

Block Coordinate DC Programming

Hoomaan Maskan
Umeå University, Sweden

hoomaan.maskan@umu.se

Paniz Halvachi
Massachusetts Institute of Technology, United States

panizh@mit.edu

Suvrit Sra
Technical University of Munich, Germany

s.sra@tum.de

Alp Yurtsever
Umeå University, Sweden

alp.yurtsever@umu.se

Abstract

We introduce an extension of the Difference of Convex Algorithm (DCA) in the form of a block coordinate approach for problems with separable structure. For n coordinate-blocks and k iterations, our main result proves a non-asymptotic convergence rate of $O(n/k)$ for the proposed method. Furthermore, leveraging the connection between DCA and Expectation Maximization (EM), we propose a block coordinate EM algorithm.

1 Introduction

In nonconvex optimization, many problems find their abstraction through the lens of Difference of Convex (DC) programming, which involves minimizing an objective function represented as the difference of convex functions. This paper focuses on the following DC programming template:

$$\min_{x \in \mathcal{M}} \phi(x) := f(x) + g(x) - h(x), \quad (1)$$

where f , g , and h are lower semi-continuous convex functions defined on the closed convex set $\mathcal{M} \subseteq \mathbb{R}^m$. We assume that the function f is smooth, and that the function g and the set \mathcal{M} exhibit a block-separable structure in relation to the coordinates of x (see [Section 3](#)). We do not impose the smoothness assumption on the terms g and h .

A large and popular class of algorithms for solving problem (1), introduced by [Tao and Souad \(1986\)](#), comprises instances of the DCA method. These algorithms rely on the simple yet effective idea of approximating $\phi(x)$ with the convex surrogate function obtained by using the first-order Taylor approximation of the concave component $-h(x)$. Starting from a feasible initial point $x^0 \in \mathcal{M}$, DCA solves a sequence of convex optimization subproblems defined as

$$\min_{x \in \mathcal{M}} \hat{\phi}(x, x^k) := f(x) + g(x) - h(x^k) - \langle \nabla h(x^k), x - x^k \rangle, \quad (2)$$

where x^k represents the current estimate. Observe that function $\hat{\phi}$ is a global upper bound on ϕ since h is convex; as a result, DCA can be viewed as a variant of the Majorization–Minimization (MM) algorithm ([Lange, 2016](#)).

Although we utilize the gradient of h in (2), suggesting differentiability of the function, it is worth noting that our discussions and results remain valid even when h is non-differentiable, as long as we substitute the gradient with a subgradient of h at x^k . The general template of DCA does not specify how, or to what degree of precision, we should solve these subproblems, and various instances can be derived by employing different techniques for their resolution. Motivated by its strong empirical performance, wide applicability, simplicity, and its interpretability, DCA has been extensively studied.

With the recent focus in machine learning on nonconvex optimization problems, along with the increasing complexity and scale of these problems, the demand for more efficient and scalable algorithms for nonconvex optimization continuously grows. A particular technique in addressing this challenge is the ‘coordinate descent’ (CD) approach. The main idea in CD is to minimize the multivariate objective function by updating one variable (or a *block* of variables) at every iteration, while keeping the others fixed. This approach often leads to simple subproblems that can be solved more efficiently, particularly for large-scale optimization problems commonly encountered in data science and related fields.

DC programming and DCA have been widely used in machine learning applications. Examples include kernel selection (Argyriou et al., 2006), sparse principal component analysis (Sriperumbudur et al., 2007), discrepancy estimation in domain adaptation (Awasthi et al., 2024a), training neural networks (Awasthi et al., 2024b), and many more. However, to our knowledge, there is currently no true CD variant of DCA with convergence guarantees. Our main goal in this paper is to address this gap. With this in mind, we summarize now the key contributions of this paper:

- Our primary contribution is the development of a novel variant of the DCA method that incorporates randomized CD updates. We refer to this algorithm as the Block Coordinate DC algorithm (BDCA). We analyze BDCA and demonstrate that it converges to a first-order stationary point of the problem at a non-asymptotic rate of $\mathcal{O}(n/k)$, for n blocks and k iterations. Importantly, our guarantees do not necessitate the functions g or h to be smooth.
- DCA relates to the well known Expectation Maximization (EM) algorithm when dealing with exponential family of distributions (Yuille and Rangarajan, 2003; Le Thi and Pham Dinh, 2018). Building on this connection, we introduce a block coordinate EM method, referred to as Block EM.

2 Related Work

For a comprehensive survey on the variants and applications of DCA, we refer to (Le Thi and Pham Dinh, 2018). Here, we focus specifically on the variants involving block-coordinate or alternating updates.

Related works on BDCA. Pham Dinh et al. (2022) recently proposed an alternating DC algorithm for partial DC problems involving two variables, where the objective retains a DC form when one variable is fixed. The algorithm is shown to converge to a Fréchet/Clarke critical point of the objective function under the Kurdyka-Łojasiewicz property, with non-asymptotic convergence guarantees.

More recently, Yuan (2023) proposed a block-coordinate method for DC problems; however, this method differs from DCA, in that it solves a proximal subproblem instead of the DCA subproblem in (2). Furthermore, their analysis is based on the so-called *Luo-Tseng error bound* assumption, which posits that the (sub)gradients of the objective function converge to zero in norm continuously as they approach a stationary point. This assumption severely limits the applicability of their results in non-smooth settings, as most non-smooth functions of interest exhibit discontinuous changes in subgradient norm.

DCA is also known as the convex-concave procedure (CCCP) when both the convex and concave terms are differentiable. Recently, Yurtsever and Sra (2022) established an equivalence between CCCP and the Frank-Wolfe algorithm applied to an epigraph reformulation of the original DC problem. As a result, they showed a $\mathcal{O}(1/k)$ non-asymptotic convergence rate for CCCP.

Our analysis differs from that of (Yurtsever and Sra, 2022) in several key aspects. While they transfer guarantees from the Frank-Wolfe connection, we provide a direct proof. We consider the DCA template with potentially non-smooth terms, which are managed through their subgradients. They demonstrate convergence in terms of

$$\max_{x \in \mathcal{M}} f(x^k) - f(x) - \langle \nabla h(x^k), x^k - x \rangle$$

which is not amenable to our block-coordinate analysis unless f is also separable. Consequently, we identify a new gap function (see Section 3) suitable for measuring proximity to first-order stationarity, which is an essential component of our analysis.

Related works on Block EM. Kumar and Schmidt (2017) analyzed the convergence guarantees of the EM algorithm from a majorization-minimization perspective. Assuming that the surrogate optimization problems are strongly convex, they established that EM converges at a rate of $\mathcal{O}(1/k)$ in terms of the squared gradient norm of the negative log-likelihood. However, this strong convexity assumption may not hold in general, even for exponential family distributions.

In a recent work, Kunstner et al. (2020) studied the non-asymptotic convergence of EM for exponential family distributions by drawing a connection to mirror descent, establishing a convergence rate of $\mathcal{O}(1/k)$ in Kullback-Leibler divergence and linking it to first-order stationarity through Bregman divergences.

Other works on non-convex BC methods. Block Successive Upperbound Minimization algorithm was proposed to minimize an approximated upperbound of any nonconvex and/or nonsmooth objective function (Razaviyayn et al., 2013). This method considers optimizing a sequence of approximate objective functions. They showed the asymptotic convergence of their analysis for quasi-convex functions or objectives with compact level sets. Also, they assume that each subproblem of BSUM has a unique solution.

Recent studies in nonconvex optimization using block-coordinate methods, Xu and Yin (2017) proposed block prox-linear method for smooth/nonsmooth problems. Assuming Lipschitz smoothness of the differentiable parts, they prove the convergence of a subsequence of their algorithm to a stationary point. Further, assuming the KL condition, they show that the entire sequence converges to a critical point and estimated its asymptotic convergence rate.

Aubry et al. (2018) proposes a mixture of maximum block improvement and a sequential minimization method for continuously differentiable non-convex problems with application in wireless sensor allocation with asymptotic first-order optimality analysis.

More recent studies targeted a nonasymptotic convergence analysis. In the realm of nonconvex composite optimization, Chorobura and Necoara (2023) considers objectives that include two nonconvex nonseparable terms where one of them has Lipschitz gradients while the other function is continuously differentiable. They showed a convergence rate of $\mathcal{O}(n/\sqrt{k})$ in the sense of expected gradient norm when utilizing a randomized selection of the blocks. Their cyclic scheme achieves a rate of $\mathcal{O}(n^2/\sqrt{k})$ in the same sense.

A cyclic coordinate descent algorithm with variance reduction was recently proposed (Cai et al., 2023). Their objective considers a block separable and another separable sum function whose proximal operator is efficiently computable. Their analysis assumes a nonstandard Lipschitz gradient condition and achieves a rate of $\mathcal{O}(n/k)$ in the sense of distance of the objective's subdifferential to zero.

A randomized coordinate subgradient descent was studied by Deng and Lan (2020). With additional assumption on simplicity of the structure of function g in problem (1), they propose a convergence rate of $\mathcal{O}(n/k)$ for their randomized proximal method.

3 Block Coordinate DCA

We begin by establishing notation and restating our assumptions for the model problem.

