Global Convergence to Local Minmax Equilibrium in Classes of Nonconvex Zero-Sum Games

Anonymous Author(s)
Affiliation
Address
email

Abstract

We study gradient descent-ascent learning dynamics with timescale separation in unconstrained continuous action zero-sum games where the minimizing player faces a nonconvex optimization problem and the maximizing player optimizes a Polyak-Łojasiewicz (PŁ) or strongly-concave (SC) objective. In contrast to past work, we assess convergence in relation to game-theoretic equilibria instead of only notions of stationarity. In pursuit of this goal, we prove that the only locally stable points of the continuous-time limiting system correspond to strict local minmax equilibria in each class of games. For the class of nonconvex-PŁ games, we exploit timescale separation to construct a potential function that when combined with the stability characterization and an asymptotic saddle avoidance result gives a global asymptotic almost-sure convergence guarantee to the set of strict local minmax equilibria. For the class of nonconvex-SC games, we show the surprising property that the function of the game can be made a potential with timescale separation. Combining this insight with the stability characterization allows us to generalize methods for efficiently escaping saddle points in nonconvex optimization to obtain a global finite-time convergence guarantee to local minmax equilibria.

1 Introduction

We study continuous action zero-sum games of the form

$$\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} f(x,y)$$

where $f \in C^2(\mathcal{X}\times\mathcal{Y},\mathbb{R})$ and $\mathcal{X}=\mathbb{R}^{d_1}$ and $\mathcal{Y}=\mathbb{R}^{d_2}$ denote the individual action spaces and $d=d_1+d_2$. In particular, we focus on zero-sum games in which $f(\cdot,y)$ is potentially nonconvex in $x\in\mathcal{X}$ and $f(x,\cdot)$ satisfies the Polyak-Łojasiewicz (PŁ) condition or is strongly-concave (SC) in $y\in\mathcal{Y}$. We refer to these classes of games as nonconvex-PŁ and nonconvex-SC zero-sum games.

This general formulation has a broad spectrum of applications such as fair classification [41], distributionally robust optimization [40, 43], and adversarial learning [34]. Consequently, there has been a surge of interest in recent years toward developing methods for solving these problems efficiently. So far, existing work on gradient-based learning in nonconvex-PŁ/SC zero-sum games has exclusively focused on providing global convergence rates to approximate stationary points with no attention given to the characterization in terms of game-theoretic equilibrium concepts [30, 31, 33, 41, 43, 56].

In contrast, a common theme in the study of general nonconvex-nonconcave zero-sum games is to assess the *types* of stationary points an algorithm locally converges toward in terms of their higher order structure [9, 14, 15, 21, 36, 37, 52]. The purpose of this analysis is generally to determine whether commonly deployed algorithms can guarantee local convergence to *only* game-theoretic equilibria or to design algorithms that achieve this objective.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

The goal of this paper is to close the gap between the two problem classes and determine whether gradient-based learning algorithms in nonconvex-PŁ/SC zero-sum games can be shown to globally converge to *only* game-theoretically meaningful equilibria (local minmax or Nash equilibria). We focus our attention on the canonical gradient descent-ascent learning dynamics with timescale separations between players and a stochastically perturbed variant. In these algorithms, timescale separation is manifested in the different (yet constant) learning rates of the maximizing and minimizing players. The descriptions of these systems, which we refer to as τ-GDA and τ-PGDA where $\tau > 0$ parameterizes the timescale separation between players, are provided in Algorithm 1 and Algorithm 2, respectively. Simply put, τ-GDA (Algorithm 1) corresponds to each player following their individual gradient in a noiseless setting, while τ-PGDA (Algorithm 2) describes each player following their individual (potentially stochastic) gradient with artificial noise injections.

1.1 Contributions.

We show that τ-GDA and τ-PGDA have *global* convergence guarantees in nonconvex-PŁ/SC zero-sum games to the natural game-theoretic solution concept for this problem class of strict local minmax equilibria.[1] The contributions of this work are summarized as follows.

1) In Theorem 1, we prove the only critical points that are locally stable with respect to the τ-GDA continuous-time limiting system are strict local minmax in nonconvex-PŁ/SC zero-sum games.

2) In Theorem 2, we combine Theorem 1 with a potential function construction and an asymptotic saddle avoidance result to prove that τ-GDA (Algorithm 1) globally asymptotically converges to strict local minmax equilibria almost surely in nonconvex-PŁ/SC zero-sum games. In Corollary 1, we show a corresponding local convergence rate of $\widetilde{\mathcal{O}}(\varepsilon^{-2})$ to an ε–strict local minmax equilibria.

3) In Theorem 3, for the class of nonconvex-SC zero-sum games, we provide high-probability finite-time rates showing that τ-PGDA (Algorithm 2) globally converges to ε–local minmax equilibria with a complexities of $\widetilde{\mathcal{O}}(\varepsilon^{-4})$ and $\widetilde{\mathcal{O}}(\varepsilon^{-2})$ in stochastic and deterministic problems, respectively.

To our knowledge, our results provide the broadest existing global convergence guarantee for gradient-based algorithms to *game-theoretically meaningful* equilibria in zero-sum continuous games.

1.2 Practical Motivation.

The study of nonconvex-PŁ/SC zero-sum games has often been motivated by machine learning problems formulated as games. We remark that given the problem formulations, it is natural to seek notions of minmax equilibria as solutions. Indeed, notions of stationarity are not guaranteed to reflect a meaningful solution since they could even correspond to maxmin solutions. Consequently, it is critical to give convergence guarantees to minmax equilibrium as we pursue in this work. We now provide examples of machine learning applications that belong to the class of games we consider.

Example 1. *In the fair classification problem [41], the objective is to minimize the maximum loss over multiple categories. An example formulation is*

$$\min_W \max_{i=1,\dots,m}\{\mathcal{L}_i(W)\} \tag{1}$$

where $\mathcal{L}_i$ represents the loss on category i and W denotes parameter weights of a neural network.

Example 2. *To train a robust neural network against adversarial attacks, a common approach is to formulate training as a robust min-max optimization problem of the form*

$$\min_w \sum_{i=1}^N \max_{j=1,\dots,m}\{\ell(f(\hat{x}_{ij}(w)), y_i)\} \tag{2}$$

where w is the parameter vector of the neural network and each $\hat{x}_{ij}(w)$ is the result of a targeted attack on the sample x_i seeking to change the output of the network for label j [34, 41].

Example 3. *Distributionally robust optimization and robust learning from multiple distributions is commonly formulated as a minmax optimization problem of the form*

$$\min_{x\in\mathbb{R}^d} \max_{y\in\mathcal{Y}} \sum_{i=1}^n y_i f(x) - r(y) \tag{3}$$

where $f_i(x)$ is the loss of a model x on the i-th data point, $\mathcal{Y}$ is the simplex in $\mathbb{R}^n$, and $r(y)$ is carefully selected regularizer [34, 40, 43].

