Finite-Time Last-Iterate Convergence for Learning in Multi-Player Games

Anonymous Author(s)
Affiliation
Address
email

Abstract

We study the question of last-iterate convergence rate of the *extragradient algorithm* by [Kor76] and the *optimistic gradient algorithm* by [Pop80] in multi-player games. We show that both algorithms with *constant step-size* have last-iterate convergence rate of $O(\frac{1}{\sqrt{T}})$ to a Nash equilibrium in terms of the gap function in smooth monotone games, where each player's action set is an *arbitrary convex set*. Previous results only study the unconstrained setting, where each player's action set is the entire Euclidean space. Our results address an open question raised in several recent work [HIMM19, GPD20, GPDO20], which ask for last-iterate convergence rate of either the extragradient or the optimistic gradient algorithm in the constrained setting. Our convergence rates for both algorithms are tight and match the lower bounds by [GPD20, GPDO20]. At the core of our results lies a new notion – the *tangent residual*, which we use to measure the proximity to a Nash equilibrium. We use the tangent residual (or a slight variation of the tangent residual) as the the potential function in our analysis of the extragradient algorithm (or the optimistic gradient algorithm).

1 Introduction

We consider learning in *monotone games*, a class of multi-player games introduced by [Ros65] that include many well studied games, e.g., two-player zero-sum games, convex-concave games, λ-cocoercive games [LZMJ20], zero-sum polymatrix games [BF87, DP09, CD11], and zero-sum socially-concave games [EDMN09]. We focus on the following question: *Can we obtain **last-iterate convergence rate** to a **Nash equilibrium** in monotone games when all players act according to a simple learning algorithm?*

We adopt the multi-player online learning model as introduced in [CBL06], where players interact with each other repeatedly. At every time step t, every player $i \in \{1, \ldots, N\}$ chooses an action $z_t^{(i)}$ from her action set $\mathcal{Z}^{(i)}$, which we assume to be a closed convex set in $\mathbb{R}^{n_i}$. We say the game is *unconstrained* if $\mathcal{Z}^{(i)} = \mathbb{R}^{n_i}$ for each player i. Player i's loss function $\ell_t^{(i)}(\cdot)$ is determined based on the underlying game and the actions of the other players in round t. Player i receives the loss $\ell_t^{(i)}(z_t^{(i)})$ as well as some additional feedback that informs her how to improve her decisions in the future. In this paper, we assume that each player receives the gradient feedback, i.e., player i receives the vector $\nabla \ell_t^{(i)}(z_t^{(i)})$. We make an additional mild assumption that the game is *smooth*, i.e., all players' gradients are Lipschitz. Smoothness is

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

Game class	Setting	Step size	Convergence rate
Strongly monotone	general	constant	$O(c^{-T})$ (see e.g., [Tse95] [LS19, MOP19, ZMM+21])
Cocorcive	unconstrained	constant	$O(\frac{1}{\sqrt{T}})$ [LZMJ20]
Monotone	general	constant	Asymptotic [Pop80, HIMM19]
	general	decreasing	Asymptotic * (see e.g., [ZMM+17] [ZMA+18, MZ19, HAM21])
	unconstrained	constant	$O(\frac{1}{\sqrt{T}})^{\dagger}$ [GPD20]
	general	constant	$O(\frac{1}{\sqrt{T}})$ **[This paper]**

Table 1: Last-iterate convergence for no-regret learning in smooth monotone games with perfect gradient feedback. (*) The results hold for variationally stable games. (†) The result holds under an additional second-order smoothness assumption.

a natural assumption that is satisfied in most applications and is also made in the majority of works concerning monotone games.

The standard metric to quantify an online learning algorithm's performance is the regret. Formally, the regret of player i is defined as the difference between $\sum_{t=1}^{T} \ell_t^{(i)}(z_t^{(i)})$, player i's cumulative loss, and $\min_{z \in \mathcal{Z}^{(i)}} \sum_{t=1}^{T} \ell_t^{(i)}(z)$, the loss incurred by the best fixed action in hindsight. An online learning algorithm is no-regret if, even under an adversarially chosen loss sequence $\{\ell_t^{(i)}(\cdot)\}_{t \in [T]}$, its regret at the end of round T is sublinear in T.

A vast literature on learning in games discusses the convergence to a Nash equilibrium using no-regret learning algorithms. However, most of the results concern only the *time-average* convergence, i.e., the convergence of the time average of the joint action profile, rather than the last-iterate convergence, i.e., the convergence of the joint action profile. From a game-theoretic perspective, the last-iterate convergence is more appealing compared to the time-average convergence, as only the last-iterate convergence provides a description of the evolution of the overall behavior of the players. In contrast, the trajectory of the players' joint action may be cycling around in the space perpetually while still converges in the time-average sense as demonstrated by [MPP18]. Recently, a line of work is devoted to obtain last-iterate convergence in smooth monotone games [ZMB+17, ZMM+17, ZMA+18, DP18, MOP19, HIMM19, LNPW20, GPD20, LZMJ20, ZMM+21]. However, unless the game is strongly monotone or unconstrained, only asymptotic convergence is known. Moreover, many of these results crucially rely on decreasing step-size, which, as pointed out by [LZMJ20], is unnatural from an economic point of view, because it treats newly acquired information with decreasing importance. Hence, the following question is of particular interest and is raised as an open question in [HIMM19, LZMJ20, GPD20].

*Can we establish **last-iterate rates** if all players of a **constrained** smooth monotone game act according to a no-regret learning algorithm with **constant step size**?* (*)

[LZMJ20] first realizes the importance of question (*) and takes initial steps towards addressing it. They show that if all players follow the *gradient descent* algorithm with constant step size, then for all smooth λ-cocoercive games, the joint action $(z_t^{(1)}, \ldots, z_t^{(N)})$ has last-iterate convergence rate of $O(\frac{1}{\sqrt{T}})$ to a Nash equilibrium in terms of the gap function. For smooth strongly monotone games, a subclass of λ-cocoercive games, linear last-iterate convergence rates are known [Tse95, Mal15, LS19, MOP19, ZMM+21]. Despite the generality of λ-cocoercive games, several fundamental classes of games such as two-player zero-sum games, zero-sum polymatrix games [BF87, DGP09, CD11] and its generalization zero-sum socially-concave games [EDMN09] are monotone but not λ-cocoercive. [GPD20] extends the result to smooth monotone games with an $O(\frac{1}{\sqrt{T}})$ last-iterate convergence rate using a different algorithm – the *optimistic gradient* by [Pop80] under an additional second-order smoothness assumption. Note that gradient descent has been observed to fail to converge in even two-player zero-sum games (see e.g., [DISZ17]), so a different algorithm is indeed needed. However, the results by [LZMJ20] and [GPD20] only consider *unconstrained* games, while in most game-theoretic settings, the players' actions are constrained. For example, in

any finite games, a player is restricted to choose a distribution over her finite set of actions. We summarize the results for last-iterate convergence in monotone games in Table 1.

