

# Global Convergence to Nash Equilibrium in Nonconvex General-Sum Games under the $n$ -Sided PL Condition

Yutong Chao

Technical University of Munich  
Munich, Germany  
yutong.chao@tum.de

Jalal Etesami

Technical University of Munich  
Munich Institute of Robotics and Machine Intelligence  
Munich, Germany  
j.etesami@tum.de

## ABSTRACT

We consider the problem of finding a Nash equilibrium (NE) in a general-sum game, where player  $i$ 's objective is  $f_i(x) = f_i(x_1, \dots, x_n)$ , with  $x_j \in \mathbb{R}^{d_j}$  denoting the strategy variables of player  $j$ . Our focus is on investigating first-order gradient-based algorithms and their variations, such as the block coordinate descent (BCD) algorithm, for tackling this problem. We introduce a set of conditions, called the  $n$ -sided PL condition, which extends the well-established gradient dominance condition a.k.a Polyak-Lojasiewicz (PL) condition and the concept of multi-convexity. This condition, satisfied by various classes of non-convex functions, allows us to analyze the convergence of various gradient descent (GD) algorithms. Moreover, our study delves into scenarios where the standard gradient descent methods fail to converge to NE. In such cases, we propose adapted variants of GD that converge towards NE and analyze their convergence rates. Finally, we evaluate the performance of the proposed algorithms through several experiments.

## KEYWORDS

Nonconvex optimization, General-sum games, Gradient descent

### ACM Reference Format:

Yutong Chao and Jalal Etesami. 2026. Global Convergence to Nash Equilibrium in Nonconvex General-Sum Games under the  $n$ -Sided PL Condition. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 24 pages. <https://doi.org/10.65109/>

## 1 INTRODUCTION

Optimization problems with nonconvex objectives appear in many applications from game theory to machine learning [9, 53, 74], such as training deep neural networks [20, 23, 50], or policy optimization in reinforcement learning [55, 66, 74]. In practice, the majority of these applications involve interacting agents whose optimal behavior emerges from solving nonconvex optimization problems. Despite the notable empirical advancements, there is a lack of understanding of the theoretical convergence guarantees of existing multi-agent methods. Even the fully cooperative settings remain largely underexplored and constitute one of the current frontiers in machine learning research.

From a game-theoretic perspective, we focus on a multi-agent general-sum game where the  $i$ -th player's objective function  $f_i$  captures the aggregate impact of agents' strategies, which is parameterized by  $\{x_i\}_{i=1}^n$ . Agents aim to optimize their individual objective function. The most common solution concept for the underlying optimization problem is Nash equilibrium (NE) [42]. Along this line of research, agents are said to be in a Nash equilibrium if no agent can gain by unilaterally deviating from its current strategy, assuming that all other agents maintain theirs. We are interested in finding Nash equilibrium  $x^* = (x_1^*, \dots, x_n^*) \in \mathbb{R}^d$  of the nonconvex functions  $\{f_i(x) : i \in [n]\}$ , i.e.,

$$f_i(x_i^*; x_{-i}^*) \leq f_i(y_i; x_{-i}^*), \quad \forall i \in [n], \forall y_i \in \mathbb{R}^{d_i}, \quad (1)$$

where  $x_{-i}^*$  denotes all blocks except  $i$ -th block and  $d = \sum_i d_i$ . Note that in the special case where  $f_1(x) = \dots = f_n(x)$ , the setting reduces to a potential game in which all players jointly aim to minimize a common objective function. It is known that finding the NE is PPAD-complete in general [19]. Thus, there is a need for additional structural assumptions to reduce the computational complexity of finding NE.

For example, in extensive-form games, which extend normal-form interactions to multi-stage decision processes, equilibria can often be analyzed using backward induction or dynamic programming-based techniques [11, 30]. In linear-quadratic (LQ) games, where player payoffs are quadratic functions governed by linear dynamical systems, the algebraic structure permits explicit characterizations of NE via coupled Riccati equations [22, 25, 51, 70]. However, these guarantees are inherently tied to the particular game models under consideration, and therefore do not generalize to the broader class of general-sum games with fewer structural assumptions and nonconvex payoff functions, where the existence and computability of equilibrium remain much less understood.

To tackle a nonconvex optimization problem, a straightforward approach is to introduce additional structural assumptions to achieve convergence guarantees. Within this scope, various relaxations of convexity have been proposed, for example, weak strong convexity [37], restricted secant inequality [69], error bound [10], quadratic growth [16], etc. Recently, there has been a surge of interest in analyzing nonconvex functions with block structures. Multiple assumptions have been analyzed, which are correlated to each block (coordinate) when other blocks are fixed, for example, PL-strongly-concave [24], nonconvex-PL [52], PL-PL [13, 18, 64], and multi-convex [54, 58, 59, 62]. For instance, the multi-convexity assumes the convexity of the function with respect to each block when the remaining blocks are fixed. On the other hand, the other aforementioned conditions are tailored for objective functions comprising



This work is licensed under a Creative Commons Attribution International 4.0 License.

only two blocks. They are particularly defined for min-max type optimizations rather than general minimization problems.

The nonconvex optimization realm has seen a growing interest in the gradient dominance condition a.k.a. Polyak-Łojasiewicz (PL) condition. For instance, in analyzing linear quadratic games [22], matrix decomposition [33], robust phase retrieval [57] and training neural networks [12, 26, 36]. This is due to its ability to enable sharp convergence analysis of both deterministic GD and stochastic GD algorithms while being satisfied by a wide range of nonconvex functions. More formally, a function  $f$  satisfies the PL condition if there exists a constant  $\mu > 0$  such that

$$\|\nabla f(x)\|^2 \geq 2\mu(f(x) - \min_{y \in \mathbb{R}^d} f(y)), \quad \forall x \in \mathbb{R}^d. \quad (2)$$

This was first introduced by Łojasiewicz [38], Polyak [48], who analyzed the convergence of the GD algorithm under the PL condition and showed its linear convergence to the global minimum. This condition can be perceived as a relaxation of strong convexity and as discussed in [29], it is closely related to conditions such as weak-strong convexity [43], restricted secant inequality [69] and error bound [39].

As mentioned, the PL condition has been extended and applied to optimization problems with multiple coordinates. This extension is analogous to generalizing the concept of convexity (concavity) to convex-concavity. For instance, the two-sided PL condition was introduced by Yang et al. [64] for analyzing deterministic and stochastic alternating gradient descent ascent (AGDA) in *min-max games*. They showed that under this condition, AGDA converges at linear rate to the unique NE. It is noteworthy that most literature requires convexity or PL condition to establish the last-iterate convergence rate to the NE [28, 53, 56]. A considerable relaxation is that the function satisfies the convexity or PL condition when all variables except one are fixed.

*Can similar results be achieved by extending the two-sided PL condition to accommodate problems in the form of (1), where the objectives comprise  $n$  coordinates? And is there an algorithm to guarantee convergence to a NE at a linear rate in such problems?*

Motivated by the above question, we introduce the notion of  $n$ -sided PL<sup>1</sup> condition (Definition 2.4), which is an extension to the PL condition, and show that it holds in several well-known nonconvex problems such as  $n$ -player linear quadratic game, linear residual network, etc. It is noteworthy that, unlike the two-sided PL condition, which guarantees the uniqueness of the NE in min-max optimization [13, 64], functions satisfying the  $n$ -sided PL (even 2-side PL) condition may have multiple NE points (see Section 2.1 for examples). However, as we will discuss, the set of stationary points for such functions is equivalent to their NE points. Moreover, unlike the two-sided PL condition, which ensures linear convergence of the AGDA algorithm to the NE, the BCD algorithm exhibits varying convergence rates for different functions, all satisfying the  $n$ -sided PL condition. Similar behavior has been observed with multi-convex functions [59, 63]. Therefore, additional local or global conditions

<sup>1</sup>We should emphasize that 2-sided PL and two-sided PL are slightly different conditions as the former is suitable for  $\min_{x,y} f(x,y)$  while the latter is for  $\min_x \max_y f(x,y)$ .

are required to characterize the convergence rate under the  $n$ -sided PL condition.

