On the Linear Convergence of the Approximate Proximal Splitting Method for Non-smooth Convex Optimization

Mojtaba Kadkhodaie · Maziar Sanjabi ·
Zhi-Quan Luo

Received: 16 April 2014 / Revised: 23 May 2014 / Accepted: 28 May 2014 /

Published online: 18 June 2014

© Operations Research Society of China, Periodicals Agency of Shanghai University, and Springer-Verlag Berlin Heidelberg 2014

Abstract Consider the problem of minimizing the sum of two convex functions, one being smooth and the other non-smooth. In this paper, we introduce a general class of approximate proximal splitting (APS) methods for solving such minimization problems. Methods in the APS class include many well-known algorithms such as the proximal splitting method, the block coordinate descent method (BCD), and the approximate gradient projection methods for smooth convex optimization. We establish the linear convergence of APS methods under a local error bound assumption. Since the latter is known to hold for compressive sensing and sparse group LASSO problems, our analysis implies the linear convergence of the BCD method for these problems without strong convexity assumption.

Keywords Convex optimization · Proximal splitting method · Block coordinate descent method · Convergence rate analysis · Local error bound

1 Introduction

Consider a constrained convex minimization of the form

$$\min_{\mathbf{x} \in X} F(\mathbf{x}) = f_1(\mathbf{x}) + f_2(\mathbf{x}), \quad (1.1)$$

where X is a convex closed set (can be $\mathbb{R}^n$), f_1 is a convex function (may be non-smooth), and f_2 is a smooth convex function with Lipschitz continuous gradient.

M. Kadkhodaie (✉) · M. Sanjabi · Z.-Q. Luo
University of Minnesota, Minneapolis, MN, USA
e-mail: kadkh004@umn.edu

1.1 Motivating Applications

Non-smooth convex optimization problems of the form (1.1) arise in many contemporary statistical and signal processing applications including compressive sensing, signal denoising, and sparse logistic regression. In the sequel, we outline some of the most recent applications of problem (1.1).

1.1.1 LASSO Problem

Suppose that we have a noisy observation vector $\mathbf{b} \in \mathbb{R}^m$ about an unknown sparse vector $\mathbf{x} \in \mathbb{R}^n$, where the signal model is linear and given by

$$\mathbf{b} \approx \mathbf{A}\mathbf{x},$$

for some given matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$. One of the most popular techniques to estimate the sparse vector $\mathbf{x}$ is called LASSO [29]. LASSO can be viewed as an ℓ_1 -norm regularized linear least squares problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \lambda \|\mathbf{x}\|_1, \quad (1.2)$$

where the first term $\frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$ reduces the estimation error, and the second term $\lambda \|\mathbf{x}\|_1$ promotes the sparsity of the solution. The parameter λ controls the sparsity level of the solution. The higher λ is, the fewer non-zero entries would be in the LASSO solution. Clearly, by setting $f_2(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$, $f_1(\mathbf{x}) = \lambda \|\mathbf{x}\|_1$ and $\mathbf{X} = \mathbb{R}^n$, problem (1.2) becomes a special case of problem (1.1).

1.1.2 Group LASSO Problem

In many applications, such as image denoising, the desired solution should have the so called *group sparse* structure [35], i.e., the solution $\mathbf{x}$ should admit a group separable structure $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_K]'$, where $\mathbf{x}_i \in \mathbb{R}^{n_i}$ and $\sum_{i=1}^K n_i = n$, with only a few non-zero representing groups. In these applications, the following optimization problem needs to be solved

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\|, \quad (1.3)$$

where $\mathcal{J}$ is a partition of $\{1, \dots, n\}$ and w_J is the sparsity weight of block J . Setting $f_2(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$ and $f_1(\mathbf{x}) = \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\|$, the Group LASSO problem (1.3) follows the structure of problem (1.1).

1.1.3 Group LASSO for Logistic Regression

Given a set of n -dimensional feature vectors $\mathbf{a}_i$, $i = 1, \dots, m$, and the corresponding class labels $b_i \in \{0, 1\}$, $i = 1, \dots, m$, our task is to find a linear classifier for the

vectors $\mathbf{a}_i$. Assume the probability distribution of the class label b , given a feature vector $\mathbf{a}$ is given by

$$p(b = 1|\mathbf{a}; \mathbf{x}) = \frac{\exp(\mathbf{a}^T \mathbf{x})}{1 + \exp(\mathbf{a}^T \mathbf{x})},$$

where $\mathbf{x}$ is the logistic coefficient vector. The logistic Group LASSO problem [18] can be written as

$$\min_{\mathbf{x} \in \mathbb{R}^n} \sum_{i=1}^m (\log(1 + \exp(\mathbf{a}_i^T \mathbf{x})) - b_i \mathbf{a}_i^T \mathbf{x}) + \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\|, \quad (1.4)$$

where w_J is the sparsity weight for the corresponding block $\mathbf{x}_J$. Problem (1.4) can also be interpreted as a special form of problem (1.1). We refer the readers to [2, 15] for further applications of Group LASSO, and to [3, 7, 10, 18, 27, 33, 34] for further studies on Group LASSO type of techniques in statistical problems.

1.2 First-Order Methods to Solve Convex Problem (1.1)

For the large scale optimization problems of the form (1.1), the preferred approach is to use iterative descent algorithms along gradient-related directions. Such first-order methods have a long history in optimization. For example, if we assume $X = \mathbb{R}^n$ and $f_1(\cdot)$ to be the indicator function of a closed convex set $\mathbf{C}$, then the problem in (1.1) turns out to be the smooth minimization of $F(\mathbf{x}) = f_2(\mathbf{x})$ over the set $\mathbf{C}$. A well-known first-order method to solve this problem is called gradient projection (GP) [25, 26] algorithm. In iteration k , the GP algorithm takes a gradient step of size α_k and then projects the point back into the feasible set $\mathbf{C}$,

$$\mathbf{x}^{k+1} = \text{Proj}_{\mathbf{C}}[\mathbf{x}^k - \alpha_k \nabla f_2(\mathbf{x}^k)]. \quad (1.5)$$

The convergence analysis of the GP method has been studied before [13]. It has been shown that such analysis can be generalized to approximate versions of the GP method [9, 12–14, 17] which are also known as approximate gradient projection (AGP) methods. In the framework of AGP, an error is allowed in the computation of gradient, as long as the size of the error vector is sufficiently small (see Sect. 3). In particular, the update of an AGP algorithm can be defined as

$$\mathbf{x}^{k+1} = \text{Proj}_{\mathbf{C}}[\mathbf{x}^k - \alpha_k \nabla f_2(\mathbf{x}^k) + \mathbf{e}^k], \quad (1.6)$$

where $\|\mathbf{e}^k\| \leq \kappa \|\mathbf{x}^{k+1} - \mathbf{x}^k\|$ for some $\kappa > 0$. It has been shown [13] that many well-known algorithms such as matrix splitting method [20] and extragradient method [8] are all special cases of the AGP algorithm (1.6). We will extend this approach to the non-smooth optimization (1.1); see Sect. 3.

The analog of the GP algorithm for the general non-smooth version of problem (1.1) is the so called *proximal splitting method* (PSM). In order to introduce this method, we first need to define the proximity operator. For any convex function $\varphi(\cdot)$ (possibly non-smooth), the Moreau–Yashida proximity operator $\text{prox}_{\varphi}(\cdot, X) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is defined as

$$\text{prox}_\varphi(\mathbf{x}, X) = \underset{\mathbf{y} \in X}{\operatorname{argmin}} \varphi(\mathbf{y}) + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|^2. \quad (1.7)$$

Note that since $\frac{1}{2} \|\cdot - \mathbf{x}\|^2$ is strongly convex and $\varphi(\cdot)$ is convex, the minimizer of (1.7) is unique. Furthermore, if the function φ is chosen to be the indicator function, ι_C , of the closed convex set C and $X = \mathbb{R}^n$, then the proximity operator is reduced to the projection operator onto the set C . Thus, the proximity operator is a natural extension of the projection operator. In the sequel, we will denote the proximity operator simply by $\text{prox}_\varphi(\cdot)$ and assume that its dependence on the set X is understood from the context.

