A HIGHLY EFFICIENT SEMISMOOTH NEWTON AUGMENTED LAGRANGIAN METHOD FOR SOLVING LASSO PROBLEMS*

XUDONG LI[†], DEFENG SUN[‡], AND KIM-CHUAN TOH[§]

Abstract. We develop a fast and robust algorithm for solving large-scale convex composite optimization models with an emphasis on the ℓ_1-regularized least squares regression (lasso) problems. Despite the fact that there exist a large number of solvers in the literature for the lasso problems, we found that no solver can efficiently handle difficult large-scale regression problems with real data. By leveraging on available error bound results to realize the asymptotic superlinear convergence property of the augmented Lagrangian algorithm, and by exploiting the second order sparsity of the problem through the semismooth Newton method, we are able to propose an algorithm, called Ssnal, to efficiently solve the aforementioned difficult problems. Under very mild conditions, which hold automatically for lasso problems, both the primal and the dual iteration sequences generated by Ssnal possess a fast linear convergence rate, which can even be superlinear asymptotically. Numerical comparisons between our approach and a number of state-of-the-art solvers, on real data sets, are presented to demonstrate the high efficiency and robustness of our proposed algorithm in solving difficult large-scale lasso problems.

Key words. lasso, sparse optimization, augmented Lagrangian, metric subregularity, semismoothness, Newton's method

AMS subject classifications. 65F10, 90C06, 90C25, 90C31

DOI. 10.1137/16M1097572

1. Introduction. In this paper, we aim to design a highly efficient and robust algorithm for solving convex composite optimization problems including the following ℓ_1-regularized least squares (LS) problem:

$$(1) \qquad \min\left\{ \frac{1}{2}\|\mathcal{A}x - b\|^2 + \lambda\|x\|_1 \right\},$$

where $\mathcal{A} : \mathcal{X} \to \mathcal{Y}$ is a linear map whose adjoint is denoted as $\mathcal{A}^*$, $b \in \mathcal{Y}$ and $\lambda > 0$ are given data, and $\mathcal{X}$, $\mathcal{Y}$ are two real finite-dimensional Euclidean spaces each equipped with an inner product $\langle \cdot, \cdot \rangle$ and its induced norm $\| \cdot \|$.

With the advent of convenient automated data collection technologies, the Big Data era brings new challenges in analyzing massive data due to the inherent sizes—large samples and high dimensionality [15]. In order to respond to these challenges, researchers have developed many new statistical tools to analyze such data. Among these, the ℓ_1-regularized models are arguably the most intensively studied. They are used in many applications, such as in compressive sensing, high-dimensional variable selection, and image reconstruction. Most notably, the model (1), called lasso, was

†Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544 (xudongl@princeton.edu).

‡Department of Applied Mathematics, The Hong Kong Polytechnic University, Hung Hom, Hong Kong. On leave from Department of Mathematics, National University of Singapore (defeng.sun@polyu.edu.hk).

§Department of Mathematics and Institute of Operations Research and Analytics, National University of Singapore, 10 Lower Kent Ridge Road, Singapore 119076, Singapore (mattohkc@nus.edu.sg).

proposed in [48] and has been used extensively in high-dimensional statistics and machine learning. The model (1) has also been studied in the signal processing context under the name basis pursuit denoising [9]. In addition to its own importance in statistics and machine learning, lasso problem (1) also appears as an inner subproblem of many important algorithms. For example, in recent papers [4, 1], a level set method was proposed to solve a computationally more challenging reformulation of the lasso problem, i.e.,

$$\min \left\{ \|x\|_1 \ \mid \ \frac{1}{2} \|\mathcal{A}x - b\|^2 \leq \sigma \right\}.$$

The level set method relies critically on the assumption that the optimization problem

$$\min \left\{ \frac{1}{2} \|\mathcal{A}x - b\|^2 + \delta_{B_\lambda}(x) \right\}, \quad B_\lambda = \{x \mid \| \cdot \|_1 \leq \lambda\},$$

of the same type as (1), can be efficiently solved to high accuracy. The lasso-type optimization problems also appear as subproblems in various proximal Newton methods for solving convex composite optimization problems [7, 29, 53]. Notably, in a broader sense, all these proximal Newton methods belong to the class of algorithms studied in [17].

The above-mentioned importance together with a wide range of applications of (1) has inspired many researchers to develop various algorithms for solving this problem and its equivalent reformulations. These algorithms can roughly be divided into two categories. The first category consists of algorithms that use only the gradient information, for example, the accelerated proximal gradient (APG) method [37, 2], GPSR [16], SPGL1 [4], SpaRSA [50], FPC_AS [49], and NESTA [3], to name only a few. Meanwhile, algorithms in the second category, including but not limited to mfIPM [18], SNF [36], BAS [6], SQA [7], OBA [26], and FBS-Newton [51], utilize the second order information of the underlying problem in the algorithmic design to accelerate the convergence. Nearly all of these second order information based solvers rely on certain nondegeneracy assumptions to guarantee the nonsingularity of the corresponding inner linear systems. The only exception is the inexact interior-point-algorithm based solver mfIPM, which does not rely on the nondegeneracy assumption but requires the availability of appropriate preconditioners to ameliorate the extreme ill-conditioning in the linear systems of the subproblems. For nondegenerate problems, the solvers in the second category generally work quite well and usually outperform the algorithms in the first category when high accuracy solutions are sought. In this paper, we also aim to solve the lasso problems by making use of the second order information. The novelty of our approach is that we do not need any nondegeneracy assumption in our theory or computations. The core idea is to analyze the fundamental nonsmooth structures in the problems to formulate and solve specific semismooth equations with well conditioned symmetric and positive definite generalized Jacobian matrices, which consequently play a critical role in our algorithmic design. When applied to solving difficult large-scale sparse optimization problems, even for degenerate ones, our approach can outperform the first order algorithms by a huge margin regardless of whether low- or high-accuracy solutions are sought. This is in sharp contrast to most of the other second order based solvers mentioned above, where their competitive advantages over first order methods only become apparent when high-accuracy solutions are sought.

Our proposed algorithm is a semismooth Newton augmented Lagrangian (in short, Ssnal) method for solving the dual of problem (1), where the sparsity property of the

second order generalized Hessian is wisely exploited. This algorithmic framework is adapted from those that appeared in [54, 25, 52, 31] for solving semidefinite programming problems where impressive numerical results have been reported. Specialized to the vector case, our SSNAL possesses unique features that are not available in the semidefinite programming case. It is these combined unique features that allow our algorithm to converge at a very fast speed with very low computational costs at each step. Indeed, for large-scale sparse lasso problems, our numerical experiments show that the proposed algorithm needs at most a few dozens of outer iterations to reach solutions with the desired accuracy while all the inner semismooth Newton subproblems can be solved very cheaply. One reason for this impressive performance is that the piecewise linear-quadratic structure of the lasso problem (1) guarantees the asymptotic superlinear convergence of the augmented Lagrangian algorithm. Beyond the piecewise linear-quadratic case, we also study more general functions to guarantee this fast convergence rate. More importantly, since there are several desirable properties including the strong convexity of the objective function in the inner subproblems, in each outer iteration we only need to execute a few (usually one to four) semismooth Newton steps to solve the underlying subproblem. As will be shown later, for lasso problems with sparse optimal solutions, the computational costs of performing these semismooth Newton steps can be made to be extremely cheap compared to other costs. This seems to be counterintuitive as normally one would expect a second order method to be computationally much more expensive than the first order methods at each step. Here, we make this counterintuitive achievement possible by carefully exploiting the second order sparsity in the augmented Lagrangian functions. Notably, our algorithmic framework not only works for models such as lasso, adaptive lasso [56], and elastic net [57], but can also be applied to more general convex composite optimization problems. The high performance of our algorithm also serves to show that the second order information, more specifically the nonsmooth second order information, can be and should be incorporated intelligently into the algorithmic design for large-scale optimization problems.

The remaining parts of this paper are organized as follows. In the next section, we introduce some definitions and present preliminary results on the metric subregularity of multivalued mappings. We should emphasize here that these stability results play a pivotal role in the analysis of the convergence rate of our algorithm. In section 3, we propose an augmented Lagrangian algorithm to solve the general convex composite optimization model and analyze its asymptotic superlinear convergence. The semismooth Newton algorithm for solving the inner subproblems and the efficient implementation of the algorithm are also presented in this section. In section 4, we conduct extensive numerical experiments to evaluate the performance of SSNAL in solving various lasso problems. We conclude our paper in sections.

2. Preliminaries. We discuss in this section some stability properties of convex composite optimization problems. It will become apparent later that these stability properties are the key ingredients for establishing the fast convergence of our augmented Lagrangian method.

Recall that $\mathcal{X}$ and $\mathcal{Y}$ are two real finite-dimensional Euclidean spaces. For a given closed proper convex function $p : \mathcal{X} \to (-\infty, +\infty]$, the proximal mapping $\mathrm{Prox}_p(\cdot)$ associated with p is defined by

$$\mathrm{Prox}_p(x) := \arg\min_{u \in \mathcal{X}} \left\{ p(x) + \frac{1}{2} \|u - x\|^2 \right\} \quad \forall x \in \mathcal{X}.$$

We will often make use of the Moreau identity $\text{Prox}_{tp}(x) + t\text{Prox}_{p^*/t}(x/t) = x$, where $t > 0$ is a given parameter. Denote $\text{dist}(x, C) = \inf_{x' \in C} \|x - x'\|$ for any $x \in \mathcal{X}$ and any set $C \subset \mathcal{X}$.

Let $F : \mathcal{X} \rightrightarrows \mathcal{Y}$ be a multivalued mapping. We define the graph of F to be the set

$$\text{gph}F := \{(x, y) \in \mathcal{X} \times \mathcal{Y} \mid y \in F(x)\}.$$

F^{-1}, the inverse of F, is the multivalued mapping from $\mathcal{Y}$ to $\mathcal{X}$ whose graph is $\{(y, x) \mid (x, y) \in \text{gph}F\}$.

DEFINITION 2.1 (error bound). *Let $F : \mathcal{X} \rightrightarrows \mathcal{Y}$ be a multivalued mapping and $y \in \mathcal{Y}$ satisfy $F^{-1}(y) \neq \emptyset$. F is said to satisfy the error bound condition for the point y with modulus $\kappa \geq 0$ if there exists $\varepsilon > 0$ such that if $x \in \mathcal{X}$ with $\text{dist}(y, F(x)) \leq \varepsilon$, then*

$$(2) \qquad\qquad \text{dist}\left(x, F^{-1}(y)\right) \leq \kappa \, \text{dist}(y, F(x)).$$

The above error bound condition was called the growth condition in [34] and was used to analyze the local linear convergence properties of the proximal point algorithm. Recall that $F : \mathcal{X} \rightrightarrows \mathcal{Y}$ is called a polyhedral multifunction if its graph is the union of finitely many polyhedral convex sets. In [39], Robinson established the following celebrated proposition on the error bound result for polyhedral multifunctions.

PROPOSITION 2.2. *Let F be a polyhedral multifunction from $\mathcal{X}$ to $\mathcal{Y}$. Then, F satisfies the error bound condition (2) for any point $y \in \mathcal{Y}$ satisfying $F^{-1}(y) \neq \emptyset$ with a common modulus $\kappa \geq 0$.*

For later use, we present the following definition of metric subregularity from Chapter 3 in [13].

DEFINITION 2.3 (metric subregularity). *Let $F : \mathcal{X} \rightrightarrows \mathcal{Y}$ be a multivalued mapping and $(\bar{x}, \bar{y}) \in \text{gph}F$. F is said to be metrically subregular at $\bar{x}$ for $\bar{y}$ with modulus $\kappa \geq 0$ if there exist neighborhoods U of $\bar{x}$ and V of $\bar{y}$ such that*

$$\text{dist}(x, F^{-1}(\bar{y})) \leq \kappa \, \text{dist}(\bar{y}, F(x) \cap V) \quad \forall x \in U.$$

From the above definition, we see that if $F : \mathcal{X} \rightrightarrows \mathcal{Y}$ satisfies the error bound condition (2) for $\bar{y}$ with the modulus κ, then it is metrically subregular at $\bar{x}$ for $\bar{y}$ with the same modulus κ for any $\bar{x} \in F^{-1}(\bar{y})$.

The following definition on essential smoothness is taken from [40, section 26].

DEFINITION 2.4 (essential smoothness). *A proper convex function f on $\mathcal{X}$ is essentially smooth if f is differentiable on $\text{int}\,(\text{dom}\,f) \neq \emptyset$ and $\lim_{k \to \infty} \|\nabla f(x^k)\| = +\infty$ whenever $\{x^k\}$ is a sequence in $\text{int}\,(\text{dom}\,f)$ converging to a boundary point x of $\text{int}\,(\text{dom}\,f)$.*

In particular, a smooth convex function on $\mathcal{X}$ is essentially smooth. Moreover, if a closed proper convex function f is strictly convex on $\text{dom}\,f$, then its conjugate f^* is essentially smooth [40, Theorem 26.3].

