

[Mingyi Hong,
Meisam Razaviyayn,
Zhi-Quan Luo, and
Jong-Shi Pang]

A Unified Algorithmic Framework for Block-Structured Optimization Involving Big Data

[With applications in machine learning and signal processing]

This article presents a powerful algorithmic framework for big data optimization, called the *block successive upper-bound minimization (BSUM)*. The BSUM includes as special cases many well-known methods for analyzing massive data sets, such as the block coordinate descent (BCD) method, the convex-concave procedure

(CCCP) method, the block coordinate proximal gradient (BCPG) method, the nonnegative matrix factorization (NMF) method, the expectation maximization (EM) method, etc. In this article, various features and properties of the BSUM are discussed from the viewpoint of design flexibility, computational efficiency, parallel/distributed implementation, and the required communication overhead. Illustrative examples from networking, signal processing, and machine learning are presented to demonstrate the practical performance of the BSUM framework.

Digital Object Identifier 10.1109/MSP.2015.2481563
Date of publication: 29 December 2015

INTRODUCTION


OVERVIEW OF OPTIMIZATION FOR BIG DATA

With advances in sensor, communication, and storage technologies, data acquisition is now more ubiquitous than at any other time. This has made available big data sets in many areas of engineering, biological, social, and physical sciences. While the proper modeling and analysis of such data sets can yield valuable information for inference, estimation, tracking, learning, and decision making, their size and complexity present great challenges in algorithm design and implementation.

Due to its central role in big data analytics, large-scale optimization has recently attracted significant attention not only from the optimization community, but also from the machine-learning, statistics, and signal processing communities. For example, emerging problems in image processing, social network analysis, and computational biology can easily exceed millions or billions of variables, and significant research is underway to enable fast accurate solutions to these problems [1]–[4]. Also, problems related to the design and provision of large-scale smart infrastructures such as wireless communication networks require real-time efficient resource allocation decisions to ensure optimal network performance. Traditional general-purpose optimization tools are inadequate for these problems due to the complexity of the model, the heterogeneity of the data, and, most importantly, the sheer data size [5]–[9]. Modern large-scale optimization algorithms, especially those that are capable of exploiting problem structures; dealing with distributed, time-varying, and incomplete data sets; and utilizing massively parallel computing and storage infrastructures, have become the workhorse in the big data era.

To be efficient for big data applications, optimization algorithms must have certain properties:

- 1) Each of their computational steps must be simple and easy to perform.
- 2) The intermediate results are easily stored.



[FIG1] The BCD method for a two-dimensional problem. The dashed curves represent the contours of the objective function, the solid lines represent the progress of the algorithm.

3) They can be implemented in a distributed and/or parallel manner so as to exploit the modern multicore and cluster computing architectures.

4) A high-quality solution can be found using a small number of iterations.

These requirements preclude the use of high-order information about the problem (i.e., the Hessian matrix of the objective), which is usually too expensive to obtain, even for modest-sized problems.

THE BCD METHOD

A very popular family of optimization algorithms that satisfies most of the aforementioned properties is the BCD method, sometimes known as the *alternating minimization/maximization (AM)* algorithm. The basic steps of the BCD are simple: 1) partition the entire optimization variables into small blocks and 2) optimize one block variable (or few blocks of variables) at each iteration while holding the remaining variables fixed. More specifically, consider the following block-structured optimization problem

$$\underset{x}{\text{minimize}} \quad f(x_1, x_2, \dots, x_n), \quad \text{s.t. } x_i \in \mathcal{X}_i, i = 1, 2, \dots, n, \quad (1)$$

where $f(\cdot)$ is a continuous function (possibly nonconvex, nonsmooth), each $\mathcal{X}_i \subseteq \mathbb{R}^{m_i}$ is a closed convex set, and each x_i is a block variable, $i = 1, 2, \dots, n$. Define $x := (x_1, \dots, x_n) \in \mathbb{R}^m$, and let $\mathcal{X} := \mathcal{X}_1 \times \dots \times \mathcal{X}_n \subseteq \mathbb{R}^m$. When applying the classical BCD method to solve (1), at every iteration r , a single block of variables, say $i = (r \bmod n) + 1$, is optimized by solving the following problem

$$x_i^r \in \underset{x_i \in \mathcal{X}_i}{\text{argmin}} \quad f(x_i, x_{-i}^{r-1}), \quad (2)$$

where we have defined $x_{-i}^{r-1} := (x_1^{r-1}, \dots, x_{i-1}^{r-1}, x_{i+1}^{r-1}, \dots, x_n^{r-1})$; for the rest of the variables $j \neq i$, let $x_j^r = x_j^{r-1}$. See Figure 1 for a graphical illustration of the algorithm.

The BCD algorithm is intuitively appealing and very simple to implement, yet it is extremely powerful. It enjoys tremendous popularity in a wide range of applications from signal processing communications to machine learning. Representative examples include image deblurring [10], statistical learning [11], wireless communications [12], etc. In recent years, there has been a surge of renewed interests to extend and generalize the BCD type of algorithms due to its applications in modern big data optimization problems. Theoretically, the BCD algorithm and its variants have been generalized significantly to accommodate various coordinate update rules [5], [13], [14] and have been made suitable for implementation on modern parallel processing architecture [15]–[19]. It can handle a wide range of nonsmooth, nonconvex cost functions [20]–[22]. Practically, it has been used in emerging large-scale signal processing and machine-learning applications, such as compressive sensing/sparse signal recovery [23], [24], matrix completion [25], [26], and tensor decomposition [27], [28], to name just a few. A recent survey of this algorithm can be found in [29].

THE BSUM METHOD

In this article, we introduce a unifying framework, the BSUM method, which generalizes the BCD type of algorithms [21]. Simply put, the BSUM framework includes algorithms that successively optimize certain upper bounds or surrogate functions of the original objectives, possibly in a block-by-block manner. The BSUM framework significantly expands the application domain of the traditional BCD algorithms. For example, it covers many classical statistics and machine-learning algorithms such as the EM method [30], the CCCP [31], and the multiplicative NMF [32]. It also includes as special cases many well-known signal processing algorithms such as the family of interference pricing algorithms [33], [34] and the weighted minimum mean square error algorithms [12], [35] for interference management in large-scale wireless systems.

The generality and flexibility of the BSUM offers an ideal platform to explore optimization algorithms for big data. Through the lens of the BSUM, one obtains not only a thorough understanding of the variety of algorithms and applications that are being covered, but more importantly the principle for designing a good algorithm suitable for big data. To this end, this article will first provide a concise overview of a few key theoretical characterizations of the algorithms that fall under the BSUM framework, BCD included. Our emphasis will be given on providing intuitive understanding as to when and where the BSUM framework should (or should not) work, and how its performance can be characterized. The second part of this article offers a detailed account of many existing large-scale optimization algorithms that fall under the BSUM framework, together with a few big data related applications that are of significant interests to the signal processing community. The last part of the article outlines some interesting extensions of the BSUM that further help expand its application domains. Throughout this article, special emphasis will be given on computational issues arising from big data optimization problems such as algorithm design for parallel and distributed computation, algorithm implementation

[TABLE 1] A PSEUDOCODE OF THE BSUM ALGORITHM.

```

1 FIND A FEASIBLE POINT  $x^0 \in \mathcal{X}$  AND SET  $r = 0$ 
2 REPEAT
3   PICK INDEX SET  $\mathcal{I}^r$ 
4   LET  $x_i^r \in \operatorname{argmin}_{x_i \in \mathcal{X}_i} u_i(x_i, x^{r-1}), \forall i \in \mathcal{I}^r$ 
5   SET  $x_k^r = x_k^{r-1}, \forall k \notin \mathcal{I}^r$ 
6    $r = r + 1$ ,
7 UNTIL SOME CONVERGENCE CRITERION IS MET

```

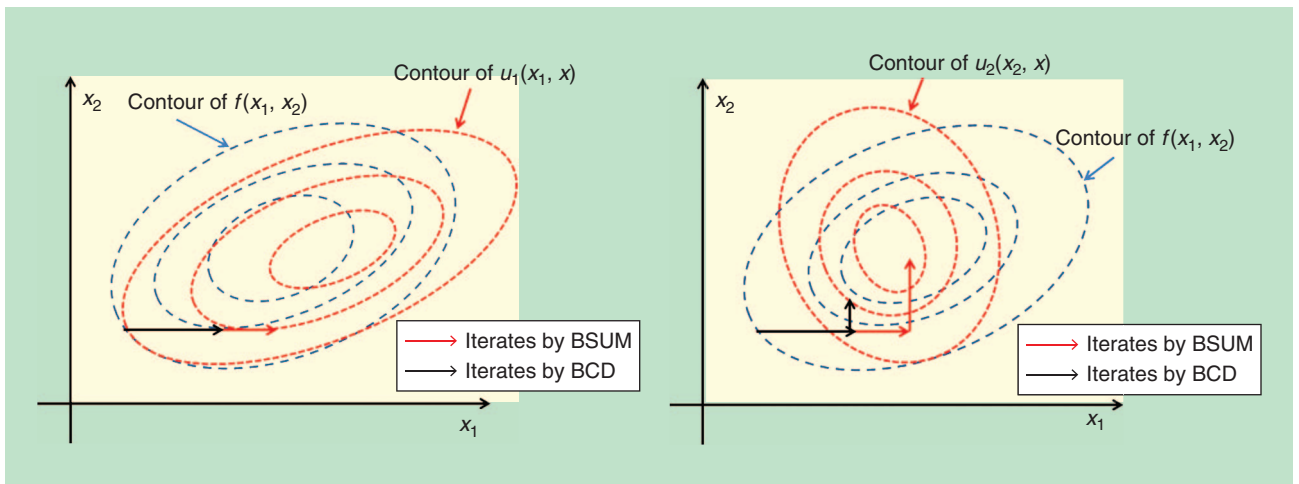
on multicore computing platforms, and distributed storage of the data. In particular, we will discuss how these issues can be addressed in the BSUM framework by providing theoretical insights and examples from real practical problems.

BSUM AND ITS THEORETICAL PROPERTIES

THE MAIN IDEA

We start by providing a high-level description of the BSUM method and some of its theoretical properties. In practice, one of the main problems of directly applying BCD to solve (1) is that each of its subproblems (2) is often difficult to solve exactly, especially when $f(x)$ is nonconvex. Moreover, even such exact minimization can be performed, the BCD may not necessarily converge. One of the key insights to be offered by the BSUM framework is that, for both practical and theoretical considerations, obtaining an approximate solution of (2) is good enough to keep the algorithm going. To be more specific, let us introduce $u_i(x_i, z): \mathcal{X}_i \rightarrow \mathbb{R}$ as an approximation function of $f(x_i, z_{-i})$ for each coordinate i at a given feasible point $z \in \mathcal{X}$. Let us define a set \mathcal{I}^r (possibly with $|\mathcal{I}^r| > 1$) as the block-variable indices to be picked at iteration r . Then at each iteration r , the BSUM method performs the following simple update

$$\begin{cases} x_i^r \in \operatorname{argmin}_{x_i \in \mathcal{X}_i} u_i(x_i, x^{r-1}), \forall i \in \mathcal{I}^r \\ x_k^r = x_k^{r-1}, \forall k \notin \mathcal{I}^r \end{cases}. \quad (3)$$



[FIG2] A comparison of BSUM and BCD methods for solving a two-dimensional problem. Each time a single coordinate is picked for update. The blue dashed curves represent the contours of the objective function $f(x_1, x_2)$, the red dotted curves represent the contours of the upper-bound functions $u_1(x_1, x)$ and $u_2(x_2, x)$; the black (resp. red) solid lines represent the progress of the BCD (resp. BSUM).

[TABLE 2] THE COMMONLY USED COORDINATE SELECTION RULES.

AT EACH ITERATION r , DEFINE A SET OF AUXILIARY VARIABLES $\{\tilde{x}_i^r\}_{i=1}^n$ AS:

$$\tilde{x}_i^r \in \operatorname{argmin}_{x_i \in X_i} u_i(x_i, x^{r-1}), i = 1, \dots, n.$$

THEN WE HAVE THE FOLLOWING COMMONLY USED COORDINATE SELECTION RULES.

- **CYCLIC RULE:** THE COORDINATES ARE CHOSEN CYCLICALLY, I.E., IN THE ORDER OF $1, 2, \dots, n, 1, 2, \dots$.
- **ESSENTIALLY CYCLIC (E-C) RULE:** THERE EXISTS A GIVEN PERIOD $T \geq 1$ DURING WHICH EACH BLOCK IS UPDATED AT LEAST ONCE, I.E.,

$$\bigcup_{i=1}^T \mathcal{I}^{r+i} = \{1, \dots, n\}, \forall r.$$

- **GAUSS-SOUTHWELL (G-SO) RULE:** AT EACH ITERATION r , \mathcal{I}^r CONTAINS A SINGLE INDEX i^* THAT SATISFIES:

$$i^* \in \{i \mid \|\tilde{x}_i^r - x_i^{r-1}\| \geq q \max_j \|\tilde{x}_j^r - x_j^{r-1}\|\}$$

FOR SOME CONSTANT $q \in (0, 1]$.

- **MAXIMUM BLOCK IMPROVEMENT (MBI) RULE:** AT EACH ITERATION r , \mathcal{I}^r CONTAINS A SINGLE INDEX i^* THAT ACHIEVES THE BEST OBJECTIVE: $i^* \in \operatorname{argmin}_i f(\tilde{x}_i^r, x^{r-1})$.
- **RANDOMIZED RULE:** LET $p_{\min} \in (0, 1)$ BE A CONSTANT. AT EACH ITERATION r , THERE EXISTS A PROBABILITY VECTOR $p^r = (p_1^r, \dots, p_n^r) \in \mathbb{R}^n$ SATISFYING $\sum_{i=1}^n p_i^r = 1$ AND $p_i^r > p_{\min}, \forall i$, WITH \mathcal{I}^r CONTAINING A SINGLE RANDOM INDEX i_r^* DETERMINED BY

$$\Pr(i \in \mathcal{I}^r \mid x^{r-1}, x^{r-2}, \dots, x^0) = p_i^r, \forall i.$$

The complete description of the BSUM is given in Table 1. See Figure 2 for a graphical comparison of the iterates generated by BSUM and BCD for a two-dimensional problem. It should be clear at this point that when $\mathcal{I}^r = \{(r \bmod n) + 1\}$ and no approximation is used [i.e., $u_i(x_i, z) = f(x_i, z_{-i})$] and at each iteration a single coordinate is selected, then the BSUM reduces to the classical cyclic BCD method. In Table 2, we also present several index selection rules that are covered by the BSUM framework. For simplicity of presentation, we will use the classical cyclic index selection rule, where $\mathcal{I}^r = \{(r \bmod n) + 1\}$, in the remainder of this article unless otherwise noted.

Next, we introduce the precise definition of the approximation function. The main idea is that, for each i , the approximation $u_i(x_i, x^r)$ should be an upper bound of the original objective function at the point of x^r (hence, the BSUM name of the framework); see Figure 3 for an illustration of the upper-bound minimization process. Intuitively, picking an upper-bound approximation function is reasonable because optimizing it at least should guarantee some descent of the original objective f ; see Figure 3(c).