Our Contributions. Our first contribution is to provide an affirmative answer to question (*).

> **Contribution 1:** In Theorem 3, we show that if all players of a constrained smooth monotone game act according to the optimistic gradient algorithm, which is no-regret, with a **constant step size**, then their joint action exhibits a last-iterate convergence rate of $O\left(\frac{1}{\sqrt{T}}\right)$ in terms of the gap function (Definition 2) to a Nash equilibrium.

Our result holds in the constrained setting and does not rely on the second-order smoothness assumption made in [GPD20]. Moreover, our upper bound is tight and matches the lower bound of [GPD20].

The problem of finding a Nash equilibrium in a smooth monotone game is essentially equivalent to solving a Lipschitz and monotone variational inequality (VI) (Definition 5 in Appendix D),[1] which has been studied since the 1960s [HS66, Bro65, LS67, BS68, Sib70]. There is a vast literature on solving VIs, and we refer the reader to [FP07] for further references. The extragradient (EG) algorithm by [Kor76] and the optimistic gradient (OG) algorithm by [Pop80] are arguably the two most classical and popular methods for solving monotone VIs. Despite the long history, the last-iterates of both algorithms are only known to asymptotically converge to a solution of the monotone and Lipschitz VI,[2] but no upper bounds on the rate of convergence had been provided for the general setting.

> **Contribution 2:** We provide the first and tight last-iterate convergence rate of $O(\frac{1}{\sqrt{T}})$ in terms of the gap function for both EG and OG with constant step size for solving Lipschitz and monotone VIs (Theorem 9 and Theorem 10 in Appendix I).

Our analysis of OG for monotone games directly applies to monotone VIs. To be more consistent with our discussion of OG, we state our last-iterate convergence rate of EG in the context of online learning in monotone games (Theorem 3).[3] Prior to our work, last-iterate convergence rate of EG only exists in the unconstrained setting. [GPDO20] shows a $O(\frac{1}{\sqrt{T}})$ upper bound in terms of the gap function under an additional second-order smoothness condition. [GLG21] improves the result and shows that the same upper bound holds without the second-order smoothness condition but still require the setting to be unconstrained.

Our Analysis. As mentioned in [GPD20], "The lack of existence of a natural potential function in general monotone games is a significant challenge in establishing last-iterate convergence." Indeed, many of the natural quantities such as the gap function, the norm of the gradient, the difference of two consecutive iterates are provably non-monotone even in bilinear games. See Appendix J for more discussion and examples. We propose a novel notion that (i) measures the proximity to a Nash equilibrium, and (ii) can be used to construct natural potential functions to analyze both EG and OG in monotone games. We call the new notion the **tangent residual**, which can be viewed as the norm of the gradient projected to the tangent cone of the current iterate (Definition 4). The tangent residual plays a crucial role in our analyses for both algorithms. Unlike the other quantities mentioned above, we show that the *tangent residual is monotonically decreasing* and has *a last-iterate convergence rate of* $O(\frac{1}{\sqrt{T}})$ for EG. For OG, we prove that a small modification of the tangent residual is monotonically decreasing, which implies that the tangent residual has a last-iterate convergence rate of $O(\frac{1}{\sqrt{T}})$. Using the convergence rate of the tangent residual, we can easily derive the last-iterate convergence rate of other classical performance measures such as the gap function and the total gap function. However, we suspect these

[1] Technically, the definition of Lipschitz and monotone VI captures finding a Nash equilibrium in a smooth monotone game as a special case, but the difference has little impact on our analyses, and our results hold for Lipschitz and monotone VIs.

[2] The last-iterate asymptotic convergence of EG can be found in [Kor76] and [FP07], and the last-iterate asymptotic convergence of OG can be found in [Pop80] and [HIMM19].

[3] Although EG is not a no-regret learning algorithm (See Proposition 10 in [GPD20]), it is nevertheless a simple and natural algorithm.

rates can be challenging to obtain directly without the help of a potential function. Finally, we establish the monotonicity of our potential functions using computer-aided proofs based on **sum-of-squares (SOS) programming** [Nes00, Par00, Par03, Las01, Lau09]. Indeed, our potential function of OG is directly constructed using SOS programming. We think the tangent residual and the SOS-based analysis might be of independent interest. See Section 4.2, Appendix B, and Appendix C.1 for more discussion.

1.1 Related Work

Last-Iterate Convergence Rate for EG/OG like Algorithms. [GPDO20, GPD20] show a lower bound of $\Omega(\frac{1}{\sqrt{T}})$ for solving bilinear games using any p-SCLI algorithms, which include EG and OG. In the unconstrained setting, if we further assume that either the game is strongly monotone or the payoff matrix A in a bilinear game has all singular values bounded away from 0, linear convergence rate is known for EG, OG, and several of their variants [DISZ17, GBV+18, LS19, MOP19, PDZC20, ZY19]. The results for the constrained setting are sparser. Unless the game is strongly monotone, most results only guarantee asymptotic convergence, i.e., converge in the limit, [DP18, LNPW20]. Finally, a recent paper by [WLZL21b] provides a linear convergence rate of OG for bilinear games when the domain is a polytope. They show that there is a *problem dependent* constant $0 < c < 1$ that depends on the payoff matrix of the game as well as the constraint set, so that the error shrinks by a $1 - c$ factor per iteration. However, c may be arbitrarily close to 0, even if we assume the corresponding operator to be L-Lipschitz. Overall, their "problem-dependent" bound is incomparable and complements the worst-case view taken in this paper, where we want to derive the worst-case convergence rate for all smooth and monotone games. Our results are the first last-iterate convergence rates in this worst-case view and match the lower bounds by [GPDO20, GPD20].

Other Algorithms and Performance Measures. It is well-known that both EG [Nem04] and OG [HIMM19, MOP20] have a time-average convergence rate of $O(\frac{1}{T})$ in terms of the gap function for smooth monotone games. Other than the gap function, one can also measure the convergence using the norm of the operator if the setting is unconstrained, or the natural residual (Definition 6 in Appendix E) or similar notions if the setting is constrained. In the unconstrained setting, [Kim21], [YR21], and [LK21] provide algorithms that obtain $O(\frac{1}{T})$ convergence rate in terms of the norm of the operator, which is shown to be optimal by [YR21] for Lipschitz and monotone VIs. In the constrained setting, [Dia20] shows the same $O(\frac{1}{T})$ convergence rate under the extra assumption that the operator is cocoercive and loses an additional logarithmic factor when the operator is only monotone. Our result implies a $O(\frac{1}{\sqrt{T}})$ last-iterate convergence rate in terms of the natural residual for both EG and OG. From an optimization point of view, i.e., the goal is to *compute* a Nash equilibrium of the game as quickly as possible, we should either output the average-iterate of EG or OG if we care about the gap function, or choose one of the above faster algorithms if we care about the norm of the operator or the natural residual. However, one of our main motivation is game-theoretic, that is, we would like to view simple and no-regret learning algorithms as models of agents' behavior in online learning and understand the speed for the overall behavior to converge to a Nash equilibrium. The results above, although obtain a faster rate, either only hold in the unconstrained setting, or only in the time-average sense, or are not no-regret. From this game-theoretic view point, we believe understanding the last-iterate convergence rate of no-regret learning algorithms is an important problem.