A common approach in each problem is to transform the inner maximization problem through relaxations and regularization methods to give an unconstrained problem in the parameter space.

[1] Strict local minmax equilibria are also known as differential Stackelberg equilibria in the literature [14, 15] and we use the terms interchangeably in this work.

Table 1: The gradient complexity of gradient descent (GD), gradient descent-ascent with timescale separation (τ-GDA), and alternating (including multi-step) gradient descent-ascent (AGDA) or perturbed variants (PGD, τ-PGDA) in deterministic and stochastic nonconvex optimization, nonconvex-PŁ zero-sum games, and nonconvex-SC zero-sum games. We state the complexity in terms of the ε tolerance of the guarantee and the dimension d with the notation $\widetilde{\mathcal{O}}(\cdot)$ hiding logarithmic factors in ε and d.

Problem	Algorithm & Reference	Complexity		Guarantee
		Deterministic	Stochastic	
Nonconvex Optimization	GD [20]	$\widetilde{\mathcal{O}}(\varepsilon^{-2})$	$\widetilde{\mathcal{O}}(\varepsilon^{-4})$	ε–Stationarity
	PGDA [20]	$\widetilde{\mathcal{O}}(\varepsilon^{-2})$	$\widetilde{\mathcal{O}}(\varepsilon^{-4})$	ε–Local Min
Nonconvex SC Zero-Sum	τ-GDA, τ-AGDA [30]	$\widetilde{\mathcal{O}}(\varepsilon^{-2})$	$\widetilde{\mathcal{O}}(\varepsilon^{-4})$	ε–Stationarity
	Theorem 3 (τ-PGDA)	$\widetilde{\mathcal{O}}(\varepsilon^{-2})$	$\widetilde{\mathcal{O}}(\varepsilon^{-4})$	ε–**DSE**
Nonconvex PŁ Zero-Sum	AGDA [41], [56, Appendix D]	$\widetilde{\mathcal{O}}(\varepsilon^{-2})$	–	ε–Stationarity
	Theorem 2 (τ-GDA)	Asymptotic	–	**DSE**

2 Related Work

We now cover the most relevant related work with further discussion provided in Appendix A.

Nonconvex-Nonconcave Zero-Sum Games. A common theme in analyzing gradient descent-ascent with or without timescale separation in nonconvex-nonconcave zero-sum games has been to assess the local stability around critical points of the continuous-time limiting system and draw connections to the differential Nash and Stackelberg equilibrium notions (see Definition 4) [9, 14, 15, 21, 36, 38, 39, 57]. Importantly, unless the timescale separation is chosen very carefully, the stable critical points of gradient descent-ascent may not be game-theoretically meaningful [15, 21]. We obtain stronger stability characterizations in our analysis of the continuous-time system using the structure of nonconvex-PŁ/SC zero-sum games. Further discussion of this topic is given in Section 4.

Nonconvex-PŁ and Nonconvex-SC Zero-Sum Games. Table 1 provides a comprehensive comparison between our results and existing results for gradient descent variants in nonconvex optimization and nonconvex-PŁ/SC zero-sum games. We leave discussion of research on these classes of games with other algorithmic methods to Appendix A. The key distinction between this paper and past work is that instead of assessing convergence in terms of only reaching an approximate stationary point of the dynamics or a surrogate function, we obtain convergence guarantees in regards to differential Stackelberg equilibria (equivalently strict local minmax equilibria). Despite this being a much stricter and meaningful notion of solving the problem, we show that global asymptotic convergence guarantees remain obtainable in nonconvex-PŁ zero-sum games (Theorem 2) and that the local convergence rate (Corollary 1) is comparable to existing results. Moreover, for the subclass of nonconvex-SC zero-sum games, we provide novel global finite-time convergence guarantees to differential Stackelberg equilibria with rates comparable to existing results for finding a local minimum in nonconvex optimization or any stationary point in this class of games (Theorem 3).

Escaping Saddle Points in Nonconvex Optimization. A key contribution of this work is in regards to escaping saddles of the dynamics in classes of nonconvex zero-sum games. We build on methods from nonconvex optimization to obtain analogous results. In general, saddle avoidance results for variants of gradient descent in nonconvex optimization are asymptotic or finite-time. The former states that almost surely the algorithm does not converge to saddles points [26, 27], while the latter gives rates of escape to conclude convergence to approximate local minimum [16, 19, 20]. A key assumption in the aforementioned works is what is known as the strict saddle property, which ensures directions of escape exists from a saddle point. We make an analogous assumption refined for games. Further discussion of how we extend the methods from nonconvex optimization to the zero-sum game problem along with the challenges is included in Section 6.

3 Preliminaries

In the zero-sum games we study, we refer to the minimizing player controlling x as player 1 and the maximizing player controlling y as player 2. We consider objective functions $f \in C^2(\mathcal{Z}, \mathbb{R})$

where the joint strategy space is denoted by $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ where $\mathcal{X} = \mathbb{R}^{d_1}$ and $\mathcal{Y} = \mathbb{R}^{d_2}$ denote the individual action spaces and $d = d_1 + d_2$. We often denote a joint strategy by $z = (x, y) \in \mathcal{Z}$.

Notation. We denote ∇f as the total derivative of f, $\nabla_i f$ as the derivative of f with respect to the choice variable of player i, $\nabla_{ij} f$ as the partial derivative of $\nabla_i f$ with respect to the choice variable of player j, and $\nabla_i^2 f$ as the partial derivative of $\nabla_i f$ with respect to the choice variable of player i. We let $\|\cdot\|$ denote the 2-norm of vectors unless otherwise specified, $\mathrm{spec}(\cdot)$ denote the set of eigenvalues of a matrix, $\mathrm{Re}(\cdot)$ denote the real part of a complex number, and $\mathbb{C}_-^\circ$ and $\mathbb{C}_+^\circ$ denote the open left-half and right-half complex plane, respectively. Let $\lambda_{\min}(A)$ denote the eigenvalue of A with the minimum real part, and $\lambda_{\max}(A)$ the eigenvalue of A with the maximum real part.

Classes of Games. We study and analyze both nonconvex-PŁ and nonconvex-SC zero-sum games. Throughout, we make the following assumptions on the class of functions that define the games.