We consider problem (1), and throughout, we assume that a finite solution $x_* \in \mathcal{M}$ exists, satisfying $\phi_* := \phi(x_*) \leq \phi(x)$ for all $x \in \mathcal{M}$. We assume that the function f is L -smooth, meaning that its gradient is Lipschitz continuous with a constant $L \geq 0$. In addition, we assume that the function g and set \mathcal{M} exhibit a block separable structure in relation to the coordinates of x :

$$g(x) = \sum_{i=1}^n g_i(D_i x),$$

where D_i represents $(m_i \times m)$ dimensional row subsets of the identity matrix I , serving as the selection operator for a non-overlapping partition of the coordinates into n blocks:

$$\sum_{i=1}^n D_i^\top D_i = I \quad \text{and} \quad \sum_{i=1}^n m_i = m.$$

Similarly, \mathcal{M} can be decomposed as $\mathcal{M}_1 \times \cdots \times \mathcal{M}_n$, where the components $\mathcal{M}_i \in \mathbb{R}^{m_i}$, such that

$$\mathbf{x} \in \mathcal{M} \iff D_i \mathbf{x} \in \mathcal{M}_i \quad \text{for } i = 1, \dots, n.$$

For the ease of presentation, we define the following coordinate vector notation:

$$\mathbf{x}_i = D_i \mathbf{x}, \quad \mathbf{x}_i = D_i^\top D_i \mathbf{x}, \quad \text{and} \quad \bar{\mathbf{x}}_i = (I - D_i^\top D_i) \mathbf{x}.$$

Here, $\mathbf{x}_i \in \mathbb{R}^{m_i}$ represents the i^{th} coordinate block of $\mathbf{x} \in \mathbb{R}^m$. The vector \mathbf{x}_i is an extension of x_i to \mathbb{R}^m with all other blocks padded with zeros. Conversely, $\bar{\mathbf{x}}_i$ is the complement of \mathbf{x}_i , containing zeros in the i^{th} block and ensuring that $\mathbf{x}_i + \bar{\mathbf{x}}_i = \mathbf{x}$.

The foundation of DCA is the surrogate convex objective obtained by linearizing the concave term in the vicinity of current estimate \mathbf{x}^k , as defined in (2). To derive a block coordinate variant of DCA, we minimize this surrogate function along specific directions determined by randomly chosen coordinate blocks. The algorithm operates as follows:

1. Start from a feasible initial point $\mathbf{x}^0 \in \mathcal{M}$.
2. For $k = 1, \dots, K$, update the estimate as follows:
 - (a) Choose a coordinate block i_k uniformly at random from the set $\{1, \dots, n\}$.
 - (b) Keep the values fixed for all coordinates except i_k , resulting in $\bar{\mathbf{x}}_{i_k}^{k+1} = \bar{\mathbf{x}}_{i_k}^k$.
 - (c) Update the i_k^{th} block by minimizing the surrogate objective (2) along these coordinates.

Step 2c in this algorithm description amounts to solving the following subproblem:

$$\min_{\mathbf{x} \in \mathcal{M}} \hat{\phi}(\mathbf{x}, \mathbf{x}^k) \quad \text{subj.to} \quad \mathbf{x} = \mathbf{x}_{i_k} + \bar{\mathbf{x}}_{i_k}^k,$$

which is clearly equivalent to solving the following:

$$\min_{\mathbf{x}_{i_k} \in \mathcal{M}_{i_k}} \hat{\phi}(D_{i_k}^\top \mathbf{x}_{i_k} + \bar{\mathbf{x}}_{i_k}^k, \mathbf{x}^k).$$

To simplify the presentation, we introduce the notation $\hat{\phi}_i : \mathbb{R}^{m_i} \times \mathbb{R}^m \rightarrow \mathbb{R}$ for the surrogate objective restricted to the i^{th} coordinate block:

$$\begin{aligned} \hat{\phi}_i(\mathbf{x}_i, \mathbf{x}^k) &:= \hat{\phi}(D_i^\top \mathbf{x}_i + \bar{\mathbf{x}}_i^k, \mathbf{x}^k) \\ &= \hat{\phi}(\mathbf{x}_i + \bar{\mathbf{x}}_i^k, \mathbf{x}^k) \\ &= f(\mathbf{x}_i + \bar{\mathbf{x}}_i^k) + g_i(\mathbf{x}_i) - g_i(\mathbf{x}_i^k) + g(\mathbf{x}^k) - h(\mathbf{x}^k) - \langle \nabla_i h(\mathbf{x}^k), \mathbf{x}_i - \mathbf{x}_i^k \rangle. \end{aligned}$$

Here, only the terms in blue are relevant for solving the subproblem, as the remaining terms are constant with respect to \mathbf{x}_i . Building on this discussion, we propose Block Coordinate DC Algorithm (BDCA), as outlined in [Algorithm 1](#).

3.1 Convergence Guarantees for BDCA

In this section we analyze the convergence guarantees of BDCA. Specifically, we establish an upper bound on a stationarity measure for BDCA and use it to prove the convergence rate of [Algorithm 1](#).

Definition 1. We use the following ‘gap’ function as a measure of closeness to a first order stationary point:

$$\text{gap}_{\mathcal{M}}^L(\mathbf{y}) = \max_{\mathbf{x} \in \mathcal{M}} \left\{ \langle \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}), \mathbf{y} - \mathbf{x} \rangle + g(\mathbf{y}) - g(\mathbf{x}) - \frac{L}{2} \|\mathbf{x} - \mathbf{y}\|^2 \right\}.$$

Algorithm 1 Block Coordinate DC Algorithm (BDCA)

Input: Starting point $\mathbf{x}^1 \subseteq \mathcal{M}$, and total number of iterations K .

for $k = 1$ to K **do**

 Select a subgradient $\mathbf{v}^k \in \partial h(\mathbf{x}^k)$

 Choose i_k from $\{1, \dots, n\}$ uniformly at random

 Find $\mathbf{x}_{i_k}^{k+1} \in \operatorname{argmin}_{\mathbf{x}_{i_k} \in \mathcal{M}_{i_k}} f(\mathbf{x}_{i_k} + \tilde{\mathbf{x}}_{i_k}^k) + g_{i_k}(\mathbf{x}_{i_k}) - \langle \mathbf{v}_{i_k}^k, \mathbf{x}_{i_k} \rangle$

 Set $\mathbf{x}^{k+1} = \mathbf{x}_{i_k}^{k+1} + \tilde{\mathbf{x}}_{i_k}^k$

end for

Output: \mathbf{x}^{k+1}

The next lemma shows that this is a suitable measure for first-order stationarity.

Lemma 1. *Let f, g and h be convex functions, with f being L -smooth. Then, $\operatorname{gap}_{\mathcal{M}}^L(\mathbf{y}) \geq 0$ for all $\mathbf{y} \in \mathcal{M}$, and $\operatorname{gap}_{\mathcal{M}}^L(\mathbf{y}) = 0$ if and only if \mathbf{y} is a first-order stationary point of Problem (1).*

Proof. The first statement is straightforward, since the expression being maximized becomes zero when we set $\mathbf{x} = \mathbf{y}$, and $\mathbf{y} \in \mathcal{M}$.

For the second statement, let us recall the definition of first-order stationarity. A point \mathbf{y} is said to be a first-order stationary point of Problem (1) if the following condition holds:

$$0 \in \nabla f(\mathbf{y}) + \partial g(\mathbf{y}) - \nabla h(\mathbf{y}) + \mathcal{N}_{\mathcal{M}}(\mathbf{y}). \quad (3)$$

Here, ∂g is the subdifferential of g and $\mathcal{N}_{\mathcal{M}}$ represents the normal cone of \mathcal{M} .

If (3) holds, it implies there exists $\mathbf{u} \in \partial g(\mathbf{y})$ such that

$$\begin{aligned} & -\nabla f(\mathbf{y}) - \mathbf{u} + \nabla h(\mathbf{y}) \in \mathcal{N}_{\mathcal{M}}(\mathbf{y}) \\ \iff & \langle \nabla f(\mathbf{y}) + \mathbf{u} - \nabla h(\mathbf{y}), \mathbf{y} - \mathbf{x} \rangle \leq 0, \quad \forall \mathbf{x} \in \mathcal{M} \\ \implies & \langle \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}), \mathbf{y} - \mathbf{x} \rangle + g(\mathbf{y}) - g(\mathbf{x}) \leq 0, \\ \implies & \langle \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}), \mathbf{y} - \mathbf{x} \rangle + g(\mathbf{y}) - g(\mathbf{x}) - \frac{L}{2} \|\mathbf{x} - \mathbf{y}\|^2 \leq 0, \\ \iff & \operatorname{gap}_{\mathcal{M}}^L(\mathbf{y}) \leq 0, \end{aligned} \quad (4)$$

for any $\mathbf{x} \in \mathcal{M}$. Combining this with the first statement, which asserts that $\operatorname{gap}_{\mathcal{M}}^L(\mathbf{y}) \geq 0$ for all $\mathbf{y} \in \mathcal{M}$, we conclude that $\operatorname{gap}_{\mathcal{M}}^L(\mathbf{y}) = 0$ if \mathbf{y} is a first-order stationary point.