Stochastic Gradient Feedback. Several recent papers consider the case where the gradient feedback is noisy (see e.g.,[HIMM19, LZMJ20]). This is a very interesting and realistic model. We think our techniques will be useful in obtaining results in this model in future work.

2 Preliminaries

We consider the Euclidean Space $(\mathbb{R}^n, \|\cdot\|)$, where $\|\cdot\|$ is the ℓ_2 norm and $\langle\cdot,\cdot\rangle$ denotes inner product on $\mathbb{R}^n$. A **continuous game** $\mathcal{G}$ is denoted as $(\mathcal{N}, \{\mathcal{Z}^{(i)}\}_{i\in[N]}, \{f^{(i)}\}_{i\in[N]})$ where

there are N players $\mathcal{N} = \{1, \cdots, N\}$. Player $i \in \mathcal{N}$ chooses action from a closed convex set $\mathcal{Z}^{(i)} \subseteq \mathbb{R}^{n_i}$ such that $\mathcal{Z}_{\mathcal{G}} := \Pi_{i \in \mathcal{N}} \mathcal{Z}^{(i)} \subseteq \mathbb{R}^n$ and wants to minimize her cost function $f^{(i)} : \mathcal{Z}_{\mathcal{G}} \to \mathbb{R}$. For each player i, we denote by $z^{(-i)}$ the vector of actions of all the other players and by $z^{(i)}$ the action of player i. When players play according to action profile $z \in \mathcal{Z}_{\mathcal{G}}$, player i receives gradient feedback $\nabla_{z^{(i)}} f^{(i)}(z^{(1)}, \ldots, z^{(N)})$. A *Nash equilibrium* of game $\mathcal{G}$ is an action profile $z^* \in \mathcal{Z}_{\mathcal{G}}$ such that $f^{(i)}(z^*) \leq f^{(i)}(z'^{(i)}, z^{*(-i)})$ for any $z'^{(i)} \in \mathcal{Z}^{(i)}$. Let $F_{\mathcal{G}}(\cdot) = (\nabla_{z^{(1)}} f^{(1)}(\cdot), \cdots, \nabla_{z^{(N)}} f^{(N)}(\cdot))$ be an operator that maps any joint action in $\mathcal{Z}$ to the corresponding joint gradient feedback vector in $\mathbb{R}^n$. When $\mathcal{G}$ is clear from context, we omit the subscript and write write $F_{\mathcal{G}}(z)$ ($\mathcal{Z}_{\mathcal{G}}$ resp.) as $F(z)$ ($\mathcal{Z}$ resp.).

Nash Equilibria of Monotone Games. Throughout this paper, we focus on smooth monotone games (Definition 1). It is well known that finding a Nash equilibrium of a monotone game is exactly the same as finding a solution to the variational inequality with monotone operator $F_{\mathcal{G}}$ (Lemma 1).

Definition 1 ([Ros65]). *We say game $\mathcal{G}$ is L-*smooth *and* monotone *if the operator $F_{\mathcal{G}}$ is monotone, that is, $\forall z, z' \in \mathcal{Z}$, $\langle F_{\mathcal{G}}(z) - F_{\mathcal{G}}(z'), z - z' \rangle \geq 0$ and L-Lipschitz, i.e., $\forall z, z' \in \mathcal{Z}$, $L \cdot \|z - z'\| \geq \|F_{\mathcal{G}}(z) - F_{\mathcal{G}}(z')\|$.*

Lemma 1 (1.4.2 Proposition [FP07]). *For a monotone game $\mathcal{G}$, an action profile z^* is a Nash equilibrium iff $\langle F_{\mathcal{G}}(z^*), z^* - z \rangle \leq 0$, $\forall z \in \mathcal{Z}$.*

Remark 1. *One sufficient condition for a Nash equilibrium z^* to exist is when the set $\mathcal{Z}$ is bounded, but there are also other sufficient conditions that apply to unbounded $\mathcal{Z}$. See [FP07] for more details. Throughout this paper, we only consider monotone games that have a Nash equilibrium.*

Definition 2 (Gap and Total Gap Function). *For a monotone game $\mathcal{G}$, two standard ways to measure the proximity of an action profile $z \in \mathcal{Z}$ to Nash equilibrium, are by its gap function and total gap function. Let D be a fixed parameter. The gap function is defined as $\textsc{Gap}_{\mathcal{G},D}(z) = \max_{z' \in \mathcal{Z} \cap \mathcal{B}(z,D)} \langle F_{\mathcal{G}}(z), z - z' \rangle$, where $\mathcal{B}(z, D)$ is a ball with radius D centered at z.[4] The total gap function is defined as $\textsc{TGap}_{\mathcal{G},D}(z) = \sum_{i \in \mathcal{N}} \left(f^{(i)}(z) - \min_{z'^{(i)} \in Z^{(i)} \cap \mathcal{B}(z^{(i)}, D)} f(z'^{(i)}, z^{(-i)}) \right)$.*

When $\mathcal{G}$ and D are clear from context, we omit subscripts and write the gap function (total gap function resp.) at z as $\textsc{Gap}(z)$ ($\textsc{TGap}(z)$ resp.).

The Optimistic Gradient Algorithm. Let w_k be the action profile played at day k and assume that player i updates her action according to the OG algorithm. For arbitrary $z_0^{(i)}$ and $w_0^{(i)}$, player's i action at day $k+1$ is $w_{k+1}^{(i)}$ such that $z_k^{(i)} = \Pi_{\mathcal{Z}^{(i)}} \left[z_{k-1}^{(i)} - \eta \nabla_{z^{(i)}} f^{(i)}(w_k) \right]$ and $w_{k+1}^{(i)} = \Pi_{\mathcal{Z}^{(i)}} \left[z_k^{(i)} - \eta \nabla_{z^{(i)}} f^{(i)}(w_k) \right]$.

Note that action $z_k^{(i)}$ is not being played by player i and is only used to compute action $w_{k+1}^{(i)}$. When all the players update their actions according to OG with step-size η, let $w_0 = (w_0^{(1)}, \ldots, w_0^{(N)})$ and $z_0 = (z_0^{(1)}, \ldots, z_0^{(N)})$. Then at day $k+1$, players pick action profile w_{k+1}, where:

$$z_k = \Pi_{\mathcal{Z}} \left[z_{k-1} - \eta F(w_k) \right], \qquad w_{k+1} = \Pi_{\mathcal{Z}} \left[z_k - \eta F(w_k) \right] \tag{1}$$

Clearly, the OG update rule is well defined for any operator F and any closed convex set $\mathcal{Z}$. Indeed, this is how the OG algorithm is defined for variational inequalities.