Assumption 1. *Given a zero-sum game $(f, -f)$ defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$, $\nabla_1 f(\cdot, y)$, $\nabla_2 f(\cdot, y)$, $\nabla_2 f(x, \cdot)$, and $\nabla_1 f(x, \cdot)$ are L_1, L_{12}, L_2, and L_{21} Lipschitz, respectively. That is, for all $x, x', x'' \in \mathcal{X}$ and $y, y', y'' \in \mathcal{Y}$, we have*

$$\|\nabla_1 f(x, y'') - \nabla_1 f(x', y'')\| \le L_1 \|x - x'\|, \quad \|\nabla_2 f(x'', y) - \nabla_2 f(x'', y')\| \le L_2 \|y - y'\|$$
$$\|\nabla_2 f(x, y'') - \nabla_2 f(x', y'')\| \le L_{12} \|x - x'\|, \quad \|\nabla_1 f(x'', y) - \nabla_1 f(x'', y')\| \le L_{21} \|y - y'\|.$$

This assumption immediately implies that the vector of individual gradients given by $g(x, y) = (\nabla_1 f(x, y), -\nabla_2 f(x, y))$ is also Lipschitz with parameter $L \le \ell := L_1 + L_2$.

Assumption 2. *Given a zero-sum game $(f, -f)$ defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$, $\nabla^2 f(\cdot, \cdot)$ is β-Lipschitz with respect to the induced 2-norm. That is, for all $z, z' \in \mathcal{Z}$, $\|\nabla^2 f(z) - \nabla^2 f(z')\| \le \beta \|z - z'\|$.*

We now define a nonconvex-PŁ zero-sum game. This class of games allows for the objective function to be nonconvex in $x \in \mathcal{X}$, but it needs to satisfy the Polyak-Łojasiewicz (PŁ) condition in $y \in \mathcal{Y}$.

Definition 1 (Nonconvex-PŁ Game)**.** *Consider a zero-sum game $(f, -f)$ defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$. The game is called nonconvex-PŁ if $f(x, \cdot)$ is μ-PŁ with respect to the argument $y \in \mathcal{Y}$. That is, for $\mu > 0$ and for all $(x, y) \in \mathcal{Z}$,*

$$\|\nabla_2 f(x, y)\|^2 \ge 2\mu(\max_{y' \in \mathcal{Y}} f(x, y') - f(x, y)).$$

The following provides a definition for what we call a nonconvex-SC zero-sum game. In this class of games, the function that defines the game may be nonconvex in $x \in \mathcal{X}$, but it must be SC in $y \in \mathcal{Y}$.

Definition 2 (Nonconvex-SC Game)**.** *Consider a zero-sum game $(f, -f)$ defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$. The game is nonconvex-SC if $f(x, \cdot)$ is μ-strongly-concave with respect to the argument $y \in \mathcal{Y}$. That is, given any $x \in \mathcal{X}$ and for all $y, y' \in \mathcal{Y}$,*

$$f(x, y') \le f(x, y) + \langle \nabla_1 f(x, y), y' - y \rangle - \tfrac{\mu}{2} \|y - y'\|_2^2.$$

It is worth noting that nonconvex-SC zero-sum games are nonconvex-PŁ zero-sum games, but the converse does not necessarily hold as a result of the relationship between PŁ and SC functions [23].

Learning Dynamics. We study gradient descent-ascent with timescale separation and a perturbed variant. Let γ be the learning rate of player one and $\tau > 0$ be the timescale parameterization. The deterministic τ-`GDA` dynamics we study are presented in Algorithm 1 and the potentially stochastic variant with injected noise perturbations called τ-`PGDA` is given in Algorithm 2 (see Section 6 for the relevant notation). The key distinction between this study of gradient descent-ascent with timescale separation and past work is how we assess convergence as we now begin to formalize.

Stationarity Notions. A critical point corresponds to a joint strategy profile at which the individual gradient of each player is equal to zero.

Definition 3 (Critical Point)**.** *$(x, y) \in \mathcal{Z}$ is a critical point if $\nabla_1 f(x, y) = 0$ and $\nabla_2 f(x, y) = 0$.*

Critical points correspond to stationary points of the τ-`GDA` dynamics. The nonconvex-PŁ/SC zero-sum game literature has generally assessed convergence in terms of the complexity of finding an approximate stationary point (see, e.g., [41, 56]). That is, joint strategies $(x, y) \in \mathcal{Z}$ such that $\|\nabla_1 f(x, y)\| \le \varepsilon$ and $\|\nabla_2 f(x, y)\| \le \varepsilon$. A related and common notion of convergence in this body of work (see e.g., [30]) is that of finding a stationary point of the function $\max_{y \in \mathcal{Y}} f(\cdot, y)$. This criterion amounts to seeking to achieve the condition $\|\nabla_1 \max_{y \in \mathcal{Y}} f(x, y)\| \le \varepsilon$.

Algorithm 1 τ-GDA

Input: $x_0 \in \mathbb{R}^{d_1}$, $y_0 \in \mathbb{R}^{d_2}$
for $k = 0, 1, \ldots$ **do**
  $x_{k+1} \leftarrow x_k - \gamma \nabla_1 f(x_k, y_k)$
  $y_{k+1} \leftarrow y_k + \gamma\tau \nabla_2 f(x_k, y_k)$
end for

Algorithm 2 τ-PGDA

Input: $x_0 \in \mathbb{R}^{d_1}$, $y_0 \in \mathbb{R}^{d_2}$
for $k = 0, 1, \ldots$ **do**
  Sample $\theta_{i,k} \sim \mathcal{D}_i$, $i = 1, 2$; $(\xi_{1,k}, \xi_{2,k}) \sim \mathcal{N}(0, (r^2/d)I)$
  $x_{k+1} \leftarrow x_k - \gamma(g_1(x_k, y_k; \theta_{1,k}) + \xi_{1,k})$,
  $y_{k+1} \leftarrow y_k + \gamma\tau(g_2(x_k, y_k; \theta_{2,k}) + \tau^{-1}\xi_{2,k})$
end for

In contrast, we assess convergence with connections to the equilibrium notions that are commonly studied in the nonconvex-nonconcave zero-sum game literature. Since either stationarity notion may lack any game-theoretic meaning, we consider a strictly harder notion of solving a game.

Equilibrium Notions. The typical solution concept in game theory when an implicit or explicit order of play is present in the structure of the game is the (local) Stackelberg (equivalently minmax in zero-sum games) equilibrium concept [5].[2] Informally, in nonconvex-PŁ/SC zero-sum games, a local minmax equilibrium corresponds to a strategy pair $(x^*, y^*) \in \mathcal{Z}$ such that x^* is a local minimum of the function $f(x, y_*(x))$ where $y_*(x) \in \operatorname{argmax}_{y \in \mathcal{Y}} f(x, y)$ and y^* is a local maximum of the function $f(x^*, y)$. When the function f is bounded or when $f(\cdot, y)$ is bounded and $f(x, \cdot)$ is strongly concave, a minmax equilibrium is guaranteed to exist.