In this work, we study the convergence of first-order GD-based algorithms such as the BCD, and propose different variants of BCD that are more suitable for the class of nonconvex functions satisfying  $n$ -sided PL condition. We show the convergence to NE and introduce additional local conditions under which linear convergence can be guaranteed.

## 1.1 Related Work

**Block Coordinate Descent and its variants:** Block coordinate descent (BCD) is an efficient and reliable gradient-based method for optimization problems, which has been used extensively in machine learning [3, 41, 44, 68, 73]. Numerous existing works have studied the convergence of BCD and its variants applied to the special case of potential games. Most of them require the assumptions of convexity, PL condition, and their extensions [5, 14, 15, 27, 35]. For instance, Xu and Yin [62, 63] studied the convergence of BCD for the regularized block multiconvex optimization. They established the last iterate convergence under Kurdyka-Łojasiewicz, which might not hold for many functions globally. The authors in [35] considered the generalized Minty variational problem and applied cyclic coordinate dual averaging with extrapolation to find its solution. Their algorithm is independent of the dimension of the number of coordinates. However, their results rely on assuming the monotonicity of the operators, which is rarely satisfied in practice. Cai et al. [8] considered composite nonconvex optimization and applied cyclic block coordinate descent with PAGE-type variance reduced method. They proved linear and non-asymptotic convergence when the PL condition holds, which is not valid for functions with multiple local minima. However, none of these work consider the general-sum setting.

**PL condition in optimization:** The PL condition was originally proposed to relax the strong convexity in the minimization problem sufficient for achieving the global convergence for first-order methods. For example, Karimi et al. [29] showed that the standard GD algorithm admits a linear convergence to minimize an  $L$ -smooth and  $\mu$ -PL function. To be specific, in order to find an  $\epsilon$ -approximate optimal solution  $\hat{x}$  such that  $f(\hat{x}) - f^* \leq \epsilon$ , GD requires the computational complexity of the order  $O(\frac{L}{\mu} \log \frac{1}{\epsilon})$ . Besides this, different proposed methods, such as the heavy ball method and its accelerated version, have been analyzed [17, 60]. The authors in [67] proved the optimality of GD by showing that any first-order method requires at least  $\Omega(\frac{L}{\mu} \log \frac{1}{\epsilon})$  gradient costs to find an  $\epsilon$ -approximation of the optimal solution. Furthermore, many studies focus on the complexity when the objective function has a finite-sum structure, i.e.,  $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$ , e.g., [4, 32, 34, 49, 61]. It is important to emphasize that the aforementioned studies do not generalize to optimization problems of the form given in (1).

In addition to the minimization problem, extensions of the PL condition, such as two-sided conditions, have been proposed to provide convergence guarantees to saddle points for gradient-based algorithms when addressing minimax optimization problems. For example, the two-sided PL holds when both functions  $h_y(x) := f(x,y)$  and  $h_x(y) := -f(x,y)$  satisfy the PL condition [13, 64], or one-sided PL condition holds when only  $h_y(x)$  satisfies the PL

condition [24, 65]. Various types of first-order methods have been applied to such problems, for example, SPIDER-GDA [13], AGDA [64], Multi-step GDA [46, 52]. For additional information on the sample complexity of the methods mentioned earlier and their comparisons, see [13] and [4].

## 2 $n$ -SIDED PL CONDITION

**Notations:** Throughout this work, we use  $\|\cdot\|$  to denote the Euclidean norm and the lowercase letters to denote a column vector. In particular, we use  $x_{-i}$  to denote the vector  $x$  without its  $i$ -th block, where  $i \in [n] := \{1, \dots, n\}$ . The partial derivative of general function  $f(x)$  with respect to the variables in its  $i$ -th block is denoted as  $\nabla_i f(x) := \frac{\partial}{\partial x_i} f(x_i, x_{-i})$  and the full gradient is denoted as  $\nabla f(x)$  that is  $(\nabla_1 f(x), \dots, \nabla_n f(x))$ .

### 2.1 Definitions and Assumptions

In this paper, we assume all the functions  $f_i(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  belong to  $C^1$ , i.e., they are continuously differentiable. Furthermore, we assume they have Lipschitz gradient.

**Assumption 2.1** (Smoothness). We assume the  $L$ -Lipschitz continuity of the derivative  $\nabla f_i(x)$ ,

$$\|\nabla f_i(x) - \nabla f_i(y)\| \leq L_i \|x - y\|, \quad \forall i \in [n], \forall x, y,$$

where  $L_i > 0$  are constants and  $L := \max_i L_i$ . In this case,  $f_i(x)$  is also called  $L$ -smooth.

We now define two notions of optimality; Nash equilibrium (NE) and Stationary point.

**Definition 2.2** (Nash Equilibrium (NE)). Point  $x^* = (x_1^*, \dots, x_n^*)$  is called a Nash Equilibrium of functions  $\{f_i(x)\}$  if

$$f_i(x_i^*, x_{-i}^*) \leq f_i(x_i, x_{-i}^*), \quad \forall i \in [n], \forall x_i \in \mathbb{R}^{d_i}.$$

The set of all NE points of functions  $\{f_i(x)\}$  is denoted by  $\mathcal{N}(f_1, \dots, f_n)$ .

The other notion, the stationary point, is related to the first-order condition of optimality and also relevant for studying gradient-based algorithms.


**Definition 2.3** ( $\varepsilon$ -Partial Stationary point). Point  $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_n)$  is called an  $\varepsilon$ -stationary point of functions  $\{f_i(x)\}$  if

$$\|\nabla_i f_i(\tilde{x})\| \leq \varepsilon, \quad \forall i \in [n].$$

When  $\varepsilon = 0$ ,  $\tilde{x}$  is called a partial stationary point. We denote the set of all  $\varepsilon$ -stationary points by  $\mathcal{S}_\varepsilon(f_1, \dots, f_n)$ , respectively.

It is important to note that the notion of partial stationarity in general-sum games differs from the standard concept of stationarity commonly used in potential games. In the former, the partial derivative of each player's objective function with respect to its own variable is zero, whereas in the latter, the entire gradient of the potential function vanishes. Thus, it is possible that, at a partial stationary point  $x \in \mathcal{S}_\varepsilon(f_1, \dots, f_n)$ , the full gradient  $\nabla f_j(x) \neq 0$  for some players  $j \in [n]$ .

For general nonconvex minimization problems, the above two notions are not necessarily equivalent, i.e., a stationary point may not be a NE. Nevertheless, for the remainder of this work, we assume that the objective functions have at least one NE, i.e.,  $\mathcal{N}(f_1, \dots, f_n) \neq \emptyset$ . We also assume that  $\arg \min_{x_i \in \mathbb{R}^{d_i}} f_i(x_i, x_{-i})$  is non-empty for



**Figure 1:** Left is function  $f^{(1)}(x_1, x_2)$  and right is  $f^{(2)}(x_1, x_2)$ .

any  $i \in [n]$  and  $x_{-i}$ , i.e., there exists a best response to every  $x_{-i}$ . Below, we formally introduce the  $n$ -sided PL condition for the functions  $\{f_i(x)\}$ .