To be more precise, let us first define the directional derivative of a given function $f(x): X \rightarrow \mathbb{R}$ at a point $x \in X$ in direction d :

$$f'(x; d) \triangleq \liminf_{\lambda \downarrow 0} \frac{f(x + \lambda d) - f(x)}{\lambda}.$$

Using this definition, we make the following assumptions on the u_i s.

Assumption A


$$u_i(x_i, x) = f(x), \quad \forall x \in X, \forall i \quad (\text{A1})$$

$$u_i(x_i, z) \geq f(x_i, z_{-i}), \quad \forall x_i \in X_i, \forall z \in X, \forall i \quad (\text{A2})$$

$$u_i'(x_i, z; d_i) \mid_{x_i=z_i} = f'(z; d), \\ \forall d = (0, \dots, d_i, \dots, 0) \text{ s.t. } z_i + d_i \in X_i, \forall i \quad (\text{A3})$$

$$u_i(x_i, z) \text{ is continuous in } (x_i, z), \quad \forall i. \quad (\text{A4})$$

Intuitively, Assumptions (A1) and (A2) imply that the approximation function is a global upper bound of $f(x)$; while the assumption (A3) guarantees that the first-order behaviors of the



[FIG3] The upper-bound minimization step of the BSUM method is shown. Here we assume that coordinate i is updated at iteration $r + 1$. It is clear from the figure that after solving the BSUM subproblem (3), $f(x_i^{r+1}, x_{-i}^r) < f(x_i^r, x_{-i}^r)$, that is, the objective function is strictly decreased.

[TABLE 3] COMMONLY USED UPPER BOUNDS SATISFYING ASSUMPTION A.

- **PROXIMAL UPPER BOUND:** GIVEN A CONSTANT $\gamma > 0$, ONE CAN CONSTRUCT A BOUND BY ADDING A QUADRATIC PENALIZATION (I.E., THE PROXIMAL TERM)

$$u_i(x_i, z) := f(x_i, z_{-i}) + \frac{\gamma}{2} \|x_i - z_i\|^2.$$

- **QUADRATIC UPPER BOUND:** SUPPOSE $f(x) = g(x_1, \dots, x_n) + h(x_1, \dots, x_n)$, WHERE g IS SMOOTH WITH \mathbf{H}_i AS THE HESSIAN MATRIX FOR THE i TH BLOCK. THEN ONE CAN CONSTRUCT THE FOLLOWING BOUND

$$u_i(x_i, z) := g(z_i, z_{-i}) + h(x_i, z_{-i}) + \langle \nabla_i g(z_i, z_{-i}), x_i - z_i \rangle + \frac{1}{2} (x_i - z_i)^T \Phi_i (x_i - z_i),$$

WHERE BOTH Φ_i AND $\Phi_i - \mathbf{H}_i$ ARE POSITIVE SEMIDEFINITE MATRICES.

- **LINEAR UPPER BOUND:** SUPPOSE f IS DIFFERENTIABLE AND CONCAVE, THEN ONE CAN CONSTRUCT

$$u_i(x_i, z) := f(z_i, z_{-i}) + \langle \nabla_i f(z_i, z_{-i}), x_i - z_i \rangle.$$

- **JENSEN'S UPPER BOUND:** SUPPOSE $f(x) := f(a_1^T x_1, \dots, a_n^T x_n)$, WHERE $a_i \in \mathbb{R}^{m_i}$ IS A COEFFICIENT VECTOR, AND f IS CONVEX WITH RESPECT TO EACH $a_i^T x_i$. LET $w_i \in \mathbb{R}^{m_i}$ DENOTE A WEIGHT VECTOR WITH $\|w_i\|_1 = 1$. THEN ONE CAN USE JENSEN'S INEQUALITY AND CONSTRUCT

$$u_i(x_i, z) := \sum_{j=1}^{m_i} w_i(j) f\left(\frac{a_i(j)}{w_i(j)} (x_i(j) - z_i(j)) + a_i^T z_{-i}\right),$$

WHERE $w_i(j)$ REPRESENTS THE j TH ELEMENT IN VECTOR w_i .

objective function and the approximation function are the same at the point of approximation (cf. Figure 3). In Table 3, we provide a few commonly used upper bounds that satisfy Assumption A; see also [21] and [22] for additional examples. More discussion will be given in subsequent sections on how these bounds are used in practice.

For a popular subclass of problem (1), Assumption A can be further simplified; see the following example [21, Proposition 2].

Example 1: Consider the following special form of problem (1):

$$\begin{aligned} \min f(x) := & g(x_1, \dots, x_n) \\ & + h(x_1, \dots, x_n), \quad \text{s.t. } x_i \in X_i, i = 1, \dots, n, \end{aligned} \quad (4)$$

where $g: X \rightarrow \mathbb{R}$ is a smooth function and $h: X \rightarrow \mathbb{R}$ is a possibly nonsmooth function whose directional derivative exists at every point $x \in X$. Consider $u_i(x_i, z) = \hat{u}_i(x_i, z) + h(x)$, where $\hat{u}_i(x_i, z)$ approximates the smooth function g in the objective. Then assumption (A3) is implied by (A1) and (A2) and is therefore no longer needed.

Now that we have seen the main steps of the algorithm, we describe its theoretical properties. We address the following questions related to the convergence of the BSUM: When does the BSUM converge? How fast does it converge? When does the BSUM fail and why? The answers to these questions will be instrumental in understanding the framework as well as evaluating the performance of various related algorithms.

WHEN DOES THE BSUM CONVERGE?

To discuss the convergence property of the algorithm, we first investigate the optimality conditions for (1), which characterize the set of solutions that we would like our algorithm to reach. To this end, we introduce two related notions, one is called the *stationary solution* and the other is the *coordinatewise minimum solution*; see Table 4 for their precise definitions. Roughly speaking, the coordinatewise minimum \hat{x} is the point where no single block \hat{x}_i , $i = 1, \dots, n$ has a better direction to move to, while at a stationary point x^* the entire vector cannot move to a better

direction. Further, a stationary solution must be a coordinatewise minimum, but the reverse direction is generally not true; see the example next.


Example 2: Consider the convex function $f(z) = \|Az\|$, where $A = [3 \ 4; 2 \ 1] \in \mathbb{R}^{2 \times 2}$. This function has a unique stationary solution $(z_1, z_2) = (0, 0)$, which is also the global optimal solution. Further, the point $\hat{z} = (-4, 3)$ is a coordinatewise minimum with respect to the two standard coordinates since $f'(\hat{z}; d) \geq 0, \forall d \in \{(d_1, d_2) \in \mathbb{R}^2 \mid d_1 d_2 = 0\}$; but this point is not a stationary solution as $f'(\hat{z}; \hat{d}) < 0$ for $\hat{d} = (4, -3)$. This fact can also be observed in the contour plot of the function in Figure 4.

This example confirms that the coordinatewise minimum can be much inferior to the stationary solution. Therefore in the subsequent discussion we will mainly focus on finding the stationary solutions rather than the coordinatewise minimum. An immediate question is: can one easily distinguish these two types of solutions, or for that matter, when does a coordinatewise minimum become a stationary solution? Let us define a regular point $x \in X$ as a point that satisfies the following statement: if x is coordinatewise minimum, then it is a stationary solution. For a large and popular subclass of (1) expressed below in (5), where the nonsmooth function is separable across the blocks, all feasible points $x \in X$ are regular

$$\min f(x) := g(x_1, \dots, x_n) + \sum_{i=1}^n h_i(x_i), \quad \text{s.t. } x_i \in X_i, i = 1, \dots, n. \quad (5)$$

[TABLE 4] OPTIMALITY CONDITIONS.

- **STATIONARY SOLUTIONS:** THE POINT x^* IS A STATIONARY SOLUTION OF (1) IF $f'(x^*; d) \geq 0$ FOR ALL d SUCH THAT $x + d \in X$. LET X^* DENOTE THE SET OF STATIONARY SOLUTIONS.
- **COORDINATEWISE MINIMUM SOLUTIONS:** $\hat{x} \in X$ IS COORDINATEWISE MINIMUM OF PROBLEM (1) WITH RESPECT TO THE COORDINATES IN $\mathbb{R}^{m_1}, \mathbb{R}^{m_2}, \dots, \mathbb{R}^{m_n}$, IF $f(\hat{x} + d_k^0) \geq f(\hat{x}), \forall d_k \in \mathbb{R}^{m_k}$ WITH $\hat{x} + d_k^0 \in X, \forall k = 1, 2, \dots, n$, WHERE $d_k^0 = (0, \dots, d_k, \dots, 0)$.



[FIG4] The contour plot of the function $f(z) = \|Az\|_1$, where $A = \begin{bmatrix} 3 & 4 \\ 2 & 1 \end{bmatrix} \in \mathbb{R}^{2 \times 2}$. The point $\hat{z} = (-4, 3)$ is a coordinatewise minimum but not a stationary solution [21]. (Figure reprinted with permission from the Society for Industrial and Applied Mathematics.)

The main convergence result for the BSUM method is given below, which is adapted from [21, Th. 2]

Theorem 1: Suppose the cyclic coordinate selection rule is chosen, i.e., $I^r = \{(r \bmod n) + 1\}$. Let $\{x^r\}_{r=1}^\infty$ be a sequence generated by the BSUM algorithm. Suppose Assumption A holds, and that each x^r is regular. Then the following is true:

- Suppose that the function $u_i(x_i, y)$ is quasi-convex in x_i for $i = 1, \dots, n$. Furthermore, assume that the subproblem (3) has a unique solution for any point $x^{r-1} \in X$. Then every limit point x^* of $\{x^r\}$ is a stationary point of (1).
- Suppose the level set $X^0 = \{x \mid f(x) \leq f(x^0)\}$ is compact. Furthermore, assume that the subproblem (3) has a unique solution for any point $x^{r-1} \in X$, $r \geq 1$ for at least $n - 1$ blocks. Then $\{x^r\}$ converge to the set of stationary points, i.e.,

$$\lim_{r \rightarrow \infty} d(x^r, X^*) = 0,$$

where $d(x, X^*) \triangleq \min_{x^* \in X^*} \|x - x^*\|$ and X^* is the set of stationary points.

Here, the first part of the result deals with the possibility of an unbounded sequence, whereas the second part assumes that the sequence lies in a compact set, therefore it has a slightly stronger claim. Note that Theorem 1 can be easily extended to all the coordinate selection rules given in Table 2, and for the randomized version the convergence is with probability 1.

Special attention should be given to the conditions set forth by Theorem 1. First, it says that the upper-bound function needs to be picked carefully to satisfy both Assumption A and the requirement that at least $n - 1$ subproblems (3) have a unique solution. However, when the objective function $f(x)$ is convex, the uniqueness requirement of the per-block solution can be dropped; see [36]. Also, Theorem 1 requires the problem to be well-defined so that coordinatewise optimal solutions are equivalent to the

stationary solution (which precludes objective function $\|Ax\|_1$ in Example 2). Next we will demonstrate, through a couple of concrete examples, that relaxing some of these conditions indeed leads to the divergence of the BSUM.

WHEN DOES BSUM FAIL?

Here we provide a few examples in which BSUM fails to converge to any stationary solutions. The examples in this section serve as reminders that in practice, to avoid those pitfalls, extreme care must be exercised when designing large-scale optimization algorithms.

Our first example comes from Example 2. It demonstrates the consequence of lacking the regularity condition.

Example 3: Consider the following unconstrained convex optimization problem $\min_x \left\| \sum_{i=1}^2 A_i x_i \right\|_1$, where $A_1 = [3 \ 2]^T$ and $A_2 = [4 \ 1]^T$. Clearly by defining $A = [3 \ 4; 2 \ 1]$, the objective function can be rewritten as $\|Ax\|_1$, which is not regular at point $(x_1, x_2) = (-4, 3)$ (cf. Example 2). It follows that, by setting $u_i(x_i, z) = \left\| \sum_{i=1}^2 A_i x_i \right\|_1$ (no approximation is used) and letting $(x_1^0, x_2^0) = (-4, 3)$, the BSUM algorithm will not be able to move forward for either x_1 or x_2 , thus it will be stuck at the non-interesting point $(-4, 3)$, far away from the only stationary solution $(0, 0)$.

The next example shows that BSUM fails to converge if the feasible set X is no longer a Cartesian product of feasible sets X_1, \dots, X_n , a fact that we have taken for granted so far.

Example 4: Consider the following simple quadratic problem:

$$\min x_1^2 + x_2^2, \text{ s.t. } x_1 + x_2 = 2.$$

The optimal objective value is two. Consider the BSUM algorithm with an arbitrary approximation function satisfying Assumption A, but initiated at the point $(x_1^0, x_2^0) = (0, 2)$. Then the BSUM method will be stuck at this noninteresting point without making any progress because it is not possible to change a single block without violating the coupling constraint.

Our next example shows that the BSUM could diverge if the approximation function u_i violates Assumption A.

Example 5: Consider the simple quadratic problem

$$\min x_1^2 + x_2^2 + 2x_1x_2, \text{ s.t. } -1 \leq x_1, x_2 \leq 1.$$

The optimal objective value is 0, with $(x_1^*, x_2^*) = (0, 0)$. Consider using the BSUM algorithm with a linear approximation function, which violates Assumption (A2). More specifically, for a given feasible tuple $(x_1, x_2) = (\hat{x}_1, \hat{x}_2)$, the x_1 's subproblem becomes

$$\min_{x_1} \langle \hat{x}_1 + \hat{x}_2, x_1 \rangle, \text{ s.t. } -1 \leq x_1 \leq 1.$$

Clearly, the optimal solution is either $x_1 = -1$ or $x_1 = 1$. The same happens when solving the subproblem for x_2 . Therefore the algorithm will never reach the optimal solution $(x_1^*, x_2^*) = (0, 0)$. Further, if the feasible sets of x_1 and x_2 are unbounded, then the BSUM can diverge.

Our last example is taken from Powell [37]. It shows that without the uniqueness assumption of each subproblem (3), the BSUM algorithm is not necessarily convergent.

Example 6: Consider the following unconstrained problem

$$\min f(x) := -x_1x_2 - x_2x_3 - x_3x_1 + (x_1 - 1)_+^2 + (-x_1 - 1)_+^2 + (x_2 - 1)_+^2 + (-x_2 - 1)_+^2 + (x_3 - 1)_+^2 + (-x_3 - 1)_+^2,$$

where the notation $(z)_+^2$ means $(\max\{0, z\})^2$. In this case, fixing (x_2, x_3) and optimizing over x_1 yields the following solution

$$x_1 = \begin{cases} (1 + \frac{1}{2}|x_2 + x_3|) \text{sign}(x_2 + x_3), & \text{if } x_2 + x_3 \neq 0, \\ [-1, 1], & \text{otherwise} \end{cases} \quad (6)$$

A similar solution can be obtained for x_2 and x_3 as well. Suppose we set $u_i(x_i, z) = f(x)$ for all i (no approximation is used) and let $(x_1^0, x_2^0, x_3^0) = (-1 - \epsilon, 1 + 1/2\epsilon, -1 - 1/4\epsilon)$ for some $\epsilon > 0$. Then it can be shown the applying the cyclic version of the BSUM algorithm, the iterates will be cycling around six points $(1, 1, -1)$, $(1, -1, -1)$, $(1, -1, 1)$, $(-1, -1, 1)$, $(-1, 1, 1)$, $(-1, 1, -1)$, and none of these six points is a stationary solution of the original problem. The reason for such pathological behavior is that, in any one of the six limit points above, there are at least two subproblems that have multiple optimal solutions. For example, at $(1, 1, -1)$, and fixing x_2, x_3 (resp. x_1, x_3), the optimal solution for x_1 (resp. x_2) is any element in the interval $[-1, 1]$; cf. (6).