The Extragradient Algorithm. Let z_k be the action profile played at day $2 \cdot k$ and assume that all players update their actions according to EG with step-size η. Then players pick action profile $z_{k+\frac{1}{2}}$ at day $2 \cdot k + 1$ and action profile z_{k+1} at day $2 \cdot k + 2$, where:

$$z_{k+\frac{1}{2}} = \Pi_{\mathcal{Z}} \left[z_k - \eta F(z_k) \right], \qquad z_{k+1} = \Pi_{\mathcal{Z}} \left[z_k - \eta F(z_{k+\frac{1}{2}}) \right] \tag{2}$$

[4] When $\mathcal{Z}$ is bounded, we choose D to be the diameter of $\mathcal{Z}$. When $\mathcal{Z}$ is unbounded, the standard choice is to choose D so that all the iterates of the algorithms are guaranteed to maintain in the ball $\mathcal{B}(z, D)$. See Appendix D for more details.

Similarly, the EG update rule is well defined for any operator F and any closed convex set $\mathcal{Z}$. For the rest of the paper, we only consider the case where all players use the OG (or EG) algorithm with constant step-size η. We omit superscripts that denote players' identity and use Expression (1) (or Expression (2)) for the update rule when players use OG (or EG).

3 The Tangent Residual and Its Properties

We formally introduce our key performance measure, the tangent residual . We define the tangent residual over operators rather than games, as it will be easier to provide intuition behind its formulation.

Definition 3 (Unit Normal Cone)**.** *Given a closed convex set $\mathcal{Z} \subseteq \mathbb{R}^n$ and a point $z \in \mathcal{Z}$, we denote by $N_{\mathcal{Z}}(z) = \langle v \in \mathbb{R}^n : \langle v, z' - z\rangle \leq 0, \forall z' \in \mathcal{Z}\rangle$ the normal cone of $\mathcal{Z}$ at point z and by $\widehat{N}_{\mathcal{Z}}(z) = \{v \in N_{\mathcal{Z}}(z) : \|v\| \leq 1\}$ the intersection of the unit ball with the the normal cone of $\mathcal{Z}$ at z. Note that $\widehat{N}_{\mathcal{Z}}(z)$ is nonempty and compact for any $z \in \mathcal{Z}$, as $(0, \ldots, 0) \in \widehat{N}_{\mathcal{Z}}(z)$.*

Definition 4 (Tangent Residual)**.** *Given an operator $F : \mathcal{Z} \to \mathbb{R}^n$ and a closed convex set $\mathcal{Z}$, let $T_{\mathcal{Z}}(z) := \{z' \in \mathbb{R}^n : \langle z', a\rangle \leq 0, \forall a \in N_{\mathcal{Z}}(z)\}$ be the tangent cone of z, and define $J_{\mathcal{Z}}(z) := \{z\} + T_{\mathcal{Z}}(z)$. The tangent residual of F at $z \in \mathcal{Z}$ is defined as $r^{tan}_{(F,\mathcal{Z})}(z) := \|\Pi_{J_{\mathcal{Z}}(z)}[z - F(z)] - z\|$. An equivalent definition is $r^{tan}_{(F,\mathcal{Z})}(z) := \sqrt{\|F(z)\|^2 - \max_{\substack{a \in \widehat{N}_{\mathcal{Z}}(z), \\ \langle F(z), a\rangle \leq 0}} \langle a, F(z)\rangle^2}$.*

Remark 2. *We show the equivalence of the two definitions of tangent residual in Lemma 4 in Appendix D. We may use either of them depending on which one is more convenient.*

When the convex set $\mathcal{Z}$ and the operator F are clear from context, we are going to omit the subscript and denote the unit normal cone as $\widehat{N}(z) = \widehat{N}_{\mathcal{Z}}(z)$ and the tangent residual as $r^{tan}(z) = r^{tan}_{(F,\mathcal{Z})}(z)$. Although the definition is slightly technical, one can think of the tangent residual as the norm of another operator $\widehat{F}$, which is F projected to all directions that are not "blocked" by the boundary of $\mathcal{Z}$ if one takes an infinitesimally small step $\epsilon \cdot F(z)$, which is the same as projecting F to $J_{\mathcal{Z}}(z)$. Intuitively, if the tangent residual is small, then the next iterate will not be far away from the current one.

Tangent Residual and Its Connection with Games. Given a monotone game $\mathcal{G}$, we denote by $r^{nat}_{\mathcal{G}}(z) = r^{nat}_{(F_{\mathcal{G}},\mathcal{Z})}(z)$ and by $r^{tan}_{\mathcal{G}}(z) = r^{tan}_{(F_{\mathcal{G}},\mathcal{Z})}(z)$. In the next lemma, we argue that a small tangent residual implies a small gap and total gap, hence an it suffices to show that the last-iterate has a small tangent residual. The proof is postponed to Appendix E.

Lemma 2. *[Adapted from Theorem 10 in [GPDO20] and Proposition 2 in [GPD20]] Let $\mathcal{G} = (\mathcal{N}, \{\mathcal{Z}^{(i)}\}_{i\in\mathcal{N}}, \{f^{(i)}\}_{i\in[N]})$ be a monotone game where $\{\mathcal{Z}^{(i)}\}_{i\in\mathcal{N}}$ are closed convex sets. For $z \in \mathcal{Z}$, we have $\textsc{Gap}_{\mathcal{G},D}(z) \leq D \cdot r^{tan}_{\mathcal{G}}(z)$ and $\textsc{TGap}_{\mathcal{G},D}(z) \leq \sqrt{N} \cdot D \cdot r^{tan}_{\mathcal{G}}(z)$.*

4 Last-Iterate Convergence Rate for EG and OG

We prove the last-iterate convergence rate for EG and OG. We first describe our proof plan.

Proof Plan. Our analyses for both algorithms follow the same three-step procedure: (i) define a potential function that measures the proximity to a Nash equilibrium of the current iterate; (ii) prove a best-iterate convergence rate, that is, show that in T steps, there exists one iterate whose potential function is small; (iii) show that the potential function is non-increasing, so the last-iterate is the best-iterate, and the best-iterate convergence rate becomes the last-iterate convergence rate.

The first major challenge we face is to choose the appropriate potential functions. In the unconstrained case, the central quantity is the norm of the operator $F_{\mathcal{G}}$. The key component of the analyses [LZMJ20, GPDO20, GPD20, GLG21] is to establish that the norm of the operator $F_{\mathcal{G}}$ at the last iterate (also the T-th iterate) is upper bounded by $O(\frac{1}{\sqrt{T}})$, which implies a $O(\frac{1}{\sqrt{T}})$ last-iterate convergence rate for the gap function. In the constrained setting,

the norm of the operator is a poor choice to measure the proximity to a Nash equilibrium, as it can be far away from 0 even at a Nash equilibrium.