**Definition 2.4** ( $n$ -sided PL Condition). We say the set of functions  $\{f_i(x) = f_i(x_1, \dots, x_n) : i \in [n]\}$  satisfy  $n$ -sided  $\mu$ -PL condition if there exist a set of positive constants  $\mu_i > 0$  such that for all  $x \in \mathbb{R}^d, i \in [n]$ ,

$$\|\nabla_i f_i(x_i, x_{-i})\|^2 \geq 2\mu_i (f_i(x_i, x_{-i}) - f_{i,x_{-i}}^*), \quad (3)$$

where  $f_{i,x_{-i}}^* := \min_{y_i \in \mathbb{R}^{d_i}} f_i(y_i, x_{-i})$ . We also denote  $\mu := \min_i \mu_i$ .

It is worth noting that the  $n$ -sided PL condition does not imply convexity or the gradient dominance (PL) condition. It is an extension to the PL condition. In the special case of a potential game where  $f_i(x) = f(x)$  for all  $i \in [n]$  and  $f(x)$  is independent of  $x_{-i}$ , then  $f$  satisfies the PL condition. Moreover, it is considerably weaker than multi-strong convexity. Nevertheless, under the  $n$ -sided PL condition, the set of stationary points and the NE set are equivalent.

**Lemma 2.5.** If functions  $\{f_i(x_1, \dots, x_n)\}$  satisfy the  $n$ -sided PL condition, then  $\mathcal{S}_0(f_1, \dots, f_n) = \mathcal{N}(f_1, \dots, f_n)$ .


All proofs are in the Appendix D. It is noteworthy that, unlike the two-sided PL condition, even in the special case of a potential game where all players have the same objective function, under the  $n$ -sided PL condition presented in Definition 2.4, it is no longer possible to ensure that the NE is unique. In fact, there could be multiple NEs with different function values. For example, consider the following two potential games with objective functions illustrated in Figure 1,

$$\begin{aligned} f^{(1)}(x) &:= f_1^{(1)}(x) = f_2^{(1)}(x) = (x_1 - 1)^2 (x_2 + 1)^2 + (x_1 + 1)^2 (x_2 - 1)^2, \\ f^{(2)}(x) &:= f_1^{(2)}(x) = f_2^{(2)}(x) = f^{(1)}(x) + \exp(-(x_2 - 1)^2). \end{aligned}$$

As shown in Appendix C, both functions  $f^{(1)}$  and  $f^{(2)}$  are 2-sided PL, but their sets of NE and the set of minima points are not equivalent. In particular, both functions have three NE points, while  $f^{(1)}$  has two global minima and a saddle point, and  $f^{(2)}$  has a local minima, a global minima, and a saddle point.

## 3 ALGORITHMS AND CONVERGENCE ANALYSIS

In this section, we first study the BCD algorithm for finding a stationary point of (1) under the  $n$ -sided PL condition. Afterward, we propose different variants of BCD algorithms that can provably achieve better convergence rates.



**Figure 2: The result of R-BCD algorithm applied to potential setting with functions  $f^{(3)}(x_1, x_2)$  and  $f^{(4)}(x_1, x_2)$ . The y-axis is in log scale, thus the R-BCD demonstrates linear convergence for  $f^{(4)}$ .**

The BCD algorithm is a coordinate-wise approach that iteratively improves its current estimate by updating a selected block coordinate using the first-order partial derivatives until it converges. Depending on how the coordinates are chosen, various types of BCD algorithms can be devised. For example, coordinates can be selected uniformly at random for *random BCD* (R-BCD), presented in Algorithm 1 with learning rates  $\alpha$ , or in a deterministic cyclic sequence for *cyclic BCD*, presented in Algorithm 5 in the Appendix. Moreover, to update the  $i$ -th block at the  $t$ -th iteration, it employs  $\nabla_i f_i(x^{t-1})$ , where  $x^{t-1}$  denotes the latest estimated point. The next result shows that when the iterates of the BCD are bounded, the output converges to the NE set.

**Algorithm 1** Random Block Coordinate Descent (R-BCD)

```


Input: initial point  $x^0 = (x_1^0, \dots, x_n^0)$ , learning rates  $\alpha$ ,  $t = 1$ .
while not converges do
    Select  $i \in [n]$  at random
     $x_i^t = x_i^{t-1} - \alpha \nabla_i f_i(x^{t-1})$ 
     $x_j^t = x_j^{t-1}$ , for all  $j \neq i$ 
end while
    
```

Note that AGDA can be viewed as a variant of BCD applied to a two-player potential game where the first coordinate is updated via descent and the second coordinate via ascent. Unlike the two-sided PL condition that leads to linear convergence of AGDA to the min-max, the  $n$ -sided PL condition does not necessarily lead to any specific convergence rate of the BCD. To illustrate this phenomenon, we consider two settings: potential and general-sum. In each setting, we examine two different games and analyze the convergence behavior of R-BCD.

*Potential setting:* We consider the following two potential games, i.e.,  $f_1 = f_2 = f$ , in which all players share the same objective function; both games satisfy the 2-sided PL condition.

$$f^{(3)}(x_1, x_2) = \begin{cases} (x_1 + x_2)^2 + \exp\left(-\frac{1}{(x_1 - x_2)^2}\right), & x_1 \neq x_2, \\ (x_1 + x_2)^2, & \text{o.t.}, \end{cases}$$

$$f^{(4)}(x_1, x_2) = (x_1 + x_2)^2.$$



**Figure 3: The result of R-BCD algorithm applied to functions  $\{f_1^{(5)}, f_2^{(5)}\}$  and  $\{f_1^{(6)}, f_2^{(6)}\}$ . The y-axis is in log scale, thus the BCD shows linear convergence in the second game.**

We applied the BCD algorithm to both these functions with constant learning rates to find their NE points with random initializations. As it is illustrated in Figure 2, the R-BCD converges linearly for the function  $f^{(4)}$  while it converges sub-linearly for  $f^{(3)}$ .

*General-sum setting:* Here, we consider the following two-player games, both of which satisfy the 2-sided PL condition. The objectives of the first game are

$$f_1^{(5)}(x_1, x_2) = \begin{cases} (x_1 + x_2)^2 + \exp\left(-\frac{1}{(x_1 - x_2)^2}\right), & x_1 \neq x_2, \\ (x_1 + x_2)^2, & \text{o.t.}, \end{cases}$$

$$f_2^{(5)}(x_1, x_2) = (x_1 + x_2)^2,$$

and the objectives of the second game are

$$f_1^{(6)}(x_1, x_2) = x_1^2 + x_2^2, \quad f_2^{(6)}(x_1, x_2) = (x_1 + x_2)^2.$$

Figure 3 illustrates the convergence rate of R-BCD applied to the above functions, showing clearly that the convergence rates differ between them. Note that we plotted the sum of the players' objectives against the number of iterations, since in both examples  $(0, 0)$  corresponds to the global Nash equilibrium, where all players' objectives are zero.

These examples illustrate that, even under the  $n$ -sided PL condition and smoothness, characterizing the convergence rate of the BCD algorithm in general-sum games may not be feasible and further assumptions on the function class are required. Next, we study one such assumption that holds for a large class of non-convex functions and could characterize the convergence rate of R-BCD.

**3.1 Convergence of random BCD under an additional assumption**

To introduce our additional assumption, we need to define the sum of the objective functions and their best responses, denoted by  $F(x)$  and  $G_F(x)$ , respectively, which plays a central role in analyzing the convergence of coordinate-wise algorithms.

$$F(x) := \sum_{i=1}^n f_i(x), \quad F_{-i}(x) := F(x) - f_i(x), \quad (4)$$

$$G_F(x) := \sum_{i=1}^n f_i(x_i^*(x), x_{-i}), \quad (5)$$

where  $x_i^*(x)$  denotes the best response to  $x_{-i}$ , i.e.,

$$x_i^*(x) \in \arg \min_{y_i} \{ \|y_i - x_i\| : f_i(y_i, x_{-i}) \leq f_i(z_i, x_{-i}), \forall z_i \in \mathbb{R}^{d_i} \}.$$

Next result shows that the best response function is smooth when the objective functions are both smooth and  $n$ -sided PL.

