away from the means.

The effects of Stage-II of Algorithm 1 and the postprocessing step of ADM are illustrated in Figure 3.1. The figure displays values of the exact simulated vector β and of the Dantzig selector $\hat{\beta}$ approximated by each algorithm, first without performing postprocessing (the left column of Figure 3.1) and then with postprocessing (the right column of Figure 3.1) for one simulation with parameters $(n, p, s) = (720, 2560, 80)$ and noise $\sigma = 0.05$. One can clearly see that the postprocessing not only corrects the underestimated magnitudes of nonzero components of the estimates, but also eliminates unwanted nonzero components.

The results of the above simulations suggest that Algorithm 1 has less overall complexity than ADM, since the accuracy of the Dantzig selectors approximated by each method are similar yet Algorithm 1 completes much faster than ADM, even when requiring more iterations. Figure 3.2 displays the mean and standard deviation of ρ over 100 simulations for each parameter m and σ and for both Algorithm 1 and ADM. Note that the accuracy of the

Figure 3.3: A comparison of the CPU time required to recover the Dantzig selector using Algorithm 1 and ADM for noise levels $\sigma = 0.01, 0.05, 0.10$ and 0.15 as in Example 3.1. In each plot, the points along the curve represent the mean number of iterations required for each parameter m over 100 simulations, and the points on the vertical lines represent one standard deviation away from the means.

Dantzig selector approximated by the two algorithms are very similar across all parameter levels. Figure 3.3 displays the mean and standard deviation of the CPU time, and Figure 3.4 displays the mean and standard deviation of the total number of iterations performed by Algorithm 1 and the total number of iterations performed in the inner loop of ADM for 100 simulations for each parameter m and σ . From the figures, one can see that although Algorithm 1 requires more iterations than ADM, Algorithm 1 completes significantly faster.

Example 3.2. Leukemia Data Set

In this experiment, the Dantzig selectors produced by Algorithm 1 and by ADM are used with a collection of biomarker data to indicate whether a patient may be diagnosed with a specific type of cancer. The biomarker dataset, first introduced in [15] and studied in [29, 30], contains the measurements of 7128 genes related to leukemia diagnoses. The dataset is split into a training set and a testing set. The training set is sampled from 38 patients, 27 of whom were diagnosed with acute lymphocytic leukemia (ALL) and 11 with acute myelogenous leukemia (AML). The testing set is sampled from 34 patients, 20 diagnosed with ALL and 14 with AML.

Figure 3.4: A comparison of the number of iterations required to recover the Dantzig selector using Algorithm 1 and ADM for noise levels $\sigma = 0.01, 0.05, 0.10$ and 0.15 in Example 3.1. In each plot, the points along the curve represent the mean number of iterations required for each parameter m over 100 simulations, and the points on the vertical lines represent one standard deviation away from the means.

Let $X_{\text{train}} \in \mathbb{R}^{38 \times 7128}$ contain the biomarker data in the training set, where each row is all 7128 gene measurements of a single patient and each column has been normalized to have unit ℓ_2 norm. Let $y_{\text{train}} \in \mathbb{R}^{38}$ be the column vector indicating the diagnosis of each patient in the training set:

$$y_{\text{train}}(j) = \begin{cases} 0, & \text{if patient } j \text{ in the training set is diagnosed with ALL,} \\ 1, & \text{if patient } j \text{ in the training set is diagnosed with AML.} \end{cases}$$

Similarly define $X_{\text{test}} \in \mathbb{R}^{34 \times 7128}$ and $y_{\text{test}} \in \mathbb{R}^{34}$ from the data in the testing set.

This experiment has a training phase and a testing phase. In the training phase, a sparse vector $\hat{\beta}$ is found such that $X_{\text{train}}\hat{\beta} = y_{\text{train}}$. To preprocess the data, only the biomarkers with the largest variance are used to train the parameter $\hat{\beta}$. To this end, select a positive integer N , and let Λ be the N indices of columns from X_{train} with largest variance. Let $\tilde{X}_{\text{train}} \in \mathbb{R}^{38 \times N}$ be the submatrix of X_{train} with columns in Λ . Form the reduced problem

$$\hat{\beta}_{\Lambda} \in \operatorname{argmin}_{\beta \in \mathbb{R}^N} \left\{ \|\beta\|_1 : \left\| \tilde{X}_{\text{train}}^{\top} \left(\tilde{X}_{\text{train}}\beta - y_{\text{train}} \right) \right\|_{\infty} \leq \delta \right\}. \quad (17)$$

The Dantzig selector $\hat{\beta}_\Lambda \in \mathbb{R}^N$ satisfying problem (17) is computed using Algorithm 1 and ADM, then extended to form $\hat{\beta} \in \mathbb{R}^{7128}$ via

$$\begin{cases} \hat{\beta}(\Lambda(j)) = \hat{\beta}_\Lambda(j), & \text{for } j = 1 : N, \\ \hat{\beta}(k) = 0, & \text{if } k \notin \Lambda. \end{cases}$$