Potential function for EG. We use the **tangent residual** as the potential function for EG. Our starting point is to find a suitable generalization for the norm of the operator in the constrained setting. A standard generalization is the natural residual (Definition 6 in Appendix D), which takes the constraints into account and is guaranteed to converge to 0 at the Nash equilibrium. Unfortunately, we observe that the natural residual is *not monotonically decreasing* even in basic bilinear games, making it difficult to directly analyze. Similar non-monotonicity has been observed for several other natural performance measures such as $\|z_k - z_{k+1/2}\|$,[5] $\|z_k - z_{k+1}\|$, $\max_{z\in\mathcal{Z}} \langle F(z), z_k - z\rangle$ and $\max_{z\in\mathcal{Z}} \langle F(z_k), z_k - z\rangle$, leaving all these quantities unsuitable as a potential function. See more discussion in Appendix J. Indeed, tangent residual is the only natural generalization of the norm of the operator that is always monotone in our numerical experiments.

Potential function for OG. We choose the potential function $\Phi(z_k, w_k) = \|F(z_k) - F(w_k)\|^2 + r^{tan}(z_k)^2$ in our analysis for OG. The potential function can be interpreted as the tangent residual ($r^{tan}(z_k)^2$) and an extra correction term $\|F(z_k) - F(w_k)\|^2$. The potential function is discovered via SOS programming, with more details in Section 4.2.

4.1 Best-Iterate Convergence

En route to establish the last-iterate convergence of EG and OG, we first show a weaker guarantee known as the best-iterate convergence. In Lemma 3 we show that when all the players use EG for $2 \cdot T$ steps (OG for T steps), then there exists $t^* \in [T]$ where $r^{tan}(z_{t^*+1}) \le O(\frac{1}{\sqrt{T}})$ ($\Phi(z_{t^*}, w_{t^*}) \le O(\frac{1}{T})$ resp.). The proof for EG builds upon the proof for the best-iterate w.r.t. $\|z_k - z_{k+\frac{1}{2}}\|$ by [Kor76, FP07], while the proof for OG builds upon the best-iterate proof for $\|z_k - w_{k+1}\|$ by [WLZL21a, HIMM19]. The proof of Lemma 3 is in Appendix F.

Lemma 3. *Let $\mathcal{G} = (\mathcal{N}, \{\mathcal{Z}^{(i)}\}_{i\in\mathcal{N}}, \{f^{(i)}\}_{i\in[N]})$ be an L-smooth and monotone game where $\{\mathcal{Z}^{(i)}\}_{i\in\mathcal{N}}$ are closed convex sets and let z^* be a Nash Equilibrium of $\mathcal{G}$. Assume that all the players update their actions using the EG algorithm with arbitrary starting action profile z_0 and step-size $\eta \in (0, \frac{1}{L})$. Then for and any $T > 0$, there exists $t^* \in [T]$ such that:*

$$r^{tan}(z_{t^*+1}) \le \frac{1 + \eta L + (\eta L)^2}{\eta} \frac{1}{\sqrt{T}} \frac{\|z_0 - z^*\|}{\sqrt{1 - (\eta L)^2}}.$$

Assume that all the players update their actions using the OG algorithm with arbitrary starting ation profiles z_0,w_0 and step-size $\eta \in (0, \frac{1}{2L})$. Then for and any $T > 0$, there exists $t^* \in [T]$ such that:

$$\Phi(z_{t^*}, w_{t^*}) \le \frac{1}{\eta^2 \cdot T}\left(\frac{4 + 6\eta^4 L^4}{1 - 4\eta^2 L^2}\|z_0 - z^*\|^2 + \frac{16\eta^2 L^2 + 6\eta^4 L^4}{1 - 4\eta^2 L^2}\|w_0 - z_0\|^2\right).$$

4.2 Monotonicity of the Potentials

In this section, we prove that $r^{tan}(z_k)$ ($\Phi(z_k, w_k)$ resp.) is non-increasing across iterates of EG (OG resp.), which, in combination with Lemma 3, implies the last-iterate convergence rate of smooth monotone games when all players update their actions using EG (OG resp.).

SOS Programming. Suppose we want to prove that a polynomial $g(\boldsymbol{x}) \in \mathbb{R}[\boldsymbol{x}]$ is non-negative over a semialgebraic set $\mathcal{S} = \{\boldsymbol{x} : g_i(\boldsymbol{x}) \le 0, \forall i \in [M], h_i(\boldsymbol{x}) = 0, \forall i \in [N]\}$, where each $g_i(\boldsymbol{x})$ ($h_i(\boldsymbol{x})$ resp.) is also a polynomial. One way is to construct a *certificate of non-negativity*, by providing a set of nonnegative coefficients $\{p_i \ge 0\}_{i\in[M]}$ and $\{q_i\}_{i\in[N]}$ such that $g(\boldsymbol{x}) + \sum_{i\in[M]} p_i \cdot g_i(\boldsymbol{x}) + \sum_{i\in[N]} q_i \cdot h_i(\boldsymbol{x})$ is a SOS polynomial. Surprisingly, if $g(\boldsymbol{x})$ is indeed non-negative over $\mathcal{S}$, a certificate of non-negativity always exists as guaranteed by a foundational result in real algebraic geometry – the Krivine-Stengle Positivestellensatz [Kri64, Ste74], a generalization of Artin's resolution of Hilbert's 17th problem [Art27].

[5] $\|z_k - z_{k+1/2}\|$ is proportional to the norm of the operator mapping introduced in [Dia20].

The SOS programming [Nes00, Par00, Par03, Las01, Lau09] is a systematic way to search for such a certificate using semidefinite programming. See Appendix B for details.

Our approach is to apply SOS programming to search for a certificate of non-negativity for $r^{tan}(z_k)^2 - r^{tan}(z_{k+1})^2$ and $\Phi(z_k, w_k) - \Phi(z_{k+1}, w_{k+1})$ for every k, over a semialgebraic set defined by the update rule of the corresponding algorithm and set of constraints $\mathcal{Z}$.

Two important challenges with this approach is that the number of variables depends on the dimension of $\mathcal{Z}$ and there are infinitely many constraints associated with the problem (e.g. the set $\{\langle F(z) - F(z'), z - z'\rangle \geq 0, \forall z, z' \in \mathcal{Z}\}$). We provide a detailed exposition for the reduction of the number of constraints and the efficient formulation into a SOS program in Appendix G.3, where we prove the monotonicity of our potential function for EG in both the unconstrained and constrained settings.