Next, we show that if $\operatorname{gap}_{\mathcal{M}}^L(\mathbf{y}) = 0$, then \mathbf{y} is a first-order stationary point. Suppose $\operatorname{gap}_{\mathcal{M}}^L(\mathbf{y}) = 0$. Then,

$$\langle \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle + g(\mathbf{x}) - g(\mathbf{y}) + \frac{L}{2} \|\mathbf{x} - \mathbf{y}\|^2 \geq 0,$$

for all $\mathbf{x} \in \mathcal{M}$. Consider $\mathbf{x} = \mathbf{y} + \alpha \mathbf{d}$ for an arbitrary feasible direction \mathbf{d} (unit norm) and step-size $\alpha > 0$:

$$\alpha \langle \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}), \mathbf{d} \rangle + g(\mathbf{y} + \alpha \mathbf{d}) - g(\mathbf{y}) + \frac{L}{2} \alpha^2 \geq 0, \quad \forall \alpha \mathbf{d} : \mathbf{y} + \alpha \mathbf{d} \in \mathcal{M}.$$

By dividing both sides by α and taking the limit as $\alpha \rightarrow 0^+$, we obtain

$$\langle \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}), \mathbf{d} \rangle + \langle \mathbf{u}, \mathbf{d} \rangle \geq 0, \quad \forall \mathbf{d} : \lim_{\alpha \rightarrow 0^+} (\mathbf{y} + \alpha \mathbf{d}) \in \mathcal{M},$$

for some $\mathbf{u} \in \partial g(\mathbf{y})$. Since \mathcal{M} is a closed and convex set, we have $\lim_{\alpha \rightarrow 0^+} (\mathbf{y} + \alpha \mathbf{d}) \in \mathcal{M}$ for all $\mathbf{d} = \mathbf{x} - \mathbf{y}$ such that $\mathbf{x} \in \mathcal{M}$. Consequently, we have

$$\langle \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle + \langle \mathbf{u}, \mathbf{x} - \mathbf{y} \rangle \geq 0,$$

for any $\mathbf{x} \in \mathcal{M}$. This is equivalent to (4), thus concludes the proof. \square

Remark 2. We can relate the gap function in [Definition 1](#) to the Moreau-Yosida envelope: $\text{gap}_{\mathcal{M}}^L(\mathbf{y}) = 0$ if and only if \mathbf{y} is a fixed point of the proximal mapping:

$$\mathbf{y} = \text{prox}_{\frac{1}{L}g + I_{\mathcal{M}}} \left(\left(\mathbf{y} - \frac{1}{L}(\nabla f(\mathbf{y}) - \nabla h(\mathbf{y})) \right) \right)$$

where $I_{\mathcal{M}}$ is the indicator function of the set \mathcal{M} . The proof of this connection is deferred to the appendix.

Remark 3. When h is nonsmooth, we can replace $\nabla h(\mathbf{y})$ with an arbitrary subgradient $\mathbf{v} \in \partial h(\mathbf{y})$ in the definition of gap, and [Lemma 1](#) remains valid. In this setting, a first-order stationary point is characterized by the condition

$$(\nabla f(\mathbf{y}) + \partial g(\mathbf{y}) + \mathcal{N}_{\mathcal{M}}(\mathbf{y})) \cap \partial h(\mathbf{y}) \neq \emptyset. \quad (5)$$

This implies that there exists $\mathbf{u} \in \partial g(\mathbf{y})$ and $\mathbf{v} \in \partial h(\mathbf{y})$ such that $-\nabla f(\mathbf{y}) - \mathbf{u} + \mathbf{v} \in \mathcal{N}_{\mathcal{M}}(\mathbf{y})$, which is the same as [\(4\)](#). The remainder of the proof follows analogously.

We can now present the main convergence result.

Theorem 4. Suppose $\mathbf{x}^1, \dots, \mathbf{x}^K$ is a sequence generated by BDCA for solving problem [\(1\)](#). Then, the following bound holds:

$$\min_{k \in \{1, \dots, K\}} \mathbb{E} \left[\text{gap}_{\mathcal{M}}^L(\mathbf{x}^k) \right] \leq \frac{n}{K} \left(\phi(\mathbf{x}^1) - \phi^* \right).$$

Proof. We begin by noting that $\forall x_{i_k} \in \mathcal{M}_{i_k}$, we have

$$\begin{aligned} \phi(\mathbf{x}^{k+1}) &\leq \hat{\phi}_{i_k}(\mathbf{x}_{i_k}^{k+1}, \mathbf{x}^k) \\ &\leq \hat{\phi}_{i_k}(x_{i_k}, \mathbf{x}^k) \\ &= f(\bar{\mathbf{x}}_{i_k}^k + \mathbf{x}_{i_k}) + g(\mathbf{x}^k) + g_{i_k}(x_{i_k}) - g_{i_k}(\mathbf{x}_{i_k}^k) - h(\mathbf{x}^k) - \langle \nabla_{i_k} h(\mathbf{x}^k), x_{i_k} - \mathbf{x}_{i_k}^k \rangle. \end{aligned} \quad (6)$$

Here, the first inequality holds due to the convexity of h , while the second inequality follows from the definition of $x_{i_k}^{k+1}$ in the update rule of BDCA. Adding $f(\mathbf{x}^k)$ to both sides and rearranging, we get

$$f(\mathbf{x}^k) - f(\bar{\mathbf{x}}_{i_k}^k + \mathbf{x}_{i_k}) + g_{i_k}(\mathbf{x}_{i_k}^k) - g_{i_k}(x_{i_k}) + \langle \nabla_{i_k} h(\mathbf{x}^k), x_{i_k} - \mathbf{x}_{i_k}^k \rangle \leq \phi(\mathbf{x}^k) - \phi(\mathbf{x}^{k+1}). \quad (7)$$

Since the function f is L -smooth, we have

$$f(\bar{\mathbf{x}}_{i_k}^k + \mathbf{x}_{i_k}) \leq f(\mathbf{x}^k) + \langle \nabla_{i_k} f(\mathbf{x}^k), x_{i_k} - \mathbf{x}_{i_k}^k \rangle + \frac{L}{2} \|x_{i_k} - \mathbf{x}_{i_k}^k\|^2. \quad (8)$$

Combining the last two inequalities, we obtain

$$\langle \nabla_{i_k} f(\mathbf{x}^k) - \nabla_{i_k} h(\mathbf{x}^k), x_{i_k} - \mathbf{x}_{i_k}^k \rangle + g_{i_k}(\mathbf{x}_{i_k}^k) - g_{i_k}(x_{i_k}) - \frac{L}{2} \|x_{i_k} - \mathbf{x}_{i_k}^k\|^2 \leq \phi(\mathbf{x}^k) - \phi(\mathbf{x}^{k+1}). \quad (9)$$

Let us denote by \mathbb{E}_k the conditional expectation with respect to the random selection of i_k , given all the random choices in the previous iterations. Then,

$$\begin{aligned} &\mathbb{E}_k \left[\langle \nabla_{i_k} f(\mathbf{x}^k) - \nabla_{i_k} h(\mathbf{x}^k), x_{i_k} - \mathbf{x}_{i_k}^k \rangle + g_{i_k}(\mathbf{x}_{i_k}^k) - g_{i_k}(x_{i_k}) - \frac{L}{2} \|x_{i_k} - \mathbf{x}_{i_k}^k\|^2 \right] \\ &= \frac{1}{n} \sum_{i=1}^n \left(\langle \nabla_i f(\mathbf{x}^k) - \nabla_i h(\mathbf{x}^k), x_i - \mathbf{x}_i^k \rangle + g_i(\mathbf{x}_i^k) - g_i(x_i) - \frac{L}{2} \|x_i - \mathbf{x}_i^k\|^2 \right) \\ &= \frac{1}{n} \left(\langle \nabla f(\mathbf{x}^k) - \nabla h(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x} \rangle + g(\mathbf{x}^k) - g(\mathbf{x}) - \frac{L}{2} \|\mathbf{x} - \mathbf{x}^k\|^2 \right). \end{aligned}$$

Combining this with [\(9\)](#) gives the following bound, which holds for all $\mathbf{x} \in \mathcal{M}$:

$$\langle \nabla f(\mathbf{x}^k) - \nabla h(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x} \rangle + g(\mathbf{x}^k) - g(\mathbf{x}) - \frac{L}{2} \|\mathbf{x} - \mathbf{x}^k\|^2 \leq n\phi(\mathbf{x}^k) - n\mathbb{E}_k \left[\phi(\mathbf{x}^{k+1}) \right]. \quad (10)$$

Now, we maximize this inequality over $\mathbf{x} \in \mathcal{M}$ to get

$$\text{gap}_{\mathcal{M}}^L(\mathbf{x}^k) \leq n\phi(\mathbf{x}^k) - n\mathbb{E}_k \left[\phi(\mathbf{x}^{k+1}) \right]. \quad (11)$$

Then, we take the expectation of both sides over the random choices in all iterations:

$$\mathbb{E} \left[\text{gap}_{\mathcal{M}}^L(\mathbf{x}^k) \right] \leq n\mathbb{E} \left[\phi(\mathbf{x}^k) \right] - n\mathbb{E} \left[\phi(\mathbf{x}^{k+1}) \right]. \quad (12)$$

Finally, we average this inequality over $k = 1, \dots, K$:

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E} \left[\text{gap}_{\mathcal{M}}^L(\mathbf{x}^k) \right] \leq \frac{n}{K} \left(\phi(\mathbf{x}^1) - \mathbb{E} \left[\phi(\mathbf{x}^{K+1}) \right] \right) \leq \frac{n}{K} \left(\phi(\mathbf{x}^1) - \phi^* \right). \quad (13)$$

We complete the proof by noting that the minimum of $\text{gap}_{\mathcal{M}}^L(\mathbf{x}^k)$ over $k = 1, \dots, K$ is smaller than or equal to the average gap. \square

Remark 5. It is straightforward to verify that [Theorem 4](#) remains valid if the directional gradient $\nabla_{i_k} h(\mathbf{x}^k)$ is replaced by the i_k -th coordinate of a subgradient, $v_{i_k}^k := \mathbf{D}_{i_k} \mathbf{v}^k$, where $\mathbf{v}^k \in \partial h(\mathbf{x}^k)$. The proof proceeds in the same manner.