Consider the following composite convex optimization model:

$$(3) \qquad\qquad \max\left\{-\left(f(x) = h(\mathcal{A}x) - \langle c,\, x \rangle + p(x)\right)\right\},$$

where $\mathcal{A} : \mathcal{X} \to \mathcal{Y}$ is a linear map, $h : \mathcal{Y} \to \Re$ and $p : \mathcal{X} \to (-\infty, +\infty]$ are two closed proper convex functions, and $c \in \mathcal{X}$ is a given vector. The dual of (3) can be written as

$$(4) \qquad\qquad \min\left\{h^*(y) + p^*(z) \mid \mathcal{A}^*y + z = c\right\},$$

where g^* and p^* are the Fenchel conjugate functions of g and h, respectively. Throughout this section, we make the following blanket assumption on $h^*(\cdot)$ and $p^*(\cdot)$.

Assumption 2.5. $h^*(\cdot)$ is essentially smooth and $p^*(\cdot)$ is either an indicator function $\delta_P(\cdot)$ or a support function $\delta_P^*(\cdot)$ for some nonempty polyhedral convex set $P \subseteq \mathcal{X}$. Moreover, ∇h^* is locally Lipschitz continuous and directionally differentiable on $\mathrm{int}\,(\mathrm{dom}\,h^*)$.

Under Assumption 2.5, by [40, Theorem 26.1], we know that $\partial h^*(y) = \emptyset$ whenever $y \notin \mathrm{int}\,(\mathrm{dom}\,h^*)$. Denote by l the Lagrangian function for (4):

$$(5) \qquad l(y, z, x) = h^*(y) + p^*(z) - \langle x, \mathcal{A}^* y + z - c \rangle \quad \forall\,(y, z, x) \in \mathcal{Y} \times \mathcal{X} \times \mathcal{X}.$$

Corresponding to the closed proper convex function f in the objective of (3) and the convex-concave function l in (5), define the maximal monotone operators $\mathcal{T}_f$ and $\mathcal{T}_l$ [42] by

$$\mathcal{T}_f(x) := \partial f(x), \quad \mathcal{T}_l(y, z, x) := \{(y', z', x') \,|\, (y', z', -x') \in \partial l(y, z, x)\},$$

whose inverses are given, respectively, by

$$\mathcal{T}_f^{-1}(x') := \partial f^*(x'), \quad \mathcal{T}_l^{-1}(y', z', x') := \{(y, z, x) \,|\, (y', z', -x') \in \partial l(y, z, x)\}.$$

Unlike the case for $\mathcal{T}_f$ [33, 47, 58], stability results of $\mathcal{T}_l$ which correspond to the perturbations of both primal and dual solutions are very limited. Next, as a tool for studying the convergence rate of SSNAL, we shall establish a theorem which reveals the metric subregularity of $\mathcal{T}_l$ under some mild assumptions.

The KKT system associated with problem (4) is given as follows:

$$0 \in \partial h^*(y) - \mathcal{A}x, \quad 0 \in -x + \partial p^*(z), \quad 0 = \mathcal{A}^* y + z - c, \quad (x, y, z) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{X}.$$

Assume that the above KKT system has at least one solution. This existence assumption together with the essential smoothness assumption on h^* implies that the above KKT system can be equivalently rewritten as

$$(6) \qquad \begin{cases} 0 = \nabla h^*(y) - \mathcal{A}x, \\ 0 \in -x + \partial p^*(z), \quad (x, y, z) \in \mathcal{X} \times \mathrm{int}(\mathrm{dom}\,h^*) \times \mathcal{X}. \\ 0 = \mathcal{A}^* y + z - c, \end{cases}$$

Therefore, under the above assumptions, we only need to focus on $\mathrm{int}\,(\mathrm{dom}\,h^*) \times \mathcal{X}$ when solving problem (4). Let $(\bar{y}, \bar{z})$ be an optimal solution to problem (4). Then, we know that the set of the Lagrangian multipliers associated with $(\bar{y}, \bar{z})$, denoted as $\mathcal{M}(\bar{y}, \bar{z})$, is nonempty. Define the critical cone associated with (4) at $(\bar{y}, \bar{z})$ as follows:

$$(7) \qquad \mathcal{C}(\bar{y}, \bar{z}) := \left\{ (d_1, d_2) \in \mathcal{Y} \times \mathcal{X} \;\middle|\; \begin{array}{l} \mathcal{A}^* d_1 + d_2 = 0, \\ \langle \nabla h^*(\bar{y}), d_1 \rangle + (p^*)'(\bar{z}; d_2) = 0, d_2 \in \mathcal{T}_{\mathrm{dom}(p^*)}(\bar{z}) \end{array} \right\},$$

where $(p^*)'(\bar{z}; d_2)$ is the directional derivative of p^* at $\bar{z}$ with respect to d_2 and $\mathcal{T}_{\mathrm{dom}(p^*)}(\bar{z})$ is the tangent cone of $\mathrm{dom}(p^*)$ at $\bar{z}$. When the conjugate function p^* is taken to be the indicator function of a nonempty polyhedral set P, the above definition reduces directly to the standard definition of the critical cone in the nonlinear programming setting.

DEFINITION 2.6 (second order sufficient condition). *Let $(\bar{y}, \bar{z}) \in \mathcal{Y} \times \mathcal{X}$ be an optimal solution to problem (4) with $\mathcal{M}(\bar{y}, \bar{z}) \neq \emptyset$. We say that the second order sufficient condition for problem (4) holds at $(\bar{y}, \bar{z})$ if*

$$\langle d_1, (\nabla h^*)'(\bar{y}; d_1) \rangle > 0 \quad \forall 0 \neq (d_1, d_2) \in \mathcal{C}(\bar{y}, \bar{z}).$$

By building on the proof ideas from the literature on nonlinear programming problems [11, 27, 24], we are able to prove the following result on the metric subregularity of $\mathcal{T}_l$. This allows us to prove the linear and even the asymptotic superlinear convergence of the sequences generated by the SSNAL algorithm to be presented in the next section even when the objective in problem (3) does not possess the piecewise linear-quadratic structure as in the lasso problem (1).

THEOREM 2.7. *Assume that the KKT system* (6) *has at least one solution, and denote it as $(\bar{x}, \bar{y}, \bar{z})$. Suppose that Assumption 2.5 holds and that the second order sufficient condition for problem* (4) *holds at $(\bar{y}, \bar{z})$. Then, $\mathcal{T}_l$ is metrically subregular at $(\bar{y}, \bar{z}, \bar{x})$ for the origin.*

Proof. First, we claim that there exists a neighborhood $\mathcal{U}$ of $(\bar{x}, \bar{y}, \bar{z})$ such that for any $w = (u_1, u_2, v) \in \mathcal{W} := \mathcal{Y} \times \mathcal{X} \times \mathcal{X}$ with $\|w\|$ sufficiently small, any solution $(x(w), y(w), z(w)) \in \mathcal{U}$ of the perturbed KKT system

$$(8) \qquad 0 = \nabla h^*(y) - \mathcal{A}x - u_1, \quad 0 \in -x - u_2 + \partial p^*(z), \quad 0 = \mathcal{A}^* y + z - c - v$$

satisfies the following estimate:

$$(9) \qquad \qquad \|(y(w), z(w)) - (\bar{y}, \bar{z})\| = O(\|w\|).$$

For the sake of contradiction, suppose that our claim is not true. Then, there exist some sequences $\{w^k := (u_1^k, u_2^k, v^k)\}$ and $\{(x^k, y^k, z^k)\}$ such that $w^k \to 0$ and $(x^k, y^k, z^k) \to (\bar{x}, \bar{y}, \bar{z})$, for each k the point (x^k, y^k, z^k) is a solution of (8) for $w = w^k$, and

$$\|(y^k, z^k) - (\bar{y}, \bar{z})\| > \gamma_k \|w^k\|$$

with some $\gamma_k > 0$ such that $\gamma_k \to \infty$. Passing to a subsequence if necessary, we assume that $\{((y^k, z^k) - (\bar{y}, \bar{z}))/\|(y^k, z^k) - (\bar{y}, \bar{z})\|\}$ converges to some $\xi = (\xi_1, \xi_2) \in \mathcal{Y} \times \mathcal{X}$, $\|\xi\| = 1$. Then, setting $t_k = \|(y^k, z^k) - (\bar{y}, \bar{z})\|$ and passing to a subsequence further if necessary, by the local Lipschitz continuity and the directional differentiability of $\nabla h^*(\cdot)$ at $\bar{y}$, we know that for all k sufficiently large

$$\nabla h^*(y^k) - \nabla h^*(\bar{y}) = \nabla h^*(\bar{y} + t_k \xi_1) - \nabla h^*(\bar{y}) + \nabla h^*(y^k) - \nabla h^*(\bar{y} + t_k \xi_1)$$

$$= t_k (\nabla h^*)'(\bar{y}; \xi_1) + o(t_k) + O(\|y^k - \bar{y} - t_k \xi_1\|)$$

$$= t_k (\nabla h^*)'(\bar{y}; \xi_1) + o(t_k).$$

Denote, for each k, $\hat{x}^k := x^k + u_2^k$. Simple calculations show that for all k sufficiently large

$$(10) \quad 0 = \nabla h^*(y^k) - \nabla h^*(\bar{y}) - \mathcal{A}(\hat{x}^k - \bar{x}) + \mathcal{A}u_2^k - u_1^k = t_k (\nabla h^*)'(\bar{y}; \xi_1) + o(t_k) - \mathcal{A}(\hat{x}^k - \bar{x})$$

and

$$(11) \qquad 0 = \mathcal{A}(y^k - \bar{y}) + (z^k - \bar{z}) - v^k = t_k (\mathcal{A}^* \xi_1 + \xi_2) + o(t_k).$$

Dividing both sides of (11) by t_k and then taking limits, we obtain

$$\mathcal{A}^* \xi_1 + \xi_2 = 0, \tag{12}$$

which further implies that

$$\langle \nabla h^*(\bar{y}), \xi_1 \rangle = \langle \mathcal{A}\bar{x}, \xi_1 \rangle = -\langle \bar{x}, \xi_2 \rangle. \tag{13}$$

Since $\hat{x}^k \in \partial p^*(z^k)$, we know that for all k sufficiently large,

$$z^k = \bar{z} + t_k \xi_2 + o(t_k) \in \operatorname{dom} p^*.$$

That is, $\xi_2 \in \mathcal{T}_{\operatorname{dom} p^*}(\bar{z})$.

According to the structure of p^*, we separate our discussions into two cases.

Case I. There exists a nonempty polyhedral convex set P such that $p^*(z) = \delta_P^*(z)$ for all $z \in \mathcal{X}$. Then, for each k, it holds that

$$\hat{x}^k = \Pi_P(z^k + \hat{x}^k).$$

By [14, Theorem 4.1.1], we have

$$\hat{x}^k = \Pi_P(z^k + \hat{x}^k) = \Pi_P(\bar{z} + \bar{x}) + \Pi_{\mathcal{C}}(z^k - \bar{z} + \hat{x}^k - \bar{x}) = \bar{x} + \Pi_{\mathcal{C}}(z^k - \bar{z} + \hat{x}^k - \bar{x}), \tag{14}$$

where $\mathcal{C}$ is the critical cone of P at $\bar{z} + \bar{x}$, i.e.,

$$\mathcal{C} \equiv \mathcal{C}_P(\bar{z} + \bar{x}) := \mathcal{T}_P(\bar{x}) \cap \bar{z}^\perp.$$

Since $\mathcal{C}$ is a polyhedral cone, we know from [14, Proposition 4.1.4] that $\Pi_{\mathcal{C}}(\cdot)$ is a piecewise linear function, i.e., there exist a positive integer l and orthogonal projectors $B_1, \ldots, B_l$ such that for any $x \in \mathcal{X}$,

$$\Pi_{\mathcal{C}}(x) \in \{B_1 x, \ldots, B_l x\}.$$

By restricting to a subsequence if necessary, we may further assume that there exists $1 \le j' \le l$ such that for all $k \ge 1$,

$$\Pi_{\mathcal{C}}(z^k - \bar{z} + \hat{x}^k - \bar{x}) = B_{j'}(z^k - \bar{z} + \hat{x}^k - \bar{x}) = \Pi_{\operatorname{Range}(B_{j'})}(z^k - \bar{z} + \hat{x}^k - \bar{x}). \tag{15}$$

Denote $L := \operatorname{Range}(B_{j'})$. Combining (14) and (15), we get

$$\mathcal{C} \cap L \ni (\hat{x}^k - \bar{x}) \perp (z^k - \bar{z}) \in \mathcal{C}^\circ \cap L^\perp,$$

where $\mathcal{C}^\circ$ is the polar cone of $\mathcal{C}$. Since $\hat{x}^k - \bar{x} \in \mathcal{C}$, we have $\langle \hat{x}^k - \bar{x}, \bar{z} \rangle = 0$, which, together with $\langle \hat{x}^k - \bar{x}, z^k - \bar{z} \rangle = 0$, implies $\langle \hat{x}^k - \bar{x}, z^k \rangle = 0$. Thus, for all k sufficiently large,

$$\langle z^k, \hat{x}^k \rangle = \langle z^k, \bar{x} \rangle = \langle \bar{z} + t_k \xi_2 + o(t_k), \bar{x} \rangle,$$

and it follows that

$$
\begin{aligned}
(p^*)'(\bar{z}; \xi_2) &= \lim_{k \to \infty} \frac{\delta_P^*(\bar{z} + t_k \xi_2) - \delta_P^*(\bar{z})}{t_k} = \lim_{k \to \infty} \frac{\delta_P^*(z^k) - \delta_P^*(\bar{z})}{t_k} \\
&= \lim_{k \to \infty} \frac{\langle z^k, \hat{x}^k \rangle - \langle \bar{z}, \bar{x} \rangle}{t_k} = \langle \bar{x}, \xi_2 \rangle.
\end{aligned}
\tag{16}
$$