The proximity operator inherits many useful properties of the projection operator into convex sets. For instance, the proximity operator is known to be non-expansive and is therefore Lipschitz continuous,

$$\|\text{prox}_\varphi(\mathbf{x}) - \text{prox}_\varphi(\mathbf{y})\| \leq \|\mathbf{x} - \mathbf{y}\|, \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n.$$

For large scale problems, it is not always easy to compute the proximity operator, unless the function φ has some special structure, such as separability. In those cases, the proximity operator is efficiently computable (or has closed form). For instance, if the function φ is the ℓ_1 -norm, the proximity operator has a closed form solution, also known as the shrinkage operator [31].

Using the proximity operator, the optimality condition of problem (1.1) can be formulated as

$$\mathbf{x} = \text{prox}_{\alpha f_1}(\mathbf{x} - \alpha \nabla f_2(\mathbf{x})), \quad (1.8)$$

for some $\alpha > 0$. The PSM can be viewed as an iterative approach to solve the above fixed point equation

$$\mathbf{x}^{k+1} = \text{prox}_{\alpha_k f_1}(\mathbf{x}^k - \alpha_k \nabla f_2(\mathbf{x}^k)), \quad (1.9)$$

where $\alpha_k > 0$ determines the step size at iteration k . Note that PSM is identical to the GP method, if $f_1 = \iota_C$ for some convex closed set C . It is known that if the step size α_k satisfies

$$0 < \underline{\alpha} \leq \alpha_k \leq \bar{\alpha} < 1/L, \quad k = 0, 1, \dots,$$

where L is the Lipschitz constant of ∇f_2 , then every sequence generated by the PSM converges to a solution of (1.1) (see [4]).

In spite of this convergence result, the rate of convergence for the PSM is not known, except for some specific cases. For instance, if $f_1 = \iota_C$ and f_2 has a composite structure ($f_2(\mathbf{x}) = h(\mathbf{A}\mathbf{x})$, where h is strongly convex and $\mathbf{A}$ is an arbitrary $m \times n$ matrix), then it is proved by Luo and Tseng [12] that the PSM algorithm (which coincides with GP in this case) converges linearly to an optimal solution of (1.1). This result establishes linear convergence in the absence of strong convexity of the objective function. This result has been recently extended to the case where, $f_1(\mathbf{x}) = \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\|_2$ and $f_1(\mathbf{x}) = \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\|_2 + \lambda \|\mathbf{x}\|$, where w_J and λ are

some non-negative constants, and $f_2 = h(\mathbf{A}\mathbf{x})$ is still a composite function (see [31, 36]).

1.3 Our Contribution

This paper proposes a general framework of first-order methods, called *approximate proximal splitting* (APS), for the non-smooth convex optimization problem (1.1). This framework combines the existing framework of AGP with the proximal splitting technique, and as such, it includes the GP, AGP, and PSMs as special cases. Moreover, the well known block coordinate descent (BCD) algorithm is also a special case of APS.

We analyze the convergence rate of APS class of algorithms under a local error bound condition. Our result implies the linear convergence rate of BCD method for (1.1) for the LASSO or group LASSO type of problems when $f_1(\mathbf{x}) = \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\|_2$ or $f_1(\mathbf{x}) = \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\|_2 + \lambda \|\mathbf{x}\|_1$. The BCD algorithm is one of the main algorithms used to solve large scale optimization problems due to the simplicity of its updates (especially for the LASSO or group LASSO type of problems in which each step of BCD is equivalent to a shrinkage operator [36]). This linear convergence result provides theoretical proof for effectiveness of BCD in handling such problems.

Our result differs from the existing PSMs and analysis in several aspects. Among the existing work [1, 28, 36], the only one which considers an error term in the proximal splitting algorithm is [28], while the other two ([1] and [36]) are focussed on the pure proximal splitting algorithm. The result in [28] does not provide the linear convergence except in the strongly convex case which is a special case of our result, so it cannot be used to establish the linear convergence of the BCD algorithm for LASSO type problems. The reason for such difference is that we use a local error bound condition in place of the strong convexity assumption.

The result in [1] deals with the problem of linear convergence from a statistical point of view. It assumes that problem (1.1) comes from an M -estimator formulation with some probabilistic construction. It proves that iterations will converge linearly to a neighborhood around the optimal solution, but not necessarily an optimal solution. As such, this result is probabilistic and not deterministic. This is in contrast to our result which is a general convex optimization problem in the form (1.1), regardless how it is generated. That said, by utilizing the so called restricted strong convexity and restricted smoothness (see [1]) instead of an error bound, the authors have established the linear convergence of the proximal splitting algorithm for a broad range optimization problems with non-smooth regularizers such as L_1 norm or nuclear norm.

2 Preliminaries

2.1 Proximal Gradient Vector

Before introducing the APS method, we need to define some basic concepts, which will be useful in our future convergence analysis of the algorithm.

Definition 2.1 For any $\alpha > 0$, we define *proximal gradient* vector as

$$\tilde{\nabla}F(\mathbf{x}, \alpha) = \frac{1}{\alpha}[\mathbf{x} - \text{prox}_{\alpha f_1}(\mathbf{x} - \alpha \nabla f_2(\mathbf{x}))]. \quad (2.1)$$

When $\alpha = 1$, we will use the short notation

$$\tilde{\nabla}F(\mathbf{x}) = \tilde{\nabla}F(\mathbf{x}, \alpha). \quad (2.2)$$

Note that in the special case of $f_1 = 0$ and $X = \mathbb{R}^n$, the proximal gradient reduces to the standard gradient, namely, $\tilde{\nabla}F(\mathbf{x}, \alpha) = \nabla f_2(\mathbf{x}) = \nabla F(\mathbf{x})$. In another special case where $f_1 = \iota_C$ (the indicator function of a convex set C), we have

$$\tilde{\nabla}F(\mathbf{x}, \alpha) = \frac{1}{\alpha}[\mathbf{x} - \text{proj}_C(\mathbf{x} - \alpha \nabla f_2(\mathbf{x}))], \quad (2.3)$$

which is the residual of the optimality condition for the following problem

$$\min_{\mathbf{x} \in C} F(\mathbf{x}) = f_2(\mathbf{x}). \quad (2.4)$$

Hence, $\tilde{\nabla}F(\mathbf{x}, \alpha)$ can be viewed as a generalized notion of gradient for the constrained non-smooth minimization. In addition, it inherits many useful properties of gradient. For instance, $\tilde{\nabla}F(\mathbf{x}^*, \alpha) = 0$ for some $\alpha > 0$ iff $\mathbf{x}^*$ is an optimal solution of (1.1).

The optimality condition for (1.1) given in (1.8) suggests that we can define a local measure for distance to optimality by

$$\psi(\mathbf{x}) = \|\tilde{\nabla}F(\mathbf{x})\| = \|\mathbf{x} - \text{prox}_{f_1}(\mathbf{x} - \nabla f_2(\mathbf{x}))\|. \quad (2.5)$$

It is easy to see that $\psi(\mathbf{x}) = 0$ iff $\mathbf{x}$ belongs to the set of optimal solutions of (1.1), which we denote by X^* .

2.2 Error Bounds

In this section, we formally introduce the notion of error bound. As we will see it is a vital property in obtaining linear convergence rate for solving a problem via first-order methods.

For any $x \in X$, we can define

$$\varphi(\mathbf{x}) = \min_{\mathbf{y} \in \bar{X}^*} \|\mathbf{x} - \mathbf{y}\|, \quad (2.6)$$

where $\bar{X}^*$ is the closure of X^* (the set of optimal solutions of (1.1)). It is straightforward to see that $\varphi(\mathbf{x})$ can be used as a measure for distance to optimality, and $\varphi(\mathbf{x}) = 0$ iff $\mathbf{x} \in \bar{X}^*$. However, in practice it is impossible to compute $\varphi(\mathbf{x})$, due to the requirement of knowing the set of optimal solutions, $\bar{X}^*$. This is where the error bound comes into the picture. It serves as an approximated measure of the distance

to optimality. The error bound is simply a bound on $\varphi(\mathbf{x})$, based on another measure of optimality that can be computed easily (in this case, the size of the residual $\psi(\mathbf{x})$ defined by (2.5)).