Remark 6. The rate in [Theorem 4](#) is consistent with existing results for CD methods in convex optimization. For instance, [Theorem 1](#) in ([Wright, 2015](#)) shows that $\mathbb{E}[f(\mathbf{x}^k)] - f^* \leq \frac{2nL\|\mathbf{x}^1 - \mathbf{x}_*\|^2}{k}$ for the CD method in unconstrained smooth convex minimization. In terms of first-order stationarity condition, this implies

$$\mathbb{E} \left[\|\nabla f(\mathbf{x}^k)\|^2 \right] \leq 2L\mathbb{E} \left[f(\mathbf{x}^k) - f^* \right] \leq \frac{4nL^2\|\mathbf{x}^1 - \mathbf{x}_*\|^2}{k}.$$

[Theorem 4](#) yields a comparable result, since in this setting $\phi = f$ satisfies $f(\mathbf{x}^1) - f(\mathbf{x}^*) \leq \frac{L}{2}\|\mathbf{x}^1 - \mathbf{x}^*\|^2$, and the gap takes the form of $\text{gap}^L(\mathbf{x}^k) = \frac{1}{2L}\|\nabla f(\mathbf{x}^k)\|^2$. Thus, we obtain

$$\min_{k \in \{1, \dots, K\}} \mathbb{E} \left[\|\nabla f(\mathbf{x}^k)\|^2 \right] \leq \frac{nL^2\|\mathbf{x}^1 - \mathbf{x}_*\|^2}{K}.$$

Identifying the gap function in [Definition 1](#) as a measure of first-order stationarity is a key technical aspect of our analysis. Notably, standard measures are not suitable for non-smooth non-convex optimization problems, motivating recent efforts to explore and reassess alternative definitions ([Zhang et al., 2020](#); [Kornowski and Shamir, 2021](#)). For example, the distance measure

$$\text{dist}(0, \nabla f(\mathbf{y}) + \partial g(\mathbf{y}) - \nabla h(\mathbf{y}) + \mathcal{N}_{\mathcal{M}}(\mathbf{y})),$$

or the ‘Frank-Wolfe gap’ using a subgradient $\mathbf{u} \in \partial g(\mathbf{y})$

$$\max_{\mathbf{x} \in \mathcal{M}} \langle \nabla f(\mathbf{y}) + \mathbf{u} - \nabla h(\mathbf{y}), \mathbf{y} - \mathbf{x} \rangle$$

are intractable, meaning that no algorithm can guarantee finding a point satisfying ϵ -closeness in these notions within a finite number of iterations.

Our results indicate that the challenge of identifying appropriate stationarity measures stems from the non-smooth convex term, while the non-smooth concave term remains benign. Although it may initially seem surprising, this can be explained by the fact that the concave term does not invalidate the smoothness-based quadratic upper bound, since

$$\begin{aligned} f(\mathbf{y}) &\leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2}\|\mathbf{x} - \mathbf{y}\|^2 \\ -h(\mathbf{y}) &\leq -h(\mathbf{x}) - \langle \nabla h(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \end{aligned}$$

for all \mathbf{x} and \mathbf{y} . Summing these two inequalities, we see that $f(\mathbf{y}) - h(\mathbf{y})$ behaves similarly to a smooth function in the analysis, as the quadratic upper bound is the only property required from the smoothness assumption.

4 Block EM Algorithm

The Expectation-Maximization (EM) algorithm is a cornerstone in statistical analysis for finding maximum likelihood estimates in the presence of incomplete data; widely used in applications such as mixture models, latent variable models, data imputation, and clustering (Dempster et al., 1977; McLachlan and Krishnan, 2007). EM shares strong connections with DCA (Yuille and Rangarajan, 2003; Le Thi and Pham Dinh, 2018). In this section, we leverage these connections to develop a block EM algorithm based on our proposed BDCA.

Consider negative log-likelihood minimization problem:

$$\min_{\theta} \mathcal{L}(\theta) := - \sum_{x \in \mathcal{X}} \log P(x|\theta). \quad (14)$$

Note the change in notation: here, x denotes data points, while the decision variable $\theta \in \Theta \subseteq \mathbb{R}^m$ represents the parameters of the underlying distribution. The dataset \mathcal{X} has support size N , and \mathcal{L} is the negative log-likelihood function. The distribution $P(x|\theta)$ is defined through a hidden variable y :

$$P(x|\theta) = \sum_y P(x, y|\theta),$$

where $P(x, y|\theta)$ denotes the complete-data likelihood. The presence of hidden variables complicate (14), making standard optimization techniques inapplicable.

The EM algorithm, for solving this problem, iteratively alternates between two steps – Expectation (E-step) and Maximization (M-step) – to improve the likelihood of the model parameters until convergence:

$$\begin{aligned} Q(\theta; \theta^k) &= \mathbb{E}_{P(y|x, \theta^k)} [\log P(x, y|\theta)] && \text{(E-step)} \\ \theta^{k+1} &= \arg \max_{\theta \in \Theta} Q(\theta; \theta^k) && \text{(M-step)} \end{aligned}$$

Here, $P(y|x, \theta^k)$ is the posterior distribution of the latent variables given the observed data and the current parameter estimate θ^k . In the E-step, the expected value of the complete-data log-likelihood is computed with respect to the conditional distribution of the latent variables given the observed data and the current parameter estimates. In the M-step, this expected log-likelihood is maximized to update the parameters.

For the broad class of natural exponential family distributions, EM is a specific instance of DCA.

Definition 2. *The natural exponential family of distributions with a parameter $\theta \in \Theta$ is characterized by the following probability mass (or density) function:*

$$P(x, y|\theta) = \frac{\varphi(x, y)e^{\langle \theta, T(x, y) \rangle}}{\sum_{x, y} \varphi(x, y)e^{\langle \theta, T(x, y) \rangle}}, \quad (15)$$

where φ is a positive function (base measure) and T denotes the sufficient statistic.

It is easy to see that problem (14) with a natural exponential family distribution (15) can be written as a DC program, since the log-likelihood function can be decomposed as $\mathcal{L} = f - h$, where the terms are given by

$$f(\theta) = N \log \left(\sum_{x, y} \varphi(x, y)e^{\langle \theta, T(x, y) \rangle} \right), \quad \text{and} \quad h(\theta) = \sum_{x \in \mathcal{X}} \log \left(\sum_y \varphi(x, y)e^{\langle \theta, T(x, y) \rangle} \right). \quad (16)$$

Both terms are convex and smooth, owing to their log-sum-exp structure.

Applying DCA to this problem leads to the following update rule:

$$\theta^{k+1} \in \operatorname{argmin}_{\theta} N \log \left(\sum_{x, y} \varphi(x, y)e^{\langle \theta, T(x, y) \rangle} \right) - \sum_{x \in \mathcal{X}} \frac{\sum_y \varphi(x, y) \langle \theta, T(x, y) \rangle e^{\langle \theta^k, T(x, y) \rangle}}{\sum_y \varphi(x, y) e^{\langle \theta^k, T(x, y) \rangle}}. \quad (17)$$

Algorithm 2 Block EM Algorithm

Input: total iterations K , initialize the probability distribution P using vector $\theta^{(0)}$ (P belongs to the exponential family), and number of blocks n

for $k = 1$ **to** K **do**

Randomly choose i_k in $[1, \dots, n]$ with uniform distribution

$$\text{Update } \hat{P}^{k+1}(\mathbf{y}) = \frac{P(\mathbf{x}, \mathbf{y} | \theta^k)}{\sum_{\mathbf{y}} P(\mathbf{x}, \mathbf{y} | \theta^k)}$$

$$\text{Update } \theta_{i_k}^{k+1} = \arg \min_{\theta_{i_k}} - \sum_{\mathbf{x} \in \mathcal{X}, \mathbf{y}} \hat{P}^{k+1}(\mathbf{y}) \log P(\mathbf{x}, \mathbf{y} | \theta_1^k, \dots, \theta_{i_k}, \dots, \theta_n^k)$$

$$\text{Set } \theta^{k+1} = \bar{\theta}_{i_k}^k + \theta_{i_k}^{k+1}$$

end for

Output: θ^{K+1}

After algebraic simplifications, this reduces to:

$$\theta^{k+1} \in \operatorname{argmin}_{\theta} \sum_{\mathbf{x} \in \mathcal{X}} \sum_{\mathbf{y}} \frac{\varphi(\mathbf{x}, \mathbf{y}) e^{\langle \theta^k, T(\mathbf{x}, \mathbf{y}) \rangle}}{\sum_{\mathbf{y}} \varphi(\mathbf{x}, \mathbf{y}) e^{\langle \theta, T(\mathbf{x}, \mathbf{y}) \rangle}} \times \left[\log \left(\frac{\sum_{\mathbf{x}, \mathbf{y}} \varphi(\mathbf{x}, \mathbf{y}) e^{\langle \theta, T(\mathbf{x}, \mathbf{y}) \rangle}}{\varphi(\mathbf{x}, \mathbf{y}) e^{\langle \theta, T(\mathbf{x}, \mathbf{y}) \rangle}} \right) \right] \quad (18)$$

which is exactly the EM algorithm for the exponential family.