A natural question at this point is, can we make the BSUM work for these examples? The answer is affirmative, but how this can be done requires a case by case study. To handle the first two examples (i.e., Example 3 and 4), a generalized version of BSUM is needed, which will be discussed in the “Extensions” section. For the third example, one can simply pick a better upper bound to guarantee convergence. For example, if we pick the proximal upper bound (cf. Table 3): $u_i(x_i, z) = f(x) + \gamma/2 \|x_i - z_i\|^2$, then each subproblem will have a unique solution, and the algorithm will not be trapped by the noninteresting solutions given in Example 6. Notice that the use of $\gamma/2 \|x_i - z_i\|^2$ inhibits the move of x_i from its current position z_i . Thus, the main message here is that when optimizing each block, it is beneficial, at least theoretically, to be less greedy so that a conservative step is taken towards reducing the objective. The extent of the “conservativeness” for the per block update is then naturally characterized by the chosen upper bounds. Quite interestingly, the difficulty with the nonunique subproblem solution can also be resolved by using randomization. Formally, we have the following corollary to Theorem 1 [38].

Corollary 1: Suppose the level set $X^0 = \{x | f(x) \leq f(x^0)\}$ is compact. Then, under the randomized block selection rule, the iterates generated by the BSUM algorithm converge to the set of stationary points almost surely, i.e.,

$$\lim_{r \rightarrow \infty} d(x^r, X^*) = 0, \text{ almost surely.}$$

HOW FAST DOES THE BSUM CONVERGE?

Now that we have examined the convergence of the BSUM, let us proceed next to characterize the convergence speed of the

algorithm. There is no doubt that this is an important issue, especially so for big data optimization—the sheer size of the data and limited computational resource makes it difficult to optimize a problem to global optimality. Consequently, we are generally satisfied with high-quality suboptimal solutions and are mostly concerned with how quickly such solutions can be obtained.

Recently, extensive research efforts have been devoted to analyzing the convergence rate for various BSUM-type algorithms, most of which use randomized coordinate selection rules and/or quadratic upper-bound functions (cf. Table 3) to solve convex problems; for example; see [5], [8], and [39]–[43]. Since it is not possible to go over all the details of these different variations of BSUM, here we present, at a high level, a fairly general result that covers a large family of upper-bound functions satisfying Assumption A, as well as all coordinate selection rules outlined in Table 2.

To this end, let us make the following additional assumptions.

Assumption B

- B1) $f(x) = g(x) + \sum_{i=1}^n h_i(x_i)$, where $g(x)$ is a smooth convex function with Lipschitz continuous gradient, i.e., there exists a constant L such that $\|\nabla g(x) - \nabla g(y)\| \leq L\|x - y\|$, $\forall x, y \in X$. Further both g and h_i s are convex functions.
- B2) The level set $\{x | f(x) \leq f(x^0), x \in X\}$ is compact.
- B3) Each upper-bound function $u_i(x_i, z)$ is strongly convex with respect to x_i .

An ϵ -optimal solution $x^\epsilon \in X$ is defined as an $x^\epsilon \in \{x | x \in X, f(x) - f(x^*) \leq \epsilon\}$, where $f(x^*)$ is the globally optimal objective value of problem (5). Suppose both Assumptions A and B are satisfied. Then it is shown in [38] and [42] that BSUM takes at most c/ϵ iterations to find an ϵ -optimal solution, for all coordinate rules specified in Table 2, where $c > 0$ is a constant only related to the description of the problem. Such a type of convergence rate is known as *sublinear convergence*. Here, the main message is that under Assumptions A and B, the algorithm generally converges sublinearly in the order of $1/\epsilon$. Further, for different special forms of BSUM, the constant c in front of $1/\epsilon$ can be significantly refined so that it is independent of problem dimension; see [5] and [8]. It is also interesting to note here that when the objective f is either strongly convex or convex but with certain special structure, the BSUM achieves the linear rate of convergence. That is, BSUM takes at most $O(\log(c/\epsilon))$ iterations to find an ϵ -optimal solution, which is much faster than the sublinear rate; see, e.g., [44] and [45] for the related discussions.

Finally, we briefly mention that it is possible to relax certain conditions in Assumption B to obtain refined rates. For example, [42] shows that dropping the per-block strong convexity assumption in (B3) still achieves an $O(1/\epsilon)$ sublinear convergence. In [38], [46] it is shown that, when removing the convexity Assumption (B1), it is also possible to characterize the convergence rate to stationary solutions. In [42] it is shown that when there are two blocks of variables, the cyclic version of the BSUM can be accelerated to achieve an improved $O(1/\sqrt{\epsilon})$ complexity. In a few recent works [47], [48], it is shown that when randomized block selection and the quadratic upper bound are used, it is possible to accelerate the BSUM with any $n > 2$ blocks.

ALGORITHMS COVERED BY THE BSUM FRAMEWORK

Now that we have a good understanding of the theoretical properties of the family of BSUM algorithms, we demonstrate in this section the wide applicability of the BSUM framework by revealing its connection to a few well-known algorithms for high-dimensional massive data analysis. For each of the examples presented below, we pay special attention to the design of the upper-bound functions.

THE BCD ALGORITHM

The first algorithm that the BSUM covers is obviously the classic BCD described in the “The BCD Method” section. Here the upper-bound function is simply the original objective itself, i.e., $u_i(x_i, z) := f(x_i, z_{-i})$, $\forall x_i \in \mathcal{X}_i, z \in \mathcal{X}, \forall i$. We would like to mention that all the convergence and rate of convergence analysis of BSUM carries over to the BCD method. Specifically, the result in the section “How Fast Does the BSUM Converge?” implies that the BCD method (with coordinate update rules specified in Table 2) converges in a sublinear manner whenever Assumptions (B1) and (B2) are satisfied. This result by itself is interesting, as there has been limited theoretical analysis for the general form of a BCD algorithm when applied to problems satisfying Assumptions (B1) and (B2), but not (B3).

THE CCCP

Consider the following unconstrained nonconvex problem, known as the *difference of convex (DC)* program: $\min f(x) := g_1(x) - g_2(x)$ where both g_1 and g_2 are convex functions. The well-known CCCP algorithm [31] generates a sequence $\{x^r\}$ by solving

$$x^{r+1} = \arg \min_x u(x, x^r),$$

where $u(x, x^r) := g_1(x) - \langle x - x^r, \nabla g_2(x^r) \rangle - g_2(x^r)$. Clearly, the linear upper bound in Table 3 is used here, therefore CCCP is a special case of the BSUM algorithm with a single block of variables. Furthermore, the updates can also be done in a block coordinate manner.

THE MAJORIZATION-MINIMIZATION ALGORITHM

The majorization-minimization (MM) algorithm, which has been widely used in statistics [49], can be viewed as the single block version of BSUM. Consider the problem of $\min_{x \in \mathcal{X}} f(x)$ where $f(x)$ is a smooth function. The basic idea of the MM algorithm is to first construct a “majorization” function $u(x, z)$ for the original objective $f(z)$, such that

$$u(x, z) \geq f(z), \forall x, z \in \mathcal{X}, \quad u(x, x) = f(x), \forall x \in \mathcal{X}. \quad (7)$$

Then the algorithm generates a sequence of iterates by successively minimizing $u(x, x^r)$. This algorithm represents a slight generalization of the CCCP discussed previously, but nevertheless still falls in the framework of BSUM.

As another concrete example of the MM algorithm, let us consider the celebrated EM algorithm [30]. Let w be a vector

observation used for estimating the parameter x . The maximum likelihood estimate of x is given as (where $p(w|x)$ denotes the conditional probability of w given x)

$$\hat{x}_{\text{ML}} = \arg \max_x \ln p(w|x). \quad (8)$$

Let the random vector z be some hidden/unobserved variable. The EM algorithm generates a sequence $\{x^r\}$ by iteratively performing the following two steps 1) E-Step: Calculate $g(x, x^r) := \mathbb{E}_{z|w, x^r} \{\ln p(w, z|x)\}$ and 2) M-Step: $x^{r+1} = \arg \max_x g(x, x^r)$. To see why the EM algorithm is a special case of MM (hence a special case of BSUM) let us rewrite (8) as: $\min_x -\ln p(w|x)$, the objective of which can be bounded by

$$\begin{aligned} -\ln p(w|x) &= -\ln \mathbb{E}_{z|x} p(w|z, x) \\ &= -\ln \mathbb{E}_{z|x} \left[\frac{p(z|w, x^r) p(w|z, x)}{p(z|w, x^r)} \right] \\ &= -\ln \mathbb{E}_{z|w, x^r} \left[\frac{p(z|x) p(w|z, x)}{p(z|w, x^r)} \right] \\ &\leq -\mathbb{E}_{z|w, x^r} \ln \left[\frac{p(z|x) p(w|z, x)}{p(z|w, x^r)} \right] \\ &= -\mathbb{E}_{z|w, x^r} \ln p(w, z|x) + \mathbb{E}_{z|w, x^r} \ln p(z|w, x^r) \\ &:= u(x, x^r), \end{aligned}$$

where the inequality is due to the fact that a convex function f must satisfy $\mathbb{E}[f(x)] \geq f(\mathbb{E}[x])$ (by the Jensen’s inequality). Since $\mathbb{E}_{z|w, x^r} \ln p(z|w, x^r)$ is not a function of x , the M-step can be written as

$$x^{r+1} = \arg \min_x u(x, x^r).$$

Furthermore, it is not hard to see that $u(x^r, x^r) = -\ln p(w|x^r)$, therefore, both conditions in (7) are satisfied. Similarly as in the previous case, one can extend the EM to a block coordinate version; see [21] for a detailed discussion.

THE PROXIMAL POINT ALGORITHM

The classical proximal point algorithm (PPA) (see, e.g., [50, Sec. 3.4.3]) obtains a solution of the problem $\min_{x \in \mathcal{X}} f(x)$ by solving the following equivalent problem

$$\min_{x \in \mathcal{X}, y \in \mathcal{X}} f(x) + \frac{\gamma}{2} \|x - y\|^2, \quad (9)$$

where $f(\cdot)$ is a convex function, \mathcal{X} is a closed convex set, and $\gamma > 0$ is a coefficient. Clearly (9) is strongly convex in both x and y so long as $f(x)$ is convex [but not jointly strongly convex in (x, y)]. This problem can be solved by performing the following two steps alternately

$$x^r = \arg \min_{x \in \mathcal{X}} \left\{ f(x) + \frac{1}{2c} \|x - y^{r-1}\|^2 \right\}, \quad y^r = x^{r+1}. \quad (10)$$

Equivalently, the PPA algorithm can be viewed as successively minimizing the single block version of the proximal upper bound $u(x; x^r)$ given in Table 3. Note that for a problem with a

single block of variables, if $x \in \mathcal{X}$ is coordinatewise minimum, then it must be a global minimum solution. Therefore, every feasible solution $x \in \mathcal{X}$ is regular, and the convergence of PPA is covered by BSUM.

Furthermore, the PPA can be generalized to solve the multiblock problem (1), where $f(\bullet)$ is convex in each of its block components, but not necessarily strictly convex. Directly applying BCD may fail to find a stationary solution for this problem, as the per-block subproblems can contain multiple solutions (cf. Example 6). Alternatively, we can consider an alternating PPA [51], which successively solves the following:

$$\min_{x_i} f(x_i, x_{-i}^r) + \frac{\gamma}{2} \|x_i - x_i^r\|^2, \text{ s.t. } x_i \in \mathcal{X}_i.$$

Clearly this algorithm is a special form of BSUM with strongly convex proximal upper bound (cf. Table 3). It follows that each subproblem has a unique optimal solution, and by Theorem 1 it must converge to a stationary solution.

THE FORWARD-BACKWARD SPLITTING ALGORITHM

The forward-backward splitting (FBS) algorithm (also known as the *proximal splitting algorithm*; see, e.g., [52] and the references therein) for nonsmooth optimization solves the composite problem (4) with a single block of variables (i.e., $n = 1$), where h is convex and lower semicontinuous; g is convex and has Lipschitz continuous gradient, i.e., $\|\nabla g(x) - \nabla g(y)\| \leq L\|x - y\|$, $\forall x, y \in \mathcal{X}$ and for some $L > 0$.

Define the proximity operator $\text{prox}_h: \mathcal{X} \rightarrow \mathcal{X}$ as

$$\text{prox}_h(x) = \arg \min_{y \in \mathcal{X}} h(y) + \frac{1}{2} \|x - y\|^2.$$

The FBS iteration is given below [52]:

$$x^{r+1} = \underbrace{\text{prox}_{\beta h}}_{\text{backward step}} \left(\underbrace{x^r - \beta \nabla g(x^r)}_{\text{forward step}} \right), \quad (11)$$

where $\beta \in (0, 1/L]$. Define

$$u(x, x^r) := h(x) + \frac{1}{2\beta} \|x - x^r\|^2 + \langle x - x^r, \nabla g(x^r) \rangle + g(x^r), \quad (12)$$

which is the quadratic upper bound in Table 3, with $\Phi_1 = 1/\beta I$. It is easy to see that the iteration (11) is equivalent to the following iteration $x^{r+1} = \arg \min_{x \in \mathcal{X}} u(x, x^r)$, therefore, it again falls under the BSUM framework.

Similar to the previous example, we can generalize the FBS algorithm to solve multiple block problems of the form (5). The resulting algorithm, sometimes also known as the *BCPG method*, has received significant attention recently due to its efficiency for solving certain big data optimization problem such as the least absolute shrinkage and selection operator (LASSO) [11]. For recent developments and applications for BCPG, see [5], [39], and [53].

Here we make a special note that, by appealing to the general convergence rate result in the section “How Fast Does the BSUM Converge?” the BCPG method with any coordinate selection rules in Table 2 gives a sublinear convergence rate, when it is used to solve (5) that satisfies Assumption B.