By (12), (13), and (16), we know that $(\xi_1, \xi_2) \in \mathcal{C}(\bar{y}, \bar{z})$. Since $\mathcal{C} \cap \mathcal{L}$ is a polyhedral convex cone, we know from [40, Theorem 19.3] that $\mathcal{A}(\mathcal{C} \cap \mathcal{L})$ is also a polyhedral convex cone, which, together with (10), implies

$$(\nabla h^*)'(\bar{y}; \xi_1) \in \mathcal{A}(\mathcal{C} \cap L).$$

Therefore, there exists a vector $\eta \in \mathcal{C} \cap L$ such that

$$\langle \xi_1, (\nabla h^*)'(\bar{y}; \xi_1) \rangle = \langle \xi_1, \mathcal{A}\eta \rangle = -\langle \xi_2, \eta \rangle = 0,$$

where the last equality follows from the fact that $\xi_2 \in \mathcal{C}^\circ \cap L^\perp$. Note that the last inclusion holds since the polyhedral convex cone $\mathcal{C}^\circ \cap L^\perp$ is closed and

$$\xi_2 = \lim_{k \to \infty} \frac{z^k - \bar{z}}{t_k} \in \mathcal{C}^\circ \cap L^\perp.$$

As $0 \neq \xi = (\xi_1, \xi_2) \in \mathcal{C}(\bar{y}, \bar{z})$ but $\langle \xi_1, (\nabla h^*)'(\bar{y}; \xi_1) \rangle = 0$, this contradicts the assumption that the second order sufficient condition holds for (4) at $(\bar{y}, \bar{z})$. Thus, we have proved our claim for Case I.

Case II. There exists a nonempty polyhedral convex set P such that $p^*(z) = \delta_P(z)$ for all $z \in \mathcal{X}$. Then, we know that for each k,

$$z^k = \Pi_P(z^k + \hat{x}^k).$$

Since $(\delta_P)'(\bar{z}; d_2) = 0$ for all $d_2 \in \mathcal{T}_{\text{dom}(p^*)}(\bar{z})$, the critical cone in (7) now takes the following form:

$$\mathcal{C}(\bar{y}, \bar{z}) = \left\{ (d_1, d_2) \in \mathcal{Y} \times \mathcal{X} \mid \mathcal{A}^* d_1 + d_2 = 0, \langle \nabla h^*(\bar{y}), d_1 \rangle = 0, d_2 \in \mathcal{T}_{\text{dom}(p^*)}(\bar{z}) \right\}.$$

Similar to Case I, without loss of generality, we can assume that there exists a subspace L such that for all $k \geq 1$,

$$\mathcal{C} \cap L \ni (z^k - \bar{z}) \perp (\hat{x}^k - \bar{x}) \in \mathcal{C}^\circ \cap L^\perp,$$

where

$$\mathcal{C} \equiv \mathcal{C}_P(\bar{z} + \bar{x}) := \mathcal{T}_P(\bar{z}) \cap \bar{x}^\perp.$$

Since $\mathcal{C} \cap L$ is a polyhedral convex cone, we know

$$\xi_2 = \lim_{k \to \infty} \frac{z^k - \bar{z}}{t_k} \in \mathcal{C} \cap L,$$

and consequently $\langle \xi_2, \bar{x} \rangle = 0$, which, together with (12) and (13), implies $\xi = (\xi_1, \xi_2) \in \mathcal{C}(\bar{y}, \bar{z})$. By (10) and the fact that $\mathcal{A}(\mathcal{C}^\circ \cap \mathcal{L}^\perp)$ is a polyhedral convex cone, we know that

$$(\nabla h^*)'(\bar{y}; \xi_1) \in \mathcal{A}(\mathcal{C}^\circ \cap L^\perp).$$

Therefore, there exists a vector $\eta \in \mathcal{C}^\circ \cap L^\perp$ such that

$$\langle \xi_1, (\nabla h^*)'(\bar{y}; \xi_1) \rangle = \langle \xi_1, \mathcal{A}\eta \rangle = -\langle \xi_2, \eta \rangle = 0.$$

Since $\xi = (\xi_1, \xi_2) \neq 0$, we arrive at a contradiction to the assumed second order sufficient condition. So our claim is also true for this case.

In summary, we have proven that there exists a neighborhood $\mathcal{U}$ of $(\bar{x}, \bar{y}, \bar{z})$ such that for any w close enough to the origin, (9) holds for any solution $(x(w), y(w),$

$z(w)) \in \mathcal{U}$ to the perturbed KKT system (8). Next we show that $\mathcal{T}_l$ is metrically subregular at $(\bar{y}, \bar{z}, \bar{x})$ for the origin.

Define the mapping $\Theta_{KKT} : \mathcal{X} \times \mathcal{Y} \times \mathcal{X} \times \mathcal{W} \to \mathcal{Y} \times \mathcal{X} \times \mathcal{X}$ by

$$\Theta_{KKT}(x, y, z, w) := \begin{pmatrix} \nabla h^*(y) - \mathcal{A}x - u_1 \\ z - \text{Prox}_{p^*}(z + x + u_2) \\ \mathcal{A}^* y + z - c - v \end{pmatrix} \quad \forall (x, y, z, w) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{X} \times \mathcal{W}$$

and define the mapping $\theta : \mathcal{X} \to \mathcal{Y} \times \mathcal{X} \times \mathcal{X}$ as follows:

$$\theta(x) := \Theta_{KKT}(x, \bar{y}, \bar{z}, 0) \quad \forall x \in \mathcal{X}.$$

Then, we have $x \in \mathcal{M}(\bar{y}, \bar{z})$ if and only if $\theta(x) = 0$. Since $\text{Prox}_{p^*}(\cdot)$ is a piecewise affine function, $\theta(\cdot)$ is a piecewise affine function and thus a polyhedral multifunction. By using Proposition 2.2 and shrinking the neighborhood $\mathcal{U}$ if necessary, for any w close enough to the origin and any solution $(x(w), y(w), z(w)) \in \mathcal{U}$ of the perturbed KKT system (8), we have

$$\text{dist}(x(w), \mathcal{M}(\bar{y}, \bar{z})) = O(\|\theta(x(w))\|)$$

$$= O(\|\Theta_{KKT}(x(w), \bar{y}, \bar{z}, 0) - \Theta_{KKT}(x(w), y(w), z(w), w)\|)$$

$$= O(\|w\| + \|(y(w), z(w)) - (\bar{y}, \bar{z})\|),$$

which, together with (9), implies the existence of a constant $\kappa \geq 0$ such that

$$(17) \qquad \|(y(w), z(w)) - (\bar{y}, \bar{z})\| + \text{dist}(x(w), \mathcal{M}(\bar{y}, \bar{z})) \leq \kappa \|w\|.$$

Thus, by Definition 2.3, we have proven that $\mathcal{T}_l$ is metrically subregular at $(\bar{y}, \bar{z}, \bar{x})$ for the origin. The proof of the theorem is complete. $\square$

Remark 2.8. For convex piecewise linear-quadratic programming problems such as the ℓ_1 and elastic net regularized LS problem, we know from J. Sun's thesis [46] on the characterization of the subdifferentials of convex piecewise linear-quadratic functions (see also [43, Proposition 12.30]) that the corresponding operators $\mathcal{T}_l$ and $\mathcal{T}_f$ are polyhedral multifunctions and thus, by Proposition 2.2, the error bound condition holds. Here we emphasize again that the error bound condition holds for the lasso problem (1). Moreover, $\mathcal{T}_f$, associated with the ℓ_1 or elastic net regularized logistic regression model, i.e., for a given vector $b \in \Re^m$, the loss function $h : \Re^m \to \Re$ in problem (3) takes the form

$$h(y) = \sum_{i=1}^{m} \log(1 + e^{-b_i y_i}) \quad \forall y \in \Re^m,$$

also satisfies the error bound condition [33, 47]. Meanwhile, from Theorem 2.7, we know that $\mathcal{T}_l$ corresponding to the ℓ_1 or the elastic net regularized logistic regression model is metrically subregular at any solutions to the KKT system (6) for the origin.

3. An augmented Lagrangian method with asymptotic superlinear convergence. Recall the general convex composite model (3)

$$(\mathbf{P}) \quad \max\left\{-\left(f(x) = h(\mathcal{A}x) - \langle c, x \rangle + p(x)\right)\right\}$$

and its dual

$$(\mathbf{D}) \quad \min\{h^*(y) + p^*(z) \,|\, \mathcal{A}^* y + z = c\}.$$

In this section, we shall propose an asymptotically superlinearly convergent augmented Lagrangian method for solving (**P**) and (**D**). In this section, we make the following standing assumptions regarding the function h.

Assumption 3.1.
(a) $h : \mathcal{Y} \to \Re$ is a convex differentiable function whose gradient is $1/\alpha_h$-Lipschitz continuous, i.e.,

$$\|\nabla h(y') - \nabla h(y)\| \leq (1/\alpha_h)\|y' - y\| \quad \forall y', y \in \mathcal{Y}.$$

(b) h is essentially locally strongly convex [21], i.e., for any compact and convex set $K \subset \operatorname{dom} \partial h$, there exists $\beta_K > 0$ such that

$$(1-\lambda)h(y') + \lambda h(y) \geq h((1-\lambda)y' + \lambda y) + \frac{1}{2}\beta_K \lambda(1-\lambda)\|y' - y\|^2 \quad \forall y', y \in K,$$

for all $\lambda \in [0, 1]$.

Many commonly used loss functions in the machine learning literature satisfy the above mild assumptions. For example, h can be the loss function in the linear regression, logistic regression, and Poisson regression models. While the strict convexity of a convex function is closely related to the differentiability of its conjugate, the essential local strong convexity in Assumption 3.1(b) for a convex function was first proposed in [21] to obtain a characterization of the local Lipschitz continuity of the gradient of its conjugate function.

The aforementioned assumptions on h imply the following useful properties of h^*. First, by [43, Proposition 12.60], we know that h^* is strongly convex with modulus α_h. Second, by [21, Corollary 4.4], we know that h^* is essentially smooth and ∇h^* is locally Lipschitz continuous on $\operatorname{int}(\operatorname{dom} h^*)$. If the solution set to the KKT system associated with (**P**) and (**D**) is further assumed to be nonempty, similar to what we have discussed in the last section, one only needs to focus on $\operatorname{int}(\operatorname{dom} h^*) \times \mathcal{X}$ when solving (**D**). Given $\sigma > 0$, the augmented Lagrangian function associated with (**D**) is given by

$$\mathcal{L}_\sigma(y, z; x) := l(y, z, x) + \frac{\sigma}{2}\|\mathcal{A}^*y + z - c\|^2 \quad \forall (y, z, x) \in \mathcal{Y} \times \mathcal{X} \times \mathcal{X},$$

where the Lagrangian function $l(\cdot)$ is defined in (5).

3.1. SSNAL**: A semismooth Newton augmented Lagrangian algorithm for (D).** The detailed steps of our algorithm SSNAL for solving (**D**) are given as follows. Since a semismooth Newton method will be employed to solve the subproblems involved in this prototype of the inexact augmented Lagrangian method [42], it is natural for us to call our algorithm SSNAL.

Algorithm SSNAL: An inexact augmented Lagrangian method for (D).

Let $\sigma_0 > 0$ be a given parameter. Choose $(y^0, z^0, x^0) \in \operatorname{int}(\operatorname{dom} h^*) \times \operatorname{dom} p^* \times \mathcal{X}$. For $k = 0, 1, \ldots$, perform the following steps in each iteration:
Step 1. Compute

$$(18) \qquad (y^{k+1}, z^{k+1}) \approx \arg\min\{\Psi_k(y, z) := \mathcal{L}_{\sigma_k}(y, z; x^k)\}.$$

Step 2. Compute $x^{k+1} = x^k - \sigma_k(\mathcal{A}^*y^{k+1} + z^{k+1} - c)$ and update $\sigma_{k+1} \uparrow \sigma_\infty \leq \infty$.

Next, we shall adapt the results developed in [41, 42, 34] to establish the global and local superlinear convergence of our algorithm.

Since the inner problems cannot be solved exactly, we use the following standard stopping criterion studied in [41, 42] for approximately solving (18):

$$\text{(A)} \quad \Psi_k(y^{k+1}, z^{k+1}) - \inf \Psi_k \leq \varepsilon_k^2/2\sigma_k, \quad \sum_{k=0}^{\infty} \varepsilon_k < \infty.$$

Now, we can state the global convergence of Algorithm SSNAL from [41, 42] without much difficulty.