We then design the Block EM method by addressing the same problem using BDCA. The resulting algorithm is presented in [Algorithm 2](#). Notably, we can transfer the guarantees established in [Theorem 4](#) to provide non-asymptotic convergence results for Block EM.

Corollary 7. *Suppose there exists a sequence $\theta^k, k = 1 \dots, K$ generated by Block EM algorithm. Then,*

$$\min_{k \in \{1, \dots, K\}} \mathbb{E} \left[\|\nabla \mathcal{L}(\theta^k)\|^2 \right] \leq \frac{2nL}{K} (\mathcal{L}(\theta^1) - \mathcal{L}(\theta^*)).$$

Proof. In this setting, the gap function reduces to

$$\max_{\theta} \underbrace{\langle \nabla f(\theta^k) - \nabla h(\theta^k), \theta^k - \theta \rangle}_{\nabla \mathcal{L}(\theta^k)} - \frac{L}{2} \|\theta^k - \theta\|^2.$$

The solution to this max problem is $\theta = \theta^k - \frac{1}{L} \nabla \mathcal{L}(\theta^k)$, leading to $\text{gap}^L(\theta^k) = \frac{1}{2L} \|\nabla \mathcal{L}(\theta^k)\|^2$. Substituting this into [Theorem 4](#) yields the desired result. \square

Remark 8. *The results in [Corollary 7](#) are consistent with the $\mathcal{O}(1/k)$ convergence rate guarantee for the standard EM algorithm, as provided by [Kumar and Schmidt \(2017\)](#).*

Remark 9. *For the standard EM algorithm, [Kunstner et al. \(2020\)](#) has shown $\mathcal{O}(1/k)$ convergence rate also in Kullback-Leibler (KL) divergence. It is worth noting that the Lipschitz constant and the Euclidean distance generating function are not utilized in BDCA; they appear in the analysis only through the definition of the gap function. Consequently, we believe our analysis can be extended to the general setting of Bregman divergence.*

5 Conclusions

We investigated a general class of DC problems involving smooth and separable non-smooth terms, and proposed a randomized block-coordinate DC algorithm (BDCA) specifically designed to solve such problems efficiently. We established convergence guarantees for BDCA with a rate of $\mathcal{O}(n/K)$, where the algorithm performs K iteration over n coordinate blocks. Leveraging the connection between the DCA and the EM method, we proposed a block-coordinate EM algorithm (Block EM), where convergence guarantees are directly inherited from the DCA as a special case. DCA and EM have

numerous applications in data science and machine learning; hence, scalable alternatives like BDCA and Block EM have the potential for significant impact.

As part of our analysis, we identified tractable first-order stationarity conditions for non-smooth non-convex optimization problems based on a DC representation. Given the prevalence of non-smooth non-convex problems in machine learning applications (e.g., ReLU neural networks and other models with non-differentiable activation functions) and the number of open questions surrounding their non-asymptotic convergence analysis, we believe our findings should be of broad interest to the optimization community within machine learning research.

Acknowledgments

HM and AY were supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. SS acknowledges generous support from the Alexander von Humboldt Foundation.

References

- Argyriou, A., Hauser, R., Micchelli, C. A., and Pontil, M. (2006). A DC-programming algorithm for kernel selection. In *Proceedings of the 23rd international conference on Machine learning*, pages 41–48.
- Aubry, A., De Maio, A., Zappone, A., Razaviyayn, M., and Luo, Z.-Q. (2018). A new sequential optimization procedure and its applications to resource allocation for wireless systems. *IEEE Transactions on Signal Processing*, 66(24):6518–6533.
- Awasthi, P., Cortes, C., and Mohri, M. (2024a). Best-effort adaptation. *Annals of Mathematics and Artificial Intelligence*, pages 1–46.
- Awasthi, P., Mao, A., Mohri, M., and Zhong, Y. (2024b). DC-programming for neural network optimizations. *Journal of Global Optimization*, pages 1–17.
- Cai, X., Song, C., Wright, S., and Diakonikolas, J. (2023). Cyclic block coordinate descent with variance reduction for composite nonconvex optimization. In *International Conference on Machine Learning*, pages 3469–3494. PMLR.
- Chorobura, F. and Necoara, I. (2023). Random coordinate descent methods for nonseparable composite optimization. *SIAM Journal on Optimization*, 33(3):2160–2190.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1):1–22.
- Deng, Q. and Lan, C. (2020). Efficiency of coordinate descent methods for structured nonconvex optimization. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 74–89. Springer.
- Kornowski, G. and Shamir, O. (2021). Oracle complexity in nonsmooth nonconvex optimization. *Advances in Neural Information Processing Systems*, 34:324–334.
- Kumar, R. and Schmidt, M. (2017). Convergence rate of expectation-maximization. In *10th NIPS Workshop on Optimization for Machine Learning*, page 98.
- Kunstner, F., Kumar, R., and Schmidt, M. W. (2020). Homeomorphic-invariance of EM: Non-asymptotic convergence in KL divergence for exponential families via mirror descent. In *International Conference on Artificial Intelligence and Statistics*.
- Lange, K. (2016). *MM optimization algorithms*. SIAM.
- Le Thi, H. A. and Pham Dinh, T. (2018). DC programming and DCA: thirty years of developments. *Mathematical Programming*, 169.
- Li, H., Zhang, K., and Jiang, T. (2005). The regularized EM algorithm. In *AAAI*, pages 807–812.
- McLachlan, G. J. and Krishnan, T. (2007). *The EM algorithm and extensions*. John Wiley & Sons.
- Pham Dinh, T., Huynh, V. N., Le Thi, H. A., and Ho, V. T. (2022). Alternating DC algorithm for partial DC programming problems. *Journal of Global Optimization*, 82(4):897–928.

- Plabst, D., Munir, J., Mezghani, A., and Nossek, J. A. (2018). Efficient non-linear equalization for 1-bit quantized cyclic prefix-free massive MIMO systems. In *2018 15th International Symposium on Wireless Communication Systems (ISWCS)*, pages 1–6. IEEE.
- Razaviyayn, M., Hong, M., and Luo, Z.-Q. (2013). A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal on Optimization*, 23(2):1126–1153.
- Sriperumbudur, B. K., Torres, D. A., and Lanckriet, G. R. (2007). Sparse eigen methods by DC programming. In *Proceedings of the 24th international conference on Machine learning*, pages 831–838.
- Tao, P. D. and Souad, E. B. (1986). Algorithms for solving a class of nonconvex optimization problems. methods of subgradients. In Hiriart-Urruty, J.-B., editor, *Fermat Days 85: Mathematics for Optimization*, volume 129 of *North-Holland Mathematics Studies*, pages 249–271. North-Holland.
- Wright, S. J. (2015). Coordinate descent algorithms. *Mathematical programming*, 151(1):3–34.
- Xu, Y. and Yin, W. (2017). A globally convergent algorithm for nonconvex optimization based on block coordinate update. *Journal of Scientific Computing*, 72(2):700–734.
- Yi, X. and Caramanis, C. (2015). Regularized EM algorithms: A unified framework and statistical guarantees. *Advances in Neural Information Processing Systems*, 28.
- Yuan, G. (2023). Coordinate descent methods for dc minimization: Optimality conditions and global convergence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 11034–11042.
- Yuille, A. L. and Rangarajan, A. (2003). The concave-convex procedure. *Neural computation*, 15(4):915–936.
- Yurtsever, A. and Sra, S. (2022). CCCP is Frank-Wolfe in disguise. *Advances in Neural Information Processing Systems*, 35:35352–35364.
- Zhang, J., Lin, H., Jegelka, S., Sra, S., and Jadbabaie, A. (2020). Complexity of finding stationary points of nonconvex nonsmooth functions. In *International Conference on Machine Learning*, pages 11173–11182. PMLR.

A Regularized EM

In [Section 4](#), we discussed the relationship between the EM algorithm and DCA. To maintain simplicity and consistency with the standard EM algorithm, we focused on minimizing the negative log-likelihood function. However, it is common practice to solve regularized log-likelihood problems, where a regularization term is introduced to prevent overfitting, improve generalization, or promote specific structures in the solution ([Li et al., 2005](#); [Yi and Caramanis, 2015](#)). In this context, the Regularized EM (REM) algorithm seeks to optimize the penalized negative log-likelihood problem:

$$\min_{\theta \in \Theta} \hat{\mathcal{L}}(\theta) := \sum_{x \in \mathcal{X}} \sum_{y} -\log P(x, y | \theta) + \lambda R(\theta), \quad (\text{SM1})$$

where R is a convex regularization function with block-separable structure $R(\theta) = \sum_{i=1}^n R_i(\theta_i)$, and $\lambda \geq 0$ is the regularization parameter. Similarly, we assume Θ can be decomposed as $\Theta_1 \times \dots \times \Theta_n$, where the components, such that $\theta \in \Theta$ if and only if $\theta_i \in \Theta_i$ for all $i = 1, \dots, n$.