THE NMF ALGORITHM

Consider the following NMF problem:

$$\min_{W \in \mathbb{R}^{M \times K}, H \in \mathbb{R}^{K \times N}} f(W, H) := \frac{1}{2} \|V - WH\|_F^2, \text{ s.t. } W \geq 0, H \geq 0, \quad (13)$$

where $V \in \mathbb{R}_+^{M \times N}$ is given. The problem has been extensively studied since Lee and Seung’s seminal work [32], and it has wide applications in factor analysis, dictionary learning, speech analysis and so on [54]. In [32], a simple and efficient multiplicative algorithm is proposed:

$$[H^{r+1}]_{ji} = [H^r]_{ji} \frac{[(W^r)^T V]_{ji}}{[(W^r)^T W^r H^r]_{ji}}, \quad j = 1, \dots, K, i = 1, \dots, N \quad (14a)$$

$$[W^{r+1}]_{ji} = [W^r]_{ji} \frac{[V(H^r)^T]_{ji}}{[W^r H^{r+1} (H^r)^T]_{ji}}, \quad j = 1, \dots, M, i = 1, \dots, K. \quad (14b)$$

Here, $[H^{r+1}]_{ji}$ means the (j, i) th component of matrix H^{r+1} . In the following we show that when the iterates are well-defined (i.e., $[W_{ij}^r] > 0$ and $[H_{ij}^r] > 0$), the NMF iteration (14a) and (14b) is also covered by BSUM [55].

Let H_i and V_i represent the i th column of H and V , respectively. Then, at a given iterate $\{W^r, H^r\}$, the subproblem for optimizing H_i is given by

$$\min_{H_i \geq 0} f(H_i, \{W^r, H^r\}) := \frac{1}{2} \|V_i - W^r H_i\|_F^2 + \frac{1}{2} \sum_{j \neq i} \|V_j - W^r H_j^r\|_F^2. \quad (15)$$

Define the upper-bound function $u_i(H_i, \{W^r, H^r\})$ as

$$u_i(H_i, \{W^r, H^r\}) := f(H_i^r, \{W^r, H^r\}) + (H_i - H_i^r)^T \nabla_{H_i} f(H_i, \{W^r, H^r\}) + \frac{1}{2} (H_i - H_i^r)^T \Phi_i(W^r, H_i^r) (H_i - H_i^r),$$

where $\Phi_i(W^r, H_i^r)$ is a diagonal matrix given by

$$\Phi_i(W^r, H_i^r) := \text{Diag} \left(\frac{[(W^r)^T W^r H_i^r]_1}{[H_i^r]_1}, \dots, \frac{[(W^r)^T W^r H_i^r]_K}{[H_i^r]_K} \right).$$

Clearly, $\Phi_i(W^r, H_i^r) \succ 0$, and it is easy to show that $\Phi_i(W^r, H_i^r) \succ (W^r)^T W^r$, where $(W^r)^T W^r$ is the Hessian of the objective of (15) [32]. This implies that $u_i(H_i, \{W^r, H^r\})$ is the quadratic upper bound given in Table 3 of $f(H_i, \{W^r, H^r\})$. Further, one can check that the subproblem that minimizes $u_i(H_i, \{W^r, H^r\})$ has a unique solution, given by (14a). Similar analysis can be established for the W -block update rule as well. Therefore, we conclude that the iterates (14a) and (14b) are a special case of BSUM. Finally, we note that it is also possible to use different upper-bound functions to derive more efficient update rules for the NMF problem (13); see, e.g., [56], where both the concave upper bound and the Jensen’s upper bound (cf. Table 3) are used.

THE ITERATIVE REWEIGHTED LEAST SQUARES METHOD

The iterative reweighted least squares (IRLS) method is a popular algorithm used for solving big data problems such as sparse recovery [57]. Consider

$$\min_x h(x) + \sum_{j=1}^{\ell} (|A_j x + b_j|)_2, \text{ s.t. } x \in \mathbb{R}^m, \quad (16)$$

where $A_j \in \mathbb{R}^{k_i \times m}$, $b_j \in \mathbb{R}^{k_i}$, and $h(x)$ is some convex function not necessarily smooth. For a set of applications for this model, see [41, Sec. 4]. Consider the following smooth approximation of (16):

$$\min_x h(x) + g(x) := h(x) + \sum_{j=1}^{\ell} \sqrt{\|A_j x + b_j\|^2 + \eta^2}, \text{ s.t. } x \in \mathbb{R}^m, \quad (17)$$

where η is some small constant and $g(x)$ denotes the smooth part of the objective. The IRLS algorithm solves (17) by performing the following iteration:

$$x^{r+1} = \arg \min_{x \in \mathbb{R}^m} \left\{ h(x) + \frac{1}{2} \sum_{j=1}^{\ell} \frac{\|A_j x + b_j\|^2 + \eta^2}{\sqrt{\|A_j x^r + b_j\|^2 + \eta^2}} \right\}.$$

Define the following function for $g(x)$:

$$u(x, x^r) = \frac{1}{2} \left(\sum_{j=1}^{\ell} \frac{\|A_j x + b_j\|^2 + \eta^2}{\sqrt{\|A_j x^r + b_j\|^2 + \eta^2}} + \sqrt{\|A_j x^r + b_j\|^2 + \eta^2} \right). \quad (18)$$

It is clear that $g(x^r) = u(x^r, x^r)$, so Assumption (A1) is satisfied. To verify Assumption (A2), we apply the arithmetic-geometric inequality, and have

$$\begin{aligned} u(x, x^r) &= \frac{1}{2} \left(\sum_{j=1}^{\ell} \frac{\|A_j x + b_j\|^2 + \eta^2}{\sqrt{\|A_j x^r + b_j\|^2 + \eta^2}} + \sqrt{\|A_j x^r + b_j\|^2 + \eta^2} \right) \\ &\geq \sum_{j=1}^{\ell} \sqrt{\|A_j x + b_j\|^2 + \eta^2} = g(x), \forall x \in \mathbb{R}^m. \end{aligned}$$

Then according to Example 1, Assumption (A3) is automatically true, therefore we have verified that $u(x, x^r)$ defined in (18) is indeed an upper-bound function for the smooth function $g(x)$. It follows that the IRLS algorithm corresponds to a single-block BSUM algorithm. Notice that using the BSUM framework we can easily generalize the IRLS to the multiblock scenario.

APPLICATIONS OF THE BSUM FRAMEWORK

Here we briefly review a few applications of the BSUM framework in wireless communication, bioinformatics, signal processing, and machine learning.

WIRELESS COMMUNICATION AND TRANSCIVER DESIGN

Consider a multiple-input, multiple-output interference channel with K transmitter-receiver pairs. Let M (resp. N) be the number of antennas at each transmitter (resp. receiver) and each transmitter k , $k = 1, 2, \dots, K$, is interested in transmitting one data stream to its own receiver. Let $\mathbf{x}_k \in \mathbb{C}^M$ be the transmitted signal of user k ; assuming linear channel model, the received signal of user k can be written as

$$\mathbf{y}_k = \underbrace{\mathbf{H}_{kk} \mathbf{x}_k}_{\text{desired signal}} + \underbrace{\sum_{j \neq k} \mathbf{H}_{kj} \mathbf{x}_j}_{\text{multiuser interference}} + \underbrace{\mathbf{n}_k}_{\text{noise}},$$

where $\mathbf{H}_{kj} \in \mathbb{C}^{N \times M}$ is the channel from transmitter j to receiver k and $\mathbf{n}_k \in \mathbb{C}^N$ denotes the additive white Gaussian noise at the receiver k with distribution $\mathcal{CN}(0, \sigma^2 \mathbf{I})$.

When linear beamformers are employed at the transmitters and receivers, the transmitted signal and the estimated received data stream can be respectively written as

$$\mathbf{x}_k = \mathbf{v}_k s_k$$

and

$$\hat{\mathbf{s}}_k = \mathbf{u}_k^H \mathbf{y}_k,$$

where $\mathbf{v}_k \in \mathbb{C}^M$ and $\mathbf{u}_k \in \mathbb{C}^N$ are, respectively, the transmit and receive beamformers. Here the transmitted data stream and the estimated data stream at the receiver are denoted by $s_k \in \mathbb{C}$ and $\hat{s}_k \in \mathbb{C}$, respectively.

A crucial task in modern wireless networks is to design the transmit and receive beamformers \mathbf{v}_k and \mathbf{u}_k to maximize a given utility of the system. Here, for simplicity of presentation, we consider the sum rate utility function as our objective. Therefore, our goal is to solve the following optimization problem:

$$\begin{aligned} \max_{\mathbf{u}, \mathbf{v}} \quad & \sum_{k=1}^K R_k(\mathbf{u}, \mathbf{v}) \\ \text{s.t.} \quad & \|\mathbf{v}_k\|^2 \leq P_k, \quad \forall k = 1, 2, \dots, K, \end{aligned} \quad (19)$$

where P_k is the total power budget of user k and $R_k(\mathbf{u}, \mathbf{v})$, which is the communication rate of user k , is given by

$$R_k(\mathbf{u}, \mathbf{v}) = \log \left(1 + \frac{|\mathbf{u}_k^H \mathbf{H}_{kk} \mathbf{v}_k|^2}{\sigma^2 \|\mathbf{u}_k\|^2 + \sum_{j \neq k} |\mathbf{u}_k^H \mathbf{H}_{kj} \mathbf{v}_j|^2} \right).$$

Problem (19) is nonconvex and known to be NP-hard [58]. Using the well-known relation between the signal-to-interference-plus-noise ratio (SINR) and the mean square error (MSE) value, one can rewrite (19) as [12], [59]:

$$\begin{aligned} \min_{\mathbf{v}, \mathbf{u}} \quad & \sum_{k=1}^K \log(e_k(\mathbf{u}, \mathbf{v})) \\ \text{s.t.} \quad & \|\mathbf{v}_k\|^2 \leq P_k, \quad \forall k = 1, 2, \dots, K, \end{aligned} \quad (20)$$

where $e_k(\mathbf{u}, \mathbf{v})$ is the MSE value and is given by

$$e_k(\mathbf{u}, \mathbf{v}) = |\mathbf{u}_k^H \mathbf{H}_{kk} \mathbf{v}_k - 1|^2 + \sum_{j \neq k} |\mathbf{u}_k^H \mathbf{H}_{kj} \mathbf{v}_j|^2 + \sigma^2.$$

Since the $\log(\cdot)$ function is concave, it is upper bounded by its first-order approximation (i.e., the linear upper bound in Table 3). Therefore, we can define the function

$$u(\mathbf{x}, \mathbf{x}^r) = \sum_{k=1}^K (\log(e_k(\mathbf{x}^r)) + (e_k(\mathbf{x}^r))^{-1} (e_k(\mathbf{x}) - e_k(\mathbf{x}^r))), \quad (21)$$

where $\mathbf{x} \triangleq (\mathbf{u}, \mathbf{v})$ is the optimization variable and $\mathbf{x}^r \triangleq (\mathbf{u}^r, \mathbf{v}^r)$ denotes the beamformer at iteration r . It is not hard to see that the approximation function in (21) is a valid upper bound in the

BSUM framework and at each iteration r , this choice of approximation function leads to a quadratic programming problem which has closed-form solutions. The resulting algorithm, dubbed WMMSE, converges to a stationary point of the problem and, in practice, it typically converges in a few iterations [12] even for larger-size problems [60].

For more details of the algorithm and its extensions to various beamformer design scenarios and different utility functions, refer to [12], [20], [59], [61], and [62]. It is also worth noting that many other interesting transceiver design algorithms also fall into the BSUM framework; see [33] and [63]–[65] for more details.

BIOINFORMATICS AND SIGNAL PROCESSING

Here we briefly outline two interesting big data applications of the BSUM framework in bioinformatics and signal processing.

ABUNDANCE ESTIMATION IN MODERN HIGH-THROUGHPUT SEQUENCING TECHNOLOGIES

An essential step in the analysis of modern high throughput sequencing technologies of biological data is to estimate the abundance level of each transcript in the experiment. Mathematically, this problem can be stated as follows. Consider M transcript sequences $s_1, \dots, s_M \in \{A, C, G, T\}^L$ with the corresponding abundance levels ρ_1, \dots, ρ_M such that $\sum_{m=1}^M \rho_m = 1$. Let R_1, \dots, R_N be noisy sequencing reads originated from the transcript sequences, where each read R_n , $n = 1, \dots, N$, is originated from only one of the transcript sequences s_1, \dots, s_M . Given the observed reads, the likelihood of the abundance levels ρ_1, \dots, ρ_M can be written as

$$\begin{aligned} \Pr(R_1, \dots, R_N; \rho_1, \dots, \rho_M) \\ &= \prod_{n=1}^N \Pr(R_n; \rho_1, \dots, \rho_M) \\ &= \prod_{n=1}^N \left(\sum_{m=1}^M \Pr(R_n | \text{read } R_n \text{ from sequence } s_m) \Pr(s_m) \right) \\ &= \prod_{n=1}^N \left(\sum_{m=1}^M \alpha_{nm} \rho_m \right), \end{aligned}$$

where $\alpha_{nm} \triangleq \Pr(R_n | \text{read } R_n \text{ from sequence } s_m)$ can be obtained efficiently using an alignment algorithm such as the ones based on the Burrows-Wheeler transform; see, e.g., [66] and [67]. Therefore, given $\{\alpha_{nm}\}_{n,m}$, the maximum likelihood estimation of the abundance levels can be stated as

$$\begin{aligned} \hat{\rho}_{ML} &= \arg \min_{\rho} - \sum_{n=1}^N \log \left(\sum_{m=1}^M \alpha_{nm} \rho_m \right) \\ \text{s.t. } \sum_{m=1}^M \rho_m &= 1, \text{ and } \rho_m \geq 0, \forall m = 1, \dots, M. \end{aligned} \quad (22)$$

As a special case of the EM algorithm, a popular approach for solving this optimization problem is to successively minimize a local tight upper bound of the objective function. In particular, the eXpress software [4] solves the following optimization problem at the r th iteration of the algorithm:

$$\begin{aligned} \rho^{r+1} &= \arg \min_{\rho} \\ &\quad - \sum_{n=1}^N \left(\sum_{m=1}^M \left(\frac{\alpha_{nm} \rho_m^r}{\sum_{m'=1}^M \alpha_{nm'} \rho_{m'}^r} \log \left(\frac{\rho_m}{\rho_m^r} \right) \right) \right. \\ &\quad \left. + \log \left(\sum_{m=1}^M \alpha_{nm} \rho_m^r \right) \right) \\ \text{s.t. } \sum_{m=1}^M \rho_m &= 1, \text{ and } \rho_m \geq 0, \forall m = 1, \dots, M. \end{aligned} \quad (23)$$

Using Jensen's inequality, it is not hard to check that (23) is a valid upper bound of (22) in the BSUM framework. Moreover, (23) has a closed-form solution obtained by

$$\rho_m^{r+1} = \frac{1}{N} \sum_{n=1}^N \frac{\alpha_{nm} \rho_m^r}{\sum_{m'=1}^M \alpha_{nm'} \rho_{m'}^r}, \quad \forall m = 1, \dots, M,$$

which makes the algorithm computationally efficient at each step. In practice, the above algorithm for abundance estimation converges in a few iterations. Moreover, this algorithm is perfectly suitable for distributed storage and multicore machines. In particular, since the number of reads N is much larger than the number of sequences M , one can store the reads R_1, \dots, R_N in n_p different processing units. Hence, at each iteration r , the processing unit p , $p = 1, \dots, n_p$, can compute the local value

$$\hat{\rho}_{m,p}^{r+1} = \frac{1}{N} \sum_{n \in \mathcal{N}_p} \frac{\alpha_{nm} \rho_m^r}{\sum_{m'=1}^M \alpha_{nm'} \rho_{m'}^r}, \quad \forall m = 1, \dots, M,$$

where \mathcal{N}_p is the set of reads stored at processor p with $\cup_{p=1}^{n_p} \mathcal{N}_p = \{1, 2, \dots, N\}$. Then, all processors update their global abundance estimate through the consensus procedure

$$\rho_m^{r+1} = \sum_{p=1}^{n_p} \hat{\rho}_{m,p}^{r+1}, \quad \forall m = 1, \dots, M.$$

For a very recent application of BSUM algorithm in gene RNA-seq abundance estimation, see [68].