THEOREM 3.2. *Suppose that Assumption* 3.1 *holds and that the solution set to* (**P**) *is nonempty. Let* $\{(y^k, z^k, x^k)\}$ *be the infinite sequence generated by Algorithm* SSNAL *with stopping criterion* (A). *Then, the sequence* $\{x^k\}$ *is bounded and converges to an optimal solution of* (**P**). *In addition,* $\{(y^k, z^k)\}$ *is also bounded and converges to the unique optimal solution* $(y^*, z^*) \in \text{int}(\text{dom } h^*) \times \text{dom } p^*$ *of* (**D**).

Proof. Since the solution set to (**P**) is assumed to be nonempty, the optimal value of (**P**) is finite. From Assumption 3.1(a), we have that $\text{dom } h = \mathcal{Y}$ and h^* is strongly convex [43, Proposition 12.60]. Then, by Fenchel's duality theorem [40, Corollary 31.2.1], we know that the solution set to (**D**) is nonempty and the optimal value of (**D**) is finite and equal to the optimal value of (**P**). That is, the solution set to the KKT system associated with (**P**) and (**D**) is nonempty. The uniqueness of the optimal solution $(y^*, z^*) \in \text{int}(\text{dom } h^*) \times \mathcal{X}$ of (**D**) then follows directly from the strong convexity of h^*. By combining this uniqueness with [42, Theorem 4], one can easily obtain the boundedness of $\{(y^k, z^k)\}$ and other desired results readily. $\square$

We need the the following stopping criteria for the local convergence analysis:

$$\text{(B1)} \quad \Psi_k(y^{k+1}, z^{k+1}) - \inf \Psi_k \leq (\delta_k^2/2\sigma_k)\|x^{k+1} - x^k\|^2, \quad \sum_{k=0}^{\infty} \delta_k < +\infty,$$

$$\text{(B2)} \quad \text{dist}(0, \partial\Psi_k(y^{k+1}, z^{k+1})) \leq (\delta_k'/\sigma_k)\|x^{k+1} - x^k\|, \quad 0 \leq \delta_k' \to 0,$$

where $x^{k+1} := x^k + \sigma_k(\mathcal{A}^* y^{k+1} + z^{k+1} - c)$.

THEOREM 3.3. *Assume that Assumption* 3.1 *holds and that the solution set* Ω *to* (**P**) *is nonempty. Suppose that* $\mathcal{T}_f$ *satisfies the error bound condition* (2) *for the origin with modulus* a_f. *Let* $\{(y^k, z^k, x^k)\}$ *be any infinite sequence generated by Algorithm* SSNAL *with stopping criteria* (A) *and* (B1). *Then, the sequence* $\{x^k\}$ *converges to* $x^* \in \Omega$ *and for all* k *sufficiently large,*

$$\text{(19)} \qquad \text{dist}(x^{k+1}, \Omega) \leq \theta_k \text{dist}(x^k, \Omega),$$

where $\theta_k = \left(a_f(a_f^2 + \sigma_k^2)^{-1/2} + 2\delta_k\right)(1 - \delta_k)^{-1} \to \theta_\infty = a_f(a_f^2 + \sigma_\infty^2)^{-1/2} < 1$ *as* $k \to +\infty$. *Moreover, the sequence* $\{(y^k, z^k)\}$ *converges to the unique optimal solution* $(y^*, z^*) \in \text{int}(\text{dom } h^*) \times \text{dom } p^*$ *to* (**D**).

Moreover, if $\mathcal{T}_l$ *is metrically subregular at* (y^*, z^*, x^*) *for the origin with modulus* a_l *and the stopping criterion* (B2) *is also used, then for all* k *sufficiently large,*

$$\|(y^{k+1}, z^{k+1}) - (y^*, z^*)\| \leq \theta_k' \|x^{k+1} - x^k\|,$$

where $\theta_k' = a_l(1 + \delta_k')/\sigma_k$ *with* $\lim_{k \to \infty} \theta_k' = a_l/\sigma_\infty$.

Proof. The first part of the theorem follows from [34, Theorem 2.1], [42, Proposition 7 and Theorem 5], and Theorem 3.2. To prove the second part, we recall that

if $\mathcal{T}_l$ is metrically subregular at (y^*, z^*, x^*) for the origin with the modulus a_l and $(y^k, z^k, x^k) \to (y^*, z^*, x^*)$, then for all k sufficiently large,

$$\|(y^{k+1}, z^{k+1}) - (y^*, z^*)\| + \text{dist}(x^{k+1}, \Omega) \le a_l \, \text{dist}(0, \mathcal{T}_l(y^{k+1}, z^{k+1}, x^{k+1})).$$

Therefore, by the estimate (4.21) in [42] and the stopping criterion (B2), we obtain that for all k sufficiently large,

$$\|(y^{k+1}, z^{k+1}) - (y^*, z^*)\| \le a_l(1 + \delta_k')/\sigma_k \|x^{k+1} - x^k\|.$$

This completes the proof of Theorem 3.3. □

Remark 3.4. In the above theorem, when $\sigma_\infty = +\infty$, the inequality (19) directly implies that $\{x^k\}$ converges Q-superlinearly. Moreover, recent advances in [12] also established the asymptotic R-superlinear convergence of the dual iteration sequence $\{(y^k, z^k)\}$. Indeed, from [12, Proposition 4.1], under the same conditions of Theorem 3.3, we have that for k sufficiently large,

$$(20) \qquad \|(y^{k+1}, z^{k+1}) - (y^*, z^*)\| \le \theta_k' \|x^{k+1} - x^k\| \le \theta_k'(1 - \delta_k)^{-1} \text{dist}(x^k, \ \Omega),$$

where $\theta_k'(1 - \delta_k)^{-1} \to a_l/\sigma_\infty$. Thus, if $\sigma_\infty = \infty$, the Q-superlinear convergence of $\{x^k\}$ and (20) further imply that $\{(y^k, z^k)\}$ converges R-superlinearly.

We should emphasize here that by combining Remarks 2.8 and 3.4 and Theorem 3.3, our Algorithm SSNAL is guaranteed to produce an asymptotically superlinearly convergent sequence when used to solve (**D**) for many commonly used regularizers and loss functions. In particular, the SSNAL algorithm is asymptotically superlinearly convergent when applied to the dual of (1).

3.2. Solving the augmented Lagrangian subproblems. Here we shall propose an efficient semismooth Newton algorithm to solve the inner subproblems in the augmented Lagrangian method (18). That is, for some fixed $\sigma > 0$ and $\tilde{x} \in \mathcal{X}$, we aim to solve

$$(21) \qquad \min_{y,z} \ \Psi(y, z) := \mathcal{L}_\sigma(y, z; \tilde{x}).$$

Since $\Psi(\cdot, \cdot)$ is a strongly convex function, we have that, for any $\alpha \in \Re$, the level set $\mathcal{L}_\alpha := \{(y, z) \in \text{dom}\, h^* \times \text{dom}\, p^* \mid \Psi(y, z) \le \alpha\}$ is a closed and bounded convex set. Moreover, problem (21) admits a unique optimal solution denoted as $(\bar{y}, \bar{z}) \in \text{int}(\text{dom}\, h^*) \times \text{dom}\, p^*$.

Denote, for any $y \in \mathcal{Y}$,

$$\psi(y) := \inf_z \Psi(y, z) = h^*(y) + p^*(\text{Prox}_{p^*/\sigma}(\tilde{x}/\sigma - \mathcal{A}^* y + c))$$

$$+ \frac{1}{2\sigma} \|\text{Prox}_{\sigma p}(\tilde{x} - \sigma(\mathcal{A}^* y - c))\|^2 - \frac{1}{2\sigma} \|\tilde{x}\|^2.$$

Therefore, if $(\bar{y}, \bar{z}) = \arg\min \Psi(y, z)$, then $(\bar{y}, \bar{z}) \in \text{int}(\text{dom}\, h^*) \times \text{dom}\, p^*$ can be computed simultaneously by

$$\bar{y} = \arg\min \psi(y), \quad \bar{z} = \text{Prox}_{p^*/\sigma}(\tilde{x}/\sigma - \mathcal{A}^* \bar{y} + c).$$

Note that $\psi(\cdot)$ is strongly convex and continuously differentiable on $\text{int}(\text{dom}\, h^*)$ with

$$\nabla\psi(y) = \nabla h^*(y) - \mathcal{A}\,\text{Prox}_{\sigma p}(\tilde{x} - \sigma(\mathcal{A}^* y - c)) \quad \forall y \in \text{int}(\text{dom}\, h^*).$$

Thus, $\bar{y}$ can be obtained via solving the following nonsmooth equation:

$$(22) \qquad\qquad \nabla\psi(y) = 0, \quad y \in \text{int}(\text{dom}\, h^*).$$

Let $y \in \text{int}(\text{dom}\, h^*)$ be any given point. Since h^* is a convex function with a locally Lipschitz continuous gradient on $\text{int}(\text{dom}\, h^*)$, the following operator is well defined:

$$\hat{\partial}^2\psi(y) := \partial(\nabla h^*)(y) + \sigma\mathcal{A}\partial\text{Prox}_{\sigma p}(\tilde{x} - \sigma(\mathcal{A}^* y - c))\mathcal{A}^*,$$

where $\partial(\nabla h^*)(y)$ is the Clarke subdifferential of ∇h^* at y [10], and $\partial\text{Prox}_{\sigma p}(\tilde{x} - \sigma(\mathcal{A}^* y - c))$ is the Clarke subdifferential of the Lipschitz continuous mapping $\text{Prox}_{\sigma p}(\cdot)$ at $\tilde{x} - \sigma(\mathcal{A}^* y - c)$. Note that from [10, Proposition 2.3.3 and Theorem 2.6.6], we know that

$$\partial^2\psi(y)\,(d) \subseteq \hat{\partial}^2\psi(y)\,(d) \quad \forall d \in \mathcal{Y},$$

where $\partial^2\psi(y)$ denotes the generalized Hessian of ψ at y. Define

$$(23) \qquad\qquad V := H + \sigma\mathcal{A}U\mathcal{A}^*$$

with $H \in \partial^2 h^*(y)$ and $U \in \partial\text{Prox}_{\sigma p}(\tilde{x} - \sigma(\mathcal{A}^* y - c))$. Then, we have $V \in \hat{\partial}^2\psi(y)$. Since h^* is a strongly convex function, we know that H is symmetric positive definite on $\mathcal{Y}$ and thus V is also symmetric positive definite on $\mathcal{Y}$.

Under the mild assumption that ∇h^* and $\text{Prox}_{\sigma p}$ are strongly semismooth (the definition of which is given next), we can design a superlinearly convergent semismooth Newton method to solve the nonsmooth equation (22).

DEFINITION 3.5 (semismoothness [35, 38, 45]). *Let $F : \mathcal{O} \subseteq \mathcal{X} \to \mathcal{Y}$ be a locally Lipschitz continuous function on the open set $\mathcal{O}$. F is said to be semismooth at $x \in \mathcal{O}$ if F is directionally differentiable at x and for any $V \in \partial F(x + \Delta x)$ with $\Delta x \to 0$,*

$$F(x + \Delta x) - F(x) - V\Delta x = o(\|\Delta x\|).$$

F is said to be strongly semismooth at x if F is semismooth at x and

$$F(x + \Delta x) - F(x) - V\Delta x = O(\|\Delta x\|^2).$$

F is said to be a semismooth (respectively, strongly semismooth) function on $\mathcal{O}$ if it is semismooth (respectively, strongly semismooth) everywhere in $\mathcal{O}$.

Note that it is widely known in the nonsmooth optimization/equation community that continuous piecewise affine functions and twice continuously differentiable functions are all strongly semismooth everywhere. In particular, $\text{Prox}_{\|\cdot\|_1}$, as a Lipschitz continuous piecewise affine function, is strongly semismooth. See [14] for more semismooth and strongly semismooth functions.

Now, we can design a semismooth Newton (Ssn) method to solve (22) as follows and could expect to get a fast superlinear or even quadratic convergence.

Algorithm Ssn: A semismooth Newton algorithm for solving (22) (Ssn($y^0, \tilde{x}, \sigma$)).

Given $\mu \in (0, 1/2)$, $\bar{\eta} \in (0, 1)$, $\tau \in (0, 1]$, and $\delta \in (0, 1)$. Choose $y^0 \in \text{int}(\text{dom } h^*)$. Iterate the following steps for $j = 0, 1, \ldots$.

Step 1. Choose $H_j \in \partial(\nabla h^*)(y^j)$ and $U_j \in \partial\text{Prox}_{\sigma p}(\tilde{x} - \sigma(\mathcal{A}^* y^j - c))$. Let $V_j := H_j + \sigma\mathcal{A}U_j\mathcal{A}^*$. Solve the linear system

$$(24) \qquad\qquad V_j d + \nabla\psi(y^j) = 0$$

exactly or by the conjugate gradient (CG) algorithm to find d^j such that

$$\|V_j d^j + \nabla\psi(y^j)\| \leq \min(\bar{\eta}, \|\nabla\psi(y^j)\|^{1+\tau}).$$

Step 2. (Line search) Set $\alpha_j = \delta^{m_j}$, where m_j is the first nonnegative integer m for which

$$y^j + \delta^m d^j \in \text{int}(\text{dom } h^*) \quad \text{and} \quad \psi(y^j + \delta^m d^j) \leq \psi(y^j) + \mu\delta^m \langle\nabla\psi(y^j), d^j\rangle.$$

Step 3. Set $y^{j+1} = y^j + \alpha_j d^j$.