For exponential family of distributions, as defined in [Definition 2](#), we can express [\(SM1\)](#) as a specific instance of the problem template

$$\min_{\theta \in \Theta} f(\theta) + g(\theta) - h(\theta), \quad (\text{SM2})$$

where f and h are as defined in [\(16\)](#), and $g(\theta) = \lambda R(\theta)$.

Employing BDCA to solve [\(SM2\)](#) results in Block REM method shown in [Algorithm SM1](#). The following corollary, which is a direct consequence of [Theorem 4](#), formulates the convergence behaviour of the Block REM method in terms of the introduced gap function (See [Definition 1](#)).

Corollary SM1. *Suppose there exists a sequence $\theta^k, k = 1 \dots, K$ generated by Block REM algorithm. Then,*

$$\min_{k \in \{1, \dots, K\}} \mathbb{E} \left[\text{gap}_{\Theta}^L(\theta^k) \right] \leq \frac{n}{K} \left(\hat{\mathcal{L}}(\theta^1) - \hat{\mathcal{L}}^* \right).$$

Algorithm SM1 Block REM Algorithm

Input: regularization parameter λ , total iterations K , initialize the probability distribution P using vector $\theta^{(0)}$ (P belongs to the exponential family), and number of blocks n

for $k = 1$ **to** K **do**

Randomly choose i_k in $[1, \dots, n]$ with uniform distribution

$$\text{Update } \hat{P}^{k+1}(\mathbf{y}) = \frac{P(\mathbf{x}, \mathbf{y} | \theta^k)}{\sum_{\mathbf{y}} P(\mathbf{x}, \mathbf{y} | \theta^k)}$$

$$\text{Update } \theta_{i_k}^{k+1} = \arg \min_{\theta_{i_k} \in \Theta_{i_k}} \sum_{\mathbf{x} \in \mathcal{X}, \mathbf{y}} -\hat{P}^{k+1}(\mathbf{y}) \log P(\mathbf{x}, \mathbf{y} | \theta_1^k, \dots, \theta_{i_k}, \dots, \theta_n^k) + \lambda R_{i_k}(\theta_{i_k})$$

$$\text{Set } \theta^{k+1} = \bar{\theta}^k + \theta_{i_k}^{k+1}$$

end for

Output: θ^{K+1}

B BDCA Recovers Coordinate (Sub)Gradient Descent

Randomized coordinate subgradient descent method (RCSD) is a block coordinate proximal type subgradient algorithm for similar DC nonconvex problems as ours (see problem (1)). RCSD assumes an additional assumption on function g to have a simple structure (Deng and Lan, 2020). Here, we briefly present RCSD and its convergence guarantee. Later, we attempt to show how BDCA recovers RCSD.

RCSD is a block coordinate subgradient method which iteratively updates some random coordinates while keeping the rest of the coordinates fixed (similar to BDCA). To formally present the method, we need block proximal mapping given by

$$\mathcal{P}_i(\bar{x}_i, y_i, \gamma_i) = \operatorname{argmin}_{x \in \mathbb{R}^{m_i}} \left\{ \langle y_i, x \rangle + g_i(x) + \frac{\gamma_i}{2} \|\bar{x}_i - x\|_i^2 \right\}, \quad \forall y_i \in \mathbb{R}^{m_i} \quad (\text{SM3})$$

where $\gamma = [\gamma_1, \gamma_2, \dots, \gamma_n]$, $\|\cdot\|_i$ is the standard Euclidean distance on \mathbb{R}^{m_i} , and $\mathcal{P}(\bar{x}, \mathbf{y}, \gamma) = \sum_{i=1}^n \mathbf{D}_i^T \mathbf{D}_i \mathcal{P}_i(\bar{x}_i, y_i, \gamma_i)$. The RCSD method is shown in Algorithm SM2.

Next, we will present the convergence result of RCSD for uniform sampling strategy. For more detailed description of the results, see Deng and Lan (2020).

Theorem SM2 (Theorem 3, (Deng and Lan, 2020)). *Suppose the sequence $\mathbf{x}^k, k = 1, \dots, K$, is generated by Algorithm SM2. Assume $\gamma_i = L_i$ for L_i being block wise Lipschitz constants for function f . Then,*

$$\min_{k \in \{0, \dots, K\}} \mathbb{E} \left\{ \|G(\mathbf{x}^k, \nabla f(\mathbf{x}^k) - \mathbf{v}^k, \gamma)\|^2 \right\} \leq \frac{2L_{\max} n (\phi(\mathbf{x}^1) - \phi^*)}{K}, \quad (\text{SM4})$$

where $G(\mathbf{x}^k, \nabla f(\mathbf{x}^k) - \mathbf{v}^k, \gamma) = \sum_{i=1}^n \mathbf{D}_i^T \mathbf{D}_i G_i(x_i^k, \nabla_i f(\mathbf{x}^k) - \mathbf{v}_i^k, \gamma_i)$, $L_{\max} = \max_{i \in \{1, \dots, n\}} L_i$, and

$$G_i(x_i, \nabla_i f(\mathbf{x}) - \mathbf{v}_i, \gamma_i) = \gamma_i (x_i - \mathcal{P}_i(x_i, \nabla_i f(\mathbf{x}) - \mathbf{v}_i, \gamma_i)).$$

Now, we will show how BDCA relates to RCSD. This result is based on the smoothness of f .

Proposition SM3. *BDCA covers RCSD with uniform sampling strategy as a special instance.*

Proof. Due to smoothness of f we can write the objective function as

$$\phi(\mathbf{x}) = \frac{L_{\max}}{2} \|\mathbf{x}\|^2 + g(\mathbf{x}) - (h(\mathbf{x}) + \frac{L_{\max}}{2} \|\mathbf{x}\|^2 - f(\mathbf{x})).$$

Algorithm SM2 RCSD Algorithm

Input: $\mathbf{x}^{(0)}$, total iterations K , coefficients γ , and number of block n
for $k = 1$ **to** K **do**
 Randomly choose i_k in $[1, \dots, n]$ with uniform distribution
 Compute $\nabla_{i_k} f(\mathbf{x}^k)$ and $\mathbf{v}^k \in \partial h(\mathbf{x}^k)$
 Update $\mathbf{x}_{i_k}^{k+1} = \mathcal{P}_{i_k}(\mathbf{x}_{i_k}^k, \nabla_{i_k} f(\mathbf{x}^k) - \mathbf{v}_{i_k}^k, \gamma_{i_k})$
 Set $\mathbf{x}^{k+1} = \bar{\mathbf{x}}_{i_k}^k + \mathbf{x}_{i_k}^{k+1}$
end for
Output: \mathbf{x}^{K+1}

Take $\mathbf{v}_{i_k}^k \in \partial_{i_k} h(\mathbf{x}^k)$. This gives the BDCA subproblem as

$$\begin{aligned}
 \mathbf{x}_{i_k}^{k+1} &\in \operatorname{argmin}_{\mathbf{x}_{i_k} \in \mathcal{M}_{i_k}} \frac{L_{max}}{2} \|\mathbf{x}_{i_k}\|^2 + g_{i_k}(\mathbf{x}_{i_k}) - \langle \mathbf{v}_{i_k}^k - \nabla_{i_k} f(\mathbf{x}^k) + L_{max} \mathbf{x}_{i_k}^k, \mathbf{x}_{i_k} \rangle \\
 &= \operatorname{argmin}_{\mathbf{x}_{i_k} \in \mathcal{M}_{i_k}} \langle \nabla_{i_k} f(\mathbf{x}^k) - \mathbf{v}_{i_k}^k, \mathbf{x}_{i_k} \rangle + g_{i_k}(\mathbf{x}_{i_k}) + \frac{L_{max}}{2} \|\mathbf{x}_{i_k} - \mathbf{x}_{i_k}^k\|^2 \\
 &= \operatorname{argmin}_{\mathbf{x}_{i_k} \in \mathcal{M}_{i_k}} \langle \nabla_{i_k} f(\mathbf{x}^k) \mathbf{v}_{i_k}^k, \mathbf{x}_{i_k} - \mathbf{x}_{i_k}^k \rangle + g_{i_k}(\mathbf{x}_{i_k}) + \frac{L_{max}}{2} \|\mathbf{x}_{i_k} - \mathbf{x}_{i_k}^k\|^2 \\
 &= \operatorname{argmin}_{\mathbf{x}_{i_k} \in \mathcal{M}_{i_k}} \frac{L_{max}}{2} \langle \frac{2}{L_{max}} (\nabla_{i_k} f(\mathbf{x}^k) - \mathbf{v}_{i_k}^k), \mathbf{x}_{i_k} - \mathbf{x}_{i_k}^k \rangle + g_{i_k}(\mathbf{x}_{i_k}) + \frac{L_{max}}{2} \|\mathbf{x}_{i_k} - \mathbf{x}_{i_k}^k\|^2 \\
 &= \operatorname{argmin}_{\mathbf{x}_{i_k} \in \mathcal{M}_{i_k}} g_{i_k}(\mathbf{x}_{i_k}) + \frac{L_{max}}{2} \|\mathbf{x}_{i_k} - (\mathbf{x}_{i_k}^k - \frac{1}{L_{max}} (\nabla_{i_k} f(\mathbf{x}^k) - \mathbf{v}_{i_k}^k))\|^2 \\
 &= \operatorname{prox}_{\frac{1}{L_{max}} g_{i_k} + I_{\mathcal{M}_{i_k}}} (\mathbf{x}_{i_k}^k - \frac{1}{L_{max}} (\nabla_{i_k} f(\mathbf{x}^k) - \mathbf{v}_{i_k}^k))
 \end{aligned}$$

which is the block proximal operator with constant $1/L_{max}$ with $I_{\mathcal{C}}$ denoting the indicator function on the set \mathcal{C} . Application of this update for each randomly selected block i_k in BDCA gives RCSD. \square