TENSOR DECOMPOSITION

The CANDECOMP/PARAFAC (CP) decomposition has applications in different areas such as clustering [69] and compression [70]. For ease of presentation, here we only consider third-order tensors. Given a third-order tensor $\mathfrak{X} \in \mathbb{R}^{m_1 \times m_2 \times m_3}$, its rank R CP decomposition is given by $\mathfrak{X} = \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r$, where $\mathbf{a}_r \in \mathbb{R}^{m_1}$, $\mathbf{b}_r \in \mathbb{R}^{m_2}$, $\mathbf{c}_r \in \mathbb{R}^{m_3}$; and the notation “ \circ ” stands for the outer product operator.

In general, finding the CP decomposition of a given tensor is NP-hard [71]. A well-known algorithm for finding the CP decomposition is the alternating least squares (ALS) algorithm proposed in [72] and [73]. This algorithm is, in essence, the BCD algorithm on the following optimization problem:

$$\min_{\{\mathbf{a}_r, \mathbf{b}_r, \mathbf{c}_r\}_{r=1}^R} \left\| \mathfrak{X} - \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \right\|_F^2. \quad (24)$$

In the ALS algorithm, we consider three blocks of variables: $\{\mathbf{a}_r\}_{r=1}^R$, $\{\mathbf{b}_r\}_{r=1}^R$, and $\{\mathbf{c}_r\}_{r=1}^R$. At each iteration of the algorithm, two blocks are held fixed and only one block is updated by solving (24). The

block selection rule is cyclic and, therefore, one needs the uniqueness of the minimizer assumption at each iteration for theoretical convergence guarantee. Clearly, this assumption does not hold for (24), therefore convergence is not always guaranteed. In addition, another well-known drawback of the ALS algorithm is the “swamp” effect where the objective remains almost constant for many iterations and then starts decreasing again. It has been observed in the literature that the employment of proximal upper bound (see Table 3) could help reduce the swamp effect [74]. It is also suggested in [74] that decreasing the proximal coefficient (γ in Table 3) during the ALS algorithm could further improve the performance of the algorithm. Notice that these modifications in the algorithm makes the algorithm a special case of BSUM framework. Consequently, its theoretical convergence is also guaranteed by Theorem 1.

Figure 5 compares the performance of the naive ALS algorithm with the one using proximal upper bound. The figure shows that the proximal ALS algorithm has less swamp effect as compared to the naive ALS method. For more details of the algorithm, refer to [21] and [74]; and to [75] for the application of BSUM and CP decomposition in gene expression and brain imaging.

MACHINE LEARNING: SPARSE DICTIONARY LEARNING AND SPARSE LINEAR DISCRIMINANT ANALYSIS

DICTIONARY LEARNING FOR SPARSE REPRESENTATION

In compressive sensing [76], [77] problems, a given data signal is represented by sparse linear combination of the signals in a given set called a *dictionary*. In many applications even the dictionary is not known a priori, therefore, it should be learned from the data. More precisely, given a set of training signals $\{y_1, \dots, y_N \in \mathbb{R}^n\}$, the dictionary-learning task is to find a dictionary set


$\{a_1, \dots, a_k \in \mathbb{R}^n\}$ that can sparsely represent the signals in the training set. Defining the matrices $Y \triangleq [y_1, \dots, y_N]$, $A \triangleq [a_1, \dots, a_k]$, and $X \triangleq [x_1, \dots, x_N]$, the dictionary-learning problem can be written as [78], [79]

$$\begin{aligned} \min_{A, X} \quad & d(Y, A, X) \\ \text{s.t.} \quad & A \in \mathcal{A}, X \in \mathcal{X}, \end{aligned} \quad (25)$$


where the sets \mathcal{X} and \mathcal{A} are given based on the prior knowledge on the data. The function $d(\cdot, \cdot, \cdot)$ measures the goodness-of-fit of the model. For example, a popular choice of the function $d(\cdot, \cdot, \cdot)$ and the set \mathcal{A} leads to the following optimization problem [79]:

$$\begin{aligned} \min_{A, X} \quad & \|Y - AX\|^2 + \lambda \|X\|_1 \\ \text{s.t.} \quad & \|a_i\|_2^2 \leq \beta_i, \end{aligned}$$

where the first term in the objective keeps our estimated signals close to the training set and the second term forces the representation to be sparse. One popular approach in the dictionary-learning algorithm is to alternately update the dictionary A and the coefficients X [80]. However, naively updating these two variables to its global optimum requires solving a sparse recovery problem at each iteration, which is costly for larger-sized problems. Motivated by the idea of inexact steps in the BSUM framework, one can iteratively replace the objective by a locally tight upper bound, which is easier to minimize at each iteration and, hence, leads to computationally cheaper steps in the algorithm. It is not hard to see that utilizing the quadratic upper bound in Table 3 with diagonal matrices Φ_i leads to closed-form updates at each step [78]. Unlike many existing algorithms in the literature [79], [80], the resulting algorithm is guaranteed to converge theoretically to the set of stationary solutions as the result of Theorem 1.



[FIG5] A comparison of the ALS and proximal ALS algorithm [21]. The proximal ALS algorithm is the BSUM approach using the proximal upper bound; see Table 3. The “ALS with diminishing proximal” algorithm utilizes a decreasing proximal coefficient during the iterates of the algorithm. In the example $m_1 = m_2 = m_3 = R = 100$. (Figure reprinted with permission from the Society for Industrial and Applied Mathematics.)



[FIG6] Sample denoised images ($\sigma = 100$). (a) K-SVD, (b) proposed algorithm, (c) K-SVD (zoomed in), and (d) proposed algorithm (zoomed in) [78].

Figure 6 and Table 5 show the performance of the resulting algorithm for dictionary learning in an image denoising problem. The denoising is performed on the Lena image corrupted by additive Gaussian noise with various variances σ^2 . As can be seen from Table 5, the proposed algorithm results in larger PSNR values than the K-SVD method [80] when the noise level is large. Moreover, the proposed algorithm contains less visual artifacts. Furthermore, each step of the proposed algorithm is in closed form and is computationally favorable, while each step of the K-SVD method requires an inner iterative method.

SPARSE LINEAR DISCRIMINANT ANALYSIS

The linear discriminant analysis (LDA), which is closely related to analysis of variance (ANOVA) and regression analysis, is widely used in machine learning and statistics for classification and dimensionality reduction purposes; see, e.g., [81]. Let us, for the ease of presentation, focus only on the binary classification problem: Let $\mathbf{x}_i \in \mathbb{R}^p$, $i = 1, 2, \dots, N$, denote the zero-centered observations, where each observation \mathbf{x}_i belongs to one and only one of the two classes C_0 and C_1 . Given the binary classes, the standard within-class covariance estimate can be calculated by

$$\hat{\Sigma}_w = \frac{1}{N} \sum_{k \in \{0,1\}} \sum_{i \in C_k} (\mathbf{x}_i - \hat{\mu}_k)(\mathbf{x}_i - \hat{\mu}_k)^T,$$

where $\hat{\mu}_k = 1/N \sum_{i \in C_k} \mathbf{x}_i$ is observations mean in class C_k . Similarly, the standard between-class covariance estimate is given by

$$\hat{\Sigma}_b = \frac{1}{N} (N_0 \hat{\mu}_0 \hat{\mu}_0^T + N_1 \hat{\mu}_1 \hat{\mu}_1^T),$$

with N_0 (resp. N_1) being the cardinality of the set C_0 (resp. C_1). The goal of LDA is to find a lower-dimensional subspace so that the projection of the observations onto the selected subspace leads to well-separated classes. In other words, the task is to project data points into a subspace with large between-class variance relative to the within-class variance. For simplicity, consider projection onto one-dimensional subspace defined by the vector $\beta \in \mathbb{R}^p$; see [82] for details on projection to larger-dimensional subspaces. Then the inner product $\langle \beta, \mathbf{x} \rangle$ is the projection of the observation \mathbf{x} onto the selected subspace; and the within-class variance of the projected data points is given by $\hat{\sigma}_w = \beta^T \hat{\Sigma}_w \beta$; while the between

[TABLE 5] THE IMAGE DENOISING RESULT COMPARISON ON THE "LENA IMAGE" FOR DIFFERENT NOISE LEVELS. VALUES ARE AVERAGED OVER TEN MONTE CARLO SIMULATIONS [78].

σ / PSNR	DCT	K-SVD	PROPOSED ALGORITHM
20/22.11	32	32.38	30.88
60/12.57	26.59	26.86	26.37
100/8.132	24.42	24.45	24.46
140/5.208	22.96	22.93	23.11
180/3.025	21.73	21.69	21.96

class variance can be written as $\hat{\sigma}_b = \beta^T \hat{\Sigma}_b \beta$. Therefore, in the LDA problem, we are interested in solving

$$\begin{aligned} \max_{\beta} \quad & \beta^T \hat{\Sigma}_b \beta \\ \text{s.t.} \quad & \beta^T \hat{\Sigma}_w \beta \leq 1. \end{aligned} \quad (25)$$

Unfortunately, when the number of features is large relative to N , the matrix $\hat{\Sigma}_w$ is rank deficient and therefore (25) is ill posed. To resolve this issue and to have a small generalization error, [82] suggests to regularize the optimization problem with a convex penalty function $P(\cdot)$; and solve

$$\begin{aligned} \max_{\beta} \quad & \beta^T \hat{\Sigma}_b \beta - P(\beta) \\ \text{s.t.} \quad & \beta^T \hat{\Sigma}_w \beta \leq 1. \end{aligned} \quad (26)$$

Clearly, this optimization problem is nonconvex. As suggested in [82], one can linearize the first part of the objective in (26) iteratively to obtain a tight upper bound of the objective. It is not hard to see that the algorithm used in [82] is BSUM with the linear upper bound given in Table 3.


EXTENSIONS

In this section, we discuss extensions and generalizations of the BSUM framework in various settings.

STOCHASTIC OPTIMIZATION

Consider the following stochastic optimization problem:

$$\begin{aligned} \min_x \quad & f(x) \triangleq \mathbb{E}_{\xi} [g(x, \xi)] \\ \text{s.t.} \quad & x \in \mathcal{X}, \end{aligned} \quad (27)$$



[FIG7] An iteration and running time comparison of SSUM versus SAA [20].

where \mathcal{X} is a closed convex set and ξ is a random variable modeling the uncertainty in our optimization problem. A standard classical approach for solving (27) is the sample average approximation (SAA) method; see [83] and the references therein. At iteration r of the SAA method, given a new realization ξ^r of the random variable ξ , the SAA method generates a new iterate x^r by solving a problem with the sample average $1/r \sum_{i=1}^r g(x, \xi^i)$ as its objective, where $\xi^1, \xi^2, \dots, \xi^r$ are independent identically distributed realizations of the random variable ξ .

A major drawback of the SAA method is that each of its iteration can be computationally very expensive. The computational inefficiency arises from either the nonconvexity of the objective, or not having closed-form solutions at each iteration.

Motivated by the BSUM framework, the authors of [20] and [84] suggest using an inexact version of the SAA method, in which a sequence of upper bounds of the objective are minimized. In particular, at each iteration r , the optimization variable is updated by

$$x^r \in \arg \min_x \frac{1}{r} \sum_{i=1}^r \hat{g}(x, x^{i-1}, \xi^i) \\ \text{s.t. } x \in \mathcal{X},$$

where $\hat{g}(\cdot, x^{i-1}, \xi^i)$ is an upper bound of the function $g(\cdot, \xi^i)$ around the point x^{i-1} . The approximation function $\hat{g}(\cdot, x^{i-1}, \xi^i)$ is assumed to be in the form of the BSUM approximation. The resulting algorithm, named *stochastic successive upper-bound minimization* (SSUM), is guaranteed to converge to the set of stationary solutions almost surely; see [20] for more details. Further, it is shown to be capable of dealing with various practical problems in signal processing and machine learning. For example, as we will see shortly, the authors of [85] apply the SSUM framework to cope with uncertainties in channel estimation for a wireless beamformer design problem. As another example, the online sparse dictionary-learning algorithm proposed in [86] is a special case of SSUM.

The stochastic optimization framework is well suited for many modern big data applications, especially when the entire

data set is not available initially and the data points are made available over time. These problems can be considered as the previously mentioned general stochastic optimization problem; see also [87]. In addition, many statistical model fitting problems, such as the simple classical regression problem, can be cast as minimizing the following sum-cost function $\sum_{\ell=1}^L g(x, \xi_\ell)$. Typically, the number of data points L is very large, making it difficult for batch processing. Therefore, it is desirable to implement algorithms working with only one (or a few) data point(s) at each step. In these scenarios, the stochastic optimization framework is useful since the sum-cost minimization problem can be viewed as a stochastic optimization problem $\min \mathbb{E}[g(x, \xi)]$ with ξ being uniformly drawn from the set $\{\xi_1, \dots, \xi_L\}$.

As an example of the SSUM method, consider the wireless transceiver design problem discussed in the “Wireless Communication and Transceiver Design” section, where the channel coefficients $\{\mathbf{H}_{ij}\}_{i,j}$ are not exactly known. In this scenario, we can consider the channel coefficients as random variables and solve the following stochastic optimization problem:

$$\max_{\mathbf{u}, \mathbf{v}} \sum_{k=1}^K \mathbb{E}[R_k(\mathbf{u}, \mathbf{v})] \\ \text{s.t. } \|\mathbf{v}_k\|^2 \leq P_k, \quad \forall k = 1, 2, \dots, K, \quad (28)$$

which is the stochastic counterpart of the optimization problem (19). Utilizing the upper bound (21) in the SSUM algorithm leads to the stochastic WMMSE algorithm [85]. Figure 7 illustrates the numerical performance of the SSUM methods as compared with SAA. At each iteration of the SAA procedure, one should solve a nonconvex optimization problem. Two different methods are considered: the gradient descent method with random initialization and the WMMSE algorithm, which is known to converge in few iterations for this problem. As Figure 7 illustrates, the running time of the SAA algorithm is much longer than that of the SSUM.

COUPLING CONSTRAINTS

So far in this article, we have assumed that the constraints in the optimization problem is separable and convex. In other words, the constraint set \mathcal{X} in (1) is of the form $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$ with each \mathcal{X}_i being convex. A natural extension of the BSUM framework is to modify it to deal with coupling and nonconvex constraints.