Searching for the potential functions. The potential function of our analysis of OG is directly discovered using an SOS program. The program is formulated by searching over linear combinations of $||F(z_k)||^2 - ||F(z_{k+1})||^2$,$||F(w_k)||^2 - ||F(w_{k+1})||^2$,$\langle F(z_k), F(w_k)\rangle - \langle F(z_{k+1}), F(w_{k+1})\rangle$ and $r^{tan}(z_k) - r^{tan}(z_{k+1})$, under (i) the constraint that the linear combination is non-increasing,[6] and (ii) the constraints induced by properties of the operator $F(\cdot)$, the update rule of OG and the set $\mathcal{Z}$ (See Appendix G.3 for a demonstration of the induced semialgebraic set of EG algorithm). We then use the linear combination output by the SOS program as the potential function in our analysis. We believe our heuristic for finding a potential function could be useful in other settings. In general, one can first choose a collection of basis functions that may be part of a potential function, then use SOS programming to search over all linear combinations of the basis functions subject to the constraint that the linear combination is non-negative to discover the potential function.

For technical convenience, we state Theorem 1 and Theorem 2 for general monotone operators over convex sets rather than games. We turn our attention back to games in Section 4.3 where we provide formal last-iterate convergence guarantees for smooth monotone games.

4.2.1 Monotonicity of $\Phi(z_k, w_k)$ for OG in the Unconstrained Setting

To better illustrate our approach, we prove in Theorem 1 the monotonicity of $\Phi(z_k, w_k)$ for the OG algorithm in the unconstrained setting. Note that this is already a strengthening of [GPD20], where we provide a simple potential function that can be used to directly argue that $\|F(z_k)\|^2 = O(\frac{1}{k})$ without making any second-order smoothness assumption.

Theorem 1. *Let $F : \mathcal{Z} \to \mathbb{R}^n$ be a monotone and L-Lipschitz operator. For any $z_k, w_k \in \mathbb{R}^n$, the unconstrained OG algorithm with step-size $\eta \in (0, \frac{1}{2L})$ satisfies $\Phi(z_k, w_k) \geq \Phi(z_{k+1}, w_{k+1})$.* [7]

Proof. Since F is monotone and L-Lipschitz, we have $\langle F(z_{k+1}) - F(z_k), z_k - z_{k+1}\rangle \leq 0$ and $\|F(w_{k+1}) - F(z_{k+1})\|^2 - L^2\|w_{k+1} - z_{k+1}\|^2 \leq 0$. We simplify them using the update rule of OG and $\eta^2 L^2 < \frac{1}{4}$. In particular, we replace $z_k - z_{k+1}$ by $\eta F(w_{k+1})$ and $w_{k+1} - z_{k+1}$ with $\eta F(w_{k+1}) - \eta F(w_k)$.

$$\langle F(z_{k+1}) - F(z_k), F(w_{k+1})\rangle \leq 0, \tag{3}$$

$$\|F(w_{k+1}) - F(z_{k+1})\|^2 - \frac{1}{4}\|F(w_{k+1}) - F(w_k)\|^2 \leq 0. \tag{4}$$

It is not hard to verify the following identity:

$$\|F(z_k) - F(w_k)\|^2 + \|F(z_k)\|^2 - \|F(z_{k+1}) - F(w_{k+1})\|^2 - \|F(z_{k+1})\|^2$$
$$+ 2 \cdot \text{LHS of Inequality}(3) + 2 \cdot \text{LHS of Inequality}(4) = \frac{1}{2}\|F(w_k) + F(w_{k+1}) - 2F(z_k))\|^2.$$

Thus, $\|F(z_k) - F(w_k)\|^2 + \|F(z_k)\|^2 \geq \|F(z_{k+1}) - F(w_{k+1})\|^2 + \|F(z_{k+1})\|^2$. □

[6] To avoid finding the trivial linear combination, i.e., all coefficients equal to 0, we also use the objective function in the SOS program to encourage a non-trivial solution if one exists by, for example, maximizing the sum of the coefficients of the linear combination.

[7] In the unconstrained setting, $r^{tan}_{(F,\mathbb{R}^n)}(z) = \|F(z)\|^2$.

4.2.2 Constrained EG and OG

In the constrained setting, the approach is more complicated due to the projections and we postpone any exposition to Appendix G.3. We state the results in the following theorem.

Theorem 2. *Let $\mathcal{Z} \subseteq \mathbb{R}^n$ be a closed convex set and $F : \mathcal{Z} \to \mathbb{R}^n$ be a monotone and L-Lipschitz operator. For any $z_k \in \mathcal{Z}$, the EG algorithm with step-size $\eta \in (0, \frac{1}{L})$ satisfies $r^{tan}(z_k) \geq r^{tan}(z_{k+1})$. Moereover, for any $z_k, w_k \in \mathcal{Z}$, the OG algorithm with step-size $\eta \in (0, \frac{1}{2L})$ satisfies $\Phi(z_k, w_k) \geq \Phi(z_{k+1}, w_{k+1})$.*

4.3 Last-Iterate Convergence of EG and OG

In this section, we formally combine Lemma 3 and Theorem 2 to show the last-iterate convergence of EG and OG with respect to the tangent residual, gap function, and the total gap function. Recall that when all the players update their actions using the EG algorithm, then z_k is the action profile played at day $2 \cdot k$ and $z_{k+\frac{1}{2}}$ is the action played at day $2 \cdot k + 1$, while when all the players update their action profile using the OG algorithm, then w_k is the action profile they play at day k. The formal proof of Theorem 3 is postponed at Appendix F.

Theorem 3. *Let $\mathcal{G} = (\mathcal{N}, \{\mathcal{Z}^{(i)}\}_{i \in \mathcal{N}}, \{f^{(i)}\}_{i \in [N]})$ be an L-smooth and monotone game where $\{\mathcal{Z}^{(i)}\}_{i \in \mathcal{N}}$ are closed convex sets and let z^* be a Nash Equilibrium of $\mathcal{G}$. Assume that all the players update their actions using the EG algorithm with arbitrary starting action profile z_0 and step-size $\eta \in (0, \frac{1}{L})$. Let $D_0 = \frac{3||z_0 - z^*||}{\sqrt{1-(\eta L)^2}}$, then for and any $T > 0$ and $D > 0$,*

$$\max\left(r^{tan}(z_T), \frac{\textsc{Gap}(z_T)}{D}, \frac{\textsc{TGap}(z_T)}{\sqrt{N} \cdot D}\right) \leq \frac{D_0}{\eta\sqrt{T}},$$
$$\max\left(r^{tan}(z_{T+\frac{1}{2}}), \frac{\textsc{Gap}(z_{T+\frac{1}{2}})}{D}, \frac{\textsc{TGap}(z_{T+\frac{1}{2}})}{\sqrt{N} \cdot D}\right) \leq \frac{(1+\eta L) \cdot D_0}{\eta\sqrt{T}}.$$