The convergence results for the above Ssn algorithm are stated in the next theorem.

THEOREM 3.6. *Assume that $\nabla h^*(\cdot)$ and $\text{Prox}_{\sigma p}(\cdot)$ are strongly semismooth on* $\text{int}(\text{dom } h^*)$ *and $\mathcal{X}$, respectively. Let the sequence $\{y^j\}$ be generated by Algorithm* Ssn. *Then $\{y^j\}$ converges to the unique optimal solution $\bar{y} \in \text{int}(\text{dom } h^*)$ of the problem in (22) and*

$$\|y^{j+1} - \bar{y}\| = O(\|y^j - \bar{y}\|^{1+\tau}).$$

Proof. Since, by [54, Proposition 3.3], d^j is a descent direction, Algorithm Ssn is well defined. By (23), we know that for any $j \geq 0$, $V_j \in \hat{\partial}^2\psi(y^j)$. Then, we can prove the conclusion of this theorem by following the proofs of [54, Theorems 3.4 and 3.5]. We omit the details here. □

We shall now discuss the implementations of stopping criteria (A), (B1), and (B2) for Algorithm Ssn to solve the subproblem (18) in Algorithm Ssnal. Note that when Algorithm Ssn is applied to minimize $\Psi_k(\cdot)$ to find

$$y^{k+1} = \text{Ssn}(y^k, x^k, \sigma_k) \quad \text{and} \quad z^{k+1} = \text{Prox}_{p^*/\sigma_k}(x^k/\sigma_k - \mathcal{A}^* y^{k+1} + c),$$

we have, by simple calculations and the strong convexity of h^*, that

$$\Psi_k(y^{k+1}, z^{k+1}) - \inf\Psi_k = \psi_k(y^{k+1}) - \inf\psi_k \leq (1/2\alpha_h)\|\nabla\psi_k(y^{k+1})\|^2$$

and $(\nabla\psi_k(y^{k+1}), 0) \in \partial\Psi_k(y^{k+1}, z^{k+1})$, where $\psi_k(y) := \inf_z \Psi_k(y, z)$ for all $y \in \mathcal{Y}$. Therefore, the stopping criteria (A), (B1), and (B2) can be achieved by the following implementable criteria:

$$(\text{A}') \quad \|\nabla\psi_k(y^{k+1})\| \leq \sqrt{\alpha_h/\sigma_k}\,\varepsilon_k, \quad \sum_{k=0}^{\infty}\varepsilon_k < \infty,$$

$$(\text{B1}') \quad \|\nabla\psi_k(y^{k+1})\| \leq \sqrt{\alpha_h\sigma_k}\,\delta_k\|\mathcal{A}^* y^{k+1} + z^{k+1} - c\|, \quad \sum_{k=0}^{\infty}\delta_k < +\infty,$$

$$(\text{B2}') \quad \|\nabla\psi_k(y^{k+1})\| \leq \delta_k'\|\mathcal{A}^* y^{k+1} + z^{k+1} - c\|, \quad 0 \leq \delta_k' \to 0.$$

That is, the stopping criteria (A), (B1), and (B2) will be satisfied as long as $\|\nabla\psi_k(y^{k+1})\|$ is sufficiently small.

3.3. An efficient implementation of SSN for solving subproblems (18). When Algorithm SSNAL is applied to solve the general convex composite optimization model (**P**), the key part is to use Algorithm SSN to solve the subproblems (18). In this subsection, we shall discuss an efficient implementation of SSN for solving the aforementioned subproblems when the nonsmooth regularizer p is chosen to be $\lambda\|\cdot\|_1$ for some $\lambda > 0$. Clearly, in Algorithm SSN, the most important step is the computation of the search direction d^j from the linear system (24). So we shall first discuss the solving of this linear system.

Let $(\tilde{x}, y) \in \Re^n \times \Re^m$ and $\sigma > 0$ be given. We consider the following Newton linear system:

$$(25) \qquad (H + \sigma AUA^T)d = -\nabla\psi(y),$$

where $H \in \partial(\nabla h^*)(y)$, A denotes the matrix representation of $\mathcal{A}$ with respect to the standard bases of $\Re^n$ and $\Re^m$, and $U \in \partial\mathrm{Prox}_{\sigma\lambda\|\cdot\|_1}(x)$ with $x := \tilde{x} - \sigma(A^Ty - c)$. Since H is a symmetric and positive definite matrix, (25) can be equivalently rewritten as

$$\left(I_m + \sigma(L^{-1}A)U(L^{-1}A)^T\right)(L^Td) = -L^{-1}\nabla\psi(y),$$

where L is a nonsingular matrix obtained from the (sparse) Cholesky decomposition of H such that $H = LL^T$. In many applications, H is usually a sparse matrix. Indeed, when the function h in the primal objective is taken to be the squared loss or the logistic loss functions, the resulting matrices H are in fact diagonal matrices. That is, the costs of computing L and its inverse are negligible in most situations. Therefore, without loss of generality, we can consider a simplified version of (25) as follows:

$$(26) \qquad (I_m + \sigma AUA^T)d = -\nabla\psi(y),$$

which is precisely the Newton system associated with the standard lasso problem (1). Since $U \in \Re^{n\times n}$ is a diagonal matrix, at first glance, the costs of computing AUA^T and the matrix-vector multiplication AUA^Td for a given vector $d \in \Re^m$ are $\mathcal{O}(m^2n)$ and $\mathcal{O}(mn)$, respectively. These computational costs are too expensive when the dimensions of A are large and can make the commonly employed approaches such as the Cholesky factorization and the conjugate gradient method inappropriate for solving (26). Fortunately, under the sparse optimization setting, if the sparsity of U is wisely taken into the consideration, one can substantially reduce these unfavorable computational costs to a level such that they are negligible or at least insignificant compared to other costs. Next, we shall show how this can be done by taking full advantage of the sparsity of U. This sparsity will be referred to as the second order sparsity of the underlying problem.

For $x = \tilde{x} - \sigma(A^Ty - c)$, in our computations, we can always choose $U = \mathrm{Diag}(u)$, the diagonal matrix whose ith diagonal element is given by u_i with

$$u_i = \begin{cases} 0 & \text{if} \quad |x_i| \le \sigma\lambda, \\ 1 & \text{otherwise,} \end{cases} \qquad i = 1, \dots, n.$$

Since $\mathrm{Prox}_{\sigma\lambda\|\cdot\|_1}(x) = \mathrm{sign}(x) \circ \max\{|x| - \sigma\lambda, 0\}$, it is not difficult to see that $U \in \partial\mathrm{Prox}_{\sigma\lambda\|\cdot\|_1}(x)$. Let $\mathcal{J} := \{j \mid |x_j| > \sigma\lambda, j = 1, \dots, n\}$ and $r = |\mathcal{J}|$, the cardinality of $\mathcal{J}$. By taking the special 0-1 structure of U into consideration, we have that
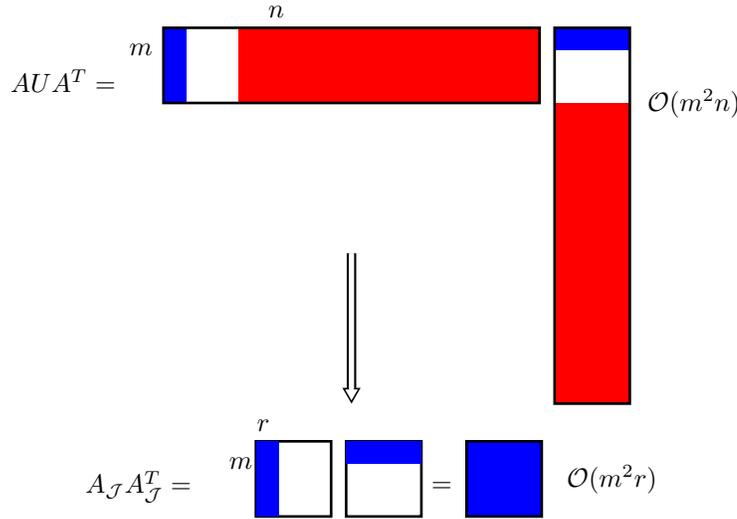
FIG. 1. *Reducing the computational costs from $\mathcal{O}(m^2 n)$ to $\mathcal{O}(m^2 r)$.*

$$(27) \qquad AUA^T = (AU)(AU)^T = A_{\mathcal{J}} A_{\mathcal{J}}^T,$$

where $A_{\mathcal{J}} \in \Re^{m \times r}$ is the submatrix of A with those columns not in $\mathcal{J}$ being removed from A. Then, by using (27), we know that now the costs of computing AUA^T and $AUA^T d$ for a given vector d are reduced to $\mathcal{O}(m^2 r)$ and $\mathcal{O}(mr)$, respectively. Due to the sparsity-promoting property of the regularizer p, the number r is usually much smaller than n. Thus, by exploring the aforementioned second order sparsity, we can greatly reduce the computational costs in solving the linear system (26) when using the Cholesky factorization. More specifically, the total computational costs of using the Cholesky factorization to solve the linear system are reduced from $\mathcal{O}(m^2(m+n))$ to $\mathcal{O}(m^2(m+r))$. See Figure 1 for an illustration on the reduction. This means that even if n happens to be extremely large (say, larger than 10^7), one can still solve the Newton linear system (26) efficiently via the Cholesky factorization as long as both m and r are moderate (say, less than 10^4). If, in addition, $r \ll m$, which is often the case when m is large and the optimal solutions to the underlying problem are sparse, instead of factorizing an $m \times m$ matrix, we can make use of the Sherman–Morrison–Woodbury formula [22] to get the inverse of $I_m + \sigma AUA^T$ by inverting a much smaller $r \times r$ matrix as follows:

$$(I_m + \sigma AUA^T)^{-1} = (I_m + \sigma A_{\mathcal{J}} A_{\mathcal{J}}^T)^{-1} = I_m - A_{\mathcal{J}}(\sigma^{-1} I_r + A_{\mathcal{J}}^T A_{\mathcal{J}})^{-1} A_{\mathcal{J}}^T.$$

See Figure 2 for an illustration on the computation of $A_{\mathcal{J}}^T A_{\mathcal{J}}$. In this case, the total computational costs for solving the Newton linear system (26) are reduced significantly further from $\mathcal{O}(m^2(m+r))$ to $\mathcal{O}(r^2(m+r))$. We should emphasize here that this dramatic reduction in computational costs results from the wise combination of the careful examination of the existing second order sparsity in the lasso-type problems and some "smart" numerical linear algebra.

From the above arguments, we can see that as long as the number of the nonzero components of $\text{Prox}_{\sigma\lambda\|\cdot\|_1}(x)$ is small, say less than $\sqrt{n}$, and $H_j \in \partial(\nabla h^*)(y^j)$ is a sparse matrix, e.g., a diagonal matrix, we can always solve the linear system (24) at very low costs. In particular, this is true for the lasso problems admitting sparse

$$A_{\mathcal{J}}^T A_{\mathcal{J}} = \quad = \quad \mathcal{O}(r^2 m)$$

FIG. 2. *Further reducing the computational costs to $\mathcal{O}(r^2 m)$.*

solutions. Similar discussions on the reduction of the computational costs can also be conducted for the case when the conjugate gradient method is applied to solve the linear systems (24). Note that one may argue that even if the original problem has only sparse solutions, at certain stages, one may still encounter the situation that the number of the nonzero components of $\text{Prox}_{\sigma\lambda\|\cdot\|_1}(x)$ is large. Our answer to this question is simple. First, this phenomenon rarely occurs in practice since we always start with a sparse feasible point, e.g., the zero vector. Second, even if at certain steps this phenomenon does occur, we just need to apply a small number of conjugate gradient iterations to the linear system (24) as in this case the parameter σ is normally small and the current point is far away from any sparse optimal solution. In summary, we have demonstrated how Algorithm Ssn can be implemented efficiently for solving sparse optimization problems of the form (18) with $p(\cdot)$ being chosen to be $\lambda\|\cdot\|_1$.

4. Numerical experiments for lasso problems. In this section, we shall evaluate the performance of our algorithm Ssnal for solving large-scale lasso problems (1). We note that the relative performance of most of the existing algorithms mentioned in the introduction has recently been well documented in the two recent papers [18, 36], which appears to suggest that for some large-scale sparse reconstruction problems, mfIPM[1] and FPC_AS[2] have mostly outperformed the other solvers. Hence, in this section we will compare our algorithm with these two popular solvers. Note that mfIPM is a specialized interior-point based second order method designed for the lasso problem (1), whereas FPC_AS is a first order method based on forward-backward operator splitting. Moreover, we also report the numerical performance of two commonly used algorithms for solving lasso problems: the accelerated proximal gradient (APG) algorithm [37, 2] as implemented by Liu, Ji, and Ye in SLEP[3] [32] and the alternating direction method of multipliers (ADMM) [19, 20]. For the purpose of comparisons, we also test the linearized ADMM (LADMM) [55]. We have implemented both ADMM and LADMM in MATLAB with the step length set to be 1.618. Although the existing solvers can perform impressively well on some easy-to-solve sparse reconstruction problems, as one will see later, they lack the ability to efficiently solve difficult problems such as the large-scale regression problems when the data $\mathcal{A}$ is badly conditioned.