In the testing phase, the trained parameter $\hat{\beta}$ is used to predict the diagnoses of patients in the testing set. The predictive indicator vector $\hat{y}_{\text{test}} \in \mathbb{R}^{34}$ is computed from $y = X_{\text{test}}\hat{\beta}$ by thresholding and clustering values near the threshold boundary. Set

$$\hat{y}_{\text{test}}(j) = \begin{cases} 0, & \text{if } y(j) < 0.49, \\ 1, & \text{if } 0.51 < y(j). \end{cases}$$

Let $y_0 = \max\{y(j) : y(j) < 0.49\}$ and $y_1 = \min\{y(j) : 0.51 < y(j)\}$. For values of j such that $0.49 \leq y(j) \leq 0.51$, set

$$\hat{y}_{\text{test}}(j) = \begin{cases} 0, & \text{if } |y(j) - y_0| \leq |y(j) - y_1|, \\ 1, & \text{if } |y(j) - y_1| < |y(j) - y_0|. \end{cases}$$

The j^{th} patient in the testing set is predicted to have a diagnosis of ALL if $\hat{y}_{\text{test}}(j) = 0$ and a diagnosis of AML if $\hat{y}_{\text{test}}(j) = 1$.

The above procedure was used to predict the diagnoses of patients in the testing set using the Dantzig selector $\hat{\beta}_\Lambda$ computed using both Algorithm 1 and ADM with parameters $N = 1000$, $\alpha = \|X_{\text{train}}^\top X_{\text{train}}\|_2^2$ and $\text{tol} = 0.1$ and stopping criteria parameters $\eta = 80$ and $\varepsilon = 10^{-4}$ for each δ in $\{0.0625, 0.125, 0.1875, 0.25, 0.3125, 0.375\}$. Figure 3.5 displays the results of these simulations regarding the accuracy of the recovered indicator vector $\hat{y}_{\text{test}}$ in predicting the leukemia diagnoses of patients in the testing set, as well as the number of iterations and CPU runtime used by Algorithm 1 and ADM. As shown in Figure 3.5(a), Algorithm 1 typically predicted the diagnoses of patients with higher accuracy than ADM. Moreover, for each parameter δ , Algorithm 1 used fewer iterations than ADM and the time used by Algorithm 1 was several orders of magnitude less than the time used by ADM, as shown in Figures 3.5(c) and (d). Figure 3.5(b) illustrates the tendency of Algorithm 1 to predict the diagnosis of patients in the testing set with higher accuracy than ADM. This plot displays the values of $y = X_{\text{test}}\hat{\beta}$ recovered using Algorithm 1 and by ADM prior to the thresholding step, along with the true values of y_{test} . Since the values recovered by Algorithm 1 tend to be more spread out, it is easier to accurately separate them into two distinct clusters.

4. Conclusion

In this paper, we developed an iterative algorithm to compute the Dantzig selector, the solution to the minimization problem (2). The algorithm is based on the proximity operator and its relationship to problem (7). The two-stage algorithm we proposed is an improvement over recently proposed methods to find the Dantzig selector, which require the use of inner loop to estimate parameters within each step of the algorithm. Additionally, our

Figure 3.5: Plots regarding the indicator vector, used to predict a leukemia diagnosis in patients in the testing set as in Example 3.2, using both Algorithm 1 and ADM. (a) The number of patients in the testing set misdiagnosed by the predicted indicator vector recovered by Algorithm 1 and ADM for various values of the parameter δ ; (b) Values of the actual diagnosis indicator vector y_{test} along with values of the predicted indicator vectors recovered by Algorithm 1 and ADM prior to separating values into classification groups for $\delta = 0.25$; (c) The number of iterations required by Algorithm 1 and ADM to recover the Dantzig selector $\hat{\beta}_\Lambda$ for various values of the parameter δ ; and (d) The CPU runtime required by Algorithm 1 and ADM to recover the Dantzig selector $\hat{\beta}_\Lambda$ for various values of the parameter δ .

proposed method uses a novel stopping criterion based upon the support of the approximated parameters.

We compared the proposed algorithm to the alternating direction method proposed in [19]. Theoretically, two methods produced results of similar quality, however each iteration of Stage-I of Algorithm 1 has less computational complexity than each iteration of the inner loop of the alternating direction method. The numerical experiments demonstrated that the proposed method and the alternating direction method typically approximated the Dantzig selectors with similar accuracy, yet Algorithm 1 produced results in significantly less time, whether it used more iterations than the alternating direction method, as in Experiment 3.1, or fewer iterations than the alternating direction method, as in Experiment 3.2.

Acknowledgements

This work is cleared for public release by WPAFB Public Affairs on 09 October 2013. Case Number: 88ABW-2013-4324. This research is supported in part by an award from National Research Council via the Air Force Office of Scientific Research and by the US National Science Foundation under grant DMS-1115523. The authors are grateful to the anonymous reviewers for their helpful comments. The authors also would like to thank Drs. X. Wang and X. Yuan for providing the MATLAB code to approximate the Dantzig selector using the alternating direction method and for sharing the real dataset used in Example 3.2.