Definition 2.2 (*Local Error Bound* [13]) Consider the optimality distance measures defined by (2.5) and (2.6). We say that problem (1.1) satisfies the local error bound property if for every $v \geq \inf_{\mathbf{x} \in X} F(\mathbf{x})$, there exist scalars $\delta > 0$ and $\tau > 0$ such that

$$\varphi(\mathbf{x}) \leq \tau \psi(\mathbf{x}), \quad (2.7)$$

for all $\mathbf{x} \in X$ with $F(\mathbf{x}) \leq v$ and $\psi(\mathbf{x}) \leq \delta$.

In other words, (2.7) says that $\varphi(\mathbf{x})$ is bounded above by the norm of the residual at $\mathbf{x}$, whenever $F(\mathbf{x})$ is bounded above and this residual is small enough. In order to gain some insight on when (2.7) holds, consider the case where $X = \mathbb{R}^n$ and $F(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x}$ for some positive definite matrix $\mathbf{A}$ and a vector $\mathbf{b} \in \mathbb{R}^n$. Then, (2.7) is equivalent to

$$\varphi(\mathbf{x}) \leq \tau \|\nabla f_2(\mathbf{x})\| = \tau \|\mathbf{A} \mathbf{x} + \mathbf{b}\|,$$

which is called the Hoffman error bound and can be easily checked to be true (using elementary linear algebra). Furthermore, the error bound holds globally for strongly convex smooth $F(\mathbf{x})$ for any closed convex set X . In case when $X = \mathbb{R}^n$ the strong convexity of $F(\mathbf{x})$ implies the existence a $\tau > 0$ such that

$$\|\mathbf{x} - \mathbf{y}\|^2 \leq \tau \langle \nabla F(\mathbf{x}) - \nabla F(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \quad \forall \mathbf{x}, \mathbf{y}.$$

Let $\hat{\mathbf{x}}$ to be a stationary point in $\bar{X}^*$ satisfying $\|\mathbf{x} - \hat{\mathbf{x}}\| = \varphi(\mathbf{x})$, then

$$\varphi(\mathbf{x})^2 \leq \tau \langle \nabla F(\mathbf{x}), \mathbf{x} - \hat{\mathbf{x}} \rangle \leq \tau \|\nabla F(\mathbf{x})\| \|\mathbf{x} - \hat{\mathbf{x}}\|.$$

Proving error bound for different problems has a long history in the literature. It was first considered by Demb and Tulowizki [5] for strongly convex quadratic functions and by Pang [21] in the context of linear complementarity problems (LCP). In the case of smooth minimization, there are numerous error bound results under different assumptions. For instance, it has been shown that the error bound hold for strongly convex functions in [22], and for quadratic functions with polyhedral constraint in [11, 12].

In the case of non-smooth optimization, the available error bound results are far fewer and only deal with problems with structured non-smooth parts. For instance, in the recent works [31, 36], it has been proved that error bounds holds for special type of non-smooth problems (Group LASSO type of problems).

Theorem 2.3 ([31, 36]) In problem (1.1) let $X = \mathbb{R}^n$, $f_1(\mathbf{x}) = \sum_{J \in \mathcal{J}} w_J \|\mathbf{x}_J\| + \lambda \|\mathbf{x}\|$ for non-negative w_J 's and λ and $f_2(\mathbf{x}) = h(\mathbf{A} \mathbf{x})$ for some strongly convex smooth function $h : \mathbb{R}^m \rightarrow \mathbb{R}$ and an $m \times n$ matrix $\mathbf{A}$. In addition if the function F is coercive, then error bound (2.7) holds for (1.1).

Corollary 2.4 The error bound condition (2.7) holds for LASSO problem (1.2), Group LASSO problem (1.3), and logistic Group LASSO problem (1.4).

We will utilize this result to establish our proof for linear convergence of APS class of methods for such problems.

3 Approximate Proximal Splitting (APS) Method

In this section, we formally introduce the APS class of methods.

Definition 3.1 (APS) An algorithm is considered in the class of APS methods if it generates a sequence of iterates $\mathbf{x}^0, \mathbf{x}^1, \dots$ in X such that

$$\mathbf{x}^{r+1} = \text{prox}_{\alpha^r f_1}(\mathbf{x}^r - \alpha^r (\nabla f_2(\mathbf{x}^r) + \mathbf{e}^r)), \quad r = 0, 1, \dots, \quad (3.1)$$

where $\{\alpha^r\}$ is a sequence of positive scalars with $\liminf \alpha^r > 0$ and $\{\mathbf{e}^r\}$ is a sequence in $\mathbb{R}^n$ with

$$\|\mathbf{e}^r\| \leq \kappa \|\mathbf{x}^r - \mathbf{x}^{r+1}\|, \quad (3.2)$$

for some non-negative scalar κ .

In Eq. (3.1), α^r and $\mathbf{e}^r$ may depend on $\mathbf{x}^r$ and can be viewed as algorithm parameters. Hence, different choices of α^r and $\mathbf{e}^r$ lead into different algorithms. For instance, the PSM algorithm, whose update rule is given by (1.9), is a special case of the APS algorithm with $\mathbf{e}^r = \mathbf{0}$. In fact, the condition (3.2) ensures that the algorithm does not deviate too much from the PSM update.

For smooth minimization which is a special case of problem (1.1) with $f_1 = \iota_C$, the AGP class of algorithms is very common. Since the proximity operator reduces to the projection operator in this case, the APS algorithm (1.6) contains the APG method (3.1) as a special case. Later we will see how the BCD algorithm is also a special case of the APS method.

4 Linear Convergence of APS

In this section, we will prove that any APS method converges linearly, if the following properties hold true.

- *Sufficient decrease.* There exists a constant $c_1 > 0$ such that $\forall r \geq 0$,

$$F(\mathbf{x}^r) - F(\mathbf{x}^{r+1}) \geq c_1 \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2. \quad (4.1)$$

- *Local error bound.* Definition 2.2.
- *Cost-to-go.* There exists a $c_2 > 0$, such that

$$F(\mathbf{x}^r) - F^* \leq c_2 (\varphi(\mathbf{x}^r)^2 + \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2), \quad \forall r, \quad (4.2)$$

where F^* is the optimal objective value of (1.1).

The local error bound property solely depends on the optimization problem. Therefore, the problem structure needs to be studied to ensure this property holds. In

the literature, this condition has been established for certain classes of optimization problems, see [13, 30, 36] and references therein. Some of the existing results are summarized in Theorem 2.3.

The section proceeds as follows. Assuming the sufficient decrease condition, we first prove that the cost-to-go will naturally follow the APS class of algorithms. Then, the sufficient decrease is proved under some assumptions on the step size α^r and the error vector $\mathbf{e}^r$. Finally, the linear convergence rate of APS methods is shown assuming the local error bound condition of the problem. The following lemma shows that the cost-to-go property follows from the sufficient decrease condition.

Lemma 4.1 *If an APS method satisfies the sufficient decrease condition (4.1), then the cost-to-go condition (4.2) also holds.*

Proof Set $\hat{\mathbf{x}}^r$ to be the point in $\bar{X}^*$, such that $\varphi(\mathbf{x}^r) = \|\mathbf{x}^r - \hat{\mathbf{x}}^r\|$. The optimality condition of $\mathbf{x}^{r+1}$ implies

$$\begin{aligned} f_1(\hat{\mathbf{x}}^r) + \langle \nabla f_2(\mathbf{x}^r) + \mathbf{e}^r, \hat{\mathbf{x}}^r - \mathbf{x}^r \rangle + \frac{1}{2\alpha^r} \|\hat{\mathbf{x}}^r - \mathbf{x}^r\|^2 &\geq \\ f_1(\mathbf{x}^{r+1}) + \langle \nabla f_2(\mathbf{x}^r) + \mathbf{e}^r, \mathbf{x}^{r+1} - \mathbf{x}^r \rangle + \frac{1}{2\alpha^r} \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2. \end{aligned}$$

This implies

$$\langle \nabla f_2(\mathbf{x}^r) + \mathbf{e}^r, \mathbf{x}^{r+1} - \hat{\mathbf{x}}^r \rangle + f_1(\mathbf{x}^{r+1}) - f_1(\hat{\mathbf{x}}^r) \leq \frac{1}{2\alpha^r} \varphi^2(\mathbf{x}^r). \quad (4.3)$$