LINEAR COUPLING

Consider the following convex problem with linear coupling constraints

$$\begin{aligned} \min_{(x_1, \dots, x_n)} \quad & f(x_1, \dots, x_n) \\ \text{s.t.} \quad & \sum_{i=1}^n A_i x_i = b \\ & x_i \in \mathcal{X}_i, \forall i = 1, 2, \dots, n, \end{aligned} \quad (29)$$

where $x_i \in \mathbb{R}^{m_i}$, $A_i \in \mathbb{R}^{k \times m_i}$, $b \in \mathbb{R}^k$ and $f(\cdot)$ is a convex function. As seen in Example 4, the direct extension of the BCD/BSUM approach does not work for this type problem. A popular approach for solving the above optimization problem is the alternating direction method of multipliers (ADMM) [88], [89]. This approach is based on finding a saddle point of the augmented Lagrangian function

$$\begin{aligned} L(x_1, \dots, x_n; \lambda) = & f(x_1, \dots, x_n) \\ & + \left\langle \lambda, \sum_{i=1}^n A_i x_i - b \right\rangle + \frac{\rho}{2} \left\| \sum_{i=1}^n A_i x_i - b \right\|^2, \end{aligned}$$

where $\lambda \in \mathbb{R}^k$ is the Lagrange multiplier corresponding to the linear constraint; $\rho > 0$ is the augmented Lagrangian coefficient; and $\langle \cdot, \cdot \rangle$ denotes the inner product operator.

At each iteration of the ADMM method, either a primal block variable x_i is updated according to

$$x_i^{r+1} \leftarrow \arg \min_{x_i \in \mathcal{X}_i} L(x_i, x_{-i}^r; \lambda^r),$$

or the dual Lagrange multiplier λ is updated according to the gradient ascent rule

$$\lambda^{r+1} \leftarrow \lambda^r + \alpha^r \left(\sum_{i=1}^n A_i x_i^r - b \right),$$

where α^r is the dual step size at iteration r . The update orders for the primal and dual variables could be either cyclic or randomized.

Similar to the BSUM framework, one can replace the augmented Lagrangian function $L(\cdot, x_{-i}^r; \lambda^r)$ with its tight upper bound $\tilde{L}_i(\cdot, x_{-i}^r; \lambda^r)$ at iteration $r + 1$, where

$$\begin{aligned} \tilde{L}_i(x_i, x_{-i}^r, \lambda^r) = & u_i(x_i, x_{-i}^r) \\ & + \left\langle \lambda^r, A_i x_i + \sum_{j \neq i} A_j x_j^r - b \right\rangle + \frac{\rho}{2} \left\| A_i x_i + \sum_{j \neq i} A_j x_j^r - b \right\|^2 \end{aligned}$$

with $u_i(\cdot, x_{-i}^r)$ being a locally tight approximation of the function $f(\cdot, x_{-i}^r)$ around the point x_i^r satisfying Assumption A. The resulting algorithm, named the *BSUM method of multipliers*

(*BSUMM*) [44], is guaranteed to converge to the global optimal of (29) under some regularity assumptions [44]. For extensions to nonconvex problems, see the recent work [90] and [91]. There are a few other interesting techniques that deal with linearly coupling constraint. For example, [92] and [93] propose to randomly pick two blocks of variables to update at each iteration, and [94] proposes new algorithms based on minimizing the augmented Lagrangian function. We refer the readers to these papers for more related works in this direction.

Let us illustrate an application of BSUMM to a multicommodity routing problem, which arises in the design of next-generation cloud-based communication networks [95]. Consider a connected wireline network $\mathcal{N} = (\mathcal{V}, \mathcal{L})$ that is controlled by $K + 1$ network controllers (NCs) as illustrated in Figure 8. Let \mathcal{V} denote the set of network nodes, which is partitioned into K subsets, i.e., $\mathcal{V} = \bigcup_{i=1}^K \mathcal{V}^i$, $\mathcal{V}^i \cap \mathcal{V}^j = \emptyset$, $\forall i \neq j$. The set of directed links that connect nodes of \mathcal{V} is denoted as $\mathcal{L} \triangleq \{l = (s_l, d_l) \mid \forall s_l, d_l \in \mathcal{V}\}$, where $l = (s_l, d_l)$ denotes the directed link from node s_l to node d_l . Each NC i controls \mathcal{V}^i and the links connecting these nodes, i.e., $\mathcal{L}^i \triangleq \{l = (s_l, d_l) \in \mathcal{L} \mid \forall s_l, d_l \in \mathcal{V}^i\}$, (cf. Figure 8). The network $\mathcal{N}^i \triangleq (\mathcal{V}^i, \mathcal{L}^i)$ is called the subnetwork i .


Our objective is to transport M data flows over the network, with each flow m being routed from a source node $s_m \in \mathcal{V}$ to the sink node $d_m \in \mathcal{V}$. We use $r_m \geq 0$ to denote flow m 's rate, and use $f_{l,m} \geq 0$ to denote its rate on link $l \in \mathcal{L}$. We also assume that a master node exists, which controls the data flow rates $\{r_m\}_{m=1}^M$. The central NC 0 controls the subnetwork \mathcal{N}^0 , consisting of the master node and the links connecting different subnetworks, i.e., $\mathcal{L}^0 = \bigcup_{i \neq j} \mathcal{L}_{ij}^0$.

We consider two types of network constraints:

1) *Link capacity constraints*. Assume each link $l \in \mathcal{L}$ has a fixed capacity denoted as C_l . The total flow rate on link l is constrained by

$$1^T \mathbf{f}_l \leq C_l, \forall l \in \mathcal{L}, \quad (30)$$

where $\mathbf{1}$ is the all-one vector and $\mathbf{f}_l \triangleq [f_{l,1}, \dots, f_{l,M}]^T$.



[FIG8] A wireline network consists of five subnetworks. Each of them is controlled by an NC, and these NCs are coordinated globally by a central NC 0 [95].

**[TABLE 6] A COMPARISON OF DIFFERENT ALGORITHMS.
M = 200.**

APPROACHES	TIME	NUMBER OF ITERATIONS
GUROBI	11.4690 s	N/A
SYNCHRONOUS BSUMM (TEN NCs, ONE CORE)	18.0679 s	264.71
SYNCHRONOUS BSUMM (TEN NCs, TEN CORE)	5.80 s	264.71
ASYNCHRONOUS BSUMM (TEN NCs, TEN CORE)	4.23 s	2109.07

2) *Flow conservation constraints.* For any node $v \in \mathcal{V}$ and data flow m , the total incoming flow should be equal to the total outgoing flow:

$$\sum_{l \in \text{In}(v)} f_{l,m} + 1_{v=s(m)} r_m = \sum_{l \in \text{Out}(v)} f_{l,m} + 1_{v=d(m)} r_m, \quad m = 1, \dots, M, \quad \forall v \in \mathcal{V}, \quad (31)$$

where $\text{In}(v) \triangleq \{l \in \mathcal{L} | d_l = v\}$ and $\text{Out}(v) \triangleq \{l \in \mathcal{L} | s_l = v\}$ denote the set of links going into and coming out of a node v respectively; $1_{v=x} = 1$ if $v = x$ otherwise $1_{v=x} = 0$.

To provide fairness to the users, we maximize the minimum rate of all data flows. The problem can be formulated as the following linear program (LP)

$$\max_{f, r} \quad r_{\min} \quad \text{s.t.} \quad f \geq 0, \quad r_m \geq r_{\min}, \quad m = 1, \dots, M \quad (32a)$$

$$(30) \text{ and } (31), \quad (32b)$$

where $f \triangleq \{f_l | l \in \mathcal{L}\}$ and $r \triangleq \{r_{\min}, r_m | m = 1 \sim M\}$. Obviously, one can use off-the-shelf optimization packages such as Gurobi [96] to solve the LP (32), but this is only viable in a centralized setting where all the flows are managed by a single controller.

To enable distributed/parallel network management across the NCs, we need to allow each NC i to independently optimize the variables belonging to the subnetwork \mathcal{N}^i . However, this task is difficult because the optimization variables of (32) is coupled (indeed each flow rate f_m appears in exactly two flow conservation constraints). To address this problem, we introduce a few sets of new variables to decouple the flow conservation constraints across different subnetworks (see [95] for the detailed reformulation). The reformulated problem (32) is given by

$$\begin{aligned} \max_x \quad & r_{\min} \\ \text{s.t.} \quad & \{r_{\min}, x_{02}\} \in \mathcal{X}_0, \quad \{x_{i1}, x_{i2}\} \in \mathcal{X}_{i1}, \quad x_{i3} \in \mathcal{X}_{i2}, \\ & \underbrace{x'_{01} = x'_{02}}_{\text{in } \mathcal{N}^0}, \quad \underbrace{x'_{i1} = x'_{01}}_{\text{in } \mathcal{N}^i \text{ and } \mathcal{N}^0}, \quad \underbrace{x'_{i2} = x_{i3}}_{\text{in } \mathcal{N}^i}, \quad i = 1 \dots K, \end{aligned}$$

where \mathcal{X}_0 , \mathcal{X}_{i1} , and \mathcal{X}_{i2} are some feasible sets, and $\{r_{\min}, x_{02}, x_{01}, x_{i1}, x_{i2}, x_{i3}\}$ are the block variables. By applying the BSUMM, we can obtain a parallel/distributed algorithm. A few remarks about the implementation of this algorithm are:

1) The replication of link/flow variables for links across different subnetworks allows each subnetwork to be considered

separately and independently. This feature makes the BSUMM subproblems solvable in parallel. The requirement of the replicated variables being the same as the original variables is enforced by the linear coupling constraints, and they can be satisfied asymptotically as the BSUMM algorithm converges.

2) The subproblems of the proposed BSUMM-based algorithm can be solved by each NC very efficiently. For example, the update of $\{r_{\min}, x_{02}\}$ can be performed by each NC in closed form; the update of $\{x_i, x'_{01}\}$ can be performed by running the well-known RELAX code [97].

3) A careful implementation of the BSUMM allows the NCs to act asynchronously, in the sense that they do not need to coordinate with each other for computation. Such asynchronous implementation has the potential of greatly improving the computational efficiency.

We illustrate the BSUMM implementation over a mesh wireline network with 126 nodes, which is randomly partitioned into nine subnetworks with 306 directed links within these subnetworks and 100 directed links connecting the subnetworks. The capacities for the links within (resp. between) the subnetworks are uniformly distributed in [50,100] megabits/second (resp. [20,50] megabits/second). All simulation results are averaged over 200 randomly selected data flow pairs and link capacity.

To demonstrate the benefit of parallelization, we also utilize a high-performance computing cluster, and make each computing node to be an NC. We compare a few different approaches for solving (32):

- 1) use Gurobi [96], a centralized LP solver
- 2) apply the synchronous BSUMM algorithm, with $K = 10$ NCs; the computation is done by either a single or by ten distributed computing cores
- 3) apply the asynchronous BSUMM with $K = 10$ NCs; the computation is done in ten distributed computing cores.

Note that the asynchrony in the network arises naturally from the per-node computational delay and network communication delay. In Table 6, we demonstrate the performance of various algorithms when $M = 200$. Clearly, the asynchronous BSUMM with a small number of NCs outperforms all the rest of the algorithms.

The numerical results suggest that appropriate network decomposition and asynchronous implementation are both critical for the fast convergence of BSUMM. In practice, we observe that the network should be decomposed following a few guidelines:

- the computation burden across the subnetworks is well balanced
- the subroutine within the network can achieve its maximum efficiency
- the total number of replicated auxiliary variables is small.

NONCONVEX CONSTRAINTS

The BSUM idea can be straightforwardly extended to a nonconvex constraint scenario for single block optimization problems. To proceed, consider the optimization problem:

$$\begin{aligned} \min_x \quad & f_0(x) \\ \text{s.t.} \quad & f_i(x) \leq 0, \quad \forall i = 1, 2, \dots, \ell, \end{aligned} \quad (33)$$

where the functions $f_i(\cdot)$ are not necessarily convex. Since dealing with nonconvex constraints is often not easy, one popular approach is to replace the functions $f_i(x), i = 1, 2, \dots, \ell$, with their locally tight upper bound $u_i(x, x^r)$ iteratively. In other words, the update rule of the iterates is given by [98]

$$\begin{aligned} x^{r+1} \leftarrow \arg \min_x \quad & u_0(x, x^r) \\ \text{s.t.} \quad & u_i(x, x^r) \leq 0, \quad \forall i = 1, 2, \dots, \ell. \end{aligned} \quad (34)$$

As illustrated in Figure 9, the iterative approximation of the constraints is a restriction of the constraints and hence the iterates remain feasible during the algorithm. If, in addition, some constraint qualification conditions are satisfied, the resulting algorithm is guaranteed to converge to the set of stationary solutions of (33); see [38, Th. 1] for detailed conditions and analysis.

PARALLEL VERSION AND EXTENSIONS TO GAME THEORY

With the recent advances in multicore and cluster computational platforms, it is desirable to design “parallel” algorithms for multiblock optimization problem where multiple cores update the block variables in parallel to optimize the objective function. A naive parallel extension of the BCD approach for solving (1) is to update all blocks (or a subset of them) in parallel by solving

$$x_i^{r+1} \leftarrow \arg \min_{x_i \in \mathcal{X}_i} f(x_i, x_i^r), \quad \forall i = 1, \dots, n.$$

Unfortunately, this naive extension of the BCD algorithm does not converge, in general, and might result in a zigzag/oscillating or divergent behavior. As an example, consider the problem


$$\begin{aligned} \min_{(x_1, x_2)} \quad & (x_1 - x_2)^2 \\ \text{s.t.} \quad & -1 \leq x_1, x_2 \leq 1. \end{aligned}$$

Clearly, this problem is convex with bounded feasible set and its optimal value is zero. However, the above naive parallel extension of the algorithm leads to the following iteration path:

$$\begin{aligned} (x_1^0, x_2^0) &= (1, -1) \rightarrow (x_1^1, x_2^1) = (-1, 1) \rightarrow (x_1^2, x_2^2) \\ &= (1, -1) \rightarrow \dots, \end{aligned}$$

which is clearly not convergent. This is caused by aggressive steps used in the algorithm. To make the algorithm convergent, it is then necessary to employ controlled steps that are also small enough. Furthermore, in the case of nonconvex objective function $f(\cdot)$ in (1), the approximation functions could be again used to obtain computationally efficient update rules. The resulting algorithm, dubbed *parallel successive convex approximation (PSCA)*, is summarized in Table 7. To see the convergence analysis of this algorithm and other related algorithms such as the flexible parallel algorithm, refer to [18], [99], [100], and the references therein.

Notice that PSCA can be viewed as a way of solving a multi-agent optimization problem where multiple agents/users try to



[FIG9] An illustration of constraint convexification.


optimize a common objective by updating their own variable iteratively. Furthermore, it can be used in a game-theoretic setting where each player in the game utilizes the best response strategy by optimizing a locally tight upper bound of its own utility function. This algorithm is guaranteed to converge for some particular class of games under some regularity assumptions on the players' utility functions. The convergence analysis presented in [101] is based on certain contraction approach as well as monotone convergence for potential games.