C Stationarity in Nonsmooth Optimization

In this section, we illustrate through simple examples that two common measures for first-order stationarity in smooth optimization, namely the distance measure

$$\operatorname{dist}(0, \nabla f(\mathbf{y}) + \partial g(\mathbf{y}) - \nabla h(\mathbf{y}) + \mathcal{N}_{\mathcal{M}}(\mathbf{y})),$$

and the Frank-Wolfe gap (but with a subgradient $\mathbf{u} \in \partial g(\mathbf{y})$)

$$\max_{\mathbf{x} \in \mathcal{M}} \langle \nabla f(\mathbf{y}) + \mathbf{u} - \nabla h(\mathbf{y}), \mathbf{y} - \mathbf{x} \rangle$$

are intractable in nonsmooth optimization. Specifically, we show that these measures can be discontinuous at the solution and remain lower bounded even when approaching the solution arbitrarily closely.

Consider the problem where $\mathcal{M} = [-1, 1]$, $f(x) = h(x) = 0$, and $g = |x|$. The problem is convex, and the only first-order stationary point is at $x = 0$, which is also the global minimum. The subdifferential of g is and the normal cone of \mathcal{M} are given by

$$\partial g(x) = \begin{cases} \{1\} & \text{if } x > 0, \\ \{-1\} & \text{if } x < 0, \\ [-1, 1] & \text{if } x = 0, \end{cases} \quad \mathcal{N}_{\mathcal{M}}(x) = \begin{cases} \{0\} & \text{if } x \in (-1, 1), \\ (-\infty, 0] & \text{if } x = -1, \\ [0, +\infty) & \text{if } x = 1. \end{cases}$$

Then, the distance measure becomes

$$\text{dist}(0, \partial g(y) + \mathcal{N}_{\mathcal{M}}(y)) = \min_{\substack{u \in \partial g(y) \\ z \in \mathcal{N}_{\mathcal{M}}(y)}} |u + z - 0| = \begin{cases} 0 & \text{if } y = 0, \\ 1 & \text{otherwise,} \end{cases}$$

indicating that the distance is 0 only at the global optimum and remains 1 at all other points, even when arbitrarily close to the global optimum.

Similarly, the Frank-Wolfe gap, even minimized over the subdifferential set, becomes

$$\min_{u \in \partial g(y)} \max_{x \in [-1,1]} \langle u, y - x \rangle = \begin{cases} 0 & \text{if } y = 0, \\ y + 1 & \text{if } y > 0, \\ 1 - y & \text{if } y < 0, \end{cases} \quad (\text{SM5})$$

which shows that, except at $y = 0$, the Frank-Wolfe gap is larger than 1 at all other points, even when arbitrarily close to the global optimum.

Notably, since the examples are convex, these findings also extend to convex nonsmooth optimization problems.

D Proof of Remark 2

We start by showing that $\text{gap}_{\mathcal{M}}^L(\mathbf{y}) = 0$ if

$$\mathbf{y} = \text{prox}_{\frac{1}{L}g + I_{\mathcal{M}}} \left(\left(\mathbf{y} - \frac{1}{L}(\nabla f(\mathbf{y}) - \nabla h(\mathbf{y})) \right) \right).$$

Using the definition of the proximal map we have

$$\begin{aligned} \mathbf{y} &= \underset{x \in \mathcal{M}}{\text{argmin}} \quad g(x) + \frac{L}{2} \|x - (\mathbf{y} - \frac{1}{L}(\nabla f(\mathbf{y}) - \nabla h(\mathbf{y}))\|^2 \\ &= \underset{x \in \mathcal{M}}{\text{argmin}} \quad g(x) + \frac{L}{2} \|x - \mathbf{y}\|^2 + \langle x - \mathbf{y}, \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}) \rangle \\ &= \underset{x \in \mathcal{M}}{\text{argmax}} \quad -g(x) - \frac{L}{2} \|x - \mathbf{y}\|^2 + \langle \mathbf{y} - x, \nabla f(\mathbf{y}) - \nabla h(\mathbf{y}) \rangle. \end{aligned} \quad (\text{SM6})$$

From (SM6) it follows that $\text{gap}_{\mathcal{M}}^L(\mathbf{y}) = 0$.

Next, we will show that \mathbf{y} is the fixed point of the proximal mapping if $\text{gap}_{\mathcal{M}}^L = 0$. We start by the reformulating the definition of the gap function:

$$\text{gap}_{\mathcal{M}}^L(\mathbf{y}) = L \max_{x \in \mathcal{M}} \left\{ \langle x + \frac{1}{L}(\nabla f(\mathbf{y}) - \nabla h(\mathbf{y})) - \mathbf{y}, \mathbf{y} - x \rangle + \frac{1}{L}g(\mathbf{y}) - \frac{1}{L}g(x) + \frac{1}{2}\|x - \mathbf{y}\|^2 \right\}.$$

By assumption we have $\text{gap}_{\mathcal{M}}^L = 0$. This implies:

$$\langle x + \frac{1}{L}(\nabla f(\mathbf{y}) - \nabla h(\mathbf{y})) - \mathbf{y}, \mathbf{y} - x \rangle + \frac{1}{L}g(\mathbf{y}) - \frac{1}{L}g(x) + \frac{1}{2}\|x - \mathbf{y}\|^2 \leq 0, \quad \forall x \in \mathcal{M}. \quad (\text{SM7})$$

Consider this inequality with $x = \text{prox}_{\frac{1}{L}g + I_{\mathcal{M}}}(u)$ and $u = \mathbf{y} - \frac{1}{L}(\nabla f(\mathbf{y}) - \nabla h(\mathbf{y}))$. Then, using non-expansiveness of the proximal operator we have

$$\begin{aligned} \langle u - x, \mathbf{y} - x \rangle &\leq \frac{1}{L}g(\mathbf{y}) - \frac{1}{L}g(x) \\ \langle u - x, \mathbf{y} - x \rangle - \frac{1}{L}g(\mathbf{y}) + \frac{1}{L}g(x) &\leq 0 \\ \langle x - u, \mathbf{y} - x \rangle + \frac{1}{L}g(\mathbf{y}) - \frac{1}{L}g(x) &\geq 0. \end{aligned} \quad (\text{SM8})$$

Utilizing (SM8) in (SM7) implies $\mathbf{y} = x = \text{prox}_{\frac{1}{L}g + I_{\mathcal{M}}}(u)$ and the proof is complete.

E Application to One-bit MIMO Signal Recovery

In this section, we will use the proposed Block EM method in one-bit MIMO signal recovery problem. First, we describe the problem. Later, simulation results are given to show the performance of our method compared to the conventional EM method.

E.1 Problem Description

Consider a MIMO uplink scenario where N transmitters (users) simultaneously send a signal of length W to the Base-Station (BS) which consists of M receiving antennas. The signal will undergo a frequency selective fading channel. Then, the received signal at the BS will have the form

$$\mathbf{R}_m = \sum_{n=1}^N \mathbf{H}_{m,n} \boldsymbol{\theta}_n + \boldsymbol{\alpha}_m \quad \text{and} \quad \mathbf{Y}_m = Q(\mathbf{R}_m), \quad \text{for } m = 1, \dots, M, \quad (\text{SM9})$$

where $\mathbf{R}_m \in \mathbb{C}^W$ is the received signal at m th antenna, $\mathbf{H}_{m,n} \in \mathbb{C}^{W \times W}$ denotes the channel circulant matrix between the n th user and m th receiver, $\boldsymbol{\theta}_n \in \mathbb{C}^W$ is the transmitted signal of the n th user, and $\boldsymbol{\alpha}_m \in \mathbb{C}^W$ is the complex additive white Gaussian noise for the m th receiver with mean $\mathbf{0}$ and covariance matrix $\sigma_\alpha^2 \mathbf{I}$. Note that each $\mathbf{H}_{m,n}$ is constructed from the channel impulse response $h_{m,n} = \{h_{m,n}^0, h_{m,n}^1, \dots, h_{m,n}^L, 0, \dots, 0\} \in \mathbb{C}^W$ between the n th user and m th receiver with L being the number of channel taps. The operator $Q(\cdot)$ serves as the sign operator on the real and imaginary parts of its argument. Throughout this section we assume the transmitted signal to be modulated with 16 QAM and $W \gg L$. This means that $\boldsymbol{\theta} \in \{-3, -1, +1, +3\}$. It is possible to recover such signal with EM method as [Plabst et al. \(2018\)](#). Here, the aim is to use the Block EM method to recover the modulated signal $\boldsymbol{\theta}$. To do so, we represent (SM9) as