Also, the mean value theorem implies

$$f_2(\mathbf{x}^{r+1}) - f_2(\hat{\mathbf{x}}^r) = \langle \nabla f_2(\xi^r), \mathbf{x}^{r+1} - \hat{\mathbf{x}}^r \rangle, \quad (4.4)$$

for some ξ^r in the line segment joining $\mathbf{x}^{r+1}$ and $\hat{\mathbf{x}}^r$. Combining the above two relations yields

$$\begin{aligned} F(\mathbf{x}^{r+1}) - F(\hat{\mathbf{x}}^r) &= f_1(\mathbf{x}^{r+1}) - f_1(\hat{\mathbf{x}}^r) + f_2(\mathbf{x}^{r+1}) - f_2(\hat{\mathbf{x}}^r) \\ &= \langle \nabla f_2(\xi^r), \mathbf{x}^{r+1} - \hat{\mathbf{x}}^r \rangle + f_1(\mathbf{x}^{r+1}) - f_1(\hat{\mathbf{x}}^r) \\ &= \langle \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \hat{\mathbf{x}}^r \rangle + \langle \nabla f_2(\xi^r) - \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \hat{\mathbf{x}}^r \rangle \\ &\quad + f_1(\mathbf{x}^{r+1}) - f_1(\hat{\mathbf{x}}^r) \\ &\leq \langle \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \hat{\mathbf{x}}^r \rangle + L \|\xi^r - \mathbf{x}^r\| \|\mathbf{x}^{r+1} - \hat{\mathbf{x}}^r\| \\ &\quad + f_1(\mathbf{x}^{r+1}) - f_1(\hat{\mathbf{x}}^r) \\ &\leq \frac{1}{2\alpha^r} \varphi^2(\mathbf{x}^r) - \langle \mathbf{e}^r, \mathbf{x}^{r+1} - \hat{\mathbf{x}}^r \rangle + L \|\xi^r - \mathbf{x}^r\| \|\mathbf{x}^{r+1} - \hat{\mathbf{x}}^r\| \\ &\leq \frac{1}{2\alpha^r} \varphi^2(\mathbf{x}^r) + L \|\xi^r - \mathbf{x}^r\| \|\mathbf{x}^{r+1} - \hat{\mathbf{x}}^r\| + \kappa \|\mathbf{x}^{r+1} - \mathbf{x}^r\| \|\mathbf{x}^{r+1} - \hat{\mathbf{x}}^r\|. \end{aligned} \quad (4.5)$$

It remains to bound the last two terms in (4.5). Using the fact that ξ^r lies in the line segment joining $\mathbf{x}^{r+1}$ and $\hat{\mathbf{x}}^r$, it follows that

$$\begin{aligned}
\|\xi^r - \mathbf{x}^r\| \|\mathbf{x}^{r+1} - \hat{\mathbf{x}}^r\| &\leq (\|\mathbf{x}^{r+1} - \mathbf{x}^r\| + \|\mathbf{x}^r - \hat{\mathbf{x}}^r\|)(\|\mathbf{x}^{r+1} - \mathbf{x}^r\| + \|\mathbf{x}^r - \hat{\mathbf{x}}^r\|) \\
&= (\|\mathbf{x}^{r+1} - \mathbf{x}^r\| + \varphi(\mathbf{x}^r))(\|\mathbf{x}^{r+1} - \mathbf{x}^r\| + \varphi(\mathbf{x}^r)) \\
&\leq 2(\|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2 + \varphi^2(\mathbf{x}^r)).
\end{aligned}$$

For the last term in (4.5) we have

$$\begin{aligned}
\|\mathbf{x}^{r+1} - \mathbf{x}^r\| \|\mathbf{x}^{r+1} - \hat{\mathbf{x}}^r\| &\leq \|\mathbf{x}^{r+1} - \mathbf{x}^r\|(\|\mathbf{x}^{r+1} - \mathbf{x}^r\| + \varphi(\mathbf{x}^r)) \\
&\leq 2\|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2 + 2\varphi^2(\mathbf{x}^r).
\end{aligned}$$

Substituting these upper bounds into the right hand side of inequality (4.5) yields

$$F(\mathbf{x}^{r+1}) - F(\hat{\mathbf{x}}^r) = O(\varphi^2(\mathbf{x}^r) + \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2). \quad (4.6)$$

This proves the desired result. $\square$

The following result establishes the sufficient decrease condition for the APS algorithm under some conditions on the error sequence $\mathbf{e}^r$ and the step size sequence α^r .

Lemma 4.2 *Consider an APS algorithm defined by (3.1)–(3.2) for some $\kappa > 0$ and some stepsizes α^r satisfying*

$$0 < \underline{\alpha} \leq \alpha^r \leq \bar{\alpha} < \frac{2}{L + 2\kappa}, \quad \text{for some } \underline{\alpha} \text{ and } \bar{\alpha}, \forall r. \quad (4.7)$$

then the sufficient decrease property (4.1) holds.

Proof By the optimality condition for $\mathbf{x}^{r+1}$, there exists a $\mathbf{g} \in \partial f_1(\mathbf{x}^{r+1})$ such that

$$\langle \alpha^r \mathbf{g} + \alpha^r \nabla f_2(\mathbf{x}^r) + \alpha^r \mathbf{e}^r + \mathbf{x}^{r+1} - \mathbf{x}^r, \mathbf{y} - \mathbf{x}^{r+1} \rangle \geq 0, \quad \forall \mathbf{y} \in X.$$

Moreover, the convexity of f_1 implies that for any $\mathbf{g} \in \partial f_1(\mathbf{x}^{r+1})$ and any $\mathbf{y} \in X$, the following inequality holds

$$f_1(\mathbf{y}) - \langle \mathbf{g}, \mathbf{y} - \mathbf{x}^{r+1} \rangle \geq f_1(\mathbf{x}^{r+1}).$$

Using the above two relations and convexity of f_2 , we obtain

$$\begin{aligned}
F(\mathbf{y}) &\geq f_1(\mathbf{y}) + f_2(\mathbf{x}^r) + \langle \nabla f_2(\mathbf{x}^r), \mathbf{y} - \mathbf{x}^r \rangle \\
&= f_1(\mathbf{y}) + f_2(\mathbf{x}^r) + \langle \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \mathbf{x}^r \rangle + \langle \nabla f_2(\mathbf{x}^r), \mathbf{y} - \mathbf{x}^{r+1} \rangle \\
&\geq f_1(\mathbf{y}) + f_2(\mathbf{x}^r) + \langle \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \mathbf{x}^r \rangle - \langle \mathbf{e}^r + \frac{1}{\alpha^r}(\mathbf{x}^{r+1} - \mathbf{x}^r) + \mathbf{g}, \mathbf{y} - \mathbf{x}^{r+1} \rangle \\
&\geq [f_1(\mathbf{x}^{r+1}) + \langle \mathbf{g}, \mathbf{y} - \mathbf{x}^{r+1} \rangle] + f_2(\mathbf{x}^r) + \langle \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \mathbf{x}^r \rangle \\
&\quad - \langle \mathbf{e}^r + \frac{1}{\alpha^r}(\mathbf{x}^{r+1} - \mathbf{x}^r) + \mathbf{g}, \mathbf{y} - \mathbf{x}^{r+1} \rangle \\
&= f_1(\mathbf{x}^{r+1}) + f_2(\mathbf{x}^r) + \langle \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \mathbf{x}^r \rangle - \langle \mathbf{e}^r + \frac{1}{\alpha^r}(\mathbf{x}^{r+1} - \mathbf{x}^r), \mathbf{y} - \mathbf{x}^{r+1} \rangle.
\end{aligned}$$

Now the Lipschitz continuity of ∇f_2 and Taylor expansion of f_2 imply that

$$\begin{aligned} F(\mathbf{y}) &\geq f_1(\mathbf{x}^{r+1}) + f_2(\mathbf{x}^r) + \langle \nabla f_2(\mathbf{x}^r), \mathbf{x}^{r+1} - \mathbf{x}^r \rangle - \langle \mathbf{e}^r + \frac{1}{\alpha^r}(\mathbf{x}^{r+1} - \mathbf{x}^r), \mathbf{y} - \mathbf{x}^{r+1} \rangle \\ &\geq f_1(\mathbf{x}^{r+1}) + f_2(\mathbf{x}^{r+1}) - \frac{L}{2} \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2 - \langle \mathbf{e}^r + \frac{1}{\alpha^r}(\mathbf{x}^{r+1} - \mathbf{x}^r), \mathbf{y} - \mathbf{x}^{r+1} \rangle \\ &= F(\mathbf{x}^{r+1}) - \frac{L}{2} \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2 - \langle \mathbf{e}^r + \frac{1}{\alpha^r}(\mathbf{x}^{r+1} - \mathbf{x}^r), \mathbf{y} - \mathbf{x}^{r+1} \rangle. \end{aligned}$$

Specializing $\mathbf{y} = \mathbf{x}^r$, and using the Cauchy–Schwartz inequality, we get

$$F(\mathbf{x}^r) - F(\mathbf{x}^{r+1}) \geq \frac{2 - 2\alpha^r\kappa - \alpha^rL}{2\alpha^r} \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2.$$

Moreover, we know that

$$\frac{2 - 2\alpha^r\kappa - \alpha^rL}{2\alpha^r} \geq \frac{2 - 2\bar{\alpha}\kappa - \bar{\alpha}L}{2\bar{\alpha}} > 0, \quad \forall r,$$

which further implies the desired result. $\square$

Finally, we state the following lemma which is needed to prove the linear convergence of the APS method. Its proof is relegated to the [Appendix](#).