Figure 10 illustrates the behavior of the cyclic and randomized parallel PSCA method as compared with their serial counterparts (i.e., the “cyclic BCD” and the “randomized BCD”) applied to the LASSO problem. The performance of the PSCA method is also illustrated for a different number of processors and various block selection rules. In Figure 10(a) and (b), parallelization can result in more efficient algorithm; however, the convergence speed does not grow linearly with the number of processing cores. Moreover, increasing the number of processors beyond certain point results in slower convergence, which can be attributed to the increased communication overhead among the nodes.

Note that the parallel update rule is very useful in dealing with distributed data sets. Consider solving the LASSO problem with the following objective: $\|Ax - b\|^2 + \lambda \|x\|_1$. Assume we have q processing cores each having their own memory. Let us partition the matrix A and vector x into q blocks: $A = [A_1, \dots, A_q]$ and $x = [x_1^T, \dots, x_q^T]^T$. If PSCA is implemented in a way that each core j is only responsible for updating

[TABLE 7] A PSEUDOCODE OF THE PSCA ALGORITHM.

- 1 FIND A FEASIBLE POINT $x^0 \in \mathcal{X}$; SET $r = 0$, AND CHOOSE A STEP-SIZE SEQUENCE $\{\gamma^r\}$
- 2 REPEAT
- 3 PICK INDEX SET \mathcal{I}^r
- 4 LET $\tilde{x}_i^r = \arg \min_{x_i \in \mathcal{X}_i} u_i(x_i, x_i^{r-1}), \quad \forall i \in \mathcal{I}^r$
- 5 SET $x_k^r = x_k^{r-1}, \quad \forall k \notin \mathcal{I}^r$
- 6 SET $x_i^r = x_i^{r-1} + \gamma^r(\tilde{x}_i^r - x_i^{r-1}), \quad \forall i \in \mathcal{I}^r$
- 7 $r = r + 1$,
- 8 UNTIL SOME CONVERGENCE CRITERION IS MET



[FIG10] A comparison of the serial BCD with PSCA method for solving the LASSO problem: $\min_x \|Ax - b\|^2 + \lambda \|x\|_1$. The matrix A and vector b are generated according to [102]; and q denotes the number of processors used in each experiment. The dimension of A is $2,000 \times 10,000$ for (a) and $1,000 \times 100,000$ for (b). The experiments are done over a computer cluster using the Message Passing Interface in C [18].

block j , then at each iteration $r + 1$, core j 's problem of interest can be written as

$$\min \|A_j x_j - b'_j\|^2 + \lambda \|x_j\|_1,$$

where $b'_j \triangleq b - \sum_{i \neq j} A_i x_i^r$. Notice that the value of b'_j can be calculated by letting each node i compute the value of $A_i x_i$ and broadcast it to other nodes. Under this architecture, each node does not need to know the complete matrix A and only local information is enough for a distributed implementation of the PSCA method.

PRACTICAL CONSIDERATIONS

There are a number of factors that we need to consider when using the BSUM framework. The first consideration is about the choice of the upper bound. What is a good bound for a given application? The answer is generally problem dependent, as we have already seen in a few examples. The general guideline is that a good upper bound should be able to ensure algorithm convergence, best exploit the problem structure, and make the subproblems easily solvable (preferably in closed form). For example, a simple proximal upper bound is not likely to perform well for the transceiver design problem discussed in the section "Wireless Communication and Transceiver Design" as the resulting subproblems will not decompose over the variables.

The second consideration is about the choice of the block update rules. As we have seen in the section "How Fast Does the BSUM Converge?" different update rules lead to quite distinct convergence behavior. For convex problems, deterministic rules such as the cyclic rule promise the worst-case convergence rates, while

the randomized rule ensures convergence rate in either averaged or high-probability sense. Further, there is barely any theoretical rate analysis for nonconvex problems, regardless of the block selection rules. Therefore, the best strategy in practice is to perform an extensive numerical study and pick the best rule for the application at hand. For example, researchers have found that the MBI rule is effective for certain tensor decomposition problems [13]; the cyclic rule can be superior to the randomized rule for certain LASSO problems, and certain G-So rules can outperform the randomized rule [43], [14].

The third consideration is about the choice of the parallelization schemes. There has been extensive research on parallelizing various special cases and variations of the BSUM type algorithm; see [17], [19], [99], [103]–[106], and the references therein. These algorithms differ in a number of implementation details and in their applicability. For example, most of the implementations use randomized block selection rules to pick the variable blocks, while [99] and [105] additionally use certain variations of the G-So rule. The majority of the schemes only work for convex problems, with the exception of [99] and [103], which work for general nonsmooth and nonconvex problems in the form of (5). When assessing whether a given problem is suitable for parallelization, it is important to know that oftentimes the number of blocks that can be updated in parallel is data dependent. For example, when solving LASSO problems, it is shown in [105] and [106] that the degree of parallelization is dependent on the maximum eigenvalue of certain submatrices of the data matrix. Some recent results [107] show that for a certain randomized coordinate descent method, such dependency could be mild. For solving general

convex nonsmooth problems, [104] shows that the step size should be carefully selected based on both the “separability” of the problem (or the sparsity of the data matrix) as well as the degree of parallelization. If the application at hand does not satisfy these conditions, the alternatives usually are: 1) exploit the problem structure and pick a good upper bound, so that the subproblems are decomposable, leading to parallel and step-size-free updates; see, for example, the NMF problem in the section “The NMF Algorithm” and the WMMSE algorithm in the section “Wireless Communication and Transceiver Design”; 2) to use the diminishing stepsize for updating the blocks; see, for example, [19], [99], and [103].

ISSUES AND OPEN RESEARCH PROBLEMS

This article presents a comprehensive algorithmic framework, BSUM, for block-structured large-scale optimization. The main strength of the BSUM framework is its strong theoretical convergence guarantee and its flexibility. As demonstrated in this article, the BSUM framework covers a number of well-known but seemingly unrelated algorithms as well as their new extensions. Moreover, it is amenable to a number of different data models as well as to parallel implementation on modern multicore computing platforms.

To close, we briefly highlight a couple of issues and open research topics related to the BSUM framework.

■ *Communication delay and overhead in parallel implementations:* As discussed in the section “Parallel Version and Extensions to Game Theory,” the convergence speed of the parallel version of the BSUM framework does not increase linearly with the number of computational nodes. In fact, after a point, increasing the number of computational nodes can lead to a slower convergence speed. As mentioned previously, this is due to the delay caused by communication among the nodes. This observation gives rise to two important research questions: First, given the maximum allowable number of computation nodes and the communication overhead of the nodes, what is the optimum choice of the number of cores for solving a given optimization problem? Answering this question requires computation/communication tradeoff analysis of the proposed optimization approach. Second, can the BSUM framework be extended and implemented in a (semi)asynchronous manner? If this is possible, then the communication overhead can be reduced significantly since the nodes are not required to wait for each other before updating the variables, making the algorithm lock-free. For recent efforts on this research direction see [17].

■ *Nonlinear coupling constraints:* As we observe in the “Extensions” section, the BSUM framework can also be used in the presence of linear coupling or nonconvex decoupled constraints. How can the BSUM framework be generalized to problems with nonlinear coupling constraints? More precisely, can the BSUM framework with block-wise update rules be applied to the optimization problem of the following form?

$$\begin{aligned} \min_x \quad & f_0(x_1, \dots, x_n) \\ \text{s.t.} \quad & f_i(x_1, \dots, x_n) \leq 0, \forall i = 1, 2, \dots, n. \end{aligned}$$

Example 4 shows that the naive extension of the BCD approach fails to find the optimal solution even in the convex setting. A popular approach to tackle the aforementioned problem is to place the constraints in the objective using Lagrange multipliers and update the multipliers iteratively. However, this approach typically leads to double-loop algorithms and requires subgradient steps in the dual space, which is known to be slow.

ACKNOWLEDGMENTS

M. Hong and M. Razaviyayn contributed equally to this work. M. Hong is supported by the National Science Foundation, grant CCF-1526078. Z.-Q. Luo is supported by the National Science Foundation, grant CCF-1526434.

AUTHORS

Mingyi Hong (mingyi@iastate.edu) received his B.E. degree in communications engineering from Zhejiang University, China, in 2005; his M.S. degree in electrical engineering from Stony Brook University, New York, in 2007; and his Ph.D. degree in systems engineering from the University of Virginia in 2011. He was previously a research assistant professor with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities. He is currently an assistant professor with the Department of Industrial and Manufacturing Systems Engineering, Iowa State University, Ames. His research interests include signal processing, wireless communications, large-scale optimization and its applications in compressive sensing, complex networks, and high-dimensional data analysis.

Meisam Razaviyayn (meisamr@stanford.edu) received his B.S. degree in electrical engineering from Isfahan University of Technology, Iran, in 2008; his M.S. degrees in mathematics and electrical engineering from the University of Minnesota in 2013; and his Ph.D. degree in electrical engineering from the University of Minnesota in 2014. He is currently a postdoctoral research fellow in the Electrical Engineering Department at Stanford University, California. His research interests include large-scale optimization, machine learning and statistics signal processing, and wireless communication and its computational issues.

Zhi-Quan Luo (luozq@umn.edu) received his B.S. degree in applied mathematics in 1984 from Peking University, Beijing, China. He received his Ph.D. degree in operations research in 1989 from the Massachusetts Institute of Technology. From 1989 to 2003, he held a faculty position with the Department of Electrical and Computer Engineering, McMaster University, Hamilton, Canada, where he eventually became the department head and held a Canada research chair in information processing. Since 2003, he has been with the Department of Electrical and Computer Engineering at the University of Minnesota, Twin Cities as a full professor and holds an endowed ADC chair in digital technology. Since 2014, he has been with the Chinese University of Hong Kong (Shenzhen) as the vice president (academic). His research interests include optimization algorithms, signal processing, and digital communication.

Jong-Shi Pang (jongship@usc.edu) received his B.S. degree in mathematics from National Taiwan University in 1973 and his

Ph.D. degree in operations research from Stanford University, California, in 1977. He is the Epstein family chair professor in the Department of Industrial and Systems Engineering, University of Southern California. He was the head and Caterpillar professor in the Department of Industrial and Enterprise Systems Engineering at the University of Illinois at Urbana-Champaign. He was previously the Margaret A. Darrin Distinguished Professor in applied mathematics at Rensselaer Polytechnic Institute in Troy, New York. He has received several awards and honors. He was the member of the inaugural 2009 class of the SIAM fellows. He is an ISI highly cited author in the mathematics category. His research interests are in continuous optimization and equilibrium programming and their applications in engineering, economics, and finance.