For testing purposes, the regularization parameter λ in the lasso problem (1) is chosen as

$$\lambda = \lambda_c \|\mathcal{A}^* b\|_\infty,$$

where $0 < \lambda_c < 1$. In our numerical experiments, we measure the accuracy of an approximate optimal solution $\tilde{x}$ for (1) by using the following relative KKT residual:

$$\eta = \frac{\|\tilde{x} - \text{prox}_{\lambda\|\cdot\|_1}(\tilde{x} - \mathcal{A}^*(\mathcal{A}\tilde{x} - b))\|}{1 + \|\tilde{x}\| + \|\mathcal{A}\tilde{x} - b\|}.$$

---

[1]http://www.maths.ed.ac.uk/ERGO/mfipmcs/
[2]http://www.caam.rice.edu/~optimization/L1/FPC_AS/
[3]http://yelab.net/software/SLEP/

For a given tolerance $\epsilon > 0$, we will stop the tested algorithms when $\eta < \epsilon$. For all the tests in this section, we set $\epsilon = 10^{-6}$. The algorithms will also be stopped when they reach the maximum number of iterations (1000 iterations for our algorithm and mfIPM, and 20000 iterations for FPC_AS, APG, ADMM, and LADMM) or the maximum computation time of 7 hours. All the parameters for mfIPM, FPC_AS, and APG are set to the default values. All our computational results are obtained by running MATLAB (version 8.4) on a Windows workstation (16-core, Intel Xeon E5-2650 @ 2.60GHz, 64 G RAM).

4.1. Numerical results for large-scale regression problems. In this subsection, we test all the algorithms with the test instances $(\mathcal{A}, b)$ obtained from large-scale regression problems in the LIBSVM data sets [8]. These data sets are collected from 10-K Corpus [28] and the UCI data repository [30]. For computational efficiency, zero columns in A (the matrix representation of $\mathcal{A}$) are removed. As suggested in [23], for the data sets **pyrim, triazines, abalone, bodyfat, housing, mpg, and space_ga**, we expand their original features by using polynomial basis functions over those features. For example, the last digit in **pyrim5** indicates that an order 5 polynomial is used to generate the basis functions. This naming convention is also used in the rest of the expanded data sets. These test instances, shown in Table 1, can be quite difficult in terms of the problem dimensions and the largest eigenvalue of $\mathcal{A}\mathcal{A}^*$, which is denoted as $\lambda_{\max}(\mathcal{A}\mathcal{A}^*)$.

Table 2 reports the detailed numerical results for SSNAL, mfIPM, FPC_AS, APG, LADMM, and ADMM in solving large-scale regression problems. In the table, m denotes the number of samples, n denotes the number of features, and "nnz" denotes the number of nonzeros in the solution x obtained by SSNAL using the following estimation:

$$\text{nnz} := \min \left\{ k \mid \sum_{i=1}^{k} |\hat{x}_i| \geq 0.999 \|x\|_1 \right\},$$

where $\hat{x}$ is obtained by sorting x such that $|\hat{x}_1| \geq \cdots \geq |\hat{x}_n|$. An entry marked as "**Error**" indicates that the algorithm breaks down due to some internal errors. The computation time is in the format of "hours:minutes:seconds," and the entry "00" in the column means that the computation time is less than 0.5 second.

One can observe from Table 2 that all the tested first order algorithms except ADMM, i.e., FPC_AS, APG, and LADMM fail to solve most of the test instances

TABLE 1
Statistics of the UCI test instances.

Probname	$m; n$	$\lambda_{\max}(\mathcal{A}\mathcal{A}^*)$
E2006.train	16087;150348	1.91e+05
log1p.E2006.train	16087;4265669	5.86e+07
E2006.test	3308;72812	4.79e+04
log1p.E2006.test	3308;1771946	1.46e+07
pyrim5	74;169911	1.22e+06
triazines4	186;557845	2.07e+07
abalone7	4177;6435	5.21e+05
bodyfat7	252;116280	5.29e+04
housing7	506;77520	3.28e+05
mpg7	392;3432	1.28e+04
space_ga9	3107;5005	4.01e+03

TABLE 2

The performance of SSNAL, mfIPM, FPC_AS, APG, LADMM, and ADMM on 11 selected regression problems (accuracy $\epsilon = 10^{-6}$). In the table, a = SSNAL, b = mfIPM, c = FPC_AS, d = APG, e = LADMM, and f = ADMM.

Probname $m;n$	λ_c	nnz	η a | b | c | d | e | f	Time a | b | c | d | e | f
E2006.train 16087;150348	10^{-3}	1	1.3-7 | 3.9-7 | 1.4-3 | 1.1-8 | 5.1-14 | 9.1-7	01 | 08 | 1:24:18 | 01 | 04 | 10:41
	10^{-4}	1	3.7-7 | 1.6-9 | 9.9-4 | 1.5-7 | 2.6-14 | 7.3-7	01 | 10 | 1:27:29 | 01 | 05 | 10:53
log1p.E2006.train 16087;4265669	10^{-3}	5	5.6-7 | 8.8-1 | 2.6-1 | 1.5-5 | 6.4-4 | 9.9-7	19 | 3:06:48 | 7:00:01 | 2:08:35 | 2:25:25 | 33:47
	10^{-4}	599	4.4-7 | 7.9-1 | 1.7-1 | 1.5-4 | 6.3-3 | 9.8-7	51 | 3:05:54 | 7:00:00 | 2:12:47 | 2:29:59 | 35:21
E2006.test 3308;72812	10^{-3}	1	1.6-9 | 4.8-7 | 3.7-4 | 5.5-8 | 4.0-14 | 6.7-7	00 | 02 | 19:25 | 00 | 01 | 18
	10^{-4}	1	2.1-10 | 4.7-7 | 2.7-4 | 3.7-7 | 2.9-10 | 6.3-7	00 | 02 | 20:21 | 00 | 01 | 21
log1p.E2006.test 3308;1771946	10^{-3}	8	9.2-7 | 3.6-7 | 6.9-1 | 1.2-4 | 9.9-7 | 9.9-7	09 | 13:36 | 4:56:49 | 39:55 | 11:52 | 3:10
	10^{-4}	1081	2.2-7 | 9.6-7 | 8.7-1 | 3.9-4 | 9.9-7 | 9.9-7	18 | 2:42:27 | 4:48:29 | 42:37 | 32:04 | 2:28
pyrim5 74;169911	10^{-3}	72	9.9-7 | 9.3-7 | 8.8-1 | 9.2-4 | 4.0-4 | 3.2-5	05 | 16:58 | 1:12:55 | 7:59 | 9:09 | 21:14
	10^{-4}	78	7.1-7 | 3.1-7 | 9.8-1 | 3.5-3 | 3.7-3 | 1.6-3	09 | 32:02 | 54:20 | 8:27 | 9:56 | 21:34
triazines4 186;557845	10^{-3}	519	8.4-7 | 8.8-1 | 9.6-1 | 2.3-3 | 2.9-3 | 3.2-4	30 | 39:20 | 7:00:05 | 1:00:03 | 1:01:59 | 2:13:00
	10^{-4}	260	9.9-7 | 8.6-1 | 9.8-1 | 1.1-2 | 1.6-2 | 1.1-3	1:35 | 37:08 | 7:47:03 | 1:06:50 | 1:05:21 | 1:55:16
abalone7 4177;6435	10^{-3}	24	5.7-7 | 3.5-7 | **Error** | 5.3-5 | 4.9-6 | 9.9-7	02 | 49 | **Error** | 10:38 | 18:04 | 9:52
	10^{-4}	59	3.7-7 | 4.7-7 | **Error** | 3.9-3 | 3.3-5 | 9.9-7	03 | 3:19 | **Error** | 10:43 | 13:58 | 9:36
bodyfat7 252;116280	10^{-3}	2	3.8-8 | 4.0-9 | 3.6-1 | 8.8-7 | 9.0-7 | 9.9-7	02 | 1:29 | 1:12:02 | 4:00 | 3:08 | 1:49
	10^{-4}	3	4.6-8 | 3.4-7 | 1.9-1 | 3.3-5 | 9.8-7 | 9.9-7	03 | 2:41 | 1:13:08 | 12:16 | 4:19 | 4:05
housing7 506;77520	10^{-3}	158	2.3-7 | 8.1-1 | 8.4-1 | 2.6-4 | 1.7-4 | 9.9-7	04 | 5:13:19 | 1:41:01 | 16:52 | 20:18 | 22:12
	10^{-4}	281	7.7-7 | 8.5-1 | 4.2-1 | 5.3-3 | 5.9-4 | 6.6-6	08 | 5:00:09 | 1:39:36 | 17:04 | 20:53 | 42:36
mpg7 392;3432	10^{-3}	47	2.0-8 | 6.5-7 | **Error** | 1.9-6 | 9.9-7 | 9.8-7	00 | 04 | **Error** | 38 | 14 | 07
	10^{-4}	128	4.1-7 | 8.5-1 | 4.5-1 | 1.2-4 | 1.3-6 | 9.9-7	00 | 13 | 4:09 | 39 | 1:01 | 11
space_ga9 3107;5005	10^{-3}	14	4.8-7 | 2.6-7 | 1.5-2 | 2.5-7 | 9.9-7 | 9.9-7	01 | 16 | 30:55 | 1:46 | 42 | 37
	10^{-4}	38	3.7-7 | 3.0-7 | 4.0-2 | 1.8-5 | 9.9-7 | 9.9-7	01 | 41 | 30:26 | 6:46 | 2:20 | 56

to the required accuracy after 20000 iterations or 7 hours. In particular, FPC_AS fails to produce a reasonably accurate solution for all the test instances. In fact, for 3 test instances, it breaks down due to some internal errors. This poor performance indicates that these first order methods cannot obtain reasonably accurate solutions when dealing with difficult large-scale problems. While ADMM can solve most of the test instances, it needs much more time than SSNAL. For example, for the instance **housing7** with $\lambda_c = 10^{-3}$, we can see that SSNAL is at least 330 times faster than ADMM. In addition, SSNAL can solve the instance **triazines4** in 30 seconds while ADMM reaches the maximum of 20000 iterations and consumes about 2 hours but only produces a rather inaccurate solution.

On the other hand, one can observe that the two second order information based methods SSNAL and mfIPM perform quite robustly despite the huge dimensions and the possibly badly conditioned data sets. More specifically, SSNAL is able to solve the instance **log1p.E2006.train** with approximately 4.3 million features in 20 seconds ($\lambda_c = 10^{-3}$). Among these two algorithms, clearly, SSNAL is far more efficient than the specialized interior-point method mfIPM for all the test instances, especially for large-scale problems where the factor can be up to 300 times faster. While SSNAL can solve all the instances to the desired accuracy, mfIPM fails on 6 instances (with high-dimensional features) out of 22. We also note that mfIPM can only reach a solution with the accuracy of 10^{-1} when it fails to compute the corresponding Newton directions. These facts indicate that the nonsmooth approach employed by SSNAL is far superior compared to the interior point method in exploiting the sparsity in the generalized Hessian. The superior numerical performance of SSNAL indicates that it is a robust, high-performance solver for high-dimensional lasso problems.

As pointed out by one referee, the polynomial expansion in our data-processing step may affect the scaling of the problems. Since first order solvers are not affine invariant, this scaling may affect their performance. Hence, we normalize each column of the matrix A (the matrix representation of $\mathcal{A}$) to have unit norm and correspondingly change the variables. This scaling step also changes the regularization parameter λ accordingly to a nonuniform weight vector. Since it is not easy to call FPC_AS when λ is not a scalar, based on the recommendation of the referee we use another popular active-set based solver PSSas[4] [44] to replace FPC_AS for the testing. All the parameters for PSSas are set to the default values. In our tests, PSSas will be stopped when it reaches the maximum number of 20000 iterations or the maximum computation time of 7 hours. We note that by default PSSas will terminate when its progress is smaller than the threshold 10^{-9}.