**Lemma 3.1.** *If  $\{f_i(x)\}$  satisfy  $n$ -sided  $\mu$ -PL and Assumption 2.1, then  $G_F(x)$  is  $nL'$ -smooth, where  $L' := L + \frac{L^2}{\mu}$ .*

It is straightforward to see that  $F(x) - G_F(x) \geq 0$  for all  $x$ . Moreover, if  $x^* \in \mathcal{N}(f_1, \dots, f_n)$ , then  $F(x^*) = G_F(x^*)$ . Conversely, if  $F(x^*) - G_F(x^*) = 0$ , then  $f_i(x^*) = \min_{x_i} f_i(x_i, x_{-i}^*)$ ,  $\forall i$ , which implies that  $x^*$  belongs to  $\mathcal{N}(f_1, \dots, f_n)$ . The next lemma summarizes these results.

**Lemma 3.2.**  *$x^*$  is a NE if  $F(x^*) - G_F(x^*) = 0$ . Moreover,  $x \in \mathcal{S}_\epsilon$  if  $F(x) - G_F(x) \leq \sqrt{2L\epsilon}$ .*

Thus,  $F(x) - G_F(x)$  serves as an indicator of whether  $x$  is a NE. Using this indicator, we show that the R-BCD algorithm achieves a linear convergence rate under an additional assumption.

**Theorem 3.3.** *Suppose  $\{f_i(x)\}$  are  $n$ -sided  $\mu$ -PL functions satisfying Assumption 2.1 such that for all  $x \in \mathbb{R}^d$*

$$\sum_{i=1}^n \langle \nabla_i(G_F(x) - F_{-i}(x)), \nabla_i f_i(x) \rangle \leq \kappa \sum_{i=1}^n \|\nabla_i f_i(x)\|^2, \quad (6)$$

where  $\kappa < 1$ , then random BCD with  $\alpha \leq \frac{1-\kappa}{n(L+L')}$  achieves linear convergence rate, i.e.,

$$\mathbb{E}[F(x^{t+1}) - G_F(x^{t+1})] \leq \left(1 - \frac{\mu\alpha(1-\kappa)}{2}\right) \mathbb{E}[F(x^t) - G_F(x^t)].$$

The expectation is taken over the randomness inherent in the procedure for selecting coordinates.

It is straightforward to verify that  $f^{(4)}$  in Figure 2 (for all  $x \in \mathbb{R}^2$ ),  $f^{(6)}$  in Figure 3 (for all  $x \in \mathbb{R}^2$ ), and  $f^{(0)}$  in Figure 1 (for all points in  $\{(x_1, x_2) : |x_1| > 0.75, |x_2| > 0.75\}$ ) satisfy the condition in (6). Notably, these sets contain all the NE points, and the R-BCD algorithm exhibits a linear convergence rate when applied to these functions. The above result further highlights the fast convergence of R-BCD for finding NEs of these functions.

To demonstrate that the additional assumption in (6) is not overly restrictive, we apply the Cauchy–Schwarz inequality together with the result of Lemma D.1 (see Appendix) to obtain the following bound, valid for all smooth functions  $\{f_i\}$  satisfying the  $n$ -sided PL condition:

$$\sum_{i=1}^n \langle \nabla_i G_F(x) - \nabla_i F_{-i}(x), \nabla_i f_i(x) \rangle \leq (\sqrt{3n} \frac{L}{\mu}) \sum_{i=1}^n \|\nabla_i f_i(x)\|^2,$$

where  $\sqrt{3n} \frac{L}{\mu} \geq 1$  (see Lemma A.3 for a proof that  $\mu \leq L$ ). This implies that for any collection of smooth functions satisfying the  $n$ -sided PL condition, there exists a constant  $c$  such that the left-hand side of (6) is bounded by  $c \sum_i \|\nabla_i f_i(x)\|^2$ . According to Theorem 3.3, R-BCD exhibits linear convergence to a NE when the corresponding constant factor is less than one. As discussed in the previous section, however, there also exist smooth,  $n$ -sided PL functions (e.g., those

depicted in Figures 2 and 3) whose constant factors exceed one but remain bounded by  $\sqrt{3n} \frac{L}{\mu}$ .

### 3.2 Adaptive random BCD with the knowledge of exact best responses

As discussed earlier, BCD algorithms may exhibit different convergence rates in general. In the remainder of this section, we therefore focus on developing variants of the randomized BCD algorithm that can achieve convergence rates close to linear. To this end, we propose algorithms that, at each iteration  $t$ , update the variables based on the best responses  $\{x_i^*(x^t)\}$ .

We initially propose an algorithm that presumes access to the exact values of the best responses at each iteration. Subsequently, we refine this assumption by integrating a sub-routine into the proposed algorithm capable of approximating the best responses. To present our results, we need the following definition.

**Definition 3.4** ( $(\theta, \nu)$ -PL condition). A function  $f$  with  $\min_x f(x) = 0$  satisfies  $(\theta, \nu)$ -PL condition if and only if there exists  $\theta \in [1, 2)$  and  $\nu > 0$  such that  $\|\nabla f(x)\|^\theta \geq (2\nu)^{\theta/2} f(x)$ .

It has been proved by Lojasiewicz [38] that for any  $C^1$  analytic function, there exists a neighborhood  $U$  around the minimizer where  $(\theta, \nu)$ -PL condition is satisfied for some  $\theta$  and  $\nu$ .

---

#### Algorithm 2 Ideal Adaptive random BCD (IA-RBCD)

---

**Input:** initial point  $x^0 = (x_1^0, \dots, x_n^0)$ ,  $T$ , step size  $\alpha$ ,  $0 \leq \gamma < 1$  and  $C > 0$

**for**  $t = 0$  **to**  $T - 1$  **do**

$$A = \sum_{i=1}^n \langle \nabla_i(G_F(x^t) - F_{-i}(x^t)), \nabla_i f_i(x^t) \rangle$$

$$B = \sum_{i=1}^n \|\nabla_i(G_F(x^t) - F_{-i}(x^t))\|^2$$

$$D = \sum_{i=1}^n \|\nabla_i f_i(x^t)\|^2$$

sample  $i^t$  uniformly at random from  $\{1, 2, \dots, n\}$

**if**  $A \leq \gamma D$  **then**

$$k^t = 0 \quad \text{:Case 1:}$$

**else if**  $(B - A)^2 \geq CA^2$  **then**

$$k^t = -2 + \frac{A}{B} \quad \text{:Case 2:}$$

**else**

$$k^t = -1 \quad \text{:Case 3:}$$

**end if**

$$x_{i^t}^{t+1} = x_{i^t}^t - \alpha \left( \nabla_{i^t} f_{i^t}(x^t) + k^t \nabla_{i^t} (G_F(x^t) - F_{-i^t}(x^t)) \right),$$

$$x_i^{t+1} = x_i^t \quad \text{for all } i \neq i^t.$$

**end for**

---

Algorithm 2 presents the steps of our modified version of the random BCD. In this algorithm, instead of updating along the direction of  $-\nabla_{i^t} f_{i^t}(x)$ , where  $i^t$  denotes the chosen coordinate at iteration  $t$ , a linear combination of  $\nabla_{i^t} f_{i^t}(x)$  and  $\nabla_{i^t} (G_F(x) - F_{-i^t}(x))$  is used to refine the updating directions. The coefficient of this linear combination,  $k^t$ , is adaptively selected based on the current estimated point. It is important to mention that  $\nabla G_F(x)$  can be computed using the gradients of  $\{f_i\}$  at the best responses,

$$\nabla G_F(x) = \sum_{i=1}^n \nabla f_i(x_i^*(x), x_{-i}). \quad (7)$$

The following theorem establishes the convergence guarantee of the algorithm presented in 2.