We characterize the local minmax (Stackelberg) equilibrium notion in terms of sufficient conditions on player costs as is typical in learning in games (see, e.g., [14, 15, 21, 36, 52]). Toward presenting this definition, we denote by $J(x, y)$ the Jacobian of the vector of individual gradients $g(x, y)$ that is given by

$$J(x, y) = \begin{bmatrix} \nabla_1^2 f(x, y) & \nabla_{12} f(x, y) \\ -\nabla_{12}^\top f(x, y) & -\nabla_2^2 f(x, y) \end{bmatrix}. \tag{4}$$

Let $\mathtt{S}_1(\cdot)$ denote the Schur complement of $(\cdot)$ with respect to the $d_2 \times d_2$ block in $(\cdot)$. The following definition is characterized by sufficient conditions for a local minmax equilibrium in zero-sum games.

Definition 4 (Differential Stackelberg/Strict Local Minmax Equilibrium [14])**.** *The joint strategy $(x^*, y^*) \in \mathcal{Z}$ is a differential Stackelberg equilibrium if the first-order conditions $\nabla f(x^*, y^*) = \nabla_1 f(x^*, y^*) - \nabla_{12} f(x^*, y^*)[\nabla_2^2 f(x^*, y^*)]^{-1} \nabla_2 f(x^*, y^*) = 0$, $\nabla_2 f(x^*, y^*) = 0$ hold and the second-order conditions $\mathtt{S}_1(J(x^*, y^*)) > 0$ and $\nabla_2^2 f(x^*, y^*) < 0$ also hold.*

A differential Stackelberg equilibrium corresponds to a joint strategy at which the minimizing player is at a local optimum with respect to its choice variable along the best response curve of the maximizing player and the maximizing player is at a local optimum with respect to its choice variable. In the next section, we also provide a definition for an ε-differential Stackelberg equilibrium.

4 Local Stability Analysis

To characterize the convergence of τ-GDA, we begin by studying its continuous-time limiting system

$$\dot{z} = -\Lambda_\tau g(z) \tag{5}$$

where $\dot{z} = (\dot{x}, \dot{y})$, $\Lambda_\tau = \operatorname{blockdiag}(I_{d_1}, \tau I_{d_2})$, and $g(z)$ is the vector of individual gradients. The Jacobian of this system is given by

$$J_\tau(z) = \Lambda_\tau J(z).$$

We analyze the stability of the continuous-time system around critical points z^* as a function of the timescale separation τ using the Jacobian $J_\tau(z^*)$ in this section toward drawing conclusions about the stability and convergence of the discrete time system τ-GDA. A critical point is said to be locally (exponentially) stable when the spectrum of $-J_\tau(z^*)$ is in the open left-half complex plane $\mathbb{C}_-^\circ$ (cf. Theorem 4, Appendix B.2). Simply put, a critical point z^* is locally exponentially stable if and only if the real parts of the eigenvalues of $-J_\tau(z^*)$ are strictly negative. Throughout, we use the broader term "stable" to mean the following.

[2] By implicit order of play, we mean the min and max order are not interchangeable.

Definition 5 (Stability)**.** *A critical point $z^* = (x^*, y^*) \in \mathcal{Z}$ is locally exponentially stable for $\dot{z} = -\Lambda_\tau g(z)$ if and only if* $\mathrm{spec}(-J_\tau(z^*)) \subset \mathbb{C}_-^\circ$ *($\equiv$ $\mathrm{spec}(J_\tau(z^*)) \subset \mathbb{C}_+^\circ$).*

Stability with respect to the continuous-time τ-GDA dynamics guarantees that the system asymptotically converges at an exponential rate to the critical point in a local neighborhood. Moreover, given a suitable choice of learning rates, equivalent insights hold for the discrete-time dynamics [7].

Stability in Nonconvex-Nonconcave Zero-Sum Games. Given the implications regarding convergence, a number of papers in the past several years study the stability of τ-GDA around critical points and the connections to differential Nash and Stackelberg equilibrium in zero-sum games [9, 14, 15, 21, 36]. However, this body of research focuses on general nonconvex-nonconcave zero-sum games. In general across the spectrum of nonconvex-nonconcave zero-sum games, the stable critical points of τ-GDA coincide with the set of differential Stackelberg equilibria only when the timescale separation $\tau \to \infty$ [21]. Given that such a choice of timescale separation requires the learning rate $\gamma \to 0$ in order to retain stability of the discrete-time system, it is not clear how to derive a practical algorithm from this insight.

Toward remedying this problem, [15] provide stability results in terms of the timescale separation concerning a given critical point, rather than across the space of nonconvex-nonconcave zero-sum games. Indeed, they prove a stability and instability result as a function of the timescale separation in the τ-GDA dynamics. The stability results say that given a differential Stackelberg equilibrium z^*, there exists a finite $\tau^* \in (0, \infty)$ that can be constructed such that z^* is stable for all $\tau \in (\tau^*, \infty)$. On the other hand, the instability results says that given a critical point which is not a differential Stackelberg equilibrium, there exists a finite $\tau_0 \in (0, \infty)$ that can be constructed such that z^* is not stable for all $\tau \in (\tau_0, \infty)$.

Stability in Nonconvex-PŁ/SC Zero-Sum Games. To our knowledge, the connection between the stability (and instability) of critical points with respect to τ-GDA dynamics and game-theoretic equilibrium notions in the semi-structured problems of nonconvex-PŁ and nonconvex-SC has not been fully characterized. We show in the following result that when a nonconvex-nonconcave game is specialized to a nonconvex-PŁ game as from Definition 1, significantly more general stability characterizations can be obtained. Notably, any critical point z^* that is not a differential Stackelberg equilibrium is unstable for all $\tau \geq 0$. Meaning that τ-GDA does not admit spurious stable points in nonconvex-PŁ game. This is in stark contrast to the known fact that 1-GDA admits spurious stable points in the more general class of nonconvex-nonconcave games as has been shown in previous literature [9, 36]. Moreover, if z^* is a differential Stackelberg equilibrium, then z^* is stable for all τ larger than the minimum τ_* for which z^* is stable and such a finite τ_* is guaranteed to exist. This result implies that in practice, one can select a finite value of τ to run τ-GDA with, and all stable critical points (if they exist) will be differential Stackelberg equilibria and if τ is scaled up and the set of stable points grows, then only differential Stackelberg equilibria can be introduced.