The detailed numerical results for SSNAL, mfIPM, PSSas, APG, LADMM, and ADMM with the normalization step in solving the large-scale regression problems are listed in Table 3. From Table 3, one can easily observe that the simple normalization technique does not change the conclusions based on Table 2. The performance of SSNAL is generally invariant with respect to the scaling and SSNAL is still much faster and more robust than other solvers. Meanwhile, after the normalization, mfIPM can solve 2 more instances to the required accuracy. On the other hand, APG and the ADMM type of solvers (i.e., LADMM and ADMM) perform worse than the unscaled case. Furthermore PSSas can only solve 6 out of 22 instances to the required accuracy. In fact, PSSas fails on all the test instances when $\lambda_c = 10^{-4}$. For the instance **triazines4**, it consumes about 6 hours but only generates a poor solution with $\eta = 3.2 \times 10^{-3}$. (Actually, we also run PSSas on these test instances without scaling

---

[4]https://www.cs.ubc.ca/schmidtm/Software/thesis.html

TABLE 3

The performance of SSNAL, *mfIPM, PSSas, APG, LADMM, and ADMM on 11 selected regression problems with* **scaling** *(accuracy* $\epsilon = 10^{-6}$*).*
In the table, a = SSNAL, *b = mfIPM, c1 = PSSas, d = APG, e = LADMM, and f = ADMM.*

Probname $m;n$	λ_c	mnz	η a | b | c1 | d | e | f	Time a | b | c1 | d | e | f
E2006.train 16087;150348	10^{-3}	1	1.3-7 | 3.9-7 | 4.6-12 | 1.8-7 | 1.4-4 | 9.9-7	01 | 1:31 | 01 | 6:32 | 21:12 | 1:18:41
	10^{-4}	1	3.6-9 | 5.5-7 | 1.1-3 | 9.1-7 | 1.2-4 | 2.9-4	02 | 2:27 | 05 | 12:01 | 19:45 | 3:19:48
log1p.E2006.train 16087;4265669	10^{-3}	5	2.6-7 | 3.6-7 | 1.1-6 | 5.5-3 | 1.2-3 | 3.3-6	32 | 59:46 | 1:34:36 | 2:33:27 | 3:04:02 | 7:00:02
	10^{-4}	599	9.0-7 | 2.3-7 | 3.3-5 | 7.9-3 | 5.4-3 | 3.4-4	2:23 | 2:42:21 | 6:06:39 | 2:48:29 | 2:58:27 | 7:00:01
E2006.test 3308;72812	10^{-3}	1	1.6-7 | 3.1-7 | 6.5-10 | 4.4-7 | 2.9-3 | 4.5-2	00 | 32 | 00 | 01 | 4:50 | 13:53
	10^{-4}	1	4.2-7 | 7.7-7 | 1.8-3 | 7.6-7 | 3.7-2 | 9.2-1	00 | 17:02 | 01 | 3:28 | 4:29 | 13:18
log1p.E2006.test 3308;1771946	10^{-3}	8	1.4-7 | 5.1-7 | 4.8-7 | 9.4-4 | 2.4-5 | 1.3-5	13 | 11:29 | 59:42 | 48:35 | 1:08:33 | 1:51:59
	10^{-4}	1081	3.3-7 | 8.8-7 | 1.0-5 | 3.9-3 | 1.0-3 | 4.7-4	1:02 | 27:08 | 3:09:16 | 51:58 | 1:01:08 | 1:36:08
pyrim5 74;169911	10^{-3}	70	3.5-7 | 8.1-7 | 4.1-3 | 2.6-3 | 1.0-3 | 1.2-4	05 | 10:42 | 12:43 | 7:57 | 12:26 | 25:50
	10^{-4}	78	5.3-7 | 7.5-7 | 7.4-2 | 6.1-3 | 5.8-3 | 3.2-3	06 | 51:24 | 30:13 | 8:41 | 12:55 | 20:01
triazines4 186;557845	10^{-3}	567	8.9-7 | 7.8-1 | 2.8-4 | 2.3-3 | 1.3-3 | 1.3-4	28 | 43:57 | 1:57:38 | 1:00:21 | 1:07:16 | 2:23:08
	10^{-4}	261	9.9-7 | 8.6-1 | 3.2-3 | 1.3-2 | 2.6-2 | 2.1-2	1:23 | 42:31 | 6:03:12 | 1:09:03 | 1:09:32 | 2:15:12
abalone7 4177;6435	10^{-3}	24	8.4-7 | 1.6-7 | 1.5-7 | 2.0-3 | 1.5-4 | 1.8-6	03 | 2:22 | 2:12 | 13:23 | 13:20 | 43:48
	10^{-4}	59	3.7-7 | 9.2-7 | 1.8-1 | 2.4-2 | 5.8-2 | 9.9-7	04 | 10:37 | 13:12 | 14:03 | 12:55 | 16:15
bodyfat7 252;116280	10^{-3}	2	1.2-8 | 1.1-7 | 2.1-15 | 1.8-2 | 2.0-1 | 9.9-7	02 | 2:01 | 2:35 | 14:05 | 16:52 | 12:32
	10^{-4}	3	7.0-8 | 1.2-1 | 1.9-5 | 3.5-2 | 3.5-1 | 9.9-7	03 | 9:33 | 14:26 | 14:29 | 16:29 | 27:11
housing7 506;77520	10^{-3}	158	8.8-7 | 8.5-7 | 8.8-7 | 2.2-4 | 5.5-5 | 9.9-7	03 | 6:02 | 8:42 | 18:26 | 22:44 | 28:05
	10^{-4}	281	8.0-7 | 5.9-1 | 3.7-5 | 6.0-3 | 3.0-3 | 1.2-5	05 | 4:50:04 | 29:16 | 21:19 | 21:55 | 53:23
mpg7 392;3432	10^{-3}	47	5.2-7 | 7.3-7 | 3.3-6 | 2.7-4 | 9.9-7 | 9.9-7	00 | 07 | 16 | 40 | 48 | 12
	10^{-4}	128	6.2-7 | 3.2-7 | 2.6-6 | 1.5-3 | 3.0-4 | 9.9-7	00 | 20 | 57 | 41 | 48 | 10
space-ga9 3107;5005	10^{-3}	14	4.3-7 | 7.9-7 | 1.7-5 | 4.5-5 | 1.1-4 | 9.9-7	01 | 39 | 04 | 7:30 | 7:37 | 3:44
	10^{-4}	38	1.6-7 | 8.7-7 | 3.1-5 | 1.7-3 | 8.6-4 | 9.9-7	01 | 1:31 | 1:16 | 7:34 | 7:03 | 7:59

and obtain similar performances. Detailed results are omitted to conserve space.) Therefore, we can safely conclude that the simple normalization technique employed here may not be suitable for general first order methods.

4.2. Numerical results for Sparco collection. In this subsection, the test instances $(\mathcal{A}, b)$ are taken from 8 real-valued sparse reconstruction problems in the Sparco collection [5]. For testing purposes, we introduce a 60dB noise to b (as in [18]) by using the MATLAB command `b = awgn(b,60,'measured')`. For these test instances, the matrix representations of the linear maps $\mathcal{A}$ are not available. Hence, ADMM will not be tested in this subsection, as it will be extremely expensive, if not impossible, to compute and factorize $\mathcal{I} + \sigma\mathcal{A}\mathcal{A}^*$.

In Table 4, we report the detailed computational results obtained by SSNAL, mfIPM, FPC_AS, APG, and LADMM in solving two large-scale instances **srcsep1** and **srcsep2** in the Sparco collection. Here, we test five choices of λ_c, i.e., $\lambda_c = 10^{-0.8}, 10^{-1}, 10^{-1.2}, 10^{-1.5}, 10^{-2}$. As one can observe, as λ_c decreases, the number of nonzeros (nnz) in the computed solution increases. In the table, we list some statistics of the test instances, including the problem dimensions (m, n) and the largest eigenvalue of $\mathcal{A}\mathcal{A}^*$ ($\lambda_{\max}(\mathcal{A}\mathcal{A}^*)$). For all the tested algorithms, we present the iteration counts, the relative KKT residuals as well as the computation times (in the format of hours:minutes:seconds). One can observe from Table 4 that all the algorithms perform very well for these test instances with small Lipschitz constants $\lambda_{\max}(\mathcal{A}\mathcal{A}^*)$. As such, SSNAL does not have a clear advantage, as shown in Table 2. Moreover, since the matrix representations of the linear maps $\mathcal{A}$ and $\mathcal{A}^*$ involved are not stored explicitly (i.e., these linear maps can only be regarded as black-box functions), the second order sparsity can hardly be fully exploited. Nevertheless, our Algorithm SSNAL is generally faster than mfIPM, APG, and LADMM, while it is comparable with the fastest algorithm, FPC_AS.

In Table 5, we report the numerical results obtained by SSNAL, mfIPM, FPC_AS, APG, and LADMM in solving various instances of the lasso problem (1). For simplicity, we only test two cases with $\lambda_c = 10^{-3}$ and 10^{-4}. We can observe that FPC_AS

TABLE 4
The performance of SSNAL, *mfIPM, FPC_AS, APG, and LADMM on srcsep1 and srcsep2 (accuracy $\epsilon = 10^{-6}$, noise 60dB). In the table, $a =$ SSNAL, $b = mfIPM$, $c = FPC\_AS$, $d = APG$, and $e = LADMM$.*

λ_c; nnz	Iteration a | b | c | d | e	η a | b | c | d | e	Time a | b | c | d | e
		srcsep1, $m = 29166$, $n = 57344$, $\lambda_{\max}(\mathcal{A}\mathcal{A}^*) = 3.56$	
0.16 ; 380	14 | 28 | 42 | 401 | 89	6.4-7| 9.0-7 | 2.7-8 | 5.3-7 | 9.5-7	07 | 09 | 05 | 15 | 07
0.10 ; 726	16 | 38 | 42 | 574 | 161	4.7-7| 5.9-7 | 2.1-8 | 2.8-7 | 9.6-7	11 | 15 | 05 | 22 | 13
0.06 ; 1402	19 | 41 | 63 | 801 | 393	1.5-7| 1.5-7 | 5.4-8 | 6.3-7 | 9.9-7	18 | 18 | 10 | 30 | 32
0.03 ; 2899	19 | 56 | 110 | 901 | 337	2.7-7| 9.4-7 | 7.3-8 | 9.3-7 | 9.9-7	28 | 53 | 16 | 33 | 28
0.01 ; 6538	17 | 88 | 223 | 1401 | 542	7.1-7| 5.7-7 | 1.1-7 | 9.9-7 | 9.9-7	1:21 | 2:15 | 34 | 53 | 45
		srcsep2, $m = 29166$, $n = 86016$, $\lambda_{\max}(\mathcal{A}\mathcal{A}^*) = 4.95$	
0.16 ; 423	15 | 29 | 42 | 501 | 127	3.2-7| 2.1-7 | 1.9-8 | 4.9-7 | 9.4-7	14 | 13 | 07 | 25 | 15
0.10 ; 805	16 | 37 | 84 | 601 | 212	8.8-7| 9.2-7 | 2.6-8 | 9.6-7 | 9.7-7	21 | 19 | 16 | 30 | 26
0.06 ; 1549	19 | 40 | 84 | 901 | 419	1.4-7| 3.1-7 | 5.4-8 | 4.7-7 | 9.9-7	32 | 28 | 18 | 44 | 50
0.03 ; 3254	20 | 69 | 128 | 901 | 488	1.3-7| 6.6-7 | 8.9-7 | 9.4-7 | 9.9-7	1:06 | 1:33 | 26 | 44 | 59
0.01 ; 7400	21 | 94 | 259 | 2201 | 837	8.8-7| 4.0-7 | 9.9-7 | 8.8-7 | 9.3-7	1:42 | 5:33 | 59 | 2:05 | 1:43

TABLE 5
The performance of SSNAL, *mfIPM, FPC_AS, APG, and LADMM on* 8 *selected sparco problems* (*accuracy* $\epsilon = 10^{-6}$, *noise* 60dB). *In the table,* $a =$ SSNAL, $b =$ *mfIPM,* $c =$ *FPC_AS,* $d =$ *APG, and* $e =$ *LADMM.*

probname $m;n$	λ_c	nnz	η a | b | c | d | e	Time a | b | c | d | e
blknheavi	10^{-3}	12	5.7-7 | 9.2-7 | *1.3-1* | *2.0-6* | 9.9-7	01 | 01 | 55 | 08 | 06
1024;1024	10^{-4}	12	9.2-8 | 8.7-7 | *4.6-3* | *8.5-5* | 9.9-7	01 | 01 | 49 | 07 | 07
srcsep1	10^{-3}	14066	1.6-7 | 7.3-7 | 9.7-7 | 8.7-7 | 9.7-7	5:41 | 42:34 | 13:25 | 1:56 | 4:16
29166;57344	10^{-4}	19306	9.8-7 | 9.5-7 | 9.9-7 | 9.9-7 | 9.5-7	9:28 | 3:31:08 | 32:28 | 2:50 | 13:06
srcsep2	10^{-3}	16502	3.9-7 | 6.8-7 | 9.9-7 | 9.7-7 | 9.8-7	9:51 | 1:01:10 | 16:27 | 2:57 | 8:49
29166;86016	10^{-4}	22315	7.9-7 | 9.5-7 | *1.0-3* | 9.6-7 | 9.5-7	19:14 | 6:40:21 | 2:01:06 | 4:56 | 16:01
srcsep3	10^{-3}	27314	6.1-7 | 9.6-7 | 9.9-7 | 9.6-7 | 9.9-7	33 | 6:24 | 8:51 | 47 | 49
196608;196608	10^{-4}	83785	9.7-7 | 9.9-7 | 9.9-7 | 9.9-7 | 9.9-7	2:03 | 1:42:59 | 3:40 | 1:15 | 3:06
soccer1	10^{-3}	4	1.8-7 | 6.3-7 | *5.2-1* | 8.4-7 | 9.9-7	01 | 03 | 13:51 | 2:35 | 02
3200;4096	10^{-4}	8	8.7-7 | 4.3-7 | *5.2-1* | *3.3-6* | 9.6-7	01 | 02 | 13:23 | 3:07 | 02
soccer2	10^{-3}	4	3.4-7 | 6.3-7 | *5.0-1* | 8.2-7 | 9.9-7	00 | 03 | 13:46 | 1:40 | 02
3200;4096	10^{-4}	8	2.1-7 | 1.4-7 | *6.8-1* | *1.8-6* | 9.1-7	01 | 03 | 13:27 | 3:07 | 02
blurrycam	10^{-3}	1694	1.9-7 | 6.5-7 | 3.6-8 | 4.1-7 | 9.4-7	03 | 09 | 03 | 02 | 07
65536;65536	10^{-4}	5630	1.0-7 | 9.7-7 | 1.3-7 | 9.7-7 | 9.9-7	05 | 1:35 | 08 | 03 | 29
blurspike	10^{-3}	1954	3.1-7 | 9.5-7 | *7.4-4* | 9.9-7 | 9.9-7	03 | 05 | 6:38 | 03 | 27
16384;16384	10^{-4}	11698	3.5-7 | 7.4-7 | *8.3-5* | 9.8-7 | 9.9-7	10 | 08 | 6:43 | 05 | 35