**Theorem 3.5.** For  $n$ -sided  $\mu$ -PL functions  $\{f_i(x)\}$  satisfying Assumption 2.1, by applying Algorithm 2,

- in Case 1 with  $\alpha \leq \frac{1-\gamma}{n(L+L')}$ , we have

$$\mathbb{E}[F(x^{t+1}) - G_F(x^{t+1}) | x^t] \leq \left(1 - \frac{\mu\alpha(1-\gamma)}{2}\right) (F(x^t) - G_F(x^t)),$$

- in Case 2 with  $\alpha \leq \min\{\frac{1}{2n(L+L')}, \frac{C}{2n(L+L')}\}$ , we have

$$\mathbb{E}[F(x^{t+1}) - G_F(x^{t+1}) | x^t] \leq \left(1 - \frac{(L+L')\mu\alpha^2}{2}\right) (F(x^t) - G_F(x^t)),$$

- in Case 3 with  $\alpha \leq \frac{1}{n(L+L')}$ ,  $\mathbb{E}[F(x^{t+1}) - G_F(x^{t+1})]$  is non-increasing. Furthermore, if  $F - G_F$  satisfies  $(\theta, \nu)$ -PL condition and case 3 are satisfied from iterates  $t$  to  $t+k$ , we have

$$\mathbb{E}[F(x^{t+k}) - G_F(x^{t+k}) | x^t] = O\left(\frac{F(x^t) - G_F(x^t)}{k^{\frac{\theta}{2-\theta}}}\right).$$

The proof in the Appendix provides the exact constant factors. It is noteworthy to highlight that BCD requires the objective function to be lower bounded to converge to the NE [62] and almost surely avoids strict saddle points [31]. However, the above result shows that IA-RBCD converges to the NE, irrespective of this assumption. Additionally, based on this result, IA-RBCD demonstrates a linear rate for two out of three cases and a sub-linear rate for the third case. Next, we show that even if the third case occurs frequently, but within a limited range, linear convergence is still guaranteed by IA-RBCD.

**Theorem 3.6.** Under the assumptions of Theorem 3.5, by applying Algorithm 2 for  $t$  iterations and denoting the number of iterations it visits Case 3 by  $b(t)$ , if  $B := \limsup_{t \rightarrow +\infty} b(t)/t < 1$ , then there exists  $0 < c_B < 1$  and  $\hat{T}_B$  such that for all  $t \geq \hat{T}_B$ ,

$$\mathbb{E}[F(x^t) - G_F(x^t)] \leq (c_B)^t (F(x^0) - G_F(x^0)).$$

In the following theorem, we demonstrate that as  $\gamma \rightarrow 1$  and  $C \rightarrow 0$ , the measure of non-NE points that fail to satisfy case 1 and case 2 converges to zero. Furthermore, when the iterates fall into case 3,  $F - G_F$  is small at those iterates which indicates that a good approximation of a NE is already achieved. As a consequence, by choosing hyperparameters  $\gamma$  and  $C$  close to one and zero, respectively, we can achieve close to linear convergence rate.

**Theorem 3.7.** Under the assumption of Theorem 3.5 and when  $F - G_F$  satisfies  $(\theta, \nu)$ -PL condition, let  $S(\gamma, C)$  be the set of non-NE points that do not satisfy case 1 and case 2, then

$$\lim_{\gamma \rightarrow 1, C \rightarrow 0} |S(\gamma, C)| = 0,$$

where  $|S(\gamma, C)|$  denotes the measure of the set  $S(\gamma, C)$ . Moreover, if  $S(\gamma, C)$  is non-empty, then

$$\lim_{\gamma \rightarrow 1, C \rightarrow 0} \max_{x \in S(\gamma, C)} F(x) - G_F(x) = 0.$$

### 3.3 Adaptive random BCD without the knowledge of exact best responses

Evaluating  $G_F$  at a given point requires knowledge of the best responses at that point. Often, these best responses are not known a priori, and they have to be computed at each iteration. Fortunately, since in our study,  $\{f_i\}$  satisfy the  $n$ -sided PL condition, the best responses can be efficiently approximated by applying the GD algorithm with the partial gradients as a sub-routine. Algorithm 3 outlines the steps of the sub-routine, which outputs an approximation of the best responses for a given step size  $\beta$  after  $T'$  iterations. These approximated best responses are then used by Algorithm 4 to find a NE point.

---

#### Algorithm 3 Approximating Best Responses (ABR)

---

**Input:** Point  $x = (x_1, \dots, x_n)$ ,  $T'$ , step size  $\beta$

**for**  $j = 1, \dots, n$  **do**

$$y_j^0 = x_j$$

**for**  $\tau = 0, \dots, T' - 1$  **do**

$$y_j^{\tau+1} = y_j^\tau - \beta \nabla_j f_j(y_j^\tau, x_{-j})$$

**end for**

**end for**

**Output:**  $\nabla \tilde{G}_F(x) = \sum_{i=1}^n \nabla f_i(y_i^{T'}, x_{-i})$

---

**Lemma 3.8.** For  $n$ -sided  $\mu$ -PL functions  $\{f_i(x)\}$  satisfying Assumption 2.1, by implementing Algorithm 3, with  $\beta \leq \frac{1}{L}$  and  $T' \geq \log\left(\frac{nL^2}{\mu^2\delta}\right) / \log\left(\frac{1}{1-\mu\beta}\right)$ , for any  $\delta > 0$ , we have

$$\|\nabla G_F(x) - \nabla \tilde{G}_F(x)\|^2 \leq \delta \sum_{i=1}^n \|\nabla_i f_i(x)\|^2.$$

Interestingly as we showed in the next result, the number of steps for approximating  $G_F$ ,  $T'$ , only depends on the function parameters, and it is independent of the final precision of  $F - G_F$ . The exact form of the step sizes are presented in the Appendix D.9.

**Theorem 3.9.** For  $n$ -sided  $\mu$ -PL functions  $\{f_i(x)\}$  satisfying Assumption 2.1, by implementing Algorithm 4 with  $\alpha$  small enough,  $\beta \leq \frac{1}{L}$  and  $T' \geq C' \log\left(\frac{1}{\alpha}\right) / \log\left(\frac{1}{1-\mu\beta}\right)$  where  $C'$  only depends on the function parameters,

- in Case 1, we have

$$\mathbb{E}[F(x^{t+1}) - G_F(x^{t+1}) | x^t] \leq \left(1 - \frac{\mu\alpha(1-\gamma)}{2}\right) (F(x^t) - G_F(x^t)),$$

- in Case 2, we have

$$\mathbb{E}[F(x^{t+1}) - G_F(x^{t+1}) | x^t] \leq \left(1 - \frac{(L+L')\mu\alpha^2}{4}\right) (F(x^t) - G_F(x^t)),$$

- in Case 3,  $\mathbb{E}[F(x^t) - G_F(x^t)]$  is non-increasing. Furthermore, if  $F - G_F$  satisfies  $(\theta, \nu)$ -PL condition and case 3 occurs from iterates  $t$  to  $t+k$ , then

$$\mathbb{E}[F(x^{t+k}) - G_F(x^{t+k}) | x^t] = O\left(\frac{F(x^t) - G_F(x^t)}{k^{\frac{\theta}{2-\theta}}}\right).$$


---

**Algorithm 4** Adaptive random BCD (A-RBCD)

---

**Input:** initial point  $x^0 = (x_1^0, \dots, x_n^0)$ ,  $T, T'$ , step sizes  $\alpha, \beta$ ,  $0 < \gamma < 1$  and  $C > 0$

**for**  $t = 0$  **to**  $T - 1$  **do**

  sample  $i^t$  uniformly from  $\{1, 2, \dots, n\}$

$\nabla \tilde{G}_F(x^t) = \text{ABR}(x^t, \beta, T^t)$  :Algorithm 3

$\tilde{A} = \sum_{i=1}^n \langle \nabla_i(\tilde{G}_F(x^t) - F_{-i}(x^t)), \nabla_i f_i(x^t) \rangle$