Lemma 4.3 *For $\alpha > 0$, we have*

- (1) *The function $\alpha \|\tilde{\nabla} f(\mathbf{x}, \alpha)\|$ is monotonically increasing with α .*
- (2) *The function $\|\tilde{\nabla} f(\mathbf{x}, \alpha)\|$ is monotonically decreasing with α .*

Now we are ready to state and prove the linear convergence of APS class of algorithms.

Theorem 4.4 *Assume that the error bound property holds for problem (1.1). Consider an APS algorithm defined by (3.1)–(3.2) for some $\kappa > 0$ and some stepsizes α^r satisfying the sufficient decrease condition (4.1). Then the sequence of iterations generated by the algorithm converges R-linearly to an optimal solution of (1.1).*

Proof The sufficient decrease condition (4.1) implies

$$\|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2 \rightarrow 0.$$

Moreover, we have

$$\begin{aligned} \|\mathbf{x}^r - \text{prox}_{\alpha^r f_1}[\mathbf{x}^r - \alpha^r \nabla f_2(\mathbf{x}^r)]\| &\leq \|\mathbf{x}^r - \mathbf{x}^{r+1}\| + \|\mathbf{x}^{r+1} - \text{prox}_{\alpha^r f_1}[\mathbf{x}^r - \alpha^r \nabla f_2(\mathbf{x}^r)]\| \\ &\leq \|\mathbf{x}^r - \mathbf{x}^{r+1}\| + \alpha^r \|\mathbf{e}^r\| \\ &\leq (\bar{\alpha}\kappa + 1) \|\mathbf{x}^r - \mathbf{x}^{r+1}\|, \end{aligned}$$

where the second inequality is due to the non-expansiveness property of the proximity operator.

Since $\alpha^r \geq \underline{\alpha}$ for all r , we obtain from Lemma 4.3 that

$$\begin{aligned}\psi(\mathbf{x}^r) &= \|\mathbf{x}^r - \text{prox}_{f_1}[\mathbf{x}^r - \nabla f_2(\mathbf{x}^r)]\| \\ &\leq \frac{1}{\min\{1, \underline{\alpha}\}} \|\mathbf{x}^r - \text{prox}_{\alpha^r f_1}[\mathbf{x}^r - \alpha^r \nabla f_2(\mathbf{x}^r)]\| \\ &\leq \frac{(\bar{\alpha}\kappa + 1)}{\min\{1, \underline{\alpha}\}} \|\mathbf{x}^r - \mathbf{x}^{r+1}\| \rightarrow 0.\end{aligned}$$

Hence,

$$\psi(\mathbf{x}^r) \rightarrow 0.$$

Using the local error bound condition, for sufficiently large r , there exists a constant τ such that

$$\varphi(\mathbf{x}^r) \leq \tau \psi(\mathbf{x}^r) \rightarrow 0, \quad (4.8)$$

which further implies $\varphi(\mathbf{x}^r) \rightarrow 0$. Notice that by Lemma 4.1, the cost-to-go estimate (4.2) holds, which together with the fact that $\varphi(\mathbf{x}^r) \rightarrow 0$ further implies

$$F(\mathbf{x}^r) \rightarrow F^*.$$

Now we use the local error bound condition together with the cost-to-go estimate to obtain

$$\begin{aligned}F(\mathbf{x}^{r+1}) - F^* &\leq c_2(\varphi^2(\mathbf{x}^r) + \|\mathbf{x}^r - \mathbf{x}^{r+1}\|^2) \\ &\leq c_2\left(\tau^2 \|\mathbf{x}^r - \text{prox}_{f_1}[\mathbf{x}^r - \nabla f_2(\mathbf{x}^r)]\|^2 + \|\mathbf{x}^r - \mathbf{x}^{r+1}\|^2\right) \\ &\leq \frac{c_2 \tau^2}{\min\{1, \underline{\alpha}^2\}} \left(\|\mathbf{x}^r - \text{prox}_{\alpha^r f_1}[\mathbf{x}^r - \alpha^r \nabla f_2(\mathbf{x}^r)]\|^2\right) + c_2 \|\mathbf{x}^r - \mathbf{x}^{r+1}\|^2.\end{aligned}$$

Next we use (3.2) and non-expansiveness of the proximity operator to bound

$$\begin{aligned}\|\mathbf{x}^r - \text{prox}_{\alpha^r f_1}[\mathbf{x}^r - \alpha^r \nabla f_2(\mathbf{x}^r)]\|^2 &\leq 2(\|\mathbf{x}^r - \mathbf{x}^{r+1}\|^2 + \|\mathbf{x}^{r+1} \\ &\quad - \text{prox}_{\alpha^r f_1}[\mathbf{x}^r - \alpha^r \nabla f_2(\mathbf{x}^r)]\|^2) \\ &\leq 2(\bar{\alpha}^2 \kappa^2 + 1) \|\mathbf{x}^{r+1} - \mathbf{x}^r\|^2.\end{aligned}$$

This further implies

$$\begin{aligned}F(\mathbf{x}^{r+1}) - F^* &\leq c_2 \left(\frac{2\tau^2(1 + \bar{\alpha}^2 \kappa^2)}{\min\{1, \underline{\alpha}^2\}} + 1 \right) \|\mathbf{x}^r - \mathbf{x}^{r+1}\|^2 \\ &\leq \frac{c_2}{c_1} \left(\frac{2\tau^2(1 + \bar{\alpha}^2 \kappa^2)}{\min\{1, \underline{\alpha}^2\}} + 1 \right) (F(\mathbf{x}^r) - F(\mathbf{x}^{r+1})),\end{aligned}$$

where the last step is due to the sufficient decrease condition (4.1). This establishes the Q-linear convergence of $F(\mathbf{x}^r) \rightarrow F^*$. In light of the sufficient decrease

condition (4.1), this further implies the R-linear convergence of $\{\mathbf{x}^r\}$ to an optimal solution. The proof of Theorem 4.4 is complete. $\square$

An immediate corollary of Theorem 4.4 is that any APS algorithm with stepsizes generated according to (4.7) converges linearly, provided that the local error bound condition (2.7) holds.

5 Block Coordinate Descent Method as an APS Method

Let $\mathbf{x} \in \mathbb{R}^n$ have the block form of $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K)'$, where $\mathbf{x}_k \in \mathbb{R}^{i_k}$ and $\sum_{k=1}^K i_k = n$. Consider the minimization problem (1.1) in which f_1 is separable over the blocks, namely,

$$f_1(\mathbf{x}) = d_1(\mathbf{x}_1) + \dots + d_K(\mathbf{x}_K), \quad (5.1)$$

where $d_k, k = 1, \dots, K$ are all convex (but not necessarily smooth) functions. Furthermore, $\mathbf{X}$ is a closed convex set in $\mathbb{R}^n$ which is in the following Cartesian product form

$$\mathbf{X} = \mathbf{X}_1 \times \mathbf{X}_2 \times \dots \times \mathbf{X}_K, \quad (5.2)$$

where $\mathbf{X}_k$ is a closed convex subset of $\mathbb{R}^{i_k}$. Note that the LASSO problem (1.2), the group LASSO problem (1.3), and the logestic group LASSO problem (1.4) admit the decomposition specified by (5.1) and (5.2).