Theorem 1. *Consider a nonconvex-PŁ zero-sum game $(f, -f)$ where $f \in C^2(\mathcal{Z}, \mathbb{R})$. Then, the following hold: 1) Any critical point z^* that is not a differential Stackelberg equilibrium is unstable for all $\tau \in (0, \infty)$; 2) If z^* is a differential Stackelberg equilibrium, then* $\mathrm{spec}(-J_\tau(z^*)) \subset \mathbb{C}_-^\circ$ *for all $\tau \in [\tau_*, \infty)$ where τ_* is the minimum $\tau \in (0, \infty)$ such that* $\mathrm{spec}(-J_\tau(z^*)) \subset \mathbb{C}_-^\circ$ *and a finite τ_* is guaranteed to exist.*

In comparison to the stability results for nonconvex-nonconcave zero-sum games of [15], we obtain the stronger results in nonconvex-PŁ zero-sum games that (i) $\tau_0 = 0$ for any critical point that is not a differential Stackelberg equilibria and (ii) a differential Stackelberg equilibria is never unstable after it becomes stable as a function of the timescale separation τ.

The stability characterization also allows us to define a natural approximate differential Stackelberg equilibrium notion. The first-order conditions in Definition 4 can equivalently be reformulated as $\nabla_1 f(x^*, y^*) = 0$ and $\nabla_2^2 f(x^*, y^*) = 0$. We presented the differential Stackelberg equilibrium using the total derivative for the minimizing player in the first-order condition to mirror the presentation of proper approximate notions. In particular, let $\bar{g}(x, y) = (\nabla f(x, y), -\nabla_2 f(x, y))$ denote the vector containing the total derivative for the minimizing player and the individual derivative for the maximizing player. We have the following ε–differential Stackelberg equilibrium.

Definition 6. *A point $z^* = (x^*, y^*)$ is an ε-differential Stackelberg equilibrium (ε-local minmax) for a nonconvex–PŁ zero-sum game $f \in C^2(\mathcal{Z}, \mathbb{R})$ satisfying Assumptions 1 and 2 if $\|\bar{g}(z^*)\| \leq \varepsilon$ and* $\mathrm{Re}(\lambda_{\min}(J_\tau(z^*))) \geq -\sqrt{\varepsilon\beta}$.

5 Global Asymptotic Convergence Analysis

Theorem 1 in the previous section completely characterizes the behavior of the continuous-time τ-GDA dynamics. However, ultimately we are concerned with the convergence properties of the discrete-time τ-GDA dynamics. We begin our study of this system by providing a global asymptotic analysis in this section of the deterministic τ-GDA dynamics presented in Algorithm 1.

We prove in Theorem 2 of this section that the deterministic τ-GDA (Algorithm 1) almost surely converges to a differential Stackelberg equilibrium in nonconvex-PŁ zero-sum games, a class of games that subsumes nonconvex-SC zero-sum games. Despite this result being asymptotic in nature, to our knowledge it is the most general class of zero-sum games in which a global convergence guarantee to established game-theoretic equilibria has been given.

To begin, observe that the τ-GDA dynamics do not necessarily correspond to a gradient flow a function, which can be observed from the fact that the Jacobian J_τ is not guaranteed to be symmetric. As a consequence, there may exist complex limiting behavior beyond convergence to a critical point such as non-trivial limit cycles or periodic orbits. To rule out this phenomenon, we prove that there exists potential function that decreases along the iterates of τ-GDA. In order to construct a potential function for nonconvex, μ-PŁ zero-sum games we also show that the best-response map of the maximizing player, $y_*(x) \in \operatorname{argmax}_y f(x,y)$, defined implicitly by $\nabla_2 f(x,y) = 0$ is L_3-Lipschitz.[3] The best response map is well-defined by the implicit mapping theorem since the PŁ condition on the maximizing player's problem implies quadratic growth [23] which in turn implies that at critical points not only is $\nabla_2^2 f$ is non-degenerate, but also it is negative definite.

Lemma 1. *Consider a non-convex, μ-PŁ zero sum game defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$ which satisfies Assumptions 1 and 2 and has condition number $\kappa = L_2/\mu$. Suppose that $\tau > 7\kappa^2$, and $\gamma < \min\{\frac{1}{3L_2\tau}, \frac{1}{2(L_1+L_3L_2)}\}$ then, for any $\Gamma \in (0, 1/7]$, $\Phi(x,y) = f(x, y_*(x)) - \Gamma f(x,y)$ is a potential function for τ-GDA.*

The function $f(x, y_*(x))$ can be seen as the function the x player would minimize if the y player was playing a best-response $y_*(x)$. The potential function $\Phi(x,y)$ essentially captures that along trajectories of τ-GDA, the function $f(x,y)$ should either decrease the value of $f(x, y_*(x))$, or decrease the value of $f(x, y_*(x)) - f(x,y)$ since the y-player converges at a fast rate to $y_*(x)$ given the time-scale separation. Indeed, this potential function implicitly guarantees that the maximizing player *tracks* the best response set, and that the minimizing player essentially ends up minimizing $f(x, y_*(x))$ as desired. The choice of τ and the learning rate γ allow us to guarantee that this occurs.

Given the potential function in Lemma 1, we can conclude that Algorithm 1 converges to critical points. The critical points may correspond to stable points of the dynamics (Definition 5) or saddle points of the dynamics, which are defined as follows.[4]

Definition 7 (Saddle Point)**.** *The critical point $z^* = (x^*, y^*) \in \mathcal{Z}$ is a saddle point of the dynamics $\dot{z} = -\Lambda_\tau g(z)$ if $\operatorname{Re}(\lambda_{\max}(-J_\tau(z^*))) = 0$ and a strict saddle if $\operatorname{Re}(\lambda_{\max}(-J_\tau(z^*))) > 0$.*

Recall that Theorem 1 indicates that the only stable points of the continuous-time τ-GDA dynamics are differential Stackelberg equilibria. Thus, if we can show that the discrete-time τ-GDA dynamics avoid saddle points of the continuous-time τ-GDA dynamics, then we can conclude that Algorithm 1 converges to only differential Stackelberg equilibria. The following result of [36] states that the discrete-time τ-GDA dynamics do indeed avoid saddle points of the continuous-time system.[5]

Lemma 2 (Theorem 4.1 [36])**.** *Consider a zero sum game defined by the non-convex, non-concave function $f \in C^2(\mathcal{Z}, \mathbb{R})$. The set of initial conditions $z \in \mathcal{Z}$ from which τ-GDA converges to strict saddle points is of measure zero.*

Lemma 2 can be combined with the existence of a potential function (Lemma 1) and the guarantee that the only stable points of τ-GDA are strict local minmax, we obtain the following global convergence guarantee for non-convex, PŁ-games. To obtain this result, we need the additional commonly assumed strict saddle point property [16, 19, 36].