$\tilde{B} = \sum_{i=1}^n \|\nabla_i(\tilde{G}_F(x^t) - F_{-i}(x^t))\|^2$

$\tilde{D} = \sum_{i=1}^n \|\nabla_i f_i(x^t)\|^2$

**if**  $\tilde{A} \leq \gamma \tilde{D}$  **then**

$\tilde{k}^t = 0$  :Case 1:

**else if**  $(\tilde{B} - \tilde{A})^2 \geq 2C\tilde{A}^2$  **then**

$\tilde{k}^t = -2 + \frac{\tilde{A}}{\tilde{B}}$  :Case 2:

**else**

$\tilde{k}^t = -1$  :Case 3:


**end if**

$x_{i^t}^{t+1} = x_{i^t}^t - \alpha(\nabla_{i^t} f_{i^t}(x^t) + \tilde{k}^t \nabla_{i^t}(\tilde{G}_F(x^t) - F_{-i^t}(x^t)))$

$x_i^{t+1} = x_i^t$ , if  $i \neq i^t$

**end for**

---



**Figure 4: Performances of A-RBCD, BM1, and BM2 applied to functions  $f_1$  and  $f_2$ .**

## 4 APPLICATIONS

In this section and in Appendix E, we evaluate the performance of our algorithms on a set of well-known nonconvex problems that satisfy the  $n$ -sided PL condition.

**Benchmarks.** We compared our algorithms against the following state-of-the-art algorithms: **BM1**, which corresponds to the R-BCD algorithm described in 1 and **BM2**, an empirical variant of Algorithm A-RBCD in which updates are performed exclusively via case 3 while ignoring the first two cases, i.e.,  $\tilde{k}^t = -1$ .

**Function with only strict saddle point:** To illustrate an example where our A-RBCD algorithm converges to a NE point while the random BCD (BM1) fails, we considered a two-player general-sum game with the objective functions  $f_1(x_1, x_2) = (x_1 - 1)^2 + 4(x_1 + 0.1 \cos(x_1))x_2 + (x_2 + 0.1 \sin(x_2))^2$  and  $f_2(x_1, x_2) = (x_1 - 1)^2 + 4(x_1 - 0.1 \cos(x_1))x_2 + (x_2 - 0.1 \sin(x_2))^2$ . The problem aims at finding the NE  $(x_1^*, x_2^*)$ , i.e.,

$$f_1(x_1^*, x_2^*) \leq f_1(x_1, x_2^*), \quad f_2(x_1^*, x_2^*) \leq f_2(x_1^*, x_2), \quad \forall x_1, x_2.$$

Figure 4 illustrates the convergence results with different initializations. The iterates of A-RBCD always converge to the NE at a linear rate. Note that the NE is a strict saddle point.

**Cournot Competition:** The Cournot competition model, first proposed by Augustin Cournot in 1838, describes a market structure in which multiple firms compete simultaneously by choosing quantities of a homogeneous good [1, 2, 40]. Each firm  $i$  determines its output level  $q_i$  assuming that the quantities of all other firms  $q_{-i}$  are fixed, and the market price is determined by an inverse demand function denoted by  $P(Q)$  that depends on the aggregate output  $Q = \sum_i q_i$ . The strategic interdependence among firms is captured by the dependence of each firm's profit  $f_i(q_i, q_{-i}) := -(P(Q)q_i - C_i(q_i))$  on the total market output, where  $C_i(q_i)$  is the cost of producing the  $q_i$  output. A *Cournot-NE* is a fixed point in which every firm's output  $q_i^*$  is the best response to the others, i.e.,

$$f_i(q_i^*, q_{-i}^*) \leq f_i(q_i, q_{-i}^*), \quad \forall q_i \geq 0.$$

Normally, the equilibrium cannot be solved in closed form, and each firm's output decreases with its marginal cost while increasing with market demand. The Cournot framework has become a cornerstone in industrial organization [47], energy economics [72], and networked systems [6], providing a tractable yet insightful model for analyzing strategic interactions among self-interested agents under cost or capacity constraints.

In Figure 5, we compared the convergence of A-RBCD and two benchmark methods for computing the Cournot-NE under two types of inverse demand functions: linear  $P(Q) = a - bQ$  and nonlinear  $P(Q) = a - bQ^2$ . The A-RBCD algorithm demonstrates faster and more stable convergence than the benchmark methods. This result highlights that the adaptive correction term in A-RBCD effectively combines  $\nabla_i f_i$  and  $\nabla_i(G_F - F_{-i})$ , leading to improved convergence behavior.

**Infinite Horizon  $n$ -player Linear-quadratic (LQ) Game:** LQ games provide a powerful tool for analyzing the behavior of the multi-agent reinforcement learning (MARL) with continuous state and action space. Owing to their closed-form solvability and interpretability, LQ games serve as canonical benchmarks in both control theory and game-theoretic learning, bridging the gap between theoretical analysis and practical algorithm design [7, 70, 71].

In this game, the  $i$ -th player's objective function is given by

$$\mathbb{E}_{s_0 \sim \mathcal{D}} \left[ \sum_{t=0}^{+\infty} [(s_t)^T Q_i s_t + (u_t^i)^T R_i u_t^i] \right],$$

where  $s_t$  denotes the state,  $u_t^i$  is the input of  $i$ -th player at time  $t$ , and  $i \in [n]$ . The state transition of the system is characterized by  $s_{t+1} = A s_t + \sum_{i=1}^n B_i u_t^i + w_t$ , where  $A \in \mathbb{R}^{d \times d}$  and  $B \in \mathbb{R}^{d \times k_i}$ . When players apply a linear feedback strategy, i.e.,  $u_t^i = -K_i s_t$ , the objective functions become

$$f_i(K_i, K_{-i}) = \mathbb{E}_{s_0 \sim \mathcal{D}} \left[ \sum_{t=0}^{+\infty} [(s_t)^T Q_i s_t + (K_i s_t)^T R_i K_i s_t] \right].$$

If  $K_i$ s are bounded and  $\Sigma_0 = \mathbb{E}_{s_0 \sim \mathcal{D}} [s_0(s_0)^T]$  is full rank, the objective functions  $\{f_i\}$  satisfy the  $n$ -sided PL condition (see Appendix E.1 for a proof). However, as it is discussed in [22], even the




Figure 5: (a) Performances of A-RBCD, BM1, and BM2 applied to the Cournot model with  $P(Q) = a - bQ$ ; (b) with  $P(Q) = a - bQ^2$ .




Figure 6: (a) and (b) Performance of A-RBCD, BM1, and BM2 applied to the LQR game for a different set of parameters. The y-axis is on a log scale.

objective of the one-player LQ game is not convex. Subsequently, the objective function of the  $n$ -player LQ game is not multi-convex. See Appendix E.2 for examples.

Figure 6 demonstrates the performances of A-RBCD, BM1, and BM2 for different set of parameters. Note that the benchmark (BM2) reflects the algorithmic setting studied in [25], which establishes the convergence of both policy gradient descent and vanilla gradient descent in the context of  $n$ -player general-sum LQ games. Nevertheless, the theoretical guarantee provided in [25] critically depends on the assumption that the covariance of the injected noise is sufficiently large to stabilize the dynamics and it is limited to the class of LQ games which is a special case of our setting.