When all the players update their actions using the OG algorithm with arbitrary starting action profiles z_0, w_0 and step-size $\eta \in (0, \frac{1}{2L})$. Denote $D_0 := \frac{\sqrt{2}(2+\eta L)}{\sqrt{1-4\cdot(\eta L)^2}} \cdot \sqrt{(4+6\eta^4 L^4)\|z_0 - z^*\|^2 + (16\eta^2 L^2 + 6\eta^4 L^4)\|w_0 - z_0\|^2}$. Then for any $T > 0$ and $D > 0$,

$$\max\left(r^{tan}(w_{T+1}), \frac{\textsc{Gap}(w_{T+1})}{D}, \frac{\textsc{TGap}(w_{T+1})}{\sqrt{N} \cdot D}\right) \leq \frac{D_0}{\eta\sqrt{T}}.$$

For EG, choosing η to be $\frac{1}{2L}$ and $D = O(\|z_0 - z^*\|)$, then $\textsc{Gap}(z_T) = O(\frac{D^2 L}{\sqrt{T}})$; For OG, choosing η to be $\frac{1}{3L}$ and $D = O(\max\{\|z_0 - z^*\|, \|w_0 - z_0\|\})$, then $\textsc{Gap}(w_{T+1}) = O(\frac{D^2 L}{\sqrt{T}})$. Both upper bounds matches the $\Omega(\frac{D^2 L}{\sqrt{T}})$ lower bound for EG, OG, and more generally all p-SCLI algorithms [GPDO20, GPD20] in terms of the dependence on D, L, and T.

In Appendix I we show that EG and OG algorithm also have last-iterate convergence for monotone variational inequalities with respect to several standard performance measures (See Appendix D for the formal definition of monotone variational inequality).

5 Conclusion

We provide the first and tight last-iterate convergence rate of the EG and OG algorithm for smooth monotone games and Lipschitzs and monotone VIs. Our proof is based on the tangent residual, which is a new proximity measure to a Nash equilibrium, and an accompanying SOS-based analysis. We believe our techniques may be useful in the study of last-iterate convergence for other algorithms.

References

[Art27] Emil Artin. Über die zerlegung definiter funktionen in quadrate. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 5, pages 100–115. Springer, 1927. Issue: 1.

[BF87] LM Bregman and IN Fokin. Methods of Determining Equilibrium Situations in Zero-Sum Polymatrix Games. *Optimizatsia*, 40(57):70–82, 1987.

[Bro65] Felix E. Browder. Nonlinear monotone operators and convex sets in Banach spaces. *Bulletin of the American Mathematical Society*, 71(5):780–785, 1965.

[BS68] Haïm Brezis and Moïse Sibony. Méthodes d'approximation et d'itération pour les opérateurs monotones. *Archive for Rational Mechanics and Analysis*, 28(1):59–82, January 1968.

[CBL06] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[CD11] Yang Cai and Constantinos Daskalakis. On Minmax Theorems for Multiplayer Games. In *Proceedings of the 2011 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, pages 217–234. Society for Industrial and Applied Mathematics, January 2011.

[DGP09] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[Dia20] Jelena Diakonikolas. Halpern Iteration for Near-Optimal and Parameter-Free Monotone Inclusion and Strong Solutions to Variational Inequalities. *arXiv:2002.08872 [cs, math]*, April 2020. arXiv: 2002.08872.

[DISZ17] Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. *arXiv preprint arXiv:1711.00141*, 2017.

[DP09] Constantinos Daskalakis and Christos H. Papadimitriou. On a Network Generalization of the Minmax Theorem. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris Nikoletseas, and Wolfgang Thomas, editors, *Automata, Languages and Programming*, Lecture Notes in Computer Science, pages 423–434, Berlin, Heidelberg, 2009. Springer.

[DP18] Constantinos Daskalakis and Ioannis Panageas. Last-Iterate Convergence: Zero-Sum Games and Constrained Min-Max Optimization. page 18 pages, 2018. Artwork Size: 18 pages Medium: application/pdf Publisher: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany Version Number: 1.0.

[DT14] Yoel Drori and Marc Teboulle. Performance of first-order methods for smooth convex minimization: a novel approach. *Mathematical Programming*, 145(1):451–482, 2014. Publisher: Springer.

[DTdB21] Radu-Alexandru Dragomir, Adrien B. Taylor, Alexandre d'Aspremont, and Jérôme Bolte. Optimal complexity and certification of bregman first-order methods. *Mathematical Programming*, 2021.

[EDMN09] Eyal Even-Dar, Yishay Mansour, and Uri Nadav. On the convergence of regret minimization dynamics in concave games. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 523–532, 2009.

[FMP18] Mahyar Fazlyab, Manfred Morari, and Victor M Preciado. Design of first-order optimization algorithms via sum-of-squares programming. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4445–4452. IEEE, 2018.

[FP07] Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.

[GBV+18] Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. *arXiv preprint arXiv:1802.10551*, 2018.

[GLG21] Eduard Gorbunov, Nicolas Loizou, and Gauthier Gidel. Extragradient Method: $O(1/K)$ Last-Iterate Convergence for Monotone Variational Inequalities and Connections With Cocoercivity. *arXiv:2110.04261 [cs, math]*, October 2021. arXiv: 2110.04261.

[GMG+22] Baptiste Goujaud, Céline Moucer, François Glineur, Julien M. Hendrickx, Adrien B. Taylor, and Aymeric Dieuleveut. Pepit: computer-assisted worst-case analyses of first-order optimization methods in python. *CoRR*, abs/2201.04040, 2022.

[GPD20] Noah Golowich, Sarath Pattathil, and Constantinos Daskalakis. Tight last-iterate convergence rates for no-regret learning in multi-player games. *arXiv:2010.13724 [cs, math]*, October 2020. arXiv: 2010.13724.

[GPDO20] Noah Golowich, Sarath Pattathil, Constantinos Daskalakis, and Asuman Ozdaglar. Last Iterate is Slower than Averaged Iterate in Smooth Convex-Concave Saddle Point Problems. *arXiv:2002.00057 [cs, math, stat]*, July 2020. arXiv: 2002.00057.

[HAM21] Yu-Guan Hsieh, Kimon Antonakopoulos, and Panayotis Mertikopoulos. Adaptive Learning in Continuous Games: Optimal Regret Bounds and Convergence to Nash Equilibrium. *arXiv:2104.12761 [cs, math]*, April 2021. arXiv: 2104.12761.

[HIMM19] Yu-Guan Hsieh, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. On the convergence of single-call stochastic extra-gradient methods. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[HS66] Philip Hartman and Guido Stampacchia. On some non-linear elliptic differential-functional equations. *Acta Mathematica*, 115(none):271–310, January 1966. Publisher: Institut Mittag-Leffler.

[Kim21] Donghwan Kim. Accelerated proximal point method for maximally monotone operators. *Mathematical Programming*, 190(1):57–87, November 2021.

[Kor76] G. M. Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.

[Kri64] Jean-Louis Krivine. Anneaux préordonnés. *Journal d'analyse mathématique*, 12:p. 307–326, 1964. Publisher: Springer.