[3] This mapping is $\ell(\kappa+1)$-Lipschitz in the case where $f(x,\cdot)$ is strongly concave [30, Lemma 4.3].

[4] The terminology of saddle point in this paper is with respect to the dynamics and should not be conflated with the terminology of a saddle point of a function.

[5] The saddle avoidance result of [36] holds for both general-sum n-player nonconvex games. For simplicity, we only state it in the context of this work.

Assumption 3 (Strict Saddle Property)**.** *Given a zero-sum game $(f, -f)$ defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$, we assume the strict saddle property that every saddle point of the dynamics is a strict saddle point.*

We are now ready to the state main result of this section.

Theorem 2. *Consider a nonconvex-PŁ zero-sum game $(f, -f)$ defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$ that satisfies Assumptions 1–3. Then, τ-GDA with $\tau > 7\kappa^2$, and stepsize $\gamma < \min\{\frac{1}{3L_2\tau}, \frac{1}{2(L_1+L_3L_2)}\}$ asymptotically converges to the set of strict local minmax that are stable for $\dot{z} = -\Lambda_\tau g(z)$ almost surely. That is, for almost all initial conditions, τ-GDA will converge to a strict local minmax point.*

If we fix a τ satisfying the assumptions of the above theorem, then if we consider the set of stable critical points for $\dot{z} = -\Lambda_\tau(z)$, we know that by Theorem 1 this set only contains differential Stackelberg (and hence, strict local minmax) points. It is precisely this set of strict local minmax to which τ-GDA converges almost surely. Moreover, by Theorem 1, as we increase τ, no new spurious non-minmax points are introduced to the set of stable critical points; only additional strict local minmax points can be added to this set.

We complement Theorem 2 with the following characterization of the convergence rate.

Corollary 1. *Consider a nonconvex-PŁ zero-sum game $(f, -f)$ defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$ that satisfies Assumptions 1–3. Given an initialization in the region of attraction of a strict local minmax, then after $\widetilde{\mathcal{O}}(\varepsilon^{-2})$ iterations at least one iterate is an ε-differential Stackelberg equilibrium.*

Observe that this Corollary is a local convergence rate guarantee. The reason for this is that it is possible that the dynamics could get stuck around a strict saddle point for a long period of time. We deal with this problem in the next section.

6 Finite Time Convergence Analysis

In this section, we show that in nonconvex-SC games not only does gradient descent-ascent provably converge in an almost sure asymptotic sense, but also that there is a finite time escape time from saddle points. The game dynamics can get stuck in a region around a non-game theoretically meaningful saddle, and while this problem has been observed in the nonconvex single player setting, it appears more of nuanced issue in the game setting since saddles of the dynamics are not always saddles of the function defining the game and vice versa. We show the result for the stochastic case in which players have unbiased stochastic gradients. The deterministic case is immediate by taking the stochasticity to zero.

Players do not have access to their exact gradients $\nabla_1 f(x, y)$ and $\nabla_2 f(x, y)$. Instead, for any (x, y) a gradient query to an oracle return $g_1(x, y; \theta_1)$ and $g_2(x, y; \theta_2)$ to the two players respectively, where each θ_i is a random variables distribution $\mathcal{D}_i$, $i = 1, 2$.

Assumption 4. *For any $(x, y) \in \mathcal{Z}$, the stochastic gradient vector $g(x, y; \theta) = (g_1(x, y; \theta_1), g_2(x, y; \theta_2))$ satisfies $\nabla_i f(x, y) = \mathbb{E}_{\theta_i \sim \mathcal{D}_i}[g_i(x, y; \theta_i)]$ and $\forall\, t \in \mathbb{R}$ and each $i = 1, 2$:*

$$\mathbb{P}(\|g_i(x, y; \theta) - \nabla_i f(x, y)\| \geq t) \leq 2\exp(-\tfrac{t^2}{2\sigma_i^2}).$$

Assumption 5. *For each $i = 1, 2$ and any $\theta_i \in \operatorname{supp}(\mathcal{D}_i)$, $g_i(\cdot, \cdot; \theta_i)$ is $\tilde{\ell}_i$-Lipschitz: i.e., for all $(x_1, y_1), (x_2, y_2) \in \mathcal{Z}$, $\|g_i(x_1, y_1; \theta_i) - g_i(x_2, y_2; \theta_i)\| \leq \tilde{\ell}_i \|(x_1, y_1) - (x_2, y_2)\|$.*

Assumption 6. *The eigenvalue of $J_\tau(z)$ with minimum real part for any strict saddle z is simple. That is, the algebraic multiplicity $m_A(-\eta) = 1$ and geometric multiplicity of $m_G(-\eta) = 1$ where $-\eta = \operatorname{Re}(\lambda_{\min}(J_\tau(z)))$.*

Let $\ell = L_1 + L_2$ if $\kappa \leq 1/2$, and $\ell = L_1 + 2\kappa L_2$ otherwise. In either case, $\lambda_{\max}(\nabla^2 f(x, y)) \leq \ell$. In the remainder, we work in the case for which $\ell \geq \sqrt{\beta\varepsilon}$; otherwise the problem of finding ε-local minmax is straightforward since all ε-stationary points will be ε-local minmax points.

The steps in showing that τ-PGDA (Algorithm 2) escapes saddles with high probability largely follow those in [8, 16, 19, 20], with several modifications due to the fact that we consider a zero-sum game as opposed to a single player optimization problem. The challenge in the minmax (or zero-sum game) setting as compared to the single player optimization setting is that the gradient descent-ascent dynamics are not a gradient flow, and hence, the eigenvalues of the local linearization may be

complex-valued. In the key step of many of the proofs on saddle avoidance, escaping from the *stuck regions* around a saddle (with high probability) is achieved by analyzing the local linearization around the saddle point, which crucially, is symmetric. Hence, the problem reduces to analyzing the escape time along the eigendirection of the minimum eigenvalue of the Hessian which relies heavily on the orthogonality of eigenvectors of symmetric matrices. In zero-sum games, the linearization is no longer symmetric and these key results have to be substantially modified. We highlight these distinctions more clearly in Appendix E.

One of the other key distinctions is in constructing a potential function. In the existing literature on convergence in the nonconvex-SC literature [32, 53, 56], the proposed potential function has a structure closely related to the potential function in Lemma 1; in particular, it takes the form $f(x, y) + r(x, y)$ where $r(x, y)$ is a tracking term of the form $\|y - y_*(x)\|^2$ or $f(x, y_*(x))$. Yet, perhaps surprisingly, despite the fact that f depends on to sequences generated by two separate gradient updates—one from minimizing f and one from maximizing it—the descent lemma below (Lemma 3) shows that the function f can be decomposed into a component that is decreasing along trajectories of τ-PGDA (Algorithm 2) and a component that exhibits a possible increase due to randomness in the stochastic gradients and injected noise. The primary reason for the decrease is large gradients in addition to the strongly concave structure of the maximizing player's problem and the timescale separation which we exploit in the proof to ensure sufficient decrease of f along trajectories of τ-PGDA.