## 5 CONCLUSION

In this paper, we identified a subclass of nonconvex functions called  $n$ -sided PL functions and studied the convergence of GD-based algorithms, particularly the R-BCD algorithm, for finding their NEs. The  $n$ -sided PL condition is a reasonable extension of the gradient dominance condition, that holds in many existing problems and is better suited for a broader class of settings, particularly general-sum games. Furthermore, we studied the convergence rate of such first-order algorithm applied to general-sum games that satisfy the  $n$ -sided PL condition. Our analysis showed that the convergence rate depends on a local relation between the sum of

the objective functions  $F$ , and the sum of their best responses  $G_F$ . Subsequently, we proposed two novel algorithms, IA-RBCD and A-RBCD, equipped with the gradient of  $G_F$ , that provably converge to a NE almost surely with random initialization, even if the individual objective functions are not lower bounded and have strict saddle points. We hope this work can shed some light on the understanding of a broader class of nonconvex optimization problems.

## REFERENCES

- [1] Melika Abolhassani, Mohammad Hossein Bateni, MohammadTaghi Hajiaghayi, Hamid Mahini, and Anshul Sawant. 2014. Network cournot competition. In *International Conference on Web and Internet Economics*. Springer, 15–29.
- [2] Blaise Allaz and Jean-Luc Vila. 1993. Cournot competition, forward markets and efficiency. *Journal of Economic theory* 59, 1 (1993), 1–16.
- [3] Zeyuan Allen-Zhu, Zheng Qu, Peter Richtárik, and Yang Yuan. 2016. Even faster accelerated coordinate descent using non-uniform sampling. In *International Conference on Machine Learning*. PMLR, 1110–1119.
- [4] Yunyan Bai, Yuxing Liu, and Luo Luo. 2024. On the Complexity of Finite-Sum Smooth Optimization under the Polyak- $\{\setminus L\}$  ojasiewicz Condition. *arXiv preprint arXiv:2402.02569* (2024).
- [5] Amir Beck and Luba Tetruashvili. 2013. On the convergence of block coordinate descent type methods. *SIAM journal on Optimization* 23, 4 (2013), 2037–2060.
- [6] Kostas Bimpikis, Shayan Ehsani, and Rahmi Ilkılıç. 2019. Cournot competition in networked markets. *Management Science* 65, 6 (2019), 2467–2481.
- [7] Jingjing Bu, Lillian J Ratliff, and Mehran Mesbahi. 2019. Global convergence of policy gradient for sequential zero-sum linear quadratic dynamic games. *arXiv preprint arXiv:1911.04672* (2019).
- [8] Xufeng Cai, Chaobing Song, Stephen Wright, and Jelena Diakonikolas. 2023. Cyclic block coordinate descent with variance reduction for composite nonconvex optimization. In *International Conference on Machine Learning*. PMLR, 3469–3494.

- [9] Yang Cai, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. 2024. Near-optimal policy optimization for correlated equilibrium in general-sum markov games. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 3889–3897.
- [10] Loris Cannelli, Francisco Facchinei, Gesualdo Scutari, and Vyacheslav Kungurtsev. 2020. Asynchronous optimization over graphs: Linear convergence under error bound conditions. *IEEE Trans. Automat. Control* 66, 10 (2020), 4604–4619.
- [11] Andrea Celli, Alberto Marchesi, Tommaso Bianchi, and Nicola Gatti. 2019. Learning to correlate in multi-player general-sum sequential games. *Advances in Neural Information Processing Systems* 32 (2019).
- [12] Zachary Charles and Dimitris Papailiopoulos. 2018. Stability and generalization of learning algorithms that converge to global optima. In *International conference on machine learning*. PMLR, 745–754.
- [13] Lesi Chen, Boyuan Yao, and Luo Luo. 2022. Faster stochastic algorithms for minimax optimization under polyak- $\{L\}$  ojasiewicz condition. *Advances in Neural Information Processing Systems* 35 (2022), 13921–13932.
- [14] Ziang Chen, Yingzhou Li, and Jianfeng Lu. 2023. On the global convergence of randomized coordinate gradient descent for nonconvex optimization. *SIAM Journal on Optimization* 33, 2 (2023), 713–738.
- [15] Flavia Chorobura and Ion Necoara. 2023. Random coordinate descent methods for nonseparable composite optimization. *SIAM Journal on Optimization* 33, 3 (2023), 2160–2190.
- [16] Ying Cui, Chao Ding, and Xinyuan Zhao. 2017. Quadratic growth conditions for convex matrix optimization problems associated with spectral functions. *SIAM Journal on Optimization* 27, 4 (2017), 2332–2355.
- [17] Marina Danilova, Anastasiia Kulakova, and Boris Polyak. 2020. Non-monotone behavior of the heavy ball method. In *Difference Equations and Discrete Dynamical Systems with Applications: 24th ICDEA, Dresden, Germany, May 21–25, 2018*. Springer, 213–230.
- [18] Constantinos Daskalakis, Dylan J Foster, and Noah Golowich. 2020. Independent policy gradient methods for competitive reinforcement learning. *Advances in neural information processing systems* 33 (2020), 5527–5540.
- [19] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *Commun. ACM* 52, 2 (2009), 89–97.
- [20] Omer Elkabetz and Nadav Cohen. 2021. Continuous vs. discrete optimization of deep neural networks. *Advances in Neural Information Processing Systems* 34 (2021), 4947–4960.
- [21] Ilyas Fatkhullin, Jalal Etesami, Niao He, and Negar Kiyavash. 2022. Sharp analysis of stochastic optimization under global Kurdyka-Lojasiewicz inequality. *Advances in Neural Information Processing Systems* 35 (2022), 15836–15848.
- [22] Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. 2018. Global convergence of policy gradient methods for the linear quadratic regulator. In *International conference on machine learning*. PMLR, 1467–1476.
- [23] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [24] Zhishuai Guo, Yan Yan, Zhuoning Yuan, and Tianbao Yang. 2023. Fast Objective & Duality Gap Convergence for Non-Convex Strongly-Concave Min-Max Problems with PL Condition. *Journal of Machine Learning Research* 24 (2023), 1–63.
- [25] Ben Hambly, Renyuan Xu, and Huining Yang. 2023. Policy gradient methods find the Nash equilibrium in N-player general-sum linear-quadratic games. *Journal of Machine Learning Research* 24, 139 (2023).
- [26] Moritz Hardt and Tengyu Ma. 2017. Identity matters in deep learning. *5th International Conference on Learning Representations (ICLR)* (2017).
- [27] Mingyi Hong, Xiangfeng Wang, Meisam Razaviyayn, and Zhi-Quan Luo. 2017. Iteration complexity analysis of block coordinate descent methods. *Mathematical Programming* 163 (2017), 85–114.
- [28] Michael Jordan, Tianyi Lin, and Zhengyuan Zhou. 2024. Adaptive, doubly optimal no-regret learning in strongly monotone and exp-concave games with gradient feedback. *Operations Research* (2024).
- [29] Hamed Karimi, Julie Nutini, and Mark Schmidt. 2016. Linear convergence of gradient and proximal-gradient methods under the polyak-lojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19–23, 2016, Proceedings, Part I*. Springer, 795–811.
- [30] Chung-Wei Lee, Christian Kroer, and Haipeng Luo. 2021. Last-iterate convergence in extensive-form games. *Advances in Neural Information Processing Systems* 34 (2021), 14293–14305.
- [31] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. 2016. Gradient descent only converges to minimizers. In *Conference on learning theory*. PMLR, 1246–1257.
- [32] Lihua Lei, Cheng Ju, Jianbo Chen, and Michael I Jordan. 2017. Non-convex finite-sum optimization via scsg methods. *Advances in Neural Information Processing Systems* 30 (2017).
- [33] Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. 2018. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*. PMLR, 2–47.
- [34] Zhize Li, Hongyan Bao, Xiangliang Zhang, and Peter Richtárik. 2021. PAGE: A simple and optimal probabilistic gradient estimator for nonconvex optimization. In *International conference on machine learning*. PMLR, 6286–6295.
- [35] Cheuk Yin Lin, Chaobing Song, and Jelena Diakonikolas. 2023. Accelerated cyclic coordinate dual averaging with extrapolation for composite convex optimization. In *International Conference on Machine Learning*. PMLR, 21101–21126.
- [36] Chaoyue Liu, Libin Zhu, and Mikhail Belkin. 2022. Loss landscapes and optimization in over-parameterized non-linear systems and neural networks. *Applied and Computational Harmonic Analysis* 59 (2022), 85–116.
- [37] Ji Liu, Steve Wright, Christopher Ré, Victor Bittorf, and Srikrishna Sridhar. 2014. An asynchronous parallel stochastic coordinate descent algorithm. In *International Conference on Machine Learning*. PMLR, 469–477.
- [38] Stanislaw Lojasiewicz. 1963. A topological property of real analytic subsets. *Coll. du CNRS, Les équations aux dérivées partielles* 117, 87–89 (1963), 2.
- [39] Zhi-Quan Luo and Paul Tseng. 1993. Error bounds and convergence analysis of feasible descent methods: a general approach. *Annals of Operations Research* 46, 1 (1993), 157–178.
- [40] Eric Maskin and Jean Tirole. 1987. A theory of dynamic oligopoly, III: Cournot competition. *European economic review* 31, 4 (1987), 947–968.
- [41] Kensuke Nakamura, Stefano Soatto, and Byung-Woo Hong. 2021. Block-cyclic stochastic coordinate descent for deep neural networks. *Neural Networks* 139 (2021), 348–357.
- [42] John F Nash et al. 1950. The bargaining problem. *Econometrica* 18, 2 (1950), 155–162.
- [43] Ion Necoara, Yu Nesterov, and Francois Glineur. 2019. Linear convergence of first order methods for non-strongly convex optimization. *Mathematical Programming* 175 (2019), 69–107.
- [44] Yu Nesterov. 2012. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization* 22, 2 (2012), 341–362.
- [45] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. 2007. *Algorithmic game theory*. Cambridge University press.
- [46] Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. 2019. Solving a class of non-convex min-max games using iterative first order methods. *Advances in Neural Information Processing Systems* 32 (2019).
- [47] Georgia Perakis and Wei Sun. 2014. Efficiency analysis of Cournot competition in service industries with congestion. *Management Science* 60, 11 (2014), 2684–2700.
- [48] Boris T Polyak. 1963. Gradient methods for the minimisation of functionals. *U. S. S. R. Comput. Math. and Math. Phys.* 3, 4 (1963), 864–878.
- [49] Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alex Smola. 2016. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*. PMLR, 314–323.
- [50] Mohamed Reyad, Amany M Sarhan, and Mohammad Arafa. 2023. A modified Adam algorithm for deep neural network optimization. *Neural Computing and Applications* 35, 23 (2023), 17095–17112.
- [51] Masoud Roudneshin, Jalal Arabneydi, and Amir G Aghdam. 2020. Reinforcement learning in nonzero-sum Linear Quadratic deep structured games: Global convergence of policy optimization. In *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 512–517.
- [52] Maziar Sanjabi, Meisam Razaviyayn, and Jason D Lee. 2018. Solving non-convex non-concave min-max games under polyak- $\{L\}$  ojasiewicz condition. *arXiv preprint arXiv:1812.02878* (2018).
- [53] Gesualdo Scutari, Daniel P Palomar, Francisco Facchinei, and Jong-Shi Pang. 2010. Convex optimization, game theory, and variational inequality theory. *IEEE Signal Processing Magazine* 27, 3 (2010), 35–49.
- [54] Xinyue Shen, Steven Diamond, Madeleine Udell, Yuantao Gu, and Stephen Boyd. 2017. Disciplined multi-convex programming. In *2017 29th Chinese control and decision conference (CCDC)*. IEEE, 895–900.
- [55] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. 2014. Deterministic policy gradient algorithms. In *International conference on machine learning*. Pmlr, 387–395.
- [56] Mohammad Karim Sohrabi and Hossein Azgomi. 2020. A survey on the combined use of optimization methods and game theory. *Archives of Computational Methods in Engineering* 27, 1 (2020), 59–80.
- [57] Ju Sun, Qing Qu, and John Wright. 2018. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics* 18 (2018), 1131–1198.
- [58] Junxiang Wang, Hongyi Li, and Liang Zhao. 2022. Accelerated gradient-free neural network training by multi-convex alternating optimization. *Neurocomputing* 487 (2022), 130–143.
- [59] Junxiang Wang, Liang Zhao, and Lingfei Wu. 2019. Multi-convex inequality-constrained alternating direction method of multipliers. *arXiv* 2019 (2019).
- [60] Jun-Kun Wang, Chi-Heng Lin, Andre Wibisono, and Bin Hu. 2022. Provable acceleration of heavy ball beyond quadratics for a class of polyak-lojasiewicz functions when the non-convexity is averaged-out. In *International Conference on Machine Learning*.
- [61] Zhe Wang, Kaiyi Ji, Yi Zhou, Yingbin Liang, and Vahid Tarokh. 2019. Spiderboost and momentum: Faster variance reduction algorithms. *NeurIPS* 32 (2019).
- [62] Yangyang Xu and Wotao Yin. 2013. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences* 6, 3 (2013), 1758–1789.