Consider the BCD method whereby after the r -th iteration, $r \geq 0$, we choose an index $s \in \{1, 2, \dots, K\}$ and compute the new iterate $\mathbf{x}^{r+1} = (\mathbf{x}_1^{r+1}, \mathbf{x}_2^{r+1}, \dots, \mathbf{x}_K^{r+1})$ as follows

$$\begin{aligned} \mathbf{x}_s^{r+1} &= \operatorname{argmin}_{\mathbf{x}_s \in \mathbf{X}_s} F(\mathbf{x}_1^r, \mathbf{x}_2^r, \dots, \mathbf{x}_{s-1}^r, \mathbf{x}_s, \mathbf{x}_{s+1}^r, \dots, \mathbf{x}_K^r) \\ \mathbf{x}_j^{r+1} &= \mathbf{x}_j^r, \quad j \neq s. \end{aligned} \quad (5.3)$$

where $(\mathbf{x}_1^r, \mathbf{x}_2^r, \dots, \mathbf{x}_K^r)$ denotes the iterating at r -th iteration. The blocks are chosen cyclically or essentially cyclically to be updated at every iteration. The essentially cyclic update ensures that there exists an integer $N \geq K$ such that after this many iterations all the blocks are updated at least once.

It is known [30] that the BCD algorithm with cyclic update or essentially cyclic update converges to the optimal solution for the set of non-smooth optimization problems that the non-smooth part is separable as defined in (5.1). To establish linear convergence of the BCD method, we need the assumption that the smooth part f_2 is strongly convex in each block, in the sense that there exists a scalar $\gamma \geq 0$ such that, for any $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K) \in \mathbf{X}$ and any $s \in \{1, 2, \dots, K\}$,

$$f_2(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{s-1}, \mathbf{x}_s + \Delta \mathbf{x}_s, \mathbf{x}_{s+1}, \dots, \mathbf{x}_K) - f_2(\mathbf{x}) - \langle \nabla_s f_2(\mathbf{x}), \Delta \mathbf{x}_s \rangle \geq \gamma \|\Delta \mathbf{x}_s\|^2, \quad (5.4)$$

for all feasible $\Delta \mathbf{x}_s \in \mathbb{R}^{i_s}$, where $\nabla_s f_2$ denotes the vector of partial derivatives of f_2 with respect to the s -th block. It is obvious that if the function f_2 is block coordinate-

wise strongly convex, then the coordinate descent method satisfies the sufficient decrease condition (4.1) (cf. Proposition 3.4 in [13]).

For the applications described in Sect. 1, the coordinate-wise strong convexity of f_2 imposes a mild condition on the linear operator $\mathbf{A}$. For example, in LASSO problem (1.2), if each column of $\mathbf{A}$ is non-zero and we consider each element in $\mathbf{x}$ to be a block, then the problem is coordinate-wise strongly convex. Furthermore, for the group LASSO problem, f_2 is block coordinate-wise strongly convex if the columns of $\mathbf{A}$ corresponding to a block are linearly independent. A similar condition can be derived for the logistic group LASSO problem (1.4) to ensure block coordinate-wise strong convexity. The following proposition shows that the block coordinate method for the L_1 norm minimization problem and the Group LASSO minimization problem is an APS method.

Proposition 5.1 *Under the assumptions (5.1), (5.2) and (5.4), the BCD method with cyclic update is an APS method satisfying (3.1)–(3.2) with some constant $\kappa > 0$.*

Proof Let us define two different iteration counters. The outer iteration index s is a counter for the number of updating cycles of the BCD algorithm, and the inner iteration index k corresponds to the variable block being updated in a given cycle. Thus, at iteration $r = sK + k$ (with $1 \leq k \leq K$), the k -th variable block is updated in the s -th cycle. Throughout the proof, the notation $\mathbf{x}^r$ means the r -th iteration of the BCD algorithm, and $\mathbf{x}_k^r$ represents the k -th block of r -th iteration.

For simplicity, let us assume that there are no constraints. This assumption is not restricting as one can always add the indicator functions of the constraining sets to the objective. Since the feasible set is assumed to have a special structure as in (5.2), the separability of the non-smooth objective component will still be preserved after this change.

The optimality condition at the r -th iteration for BCD method is,

$$\mathbf{g} + \nabla_k f_2(\mathbf{x}^r) = \mathbf{0}, \quad (5.5)$$

for some $\mathbf{g}$ in $\partial d_k(\mathbf{x}_k^r)$ (Note that we assumed $f_1(\mathbf{x}) = d_1(\mathbf{x}_1) + \cdots + d_K(\mathbf{x}_K)$). Now in each fixed cycle s , define $r' = Ks$ and the error vector $\mathbf{e}^s = (\mathbf{e}_1^s, \dots, \mathbf{e}_K^s)$, as follows

$$\mathbf{e}_k^s = \mathbf{x}_k^{r'+k} - \mathbf{x}_k^{r'} + \nabla_k f_2(\mathbf{x}^{r'}) - \nabla_k f_2(\mathbf{x}^{r'+k}), \quad \forall k = 1, \dots, K. \quad (5.6)$$

Then, it is obvious that $\mathbf{x}^{r'+K}$ generated by BCD can also be derived from the following update rule

$$\mathbf{x}^{(s+1)K} = \text{prox}_{f_1}(\mathbf{x}^{sK} - \nabla f_2(\mathbf{x}^{sK}) + \mathbf{e}^s). \quad (5.7)$$

Now we can show that $\|\mathbf{e}^s\| \leq \kappa \|\mathbf{x}^{(s+1)K} - \mathbf{x}^{sK}\|$ for some $\kappa > 0$. Since f_2 has Lipschitz continuous gradient, it follows that

$$\begin{aligned}\|e_k^s\| &\leq \|x_k^{r'+k} - x_k^{r'}\| + \|\nabla_k f_2(x^{r'}) - \nabla_k f_2(x^{r'+k})\| \\ &\leq \|x_k^{r'+k} - x_k^{r'}\| + L\|x^{r'} - x^{r'+k}\| \\ &\leq (L+1)\|x^{r'+k} - x^{r'}\| \leq (L+1)\|x^{r'+K} - x^{r'}\| = (L+1)\|x^{(s+1)K} - x^{sK}\|,\end{aligned}$$

where the second step is due to the Lipschitz condition on ∇f_2 and the last inequality is due to the block coordinate-wise update in the algorithm. This further implies that

$$\|e^s\| \leq K(L+1)\|x^{(s+1)K} - x^{sK}\|,$$

so that the condition (3.2) holds with $\kappa = K(L+1)$. $\square$

Note that a similar argument can be used to show that the BCD algorithm with essentially cyclic update also lies in the APS framework, with an error term e satisfying (3.2).

The following result is a direct consequence of Corollary 2.4 and Proposition 5.1.

Proposition 5.2 *The BCD algorithm (with cyclic or essentially cyclic update) generates a sequence of iterates that converges R -linearly to a solution in X^* for the LASSO problem (1.2), the Group LASSO problem (1.3), and the logistic Group LASSO problem (1.4), provided that the objective function is block coordinate-wise strongly convex.*

To our knowledge, this is the first result which shows the linear convergence rate of the BCD algorithm under the local error bound condition. In two closely related works, the authors of [19, 32] established the linear convergence of the (Block) Coordinate Gradient Descent (CGD) algorithm for solving problem (1.1) under the assumptions (5.1)–(5.2) and the error bound condition. The BCD and CGD algorithms both exploit block coordinate-wise updates to solve the problem (1.1). However, unlike BCD which solves the exact subproblem in each iteration, CGD approximates the smooth component f_2 by a strictly convex quadratic function. Therefore, the analysis given in [19, 32] does not imply the linear convergence of BCD. In fact, with a similar argument to Proposition 5.1, it can be shown that CGD with (essentially) cyclic update rule lies in the APS framework with a properly defined error vector sequence. Other relevant works are [16, 23] which provide the convergence analysis for a general class of inexact BCD methods. However, their analysis does not provide convergence rate guarantees for the BCD algorithm when applied to problem (1.1). Finally, the readers can refer to [6] for a unified iteration complexity analysis for a family of BCD-type algorithms.