Lemma 3 (Descent Lemma)**.** *Consider a non-convex, μ-SC zero-sum game defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$. Under Assumptions 1–5, if $\frac{3}{2\mu\tau} < \gamma < 1/\ell$, then with probability at least $1 - \delta$ for some $\delta > 0$, $f(x_t, y_t)$ is a potential function for τ-PGDA in both the stochastic and deterministic settings.*

The descent lemma enables us to argue that with high probability, either the function value decreases a sufficient amount or the iterates remain in a small region. Using this property of τ-PGDA together with the descent lemma, we then show that if two sequences that are *coupled* (i.e., two sequences starting from the same initial condition and having equivalent noise except along the direction of escape) and are in the region around a saddle point, then with high probability, at least one of them escapes The full proof is provided in Appendix E.

Putting these results together, we have the following finite time guarantee on convergence to ε-differential Stackelberg equilibria (ε–local minmax points). Let $f^* = \min_{x,y} f(x, y)$ be the minimum value of the function f, which we assume is finite. These results provide novel convergence guarantees to not just ε-stationary points but those that are game theoretically meaningful.

Theorem 3. *Consider a non-convex, μ-strongly concave zero-sum game defined by $f \in C^2(\mathcal{Z}, \mathbb{R})$ and suppose that Assumptions 1–6 hold. For any $\varepsilon, \delta > 0$, there exists γ and τ such that, with probability $1 - \delta$ for some $\delta > 0$, starting from any $z_0 = (x_0, y_0)$, at least half the iterates of τ-PGDA will be ε-differential Stackelberg equilibria after $\widetilde{\mathcal{O}}(\varepsilon^{-4})$ iterations and $\widetilde{\mathcal{O}}(\varepsilon^{-2})$ in the stochastic and deterministic settings, respectively.*

7 Discussion

To the best of our knowledge, our results are the first to guarantee global convergence of gradient-based algorithms to game theoretically meaningful equilibria in such a general class of games. We believe that a more detailed analysis could make use of the potential function defined for nonconvex-PŁ in Section 5 to prove similar finite-time results for τ-PGDA, though the proof appears to be tedious and hence, we leave this for future work. We also believe future work can be done to obtain global guarantees in this class of games to game-theoretically meaningful equilibria without gradient information and also in the setting with constraints extending recent saddle avoidance results for the single player nonsmooth, constrained setting.

References

[1] Leonard Adolphs, Hadi Daneshmand, Aurelien Lucchi, and Thomas Hofmann. Local saddle point optimization: A curvature exploitation approach. In *International Conference on Artificial Intelligence and Statistics*, pages 486–495, 2019.

[2] Naman Agarwal, Zeyuan Allen-Zhu, Brian Bullins, Elad Hazan, and Tengyu Ma. Finding approximate local minima faster than gradient descent. In *ACM SIGACT Symposium on Theory of Computing*, pages 1195–1199, 2017.

[3] Zeyuan Allen-Zhu. Natasha 2: faster non-convex optimization than sgd. In *Advances in Neural Information Processing Systems*, pages 2680–2691, 2018.

[4] Zeyuan Allen-Zhu and Yuanzhi Li. Neon2: Finding local minima via first-order oracles. In *Advances in Neural Information Processing Systems*, pages 3716–3726, 2018.

[5] Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory*. SIAM, 1998.

[6] Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Accelerated methods for nonconvex optimization. *SIAM Journal on Optimization*, 28(2):1751–1772, 2018.

[7] Benjamin Chasnov, Lillian Ratliff, Eric Mazumdar, and Samuel Burden. Convergence analysis of gradient-based learning in continuous games. In *Conference on Uncertainty in Artificial Intelligence*, pages 935–944, 2020.

[8] Hadi Daneshmand, Jonas Kohler, Aurelien Lucchi, and Thomas Hofmann. Escaping saddles with stochastic gradients. In *International Conference on Machine Learning*, pages 1155–1164, 2018.

[9] Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *Advances in Neural Information Processing Systems*, pages 9236–9246, 2018.

[10] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. *arXiv preprint arXiv:2009.09623*, 2020.

[11] Simon S Du, Chi Jin, Jason D Lee, Michael I Jordan, Barnabás Póczos, and Aarti Singh. Gradient descent can take exponential time to escape saddle points. *Advances in Neural Information Processing Systems*, 2017:1068–1078, 2017.

[12] C Fang, CJ Li, Z Lin, and T Zhang. Spider: Near-optimal non-convex optimization via stochastic path integrated differential estimator. *Advances in Neural Information Processing Systems*, 2018.

[13] Cong Fang, Zhouchen Lin, and Tong Zhang. Sharp analysis for nonconvex sgd escaping from saddle points. In *Conference on Learning Theory*, pages 1192–1234, 2019.

[14] Tanner Fiez, Benjamin Chasnov, and Lillian Ratliff. Implicit learning dynamics in stackelberg games: Equilibria characterization, convergence analysis, and empirical study. In *International Conference on Machine Learning*, pages 3133–3144, 2020.

[15] Tanner Fiez and Lillian Ratliff. Local convergence analysis of gradient descent ascent with finite timescale separation. In *International Conference on Learning Representations*, 2021.

[16] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pages 797–842, 2015.

[17] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.

[18] Ya-Ping Hsieh, Panayotis Mertikopoulos, and Volkan Cevher. The limits of min-max optimization algorithms: Convergence to spurious non-critical sets. *arXiv preprint arXiv:2006.09065*, 2020.

[19] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In *International Conference on Machine Learning*, pages 1724–1732, 2017.

[20] Chi Jin, Praneeth Netrapalli, Rong Ge, Sham M Kakade, and Michael I Jordan. On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. *Journal of the ACM*, 68(2):1–29, 2021.

[21] Chi Jin, Praneeth Netrapalli, and Michael Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *International Conference on Machine Learning*, pages 4880–4889, 2020.

[22] Chi Jin, Praneeth Netrapalli, and Michael I Jordan. Accelerated gradient descent escapes saddle points faster than gradient descent. In *Conference On Learning Theory*, pages 1042–1085, 2018.

[23] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases*, pages 795–811. Springer International Publishing, 2016.

[24] Hassan Khalil. *Nonlinear Systems*. Prentice Hall, 3rd edition, 2002.

[25] Peter Lancaster and Miron Tismenetsky. *The theory of matrices: with applications*. Elsevier, 1985.