$$\mathbf{R} = \mathcal{H}_{cir} \boldsymbol{\theta} + \boldsymbol{\alpha} \quad \text{and} \quad \mathbf{Y} = Q(\mathbf{R}) \quad (\text{SM10})$$

where $\mathbf{R} \in \mathbb{C}^{MW \times 1}$, $\boldsymbol{\alpha} \in \mathbb{C}^{MW \times 1}$, $\mathbf{Y} \in \mathbb{C}^{MW \times 1}$, $\boldsymbol{\theta} \in \mathbb{C}^{NW \times 1}$ are the vertically stacked \mathbf{R}_m , $\boldsymbol{\alpha}_m$, \mathbf{Y}_m , $\boldsymbol{\theta}_n$ over m and n respectively. $\mathcal{H}_{cir} \in \mathbb{C}^{MW \times NW}$ is the block-circulant Toeplitz matrix made from $\mathbf{H}_{m,n}$'s. According to (SM9) we need to form the conditional probability distribution function (PDF) $p(\mathbf{R}|\mathbf{Y}, \boldsymbol{\theta})$ (note that we are treating \mathbf{y} as the observed variable and \mathbf{R} as the latent variable). Using the Gaussian assumption, the joint PDF $p(\mathbf{R}, \mathbf{Y}|\boldsymbol{\theta})$ becomes

$$p(\mathbf{R}, \mathbf{Y}|\boldsymbol{\theta}) = \frac{\mathbb{I}_{Q(\mathbf{R})}(\mathbf{Y})}{\pi \sigma_\alpha^2} e^{-\frac{(\mathbf{R} - \mathcal{H}_{cir} \boldsymbol{\theta})^2}{\sigma_\alpha^2}} \quad (\text{SM11})$$

with $\mathbb{I}_{Q(\mathbf{R})}(\mathbf{Y})$ being one when $Q(\mathbf{R}) = \mathbf{Y}$ and $(\cdot)^2$ denoting element-wise square operator. Therefore $\hat{P}(\mathbf{R}) := \hat{P}(\mathbf{R}|\mathbf{Y}, \boldsymbol{\theta})$ in Block EM becomes

$$\hat{P}(\mathbf{R}|\mathbf{Y}, \boldsymbol{\theta}) = \frac{e^{-\frac{(\mathbf{R} - \mathcal{H}_{cir} \boldsymbol{\theta})^2}{\sigma_\alpha^2}}}{\int_{\mathcal{D}} e^{-\frac{(\mathbf{R} - \mathcal{H}_{cir} \boldsymbol{\theta})^2}{\sigma_\alpha^2}} d\mathbf{R}} \quad (\text{SM12})$$

where $\mathcal{D}_i = (-\infty, 0)$ when $Y < 0$ and $\mathcal{D}_i = (0, +\infty)$ when $\mathbf{y} > 0$. The second update of Block EM has the objective $\mathbb{E}_{\hat{P}(\mathbf{R}|\mathbf{y}, \boldsymbol{\theta})} [-\log(P(\mathbf{R}, \mathbf{y}|\theta_{i_k}))]$. From the structure of $-\log(P(\mathbf{R}, \mathbf{y}|\theta_{i_k}))$ we know

$$-\log(P(\mathbf{R}, \mathbf{y}|\theta_{i_k})) = \frac{\mathbf{R}^T \mathbf{R} - \mathbf{R}^T \mathcal{H}_{cir} \boldsymbol{\theta} - \boldsymbol{\theta}^T \mathcal{H}_{cir}^T \mathbf{R} + \boldsymbol{\theta}^T \mathcal{H}_{cir}^T \mathcal{H}_{cir} \boldsymbol{\theta}}{\sigma_\alpha^2}.$$

Note that for simplicity, we considered θ_{i_k} in $\boldsymbol{\theta}$ as the optimization variable. Therefore, the first term is independent from the optimization variable and we replace it with $\hat{\mathbf{R}}^T \hat{\mathbf{R}} = \mathbb{E}_{\hat{P}(\mathbf{R}|\mathbf{y}, \boldsymbol{\theta})} [\mathbf{R}]^T \mathbb{E}_{\hat{P}(\mathbf{R}|\mathbf{y}, \boldsymbol{\theta})} [\mathbf{R}]$. Thus,

$$\mathbb{E}_{\hat{P}(\mathbf{R}|\mathbf{y}, \boldsymbol{\theta})} [-\log(P(\mathbf{R}, \mathbf{y}|\theta_{i_k}))] = \frac{\hat{\mathbf{R}}^T \hat{\mathbf{R}} - \hat{\mathbf{R}}^T \mathcal{H}_{cir} \boldsymbol{\theta} - \boldsymbol{\theta}^T \mathcal{H}_{cir}^T \hat{\mathbf{R}} + \boldsymbol{\theta}^T \mathcal{H}_{cir}^T \mathcal{H}_{cir} \boldsymbol{\theta}}{\sigma_\alpha^2} = \frac{\|\hat{\mathbf{R}} - \mathcal{H}_{cir} \boldsymbol{\theta}\|^2}{\sigma_\alpha^2} \quad (\text{SM13})$$




Figure SM1: Comparison of Block EM and classical EM in terms of BER vs. E_b/N .

Therefore, in this problem we can evaluate the expected value of the latent variables \mathbf{R} through (SM12) and then minimize the convex objective (SM13) to recover the transmitted symbols θ . The coordinate-wise calculation of the expected value of the latent variables \mathbf{R} will take the form

$$\hat{\mathbf{R}}^{k+1} := \mathbb{E}_{\hat{p}(\mathbf{R}|y, \theta^k)}[\mathbf{R}] = \sigma_\alpha^2 \left(\frac{\mathbf{Y}_{Re} \phi(\mathbf{Y}_{Re} \mathbf{Z}_{Re}^k / \sigma_\alpha)}{\Phi(\mathbf{Y}_{Re} \mathbf{Z}_{Re}^k / \sigma_\alpha)} + j \frac{\mathbf{Y}_{Im} \phi(\mathbf{Y}_{Im} \mathbf{Z}_{Im}^k / \sigma_\alpha)}{\Phi(\mathbf{Y}_{Im} \mathbf{Z}_{Im}^k / \sigma_\alpha)} \right) + \mathcal{H}_{cir}(:, i_k) \theta_{i_k}^k \quad (\text{SM14})$$

where subscripts Re, Im specify the real and the imaginary parts, $j = \sqrt{-1}$, $\mathbf{Z}^k = \mathcal{H}_{cir}(:, i_k) \theta_{i_k}^k$, $\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\{-\frac{x^2}{2}\}$, and $\Phi(x) = \int_{-\infty}^x \phi(x') dx'$. With (SM14) at hand, we can easily update $\theta_{i_k}^{k+1}$ as

$$\theta_{i_k}^{k+1} = \underset{\substack{x \in \mathbb{C}^{\|i_k\|_0} \\ x \in \{-3, -1, +1, +3\}}}{\text{argmin}} \|\hat{\mathbf{R}}^{k+1} - \mathcal{H}_{cir}(:, i_k) x\|_2^2. \quad (\text{SM15})$$

Therefore, iterative updates of (SM14) and (SM15) will recover $\theta_{i_k}^{k+1}$. Note that to project to $\{-3, -1, +1, +3\}$ we need to relax the constraint set to $x \in [-3, 3]$, and assure the convergence of the updates before projection. This will lead to a small inner loop for each i_k .

E.2 Numerical Experiments

We conducted numerical simulations to verify the performance of the Block EM in this problem. Specifically, we are aiming to study the Bit Error Rate (BER) for various bit energy to noise spectral density (E_b/N) defines as

$$\frac{E_b}{N} = \frac{P \text{Tr}(\mathbb{E}_{\mathcal{H}_{cir}}[\mathcal{H}_{cir} \mathcal{H}_{cir}^H])}{MN \sigma_\alpha^2 B}, \quad (\text{SM16})$$

where P denotes the total transmitted power and B is number of bits in the signal constellation symbol. Consider a MIMO system employing 16-QAM with cyclic prefix omitted at the transmitter, $N = 6$ users, and $M = 64$ receiving antennas. The length of signals is $W = 32$. Channel was considered with complex Gaussian noise with mean zero and covariance matrix of $\sigma^2 \mathbf{I}$. The length of the fading is $L = 5$ and signals undergo 4 paths to reach the receiver. The number of inner loops for projection was 10 and the stopping criteria was $\|\zeta^u - \zeta^{u-1}\|_2 \leq \lambda \|\zeta^{u-1}\|_2$ with $\lambda = 10^{-3}$, $\zeta^u = \frac{1}{25} \sum_{p=25(u-1)+1}^{25u} \theta^p$. This choice of ζ^u is quite flexible and any short-term criteria to understand the behaviour of the update sequence can be applied. The maximum number of iterations for the Block EM algorithm was 10^4 and 10^5 for the EM algorithm. The algorithms were initialized at zero and the block size was 10. The results are depicted in Figure SM1 for 50 Monte-Carlo simulations.

As shown in Figure SM1, the Block EM performs slightly better than EM in this setting. Note that Block EM requires less computational resources than EM to maximize the likelihood function.