REFERENCES

- [1] V. Krishnamurthy, O. Namvar, and M. Hamdi, "Signal processing of social networks: Interactive sensing and learning," *Found. Trends Signal Process.*, vol. 7, nos. 1–2, pp. 1–19, 2014.
- [2] *Special Issue on Signal Processing for Big Data, IEEE Signal Processing Mag.*, vol. 31, no. 5, Sept. 2014.
- [3] S. Sra, S. Nowozin, and S. J. Wright, *Optimization for Machine Learning*. Cambridge, MA: MIT Press, 2012.
- [4] A. Roberts and L. Pachter, "Streaming fragment assignment for real-time analysis of sequencing experiments," *Nat. Methods*, vol. 10, no. 1, pp. 71–73, 2013.
- [5] Y. Nesterov, "Efficiency of coordinate descent methods on huge-scale optimization problems," *SIAM J. Optim.*, vol. 22, no. 2, pp. 341–362, 2012.
- [6] V. Cevher, S. Becker, and M. Schmidt, "Convex optimization for big data: Scalable, randomized, and parallel algorithms for big data analytics," *IEEE Signal Process. Mag.*, vol. 31, no. 5, pp. 32–43, Sept. 2014.
- [7] N. S. Aybat, "First order methods for large-scale sparse optimization," Ph.D. dissertation, Columbia Univ., 2014.
- [8] P. Richtárik and M. Takáč, "Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function," *Math. Program.*, vol. 144, no. 1, pp. 1–38, 2014.
- [9] S. Cui, Z.-Q. Luo, A. Hero, and J. Moura, *Big Data over Networks*. Cambridge, U.K.: Cambridge Univ. Press, 2015.
- [10] Y.-L. You and M. Kaveh, "A regularization approach to joint blur identification and image restoration," *IEEE Trans. Image Processing*, vol. 5, no. 3, pp. 416–428, 1996.
- [11] J. Friedman, T. Hastie, and R. Tibshirani, "Regularization paths for generalized linear models via coordinate descent," *J. Stat. Softw.*, vol. 33, no. 1, pp. 1–22, 2010.
- [12] Q. Shi, M. Razaviyayn, Z.-Q. Luo, and C. He, "An iteratively weighted MMSE approach to distributed sum-utility maximization for a MIMO interfering broadcast channel," *IEEE Trans. Signal Processing*, vol. 59, no. 9, pp. 4331–4340, 2011.
- [13] B. Chen, S. He, Z. Li, and S. Zhang, "Maximum block improvement and polynomial optimization," *SIAM J. Optim.*, vol. 22, no. 1, pp. 87–107, 2012.
- [14] J. Nutini, M. Schmidt, I. H. Laradji, M. Friedlander, and H. Koepke, "Coordinate descent converges faster with the Gauss-Southwell rule than random selection," in *Proc. 30th Int. Conf. Machine Learning (ICML)*, 2015, pp. 1632–1641.
- [15] P. Richtárik and M. Takáč, "Parallel coordinate descent methods for big data optimization," *Math. Program.*, to be published.
- [16] F. Facchinei, S. Sagratella, and G. Scutari, "Flexible parallel algorithms for big data optimization," in *Proc. 2014 IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 7208–7212.
- [17] J. Liu, C. Re S. J. Wright, and V. Bittorf, "An asynchronous parallel stochastic coordinate descent algorithm," in *Proc. Int. Conf. Machine Learning (ICML)*, 2014.
- [18] M. Razaviyayn, M. Hong, Z.-Q. Luo, and J. S. Pang, "Parallel successive convex approximation for nonsmooth nonconvex optimization," in *Proc. Neural Information Processing (NIPS)*, 2014, pp. 1440–1448.
- [19] F. Facchinei, G. Scutari, and S. Sagratella, "Parallel selective algorithms for nonconvex big data optimization," *IEEE Trans. Signal Processing*, vol. 63, no. 7, pp. 1874–1889, 2015.
- [20] M. Razaviyayn, M. Sanjabi, and Z.-Q. Luo, "A stochastic successive minimization method for nonsmooth nonconvex optimization with applications to transceiver design in wireless communication networks," preprint, 2013.
- [21] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM J. Optim.*, vol. 23, no. 2, pp. 1126–1153, 2013.
- [22] J. Mairal, "Optimization with first-order surrogate functions," in *Proc. Int. Conf. Machine Learning (ICML)*, 2013, pp. 783–791.
- [23] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [24] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Top. Signal Processing*, vol. 1, no. 4, pp. 586–597, 2007.
- [25] E. Candès and B. Recht, "Exact matrix completion via convex optimization," *Commun. ACM*, vol. 55, no. 6, pp. 111–119, 2012.
- [26] R. Sun and Z.-Q. Luo, "Guaranteed matrix completion via nonconvex factorization," preprint, 2014.
- [27] J. Kim, Y. He, and H. Park, "Algorithms for nonnegative matrix and tensor factorizations: A unified view based on block coordinate descent framework," *J. Glob. Optim.*, vol. 58, no. 2, pp. 285–319, Feb. 2014.
- [28] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Rev.*, vol. 51, no. 3, pp. 455–500, 2009.
- [29] S. J. Wright, "Coordinate descent algorithms," *Math. Program.*, vol. 151, no. 1, pp. 3–34, 2015.
- [30] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. R. Stat. Soc. Series B*, vol. 39, no. 1, pp. 1–38, 1977.
- [31] A. L. Yuille and A. Rangarajan, "The concave-convex procedure," *Neural Comput.*, vol. 15, pp. 915–936, 2003.
- [32] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Proc. Neural Information Processing Systems (NIPS)*, 2000, pp. 556–562.
- [33] C. Shi, R. A. Berry, and M. L. Honig, "Monotonic convergence of distributed interference pricing in wireless networks," in *Proc. IEEE Int. Conf. Symp. Information Theory*, 2009, pp. 1619–1623.
- [34] M. Hong and Z.-Q. Luo, "Signal processing and optimal resource allocation for the interference channel," in *Academic Press Library in Signal Processing*. New York: Academic, 2013.
- [35] M. Hong, Q. Li, and Y.-F. Liu, "Decomposition by successive convex approximation: A unifying approach for linear transceiver design in interfering heterogeneous networks," *IEEE Trans. Wireless Commun.*, to be published.
- [36] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *J. Optim. Theory Appl.*, vol. 103, no. 9, pp. 475–494, 2001.
- [37] M. J. D. Powell, "On search directions for minimization algorithms," *Math. Program.*, vol. 4, no. 1, pp. 193–201, 1973.
- [38] M. Razaviyayn, "Successive convex approximation: Analysis and applications," Ph.D. dissertation, Univ. of Minnesota, 2014.
- [39] Z. Lu and L. Xiao, "On the complexity analysis of randomized block-coordinate descent methods," *Math. Program.*, vol. 152, no. 1, pp. 615–642, Aug. 2015.
- [40] A. Beck and L. Tetruashvili, "On the convergence of block coordinate descent type methods," *SIAM J. Optim.*, vol. 23, no. 4, pp. 2037–2060, 2013.
- [41] A. Beck, "On the convergence of alternating minimization with applications to iteratively reweighted least squares and decomposition schemes," *SIAM J. Optim.*, vol. 25, no. 1, pp. 185–209, 2015.
- [42] M. Hong, X. Wang, M. Razaviyayn, and Z.-Q. Luo, "Iteration complexity analysis of block coordinate descent methods," preprint, 2013.
- [43] A. Saha and A. Tewari, "On the nonasymptotic convergence of cyclic coordinate descent method," *SIAM J. Optim.*, vol. 23, no. 1, pp. 576–601, 2013.
- [44] M. Hong, T.-H. Chang, X. Wang, M. Razaviyayn, S. Ma, and Z.-Q. Luo, "A block successive upper bound minimization method of multipliers for linearly constrained convex optimization," preprint, 2013.
- [45] M. Kadkhodaei, M. Sanjabi, and Z.-Q. Luo, "On the linear convergence of approximate proximal splitting methods for non-smooth convex minimization," *J. Oper. Res. Soc. China*, vol. 2, no. 2, pp. 123–141, 2014.
- [46] Y. Xu and W. Yin, "A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion," *SIAM J. Imaging Sci.*, vol. 6, no. 3, pp. 1758–1789, 2013.
- [47] Q. Lin, Z. Lu, and L. Xiao, "An accelerated proximal coordinate gradient method and its application to regularized empirical risk minimization," *SIAM J. Optim.*, to be published.
- [48] O. Fercoq and P. Richtárik, "Accelerated, parallel and proximal coordinate descent," *SIAM J. Optim.*, vol. 25, no. 4, pp. 1927–2023, 2015.
- [49] D. R. Hunter and K. Lange, "A tutorial on MM algorithms," *Am. Stat.*, vol. 58, no. 1, pp. 30–37, 2004.
- [50] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, 2nd ed. Belmont, MA: Athena-Scientific, 1999.
- [51] L. Grippo and M. Sciandrone, "On the convergence of the block nonlinear Gauss-Seidel method under convex constraints," *Oper. Res. Lett.*, vol. 26, no. 3, pp. 127–136, 2000.
- [52] P. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering, Springer Optimization and Its Applications*. New York: Springer, 2011, pp. 185–212.

- [53] Z. Lu and L. Xiao, "Randomized block coordinate non-monotone gradient method for a class of nonlinear programming," preprint, 2013.
- [54] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online dictionary learning for sparse coding," in *Proc. 26th Annu. Int. Conf. Machine Learning (ICML)*, 2009, pp. 689–696.
- [55] Q. Li, "On the convergence of Lee-Seung's multiplicative update for NMF," 2012, Tech. Note.
- [56] C. Févotte and J. Idier, "Algorithms for nonnegative matrix factorization with the beta-divergence," *Neural Comput.*, vol. 23, no. 3, pp. 1–24, 2011.
- [57] I. Daubechies, R. DeVore, M. Fornasier, and C. S. Gunturk, "Iteratively reweighted least squares minimization for sparse recovery," *Commun. Pure Appl. Math.*, vol. 63, no. 1, pp. 1–38, 2010.
- [58] Z.-Q. Luo and S. Zhang, "Dynamic spectrum management: Complexity and duality," *IEEE J. Sel. Top. Signal Processing*, vol. 2, no. 1, pp. 57–73, 2008.
- [59] H. Baligh, M. Hong, W.-C. Liao, Z.-Q. Luo, M. Razaviyayn, M. Sanjabi, and R. Sun, "Cross-layer provision of future cellular networks: A WMMSE-based approach," *IEEE Signal Processing Mag.*, vol. 31, no. 6, pp. 56–68, Nov. 2014.
- [60] M. Razaviyayn, H. Baligh, A. Callard, and Z.-Q. Luo, "Joint user grouping and transceiver design in a mimo interfering broadcast channel," *IEEE Trans. Signal Processing*, vol. 62, no. 1, pp. 85–94, Jan. 2014.
- [61] S. S. Christensen, R. Agarwal, E. D. Carvalho, and J. M. Cioffio, "Weighted sum-rate maximization using weighted MMSE for MIMO-BC beamforming design," *IEEE Trans. Wireless Commun.*, vol. 7, no. 12, pp. 4792–4799, 2008.
- [62] M. Hong, R. Sun, H. Baligh, and Z.-Q. Luo, "Joint base station clustering and beamformer design for partial coordinated transmission in heterogeneous networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 2, pp. 226–240, 2013.
- [63] C. Shi, R. A. Berry, and M. L. Honig, "Distributed interference pricing with MISO channels," in *Proc. 46th Annual Allerton Conf. Communication, Control, and Computing* 2008, Sept. 2008, pp. 539–546.
- [64] S.-J. Kim and G. B. Giannakis, "Optimal resource allocation for MIMO Ad Hoc Cognitive Radio Networks," *IEEE Trans. Inform. Theory*, vol. 57, no. 5, pp. 3117–3131, 2011.
- [65] C. T. K. Ng and H. Huang, "Linear precoding in cooperative MIMO cellular networks with limited coordination clusters," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 9, pp. 1446–1454, 2010.
- [66] B. Langmead, C. Trapnell, M. Pop, and S. L. Salzberg, "Ultrafast and memory-efficient alignment of short dna sequences to the human genome," *Genome Biol.*, vol. 10, no. 3, pp. R25, 2009.
- [67] H. Li and R. Durbin, "Fast and accurate short read alignment with Burrows-Wheeler transform," *Bioinformatics*, vol. 25, no. 14, pp. 1754–1760, 2009.
- [68] N. Bray, H. Pimentel, P. Melsted, and L. Pachter, "Near-optimal RNA-Seq quantification," preprint, 2015.
- [69] A. Shashua and T. Hazan, "Non-negative tensor factorization with applications to statistics and computer vision," in *Proc. 22nd Int. Conf. Machine Learning (ICML)*, 2005, pp. 792–799.
- [70] I. Ibragimov, "Application of the three-way decomposition for matrix compression," *Numer. Linear Algebra Appl.*, vol. 9, no. 6–7, pp. 551–565, 2002.
- [71] J. Håstad, "Tensor rank is NP-complete," *J. Algorithms*, vol. 11, no. 4, pp. 644–654, 1990.
- [72] J. D. Carroll and J. J. Chang, "Analysis of individual differences in multidimensional scaling via an N-way generalization of 'Eckart-Young' decomposition," *Psychometrika*, vol. 35, no. 3, pp. 283–319, 1970.
- [73] R. A. Harshman, "Foundations of the parafac procedure: Models and conditions for an 'explanatory' multimodal factor analysis," *UCLA Working Papers in Phonetics*, 16, pp. 1–84, 1970.
- [74] C. Navasca, L. De Lathauwer, and S. Kindermann, "Swamp reducing technique for tensor decomposition," in *Proc. 16th European Signal Processing Conf. (EUSIPCO)*, Aug. 2008, pp. 1–5.
- [75] J. A. Bazerque, G. Mateos, and G. B. Giannakis, "Rank regularization and Bayesian inference for tensor completion and extrapolation," *IEEE Trans. Signal Processing*, vol. 61, no. 21–24, pp. 5689–5703, 2013.
- [76] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [77] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [78] M. Razaviyayn, H.-W. Tseng, and Z.-Q. Luo, "Dictionary learning for sparse representation: Complexity and algorithms," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 5247–5251.
- [79] H. Lee, A. Battle, R. Raina, and A. Y. Ng, "Efficient sparse coding algorithms," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2006, pp. 801–808.
- [80] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Processing*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [81] G. McLachlan, *Discriminant Analysis and Statistical Pattern Recognition*, vol. 544. Hoboken, NJ: Wiley, 2004.
- [82] D. M. Witten and R. Tibshirani, "Penalized classification using Fisher's linear discriminant," *J. R. Stat. Soc. Series B (Stat. Methodol.)*, vol. 73, no. 5, pp. 753–772, 2011.
- [83] S. Kim, R. Pasupathy, and S. G. Henderson, "A guide to sample average approximation," in *Handbook of Simulation Optimization*, M. C. Fu, Ed. vol. 216: International Series in Operations Research and Management Science. New York: Springer, 2015, pp. 207–243.
- [84] J. Mairal, "Stochastic majorization-minimization algorithms for large-scale optimization," in *Proc. Advances in Neural Information Processing Systems*, 2013, pp. 2283–2291.
- [85] M. Razaviyayn, M. Sanjabi, and Z.-Q. Luo, "A stochastic weighted MMSE approach to sum rate maximization for a MIMO interference channel," in *Proc. IEEE 14th Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, IEEE, 2013, pp. 325–329.
- [86] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, Jan. 2010.
- [87] K. Slavakis, S.-J. Kim, G. Mateos, and G. B. Giannakis, "Stochastic approximation vis-a-vis online learning for big data analytics," *IEEE Signal Processing Mag.*, vol. 31, no. 6, pp. 124–129, 2014.
- [88] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [89] R. Glowinski, *Numerical Methods for Nonlinear Variational Problems*. New York: Springer-Verlag, 1984.
- [90] M. Hong, Z.-Q. Luo, and M. Razaviyayn, "Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems," submitted for publication.
- [91] B. Ames and M. Hong, "Alternating directions method of multipliers for l1-penalized zero variance discriminant analysis and principal component analysis," submitted for publication.
- [92] I. Necoara, Y. Nesterov, and F. Glineur, "A random coordinate descent method on large-scale optimization problems with linear constraints," Univ. Politehnica Bucharest, Tech. Rep. 2011.
- [93] S. Reddi, A. Hefny, C. Downey, A. Dubey, and S. Sra, "Large-scale randomized-coordinate descent methods with non-separable linear constraints," preprint, 2014.
- [94] N. S. Aybat, G. Iyengar, and Z. Wang, "An asynchronous distributed proximal method for composite convex optimization," in *Proc. 32th Int. Conf. Machine Learning (ICML)*, 2015, pp. 2454–2462.
- [95] W.-C. Liao, M. Hong, H. Farmanbar, and Z.-Q. Luo, "Semi-asynchronous routing for large-scale hierarchical networks," in *IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2015, pp. 2894–2898.
- [96] Gurobi Optimization, Inc. *Gurobi Optimizer Reference Manual*, 2013. [Online]. Available: <http://www.gurobi.com/documentation/6.0/refman/>
- [97] D. P. Bertsekas, P. Hosein, and P. Tseng, "Relaxation methods for network flow problems with convex arc costs," *SIAM J. Control Optim.*, vol. 25, no. 5, pp. 1219–1243, Sept. 1987.
- [98] B. R. Marks and G. P. Wright, "A general inner approximation algorithm for non-convex mathematical programs," *Oper. Res.*, vol. 26, no. 4, pp. 681–683, 1978.
- [99] F. Facchinei, S. Sagratella, and G. Scutari, "Flexible parallel algorithms for big data optimization," preprint, 2013.
- [100] G. Scutari, F. Facchinei, P. Song, D. P. Palomar, and J.-S. Pang, "Decomposition by partial linearization: Parallel optimization of multi-agent systems," *IEEE Trans. Signal Processing*, vol. 63, no. 3, pp. 641–656, 2014.
- [101] J.-S. Pang and M. Razaviyayn, "A unified distributed algorithm for non-cooperative games with non-convex and non-differentiable objectives," preprint, 2014.
- [102] Yu. Nesterov, "Gradient methods for minimizing composite functions," *Math. Program.*, vol. 140, no. 1, pp. 125–161, 2013.
- [103] M. Razaviyayn, M. Hong, Z.-Q. Luo, and J.-S. Pang, "Parallel successive convex approximation for nonsmooth nonconvex optimization," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2014, pp. 1440–1448.
- [104] P. Richtárik and M. Takáč, "Distributed coordinate descent method for learning with big data," preprint, 2013.
- [105] Z. Peng, M. Yan, and W. Yin, "Parallel and distributed sparse optimization," in *Proc. IEEE 2013 Asilomar Conf. Signals, Systems and Computers*, 2013, pp. 659–666.
- [106] C. Scherrer, A. Tewari, M. Halapanavar, and D. J. Haglin, "Feature clustering for accelerating parallel coordinate descent," in *Proc. Advanced in Neural Information Processing Systems (NIPS)*, 2012, pp. 28–36.
- [107] O. Fercoq, Z. Qu, P. Richtárik, and M. Takáč, "Fast distributed coordinate descent for non-strongly convex losses," in *Proc. IEEE Int. Workshop on Machine Learning for Signal Processing (MLSP)*, 2014, pp. 1–6.