6 Concluding Remarks

In this paper, we have introduced the class of APS methods and established its linear convergence under some conditions (sufficient decrease and local error bound). This general result implies the linear convergence of the BCD algorithm for a class of

non-smooth convex problems for which the strong convexity assumption does not hold.

Appendix

Proof of Lemma 4.3

We define $h(\alpha) = \|\tilde{\nabla} f(\mathbf{x}, \alpha)\|$ for any $\alpha > 0$. Hence, the first part of the lemma is to show that $\alpha h(\alpha)$ is increasing with α .

From the definition of the proximity operator (1.7) and the proximal gradient (2.1), we have

$$\alpha h(\alpha) = \left\| \mathbf{x} - \operatorname{argmin}_{\mathbf{y} \in X} \left\{ \alpha f_1(\mathbf{y}) + \frac{1}{2} \|\mathbf{y} - (\mathbf{x} - \alpha \nabla f_2(\mathbf{x}))\|^2 \right\} \right\|.$$

By the change of variable $\mathbf{z} \triangleq \mathbf{y} - \mathbf{x}$, we have

$$\begin{aligned} \alpha h(\alpha) &= \left\| \operatorname{argmin}_{\mathbf{z} \in X'} \left\{ \alpha f_1(\mathbf{x} + \mathbf{z}) + \frac{1}{2} \|\mathbf{z} + \alpha \nabla f_2(\mathbf{x})\|^2 \right\} \right\| \\ &= \left\| \operatorname{argmin}_{\mathbf{z} \in X'} \left\{ \alpha f_1(\mathbf{x} + \mathbf{z}) + \frac{1}{2} \|\mathbf{z}\|^2 + \alpha \mathbf{z}^T \nabla f_2(\mathbf{x}) \right\} \right\|, \end{aligned}$$

where $X' = \{\mathbf{z} \mid \mathbf{z} = \mathbf{y} - \mathbf{x} \text{ for some } \mathbf{y} \in X\}$. Then, we have

$$\alpha h(\alpha) = \left\| \operatorname{argmin}_{\mathbf{z} \in X'} \left\{ \alpha g(\mathbf{z}) + \frac{1}{2} \|\mathbf{z}\|^2 \right\} \right\|, \quad (7.1)$$

where $g(\mathbf{z}) = f_1(\mathbf{x} + \mathbf{z}) + \mathbf{z}^T \nabla f_2(\mathbf{x})$ is a (non-smooth) convex function. Our goal now is to show that if $0 < \alpha_1 < \alpha_2$, then

$$\alpha_1 h(\alpha_1) = \|\mathbf{z}^*(\alpha_1)\| \leq \|\mathbf{z}^*(\alpha_2)\| = \alpha_2 h(\alpha_2)$$

where $\mathbf{z}^*(\alpha)$ denotes the optimal solution of

$$\min_{\mathbf{z} \in X'} \left\{ \alpha g(\mathbf{z}) + \frac{1}{2} \|\mathbf{z}\|^2 \right\}. \quad (7.2)$$

The optimality of $\mathbf{z}^*(\alpha_1)$ implies that

$$g(\mathbf{z}^*(\alpha_1)) + \frac{1}{2\alpha_1} \|\mathbf{z}^*(\alpha_1)\|^2 \leq g(\mathbf{z}) + \frac{1}{2\alpha_1} \|\mathbf{z}\|^2, \quad \forall \mathbf{z} \in X'$$

In particular, when $\mathbf{z}$ is set to $\mathbf{z}^*(\alpha_2)$, we have

$$g(\mathbf{z}^*(\alpha_1)) + \frac{1}{2\alpha_1} \|\mathbf{z}^*(\alpha_1)\|^2 \leq g(\mathbf{z}^*(\alpha_2)) + \frac{1}{2\alpha_1} \|\mathbf{z}^*(\alpha_2)\|^2.$$

Similarly, the optimality of $\mathbf{z}^*(\alpha_2)$ implies that

$$g(\mathbf{z}^*(\alpha_2)) + \frac{1}{2\alpha_2} \|\mathbf{z}^*(\alpha_2)\|^2 \leq g(\mathbf{z}^*(\alpha_1)) + \frac{1}{2\alpha_2} \|\mathbf{z}^*(\alpha_1)\|^2.$$

Adding up the last two inequalities yields

$$\left(\frac{\alpha_2 - \alpha_1}{2\alpha_1\alpha_2}\right) \|\mathbf{z}^*(\alpha_1)\| \leq \left(\frac{\alpha_2 - \alpha_1}{2\alpha_1\alpha_2}\right) \|\mathbf{z}^*(\alpha_2)\|.$$

Since $0 < \alpha_1 < \alpha_2$, the above inequality implies that $\|\mathbf{z}^*(\alpha_1)\| \leq \|\mathbf{z}^*(\alpha_2)\|$. Note that the convexity of f_1 (or equivalently g) was not used in this part.

Next we prove the second part of the lemma which states that $h(\alpha)$ is monotonically decreasing with α . Introducing the new variable $\mathbf{u} \triangleq \frac{1}{\alpha}\mathbf{z}$, the equation (7.1) can be rewritten as

$$h(\alpha) = \left\| \operatorname{argmin}_{\mathbf{u} \in X''} \left\{ \alpha g(\alpha \mathbf{u}) + \frac{1}{2} \alpha^2 \|\mathbf{u}\|^2 \right\} \right\|$$

or equivalently,

$$h(\alpha) = \left\| \operatorname{argmin}_{\mathbf{u} \in X''} \left\{ \frac{1}{\alpha} g(\alpha \mathbf{u}) + \frac{1}{2} \|\mathbf{u}\|^2 \right\} \right\|$$

where $X'' = \{\mathbf{u} \mid \mathbf{u} = \frac{1}{\alpha}(\mathbf{y} - \mathbf{x}), \text{ for some } \mathbf{y} \in X\}$. We define $\mathbf{u}^*(\alpha)$ as the optimal solution of

$$h(\alpha) = \min_{\mathbf{u} \in X''} \left\{ \frac{1}{\alpha} g(\alpha \mathbf{u}) + \frac{1}{2} \|\mathbf{u}\|^2 \right\}. \quad (7.3)$$

It suffices to show that

$$h(\alpha_1) = \|\mathbf{u}^*(\alpha_1)\| \geq \|\mathbf{u}^*(\alpha_2)\| = h(\alpha_2),$$

for $0 < \alpha_1 < \alpha_2$. The first-order optimality condition of (7.3) at $\mathbf{u}^*(\alpha)$ implies

$$\mathbf{v} + \mathbf{u}^*(\alpha) = 0, \text{ for some } \mathbf{v} \in \partial g(\alpha \mathbf{u}^*(\alpha)), \quad (7.4)$$

where $\partial g(\alpha \mathbf{u}^*(\alpha))$ is the sub-differential set of the function g at the point $\alpha \mathbf{u}^*(\alpha)$. Rewriting (7.4) for $\mathbf{u}^*(\alpha_1)$ and $\mathbf{u}^*(\alpha_2)$, we obtain

$$\mathbf{v}_1 + \mathbf{u}^*(\alpha_1) = 0, \text{ for some } \mathbf{v}_1 \in \partial g(\alpha_1 \mathbf{u}^*(\alpha_1)), \quad (7.5)$$

$$\mathbf{v}_2 + \mathbf{u}^*(\alpha_2) = 0, \text{ for some } \mathbf{v}_2 \in \partial g(\alpha_2 \mathbf{u}^*(\alpha_2)). \quad (7.6)$$

Since g is a convex function, ∂g is a monotone mapping [24]. Therefore, $\mathbf{v}_1 \in \partial g(\alpha_1 \mathbf{u}^*(\alpha_1))$ and $\mathbf{v}_2 \in \partial g(\alpha_2 \mathbf{u}^*(\alpha_2))$ imply

$$\langle \alpha_1 \mathbf{u}^*(\alpha_1) - \alpha_2 \mathbf{u}^*(\alpha_2), \mathbf{v}_1 - \mathbf{v}_2 \rangle \geq 0. \quad (7.7)$$

Combining (7.7) with (7.5) and (7.6) implies that

$$\langle \alpha_1 \mathbf{u}^*(\alpha_1) - \alpha_2 \mathbf{u}^*(\alpha_2), \mathbf{u}^*(\alpha_2) - \mathbf{u}^*(\alpha_1) \rangle \geq 0.$$