- [63] Yangyang Xu and Wotao Yin. 2017. A globally convergent algorithm for non-convex optimization based on block coordinate update. *Journal of Scientific Computing* 72, 2 (2017), 700–734.
- [64] Junchi Yang, Negar Kiyavash, and Niao He. 2020. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. *Advances in Neural Information Processing Systems* 33 (2020), 1153–1165.
- [65] Junchi Yang, Antonio Orvieto, Aurelien Lucchi, and Niao He. 2022. Faster single-loop algorithms for minimax optimization without strong concavity. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 5485–5517.
- [66] Long Yang, Jiaming Ji, Juntao Dai, Linrui Zhang, Binbin Zhou, Pengfei Li, Yaodong Yang, and Gang Pan. 2022. Constrained update projection approach to safe policy optimization. *Advances in Neural Information Processing Systems* 35 (2022), 9111–9124.
- [67] Pengyun Yue, Cong Fang, and Zhouchen Lin. 2023. On the lower bound of minimizing polyak-jojasiewicz functions. In *The Thirty Sixth Annual Conference on Learning Theory*. PMLR, 2948–2968.
- [68] Jinshan Zeng, Tim Tsz-Kit Lau, Shaobo Lin, and Yuan Yao. 2019. Global convergence of block coordinate descent in deep learning. In *International conference on machine learning*. PMLR, 7313–7323.
- [69] Hui Zhang and Wotao Yin. 2013. Gradient methods for convex minimization: better rates under weaker conditions. *arXiv preprint arXiv:1303.4645* (2013).
- [70] Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. 2019. Policy optimization provably converges to Nash equilibria in zero-sum linear quadratic games. *Advances in Neural Information Processing Systems* 32 (2019).
- [71] Kaiqing Zhang, Xiangyuan Zhang, Bin Hu, and Tamer Basar. 2021. Derivative-free policy optimization for linear risk-sensitive and robust control design: Implicit regularization and sample complexity. *Advances in neural information processing systems* 34 (2021), 2949–2964.
- [72] Yajun Zhang, Chenghong Gu, Xiaohe Yan, and Furong Li. 2020. Cournot oligopoly game-based local energy trading considering renewable energy uncertainty costs. *Renewable Energy* 159 (2020), 1117–1127.
- [73] Ziming Zhang and Matthew Brand. 2017. Convergent block coordinate descent for training tikhonov regularized deep neural networks. *NeurIPS* 30 (2017).
- [74] Yulai Zhao, Yuandong Tian, Jason Lee, and Simon Du. 2022. Provably efficient policy optimization for two-player zero-sum markov games. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2736–2761.