[26] Jason D. Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I. Jordan, and Benjamin Recht. First-order methods almost always avoid strict saddle points. *Mathematical Programming*, 176(1):311–337, 2019.

[27] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *Conference on Learning Theory*, pages 1246–1257, 2016.

[28] Alistair Letcher. On the impossibility of global convergence in multi-loss optimization. *International Conference on Learning Representations*, 2021.

[29] Kfir Y Levy. The power of normalization: Faster evasion of saddle points. *arXiv preprint arXiv:1611.04831*, 2016.

[30] Tianyi Lin, Chi Jin, and Michael Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pages 6083–6093, 2020.

[31] Tianyi Lin, Chi Jin, Michael Jordan, et al. Near-optimal algorithms for minimax optimization. In *Conference on Learning Theory*, pages 2738–2779, 2020.

[32] Sijia Liu, Songtao Lu, Xiangyi Chen, Yao Feng, Kaidi Xu, Abdullah Al-Dujaili, Mingyi Hong, and Una-May O'Reilly. Min-max optimization without gradients: Convergence and applications to black-box evasion and poisoning attacks. In *International Conference on Machine Learning*, pages 6282–6293, 2020.

[33] Songtao Lu, Ioannis Tsaknakis, Mingyi Hong, and Yongxin Chen. Hybrid block successive approximation for one-sided non-convex min-max problems: algorithms and applications. *IEEE Transactions on Signal Processing*, 2020.

[34] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representation*, 2018.

[35] E. Mazumdar and L. J. Ratliff. Local Nash Equilibria are Isolated, Strict Local Nash Equilibria in 'Almost All' Zero-Sum Continuous Games. In *IEEE Conference on Decision and Control*, pages 6899–6904, 2019.

[36] Eric Mazumdar, Lillian J Ratliff, and S Shankar Sastry. On gradient-based learning in continuous games. *SIAM Journal on Mathematics of Data Science*, 2(1):103–131, 2020.

[37] Eric V Mazumdar, Michael I Jordan, and S Shankar Sastry. On finding local nash equilibria (and only local nash equilibria) in zero-sum games. *arXiv preprint arXiv:1901.00838*, 2019.

[38] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International Conference on Machine Learning*, pages 3481–3490, 2018.

[39] Vaishnavh Nagarajan and J Zico Kolter. Gradient descent gan optimization is locally stable. In *Advances in Neural Information Processing Systems*, pages 5591–5600, 2017.

[40] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *Advances in Neural Information Processing Systems*, pages 2208–2216, 2016.

[41] Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems*, pages 14934–14942, 2019.

[42] Ioannis Panageas and Georgios Piliouras. Gradient descent only converges to minimizers: Non-isolated critical points and invariant regions. In *Innovations in Theoretical Computer Science Conference*, 2017.

[43] Hassan Rafique, Mingrui Liu, Qihang Lin, and Tianbao Yang. Weakly-convex–concave min–max optimization: provable algorithms and applications in machine learning. *Optimization Methods and Software*, pages 1–35, 2021.

[44] L. J. Ratliff, S. A. Burden, and S. S. Sastry. Characterization and computation of local Nash equilibria in continuous games. In *Allerton Conference on Communication, Control, and Computing*, pages 917–924, 2013.

[45] L. J. Ratliff, S. A. Burden, and S. S. Sastry. Genericity and structural stability of non-degenerate differential Nash equilibria. In *American Control Conference*, pages 3990–3995, 2014.

[46] Lillian J Ratliff, Samuel A Burden, and S Shankar Sastry. On the characterization of local Nash equilibria in continuous games. *IEEE Transactions on Automatic Control*, 61(8):2301–2307, 2016.

[47] S. Shankar Sastry. *Nonlinear Systems Theory*. Springer-Verlag, 1999.

[48] Matthew Staib, Sashank Reddi, Satyen Kale, Sanjiv Kumar, and Suvrit Sra. Escaping saddle points with adaptive gradient methods. In *International Conference on Machine Learning*, pages 5956–5965, 2019.

[49] Gerald Teschl. *Ordinary differential equations and dynamical systems*, volume 140. American Mathematical Society, 2000.

[50] Nilesh Tripuraneni, Mitchell Stern, Chi Jin, Jeffrey Regier, and Michael I Jordan. Stochastic cubic regularization for fast nonconvex optimization. *Advances in Neural Information Processing Systems*, 2018:2899–2908, 2018.

[51] Jun-Kun Wang, Chi-Heng Lin, and Jacob Abernethy. Escaping saddle points faster with stochastic momentum. In *International Conference on Learning Representations*, 2019.

[52] Yuanhao Wang, Guodong Zhang, and Jimmy Ba. On solving minimax optimization locally: A follow-the-ridge approach. In *International Conference on Learning Representations*, 2020.

[53] Zhongruo Wang, Krishnakumar Balasubramanian, Shiqian Ma, and Meisam Razaviyayn. Zeroth-order algorithms for nonconvex minimax problems with improved complexities. *arXiv preprint arXiv:2001.07819*, 2020.

[54] Tengyu Xu, Zhe Wang, Yingbin Liang, and H Vincent Poor. Enhanced first and zeroth order variance reduced algorithms for min-max optimization. *arXiv preprint arXiv:2006.09361*, 2020.

[55] Yi Xu, Rong Jin, and Tianbao Yang. First-order stochastic algorithms for escaping from saddle points in almost linear time. In *Advances in Neural Information Processing Systems*, pages 5530–5540, 2018.

[56] Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. *Advances in Neural Information Processing Systems*, 33, 2020.

[57] Guojun Zhang, Pascal Poupart, and Yaoliang Yu. Optimality and stability in non-convex-non-concave min-max optimization. *arXiv preprint arXiv:2002.11875*, 2020.

[58] Guojun Zhang, Kaiwen Wu, Pascal Poupart, and Yaoliang Yu. Newton-type methods for minimax optimization. *arXiv preprint arXiv:2006.14592*, 2020.

[59] Siqi Zhang, Junchi Yang, Cristóbal Guzmán, Negar Kiyavash, and Niao He. The complexity of nonconvex-strongly-concave minimax optimization. *arXiv preprint arXiv:2103.15888*, 2021.

[60] Dongruo Zhou, Pan Xu, and Quanquan Gu. Finding local minima via stochastic nested variance reduction. *arXiv:1806.08782*, 2018.

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] See conclusion of paper.
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (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] See the stated assumptions in the paper and the appendix.
   (b) Did you include complete proofs of all theoretical results? [Yes] See 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)? [Yes] Code for experiments in the appendix are provided in the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (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? [Yes] Code for experiments in the appendix are provided in the supplementary material.
   (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]