Define $\mathbf{d} = \mathbf{u}^*(\alpha_2) - \mathbf{u}^*(\alpha_1)$. Then, the above inequality can be written as

$$\langle (\alpha_1 - \alpha_2) \mathbf{u}^*(\alpha_1) - \alpha_2 \mathbf{d}, \mathbf{d} \rangle \geq 0$$

which yields

$$\alpha_2 \|\mathbf{d}\|^2 \leq (\alpha_1 - \alpha_2) \langle \mathbf{u}^*(\alpha_1), \mathbf{d} \rangle.$$

Since $\alpha_1 - \alpha_2 < 0$, we have

$$\langle \mathbf{u}^*(\alpha_1), \mathbf{d} \rangle \leq \frac{\alpha_2}{\alpha_1 - \alpha_2} \|\mathbf{d}\|^2. \quad (7.8)$$

Now we can write that

$$\begin{aligned} \|\mathbf{u}^*(\alpha_2)\|^2 &= \|\mathbf{u}^*(\alpha_1)\|^2 + 2\langle \mathbf{u}^*(\alpha_1), \mathbf{d} \rangle + \|\mathbf{d}\|^2 \\ &\leq \|\mathbf{u}^*(\alpha_1)\|^2 + \frac{2\alpha_2}{\alpha_1 - \alpha_2} \|\mathbf{d}\|^2 + \|\mathbf{d}\|^2 \\ &= \|\mathbf{u}^*(\alpha_1)\|^2 + \frac{\alpha_1 + \alpha_2}{\alpha_1 - \alpha_2} \|\mathbf{d}\|^2 \\ &\leq \|\mathbf{u}^*(\alpha_1)\|^2, \end{aligned}$$

where the first inequality is due to (7.8) and the second inequality is due to the fact that $0 < \alpha_1 < \alpha_2$. This completes the proof.

References

- [1] Agarwal, A., Negahban, S., Wainwright, M.: Fast global convergence of gradient methods for high-dimensional statistical recovery. *Ann. Stat.* **40**, 2452–2482 (2012)
- [2] Bach, F.: Consistency of the group lasso and multiple kernel learning. *J. Mach. Learn. Res.* **9**, 1179–1225 (2008)
- [3] Berg, E.V.D., Schmidt, M., Friedlander, M., Murphy, K.: Group sparsity via linear-time projection, Technical Report TR-2008-09. Department of Computer Science, University of British Columbia, 2008
- [4] Combettes, P.L., Wajs, V.R.: Signal recovery by proximal forward-backward splitting. *Multiscale Model. Simul.* **4**, 1168–1200 (2005)
- [5] Dembo, R.S., Tulowizki, U.: Local convergence analysis for successive inexact quadratic programming methods, Working Paper. School of Organization and Management, Yale University, New Haven, (1984)
- [6] Hong, M., Wang, X., Razaviyayn, M., Luo, Z.-Q.: Iteration complexity analysis of block coordinate descent methods. *arXiv preprint*, <http://arxiv.org/abs/1310.6957> arXiv:1310.6957 (2013)
- [7] Kim, D., Sra, S., Dhillon, I.: A scalable trust-region algorithm with applications to mixed-norm regression. *International Conference of Machine Learning (ICML)*, vol. 1, 2010

- [8] Korpelevich, G.M.: The extra-gradient method for finding saddle points and other problems. *Ekon. i Mat Metody* **12**, 747–756 (1976). Translated to English as Matecon
- [9] Li, W.: Remarks on the convergence of the matrix splitting algorithm for the symmetric linear complementarity problem. *SIAM J. Optim.* **3**, 155–163 (1993)
- [10] Liu, J., Ji, S., Ye, J.: Sparse learning with efficient projections. Arizona State University, (2009)
- [11] Luo, Z.-Q., Tseng, P.: Error bound and convergence analysis of matrix splitting algorithm for the affine variational inequality problem. *SIAM J. Optim.* **2**, 43–54 (1992)
- [12] Luo, Z.-Q., Tseng, P.: On the linear convergence of descent methods for convex essentially smooth minimization. *SIAM J. Control Optim.* **3**, 408–425 (1992)
- [13] Luo, Z.-Q., Tseng, P.: Error bounds and convergence analysis of feasible descent methods: a general approach. *Ann. Oper. Res.* **46**, 157–178 (1993)
- [14] Luo, Z.-Q., Tseng, P.: Analysis of an approximate gradient projection method with applications to the back-propagation algorithm. *Optim. Methods Softw.* **4**, 85–101 (1994)
- [15] Ma, S., Song, X., Huang, J.: Supervised group lasso with applications to microarray data analysis. *BMC Bioinform.* **8**, 60 (2007)
- [16] Mairal, J.: Incremental majorization-minimization optimization with application to large-scale machine learning. arXiv preprint, <http://arxiv.org/abs/1402.4419> (2014)
- [17] Mangasarian, O.L.: Convergence of iterates of an inexact matrix splitting algorithm for the symmetric monotone linear complementarity problem. *SIAM J. Optim.* **1**, 114–122 (1991)
- [18] Meier, L., Geer, S.V.D., Bühlmann, P.: The group lasso for logistic regression. *J. R. Stat. Soc. B* **70**, 53–71 (2008)
- [19] Necoara, I., Clipici, D.: Distributed random coordinate descent method for composite minimization. arXiv preprint <http://arxiv.org/abs/1312.5302> (2013)
- [20] Pang, J.-S.: On the convergence of the basic iterative method for the implicit complementarity problem. *J. Optim. Theory Appl.* **37**, 149–162 (1982)
- [21] Pang, J.-S.: Inexact newton methods for nonlinear complementarity problem. *Math. Prog.* **36**(1), 54–71 (1986)
- [22] Pang, J.-S.: A posteriori error bounds for the linearly-constrained variational inequality problem. *Math. Oper. Res.* **12**(3), 474–484 (1987)
- [23] Razaviyayn, M., Hong, M., Luo, Z.-Q.: A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM J. Optim.* **23**, 1126–1153 (2013)
- [24] Rockafellar, R.T.: *Convex Analysis*. Princeton University Press, Princeton (1970)
- [25] Rosen, J.B.: The gradient projection method for non-linear programming, part I linear constraints. *J. Soc. Ind. Appl. Math.* **8**, 181–217 (1960)
- [26] Rosen, J.B.: The gradient projection method for non-linear programming, part II non-linear constraints. *J. Soc. Ind. Appl. Math.* **9**, 514–532 (1961)
- [27] Roth, V., Fischer, B.: The group-lasso for generalized linear models: uniqueness of solutions and efficient algorithms. In: *Proceedings of the 25th International Conference on Machine Learning*, ACM, 2008, pp. 848–855
- [28] Schmidt, M., Roux, N., Bach, F.: Convergence rates of inexact proximal-gradient methods for convex optimization. arxiv preprint, <http://arxiv.org/abs/1109.2415> (2011)
- [29] Tibshirani, R.: Regression shrinkage and selection via lasso. *J. R. Stat. Soc. B* **58**, 267–288 (1996)
- [30] Tseng, P.: Convergence of block coordinate method for non-differentiable minimization. *J. Optim. Theory Appl.* **109**, 475–494 (2001)
- [31] Tseng, P.: Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Math. Program.* **125**(2), 263–295 (2009)
- [32] Tseng, P., Yun, S.: A coordinate gradient descent method for nonsmooth separable minimization. *Math. Program.* **117**, 387–423 (2009)
- [33] Wright, S., Nowak, R., Figueiredo, M.: Sparse reconstruction by separable approximation. *IEEE Trans. Signal Process.* **57**, 2479–2493 (2009)
- [34] Yang, J., Sun, D., Toh, K.-C.: A proximal point algorithm for log-determinant optimization with group lasso regularization. *SIAM J. Optim.* **23**, 857–893 (2013)
- [35] Yuan, M., Lin, Y.: Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. B* **68**, 49–67 (2006)
- [36] Zhang, H., Jiang, J., Luo, Z.-Q.: On the linear convergence of proximal splitting method for a class of nonsmooth convex minimization problems. *J. Oper. Res. Soc. China* **1**, 163–186 (2013)