performs very well when it succeeds in obtaining a solution with the desired accuracy. However, it is not robust in that it fails to solve 4 out of 8 and 5 out of 8 problems when $\lambda_c = 10^{-3}$ and 10^{-4}, respectively. For a few cases, FPC_AS can only achieve a poor accuracy (10^{-1}). The same nonrobustness also appears in the performance of APG. This nonrobustness is in fact closely related to the value of $\lambda_{\max}(\mathcal{A}\mathcal{A}^*)$. For example, both FPC_AS and APG fail to solve a rather small problem **blknheavi** ($m = n = 1024$) whose corresponding $\lambda_{\max}(\mathcal{A}\mathcal{A}^*) = 709$. On the other hand, LADMM, SSNAL, and mfIPM can solve all the test instances successfully. Nevertheless, in some cases, LADMM requires much more time than SSNAL. One can also observe that for large-scale problems, SSNAL outperforms mfIPM by a large margin (sometimes up to a factor of 100). This also demonstrates the power of SSN based augmented Lagrangian methods over interior-point methods in solving large-scale problems. Due to the high sensitivity of the first order algorithms to $\lambda_{\max}(\mathcal{A}\mathcal{A}^*)$, one can safely conclude that the first order algorithms can only be used to solve problems with relatively small $\lambda_{\max}(\mathcal{A}\mathcal{A}^*)$. Moreover, in order to obtain efficient and robust algorithms for lasso problems or more general convex composite optimization problems, it is necessary to carefully exploit the second order information in the algorithmic design.

5. Conclusion. In this paper, we have proposed an inexact augmented Lagrangian method of an asymptotic superlinear convergence rate for solving the large-scale convex composite optimization problems of the form (**P**). It is particularly well suited for solving ℓ_1-regularized least squares (LS) problems. With the intelligent

incorporation of the semismooth Newton method, our Algorithm Ssnal is able to fully exploit the second order sparsity of the problems. Numerical results have convincingly demonstrated the superior efficiency and robustness of our algorithm in solving large-scale ℓ_1-regularized LS problems. Based on extensive numerical evidence, we firmly believe that our algorithmic framework can be adapted to design robust and efficient solvers for various large-scale convex composite optimization problems.

REFERENCES

[1] A. Y. Aravkin, J. V. Burke, D. Drusvyatskiy, M. P. Friedlander, and S. Roy, *Level-set methods for convex optimization*, preprint, arXiv:1602.01506, 2016.

[2] A. Beck and M. Teboulle, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202.

[3] S. Becker, J. Bobin, and E. J. Candès, *NESTA: A fast and accurate first-order method for sparse recovery*, SIAM J. Imaging Sci., 4 (2011), pp. 1–39.

[4] E. van den Berg and M. P. Friedlander, *Probing the Pareto frontier for basis pursuit solutions*, SIAM J. Sci. Comput., 31 (2008), pp. 890–912.

[5] E. van den Berg, M.P. Friedlander, G. Hennenfent, F.J. Herrman, R. Saab, and Ö. Yilmaz, *Sparco: A testing framework for sparse reconstruction*, ACM Trans. Math. Softw., 35 (2009), pp. 1–16.

[6] R. H. Byrd, G. M. Chin, J. Nocedal, and F. Oztoprak, *A family of second-order methods for convex ℓ_1-regularized optimization*, Math. Program., 159 (2016), pp. 435–467.

[7] R. H. Byrd, J. Nocedal, and F. Oztoprak, *An inexact successive quadratic approximation method for L-1 regularized optimization*, Math. Program., 157 (2016), pp. 375–396.

[8] C.-C. Chang and C.-J. Lin, *LIBSVM: A library for support vector machines*, ACM Trans. Intel. Syst. Tech., 2 (2011), 27.

[9] S. S. Chen, D. L. Donoho, and M. A. Saunders, *Atomic decomposition by basis pursuit*, SIAM J. Sci. Comput., 20 (1998), pp. 33–61.

[10] F. Clarke, *Optimization and Nonsmooth Analysis*, John Wiley and Sons, New York, 1983.

[11] A. L. Dontchev and R. T. Rockafellar, *Characterizations of Lipschitzian stability in nonlinear programming*, in Mathematical programming with Data Perturbations, Lecture Notes in Pure and Appl. Math. 195, Dekker, New York, 1998, pp. 65–82.

[12] Y. Cui, D. F. Sun, and K.-C. Toh, *On the asymptotic superlinear convergence of the augmented Lagrangian method for semidefinite programming with multiple solutions*, preprint, arXiv:1610.00875, 2016.

[13] A. L. Dontchev and R. T. Rockafellar, *Implicit Functions and Solution Mappings*, Springer Monographs in Mathematics, Springer, New York, 2009.

[14] F. Facchinei and J.-S. Pang, *Finite-Dimensional Variational Inequalities and Complementarity Problems*, Springer, New York, 2003.

[15] J. Fan, H. Fang, and L. Han, *Challenges of big data analysis*, Nat. Sci. Rev., 1 (2014), pp. 293–314.

[16] 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. Sele. Topics Signal Process., 1 (2007), pp. 586–597.

[17] A. Fischer, *Local behavior of an iterative framework for generalized equations with nonisolated solutions*, Math. Program., 94 (2002), pp. 91–124.

[18] K. Fountoulakis, J. Gondzio, and P. Zhlobich, *Matrix-free interior point method for compressed sensing problems*, Math. Program. Comput., 6 (2014), pp. 1–31.

[19] D. Gabay and B. Mercier, *A dual algorithm for the solution of nonlinear variational problems via finite element approximations*, Comput. Math. Appl., 2 (1976), pp. 17–40.

[20] R. Glowinski and A. Marroco, *Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité, d'une classe de problèmes de Dirichlet non linéaires*, Rev. Francaise d'Automatique, Inform. Rech. Op'erationelle, 9 (1975), pp. 41–76.

[21] R. Goebel and R. T. Rockafellar, *Local strong convexity and local Lipschitz continuity of the gradient of convex functions*, J. Convex Anal., 15 (2008), pp. 263–270.

[22] G. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.

[23] L. Huang, J. Jia, B. Yu, B. G. Chun, P. Maniatis, and M. Naik, *Predicting execution time of computer programs using sparse polynomial regression*, in Advances in Neural Information Processing Systems, 2010, pp. 883–891.

[24] A. F. Izmailov, A. S. Kurennoy, and M. V. Solodov, *A note on upper Lipschitz stability, error bounds, and critical multipliers for Lipschitz-continuous KKT systems*, Math. Program., 142 (2013), pp. 591–604.

[25] K. Jiang, D. F. Sun, and K.-C. Toh, *A partial proximal point algorithm for nuclear norm regularized matrix least squares problems*, Math. Program. Comput., 6 (2014), pp. 281–325.

[26] N. Keskar, J. Nocedal, F. Oztoprak, and A. Wachter, *A second-order method for convex ℓ_1-regularized optimization with active-set prediction*, Optim. Methods Software, 31 (2016), pp. 605–621.

[27] D. Klatte, *Upper Lipschitz behavior of solutions to perturbed $C^{1,1}$ optimization problems*, Math. Program., 88 (2000), pp. 169–180.

[28] S. Kogan, D. Levin, B. R. Routledge, J. S. Sagi, and N. A. Smith, *Predicting Risk from Financial Reports with Regression*, NAACL-HLT 2009, Boulder, CO, 2009.

[29] J. D. Lee, Y. Sun, and M. A. Saunders, *Proximal Newton-type methods for minimizing composite functions*, SIAM J. Optim., 24 (2014), pp. 1420–1443.

[30] M. Lichman, *UCI Machine Learning Repository*, http://archive.ics.uci.edu/ml/datasets.html.

[31] X. D. Li, D. F. Sun, and K.-C. Toh, *QSDPNAL: A Two-Phase Proximal Augmented Lagrangian Method for Convex Quadratic Semidefinite Programming*, preprint, arXiv:1512.08872, 2015.

[32] J. Liu, S. Ji, and J. Ye, *SLEP: Sparse Learning with Efficient Projections*, Arizona State University, 2009.

[33] Z.-Q. Luo and P. Tseng, *On the linear convergence of descent methods for convex essentially smooth minimization*, SIAM J. Control Optim., 30 (1992), pp. 408–425.

[34] F. J. Luque, *Asymptotic convergence analysis of the proximal point algorithm*, SIAM J. Control Optim., 22 (1984), pp. 277–293.

[35] R. Mifflin, *Semismooth and semiconvex functions in constrained optimization*, SIAM J. Control Optim., 15 (1977), pp. 959–972.

[36] A. Milzarek and M. Ulbrich, *A semismooth Newton method with multidimensional filter globalization for ℓ_1-optimization*, SIAM J. Optim., 24 (2014), pp. 298–333.

[37] Y. Nesterov, *A method of solving a convex programming problem with convergence rate $O(1/k^2)$*, Soviet Math. Dokl., 27 (1983), pp. 372–376.

[38] L. Qi and J. Sun, *A nonsmooth version of Newton's method*, Math. Program., 58 (1993), pp. 353–367.

[39] S. M. Robinson, *Some continuity properties of polyhedral multifunctions*, in Mathematical Programming at Oberwolfach, Math. Program. Stud., Springer, Berlin, Heidelberg, 1981, pp. 206–214.

[40] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.

[41] R. T. Rockafellar, *Monotone operators and the proximal point algorithm*, SIAM J. Control Optim., 14 (1976), pp. 877–898.

[42] R. T. Rockafellar, *Augmented Lagrangians and applications of the proximal point algorithm in convex programming*, Math. Oper. Res., 1 (1976), pp. 97–116.

[43] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*, Grundlehren Math. Wiss., Springer-Verlag, Berlin, 1998.

[44] M. Schmidt, *Graphical Model Structure Learning with ℓ_1-Regularization*, Ph.D. thesis, Department of Computer Science, The University of British Columbia, 2010.

[45] D. F. Sun and J. Sun, *Semismooth matrix-valued functions*, Math. Oper. Res., 27 (2002), pp. 150–169.

[46] J. Sun, *On Monotropic Piecewise Quadratic Programming*, Ph.D. thesis, Department of Mathematics, University of Washington, 1986.

[47] P. Tseng and S. Yun, *A coordinate gradient descent method for nonsmooth separable minimization*, Math. Program., 125 (2010), pp. 387–423.

[48] R. Tibshirani, *Regression shrinkage and selection via the lasso*, J. Roy. Statist. Soc. Ser. B, 58 (1996), pp. 267–288.

[49] Z. Wen, W. Yin, D. Goldfarb, and Y. Zhang, *A fast algorithm for sparse reconstruction based on shrinkage, subspace optimization, and continuation*, SIAM J. Sci. Comput., 32 (2010), pp. 1832–1857.

[50] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, *Sparse reconstruction by separable approximation*, IEEE Trans. Signal Process., 57 (2009), pp. 2479–2493.

[51] X. Xiao, Y. Li, Z. Wen, and L. Zhang, *Semi-smooth second-order type methods for composite convex programs*, preprint, arXiv:1603.07870, 2016.

[52] L. Yang, D. F. Sun, and K.-C. Toh, *SDPNAL+: A majorized semismooth Newton-CG augmented Lagrangian method for semidefinite programming with nonnegative constraints*, Math. Program. Comput., 7 (2015), pp. 331–366.

[53] M.-C. Yue, Z. Zhou, and A. M.-C. So, *Inexact Regularized Proximal Newton Method: Provable Convergence Guarantees for Non-Smooth Convex Minimization without Strong Convexity*, preprint, arXiv:1605.07522, 2016.

[54] X. Y. Zhao, D. F. Sun, and K.-C. Toh, *A Newton-CG augmented Lagrangian method for semidefinite programming*, SIAM J. Optim., 20 (2010), pp. 1737–1765.

[55] X. Zhang, M. Burger and S. Osher, *A unified primal-dual algorithm framework based on Bregman iteration*, J. Sci. Comput., 49 (2011), pp. 20–46.

[56] H. Zhou, *The adaptive lasso and its oracle properties*, J. Amer. Statist. Assoc., 101 (2006), pp. 1418–1429.

[57] H. Zhou and T. Hastie, *Regularization and variable selection via the elastic net*, J. Roy. Statist. Soc., Ser. B, 67 (2005), pp. 301–320.

[58] Z. Zhou and A. M.-C. So, *A unified approach to error bounds for structured convex optimization problems*, Math. Program., 165 (2017), pp. 689–728.