[Las01] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on optimization*, 11(3):796–817, 2001. Publisher: SIAM.

[Lau09] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270. Springer, 2009.

[LK21] Sucheol Lee and Donghwan Kim. Fast extra gradient methods for smooth structured nonconvex-nonconcave minimax problems. In *Annual Conference on Neural Information Processing Systems*, 2021.

[LNPW20] Qi Lei, Sai Ganesh Nagarajan, Ioannis Panageas, and Xiao Wang. Last iterate convergence in no-regret learning: constrained min-max optimization for convex-concave landscapes. *arXiv:2002.06768 [cs, stat]*, February 2020. arXiv: 2002.06768.

[LRP16] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and Design of Optimization Algorithms via Integral Quadratic Constraints. *SIAM Journal on Optimization*, 26(1):57–95, January 2016. arXiv: 1408.3595.

[LS67] J. L. Lions and G. Stampacchia. Variational inequalities. *Communications on Pure and Applied Mathematics*, 20(3):493–519, 1967. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpa.3160200302.

[LS19] Tengyuan Liang and James Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 907–915. PMLR, 2019.

[LZMJ20] Tianyi Lin, Zhengyuan Zhou, Panayotis Mertikopoulos, and Michael Jordan. Finite-Time Last-Iterate Convergence for Multi-Agent Learning in Games. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6161–6171. PMLR, November 2020. ISSN: 2640-3498.

[Mal15] Yu Malitsky. Projected Reflected Gradient Methods for Monotone Variational Inequalities. *SIAM Journal on Optimization*, 25(1):502–520, January 2015. arXiv: 1502.04968.

[MOP19] Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A Unified Analysis of Extra-gradient and Optimistic Gradient Methods for Saddle Point Problems: Proximal Point Approach. *arXiv:1901.08511 [cs, math, stat]*, September 2019. arXiv: 1901.08511.

[MOP20] Aryan Mokhtari, Asuman E. Ozdaglar, and Sarath Pattathil. Convergence rate of (1/k) for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM J. Optim.*, 30(4):3230–3251, 2020.

[MPP18] Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2703–2717. SIAM, 2018.

[MZ19] Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Math. Program.*, 173(1-2):465–507, 2019.

[Nem04] Arkadi Nemirovski. Prox-method with rate of convergence O (1/t) for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004. Publisher: SIAM.

[Nes00] Yurii Nesterov. Squared functional systems and optimization problems. In *High performance optimization*, pages 405–440. Springer, 2000.

[Par00] Pablo A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology, 2000.

[Par03] Pablo A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2):293–320, May 2003.

[PDZC20] Wei Peng, Yu-Hong Dai, Hui Zhang, and Lizhi Cheng. Training GANs with centripetal acceleration. *Optimization Methods and Software*, 35(5):955–973, 2020. Publisher: Taylor & Francis.

[Pop80] Leonid Denisovich Popov. A modification of the Arrow-Hurwicz method for search of saddle points. *Mathematical notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980. Publisher: Springer.

[Ros65] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965.

[RTBG20] Ernest K. Ryu, Adrien B. Taylor, Carolina Bergeling, and Pontus Giselsson. Operator splitting performance estimation: Tight contraction factors and optimal parameter selection. *SIAM Journal on Optimization*, 30(3):2251–2271, 2020.

[Sib70] Moïse Sibony. Méthodes itératives pour les équations et inéquations aux dérivées partielles non linéaires de type monotone. *CALCOLO*, 7(1):65–183, March 1970.

[Ste74] Gilbert Stengle. A nullstellensatz and a positivstellensatz in semialgebraic geometry. *Mathematische Annalen*, 207(2):87–97, June 1974.

[THG17a] Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Exact Worst-case Performance of First-order Methods for Composite Convex Optimization. *SIAM Journal on Optimization*, 27(3):1283–1313, January 2017. arXiv: 1512.07516.

[THG17b] Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Performance estimation toolbox (PESTO): automated worst-case analysis of first-order optimization methods. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1278–1283. IEEE, 2017.

[Tse95] Paul Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1):237–252, June 1995.

[TVT21] Sandra SY Tan, Antonios Varvitsiotis, and Vincent YF Tan. Analysis of optimization algorithms via sum-of-squares. *Journal of Optimization Theory and Applications*, 190(1):56–81, 2021.

[WLZL21a] Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Last-iterate Convergence of Decentralized Optimistic Gradient Descent/Ascent in Infinite-horizon Competitive Markov Games. *arXiv:2102.04540 [cs, stat]*, July 2021. arXiv: 2102.04540.

[WLZL21b] Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear Last-iterate Convergence in Constrained Saddle-point Optimization. *arXiv:2006.09517 [cs, stat]*, March 2021. arXiv: 2006.09517.

[YR21] TaeHo Yoon and Ernest K. Ryu. Accelerated Algorithms for Smooth Convex-Concave Minimax Problems with $\mathcal{O}(1/k^2)$ Rate on Squared Gradient Norm. *arXiv:2102.07922 [math]*, June 2021. arXiv: 2102.07922.

[ZBLG21] Guodong Zhang, Xuchan Bao, Laurent Lessard, and Roger B. Grosse. A unified analysis of first-order methods for smooth games via integral quadratic constraints. *Journal of Machine Learning Research*, 22:103:1–103:39, 2021.

[ZMA+18] Zhengyuan Zhou, Panayotis Mertikopoulos, Susan Athey, Nicholas Bambos, Peter W. Glynn, and Yinyu Ye. Learning in games with lossy feedback. In *Annual Conference on Neural Information Processing Systems*, pages 5140–5150, 2018.

[ZMB+17] Zhengyuan Zhou, Panayotis Mertikopoulos, Nicholas Bambos, Peter W Glynn, and Claire Tomlin. Countering Feedback Delays in Multi-Agent Learning. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[ZMM+17] Zhengyuan Zhou, Panayotis Mertikopoulos, Aris L. Moustakas, Nicholas Bambos, and Peter Glynn. Mirror descent learning in continuous games. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017.

[ZMM+21] Zhengyuan Zhou, Panayotis Mertikopoulos, Aris L. Moustakas, Nicholas Bambos, and Peter Glynn. Robust Power Management via Learning and Game Design. *Operations Research*, 69(1):331–345, January 2021.

[ZY19] Guojun Zhang and Yaoliang Yu. Convergence of gradient methods on bilinear zero-sum games. *arXiv preprint arXiv:1908.05699*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Our results are theoretical. We state the assumptions of our results in the statements of the lemmas.
   (c) Did you discuss any potential negative societal impacts of your work? [No] Societal Impacts are discussed in Appendix A.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] All assumptions are clearly stated in the statement of the theorems (e.g. see Theorem 3).
   (b) Did you include complete proofs of all theoretical results? [Yes] We include proof for all theoretical results either in the main body or the appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]