References

- [1] H. L. BAUSCHKE AND P. L. COMBETTES, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, AMS Books in Mathematics, Springer, New York, 2011.
- [2] A. BECK AND M. TEBoulLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM Journal on Imaging Sciences, 2 (2009), pp. 183–202.
- [3] S. BECKER, E. CANDÉS, AND M. GRANT, *Templates for convex cone problems with applications to sparse signal recovery*, Mathematical Programming Computation, 3 (2010), pp. 165–218.
- [4] P. J. BICKEL, *Discussion: the Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2352–2357.
- [5] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, 2004.
- [6] T. CAI AND J. LV, *Discussion: the Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2365–2369.
- [7] E. CANDÉS AND T. TAO, *The Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2313–2351.
- [8] —, *Rejoinder: the Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2392–2404.

- [9] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, Journal of Mathematical Imaging and Vision, 40 (2011), pp. 120–145.
- [10] L. DICKER AND X. LIN, *Parallelism, uniqueness, and large-sample asymptotics for the Dantzig selector*, Canadian Journal of Statistics, 4 (2013), pp. 23–51.
- [11] B. EFRON, T. HASTIE, I. JOHNSTONE, AND R. TIBSHIRANI, *Least angle regression*, The Annals of Statistics, 32 (2004), pp. 407–451.
- [12] B. EFRON, T. HASTIE, AND R. TIBSHIRANI, *Discussion: the Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2358–2364.
- [13] E. ESSER, X. ZHANG, AND T. F. CHAN, *A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science*, SIAM Journal on Imaging Sciences, 3 (2010), pp. 1015–1046.
- [14] M. P. FRIEDLANDER AND M. A. SAUNDERS, *Discussion: the Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2385–2391.
- [15] T. R. GOLUB, D. K. SLONIM, P. TAMAYO, C. HUARD, M. GAASENBEEK, J. P. MESIROV, H. COLLIER, M. L. LOH, J. R. DOWNING, M. A. CALIGUIRI, C. D. BLOOMFIELD, AND E. S. LANDER, *Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring*, Science, 286 (1999), pp. 531–537.
- [16] G. JAMES, P. RADCHENKO, AND J. LV, *DASSO: connections between the Dantzig selector and LASSO*, Journal of the Royal Statistical Society, Series B, 71 (2009), pp. 127–142.
- [17] Q. LI, C. A. MICCHELLI, L. SHEN, AND Y. XU, *A proximity algorithm accelerated by Gauss-Seidel iterations for L1/TV denoising models*, Inverse Problems, 28 (2012), p. 095003.
- [18] Y. LI, L. DICKER, AND S. ZHAO, *The Dantzig selector for censored linear regression models*, Statistica Sinica, to appear, (2012).
- [19] Z. LU, T. P. PONG, AND Y. ZHANG, *An alternating direction for finding Dantzig selectors*, Computational Statistics and Data Analysis, 56 (2012), pp. 4037–4046.
- [20] Z. LU AND Y. ZHANG, *An augmented Lagrangian approach for sparse principle component analysis*, Mathematical Programming, 135 (2012), pp. 149–193.
- [21] N. MEINSHAUSEN AND P. BUHLMANN, *High-dimensional graphs and variable selection with the LASSO*, Annals of Statistics, 34 (2006), pp. 1436–1462.
- [22] N. MEINSHAUSEN, G. ROCHA, AND B. YU, *Discussion: the Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2373–2384.

- [23] Y. NESTEROV, *Smooth minimization of non-smooth functions*, Mathematical Programming, Series A, 103 (2005), pp. 127–152.
- [24] M. OSBORNE, B. PRESNELL, AND B. TURLACH, *On the LASSO and its dual*, Journal of Computational and Graphical Statistics, 9 (2000), pp. 319–337.
- [25] Y. RITOV, *Discussion: the Dantzig selector: Statistical estimation when p is much larger than n* , Annals of Statistics, 35 (2007), pp. 2370–2372.
- [26] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.
- [27] R. TIBSHIRANI, *Regression shrinkage and selection via the LASSO*, Journal of the Royal Statistical Society, Series B, 58 (1996), pp. 267–288.
- [28] ———, *Regression shrinkage and selection and via the LASSO: a retrospective*, Journal of the Royal Statistical Society, Series B, 73 (2011), pp. 273–282.
- [29] R. TIBSHIRANI, M. SAUNDERS, S. ROSSET, J. ZHU, AND K. KNIGHT, *Sparsity and smoothness via the fused LASSO*, Journal of the Royal Statistical Society, Series B, 67 (2005), pp. 91–108.
- [30] X. WANG AND X. YUAN, *The linearized alternating direction method of multipliers for Dantzig selector*, SIAM Journal on Scientific Computing, 34 (2012), pp. 2792–2811.
- [31] P. ZHAO AND B. YU, *On model selection consistency of LASSO*, Journal of Machine Learning Research, 7 (2006), pp. 2541–2563.