Solving Min-Max Optimization with Hidden Structure via Gradient Descent Ascent

Anonymous Author(s)
Affiliation
Address
email

Abstract

Many recent AI architectures are inspired by zero-sum games, however, the behavior of their dynamics is still not well understood. Inspired by this, we study standard gradient descent ascent (GDA) dynamics in a specific class of non-convex non-concave zero-sum games, that we call hidden zero-sum games. In this class, players control the inputs of smooth but possibly non-linear functions whose outputs are being applied as inputs to a convex-concave game. Unlike general zero-sum games, these games have a well-defined notion of solution; outcomes that implement the von-Neumann equilibrium of the "hidden" convex-concave game. We provide conditions under which vanilla GDA provably converges not merely to local Nash, but the actual von-Neumann solution. If the hidden game lacks strict convexity properties, GDA may fail to converge to any equilibrium, however, by applying standard regularization techniques we can prove convergence to a von-Neumann solution of a slightly perturbed zero-sum game. Our convergence guarantees are non-local, which as far as we know is a first-of-its-kind type of result in non-convex non-concave games. Critically, under proper assumptions we combine the Center-Stable Manifold Theorem along with novel type of initialization dependent Lyapunov functions to prove that almost all initial conditions converge to the solution. Finally, we discuss diverse applications of our framework ranging from generative adversarial networks to evolutionary biology.

1 Introduction

Traditionally, our understanding of convex-concave games revolves around von Neumann's celebrated minimax theorem, which implies the existence of saddle point solutions with a uniquely defined value. Although many learning algorithms are known to be able to compute such saddle points [12], recently there has there has been a fervor of activity in proving stronger results such as faster regret minimization rates or analysis of the day-to-day behavior [46, 16, 7, 1, 66, 18, 2, 45, 5, 24, 69, 29, 6, 48, 30, 56].

This interest has been largely triggered by the impressive successes of AI architectures inspired by min-max games such as Generative Adversarial Networks (GANS) [25], adversarial training [40] and reinforcement learning self-play in games [63]. Critically, however, all these applications are based upon *non-convex non-concave games*, our understanding of which is still nascent. Nevertheless, some important early work in the area has focused on identifying new solution concepts that are widely applicable in general min-max games, such as (local/differential) Nash equilibrium [3, 41], local minmax [17], local minimax [31], (local/differential) Stackleberg equilibrium [23], local robust point [68]. The plethora of solutions concepts is perhaps suggestive that "solving" general min-max games unequivocally may be too ambitious a task. Attraction to spurious fixed points [17], cycles [65], robustly chaotic behavior [14, 15] and computational hardness issues [19] all suggest that general min-max games might inherently involve messy, unpredictable and complex behavior.

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

Are there rich classes of non-convex non-concave games with an effectively unique game theoretic solution that is selected by standard optimization dynamics (e.g. gradient descent)?

Our class of games. We will define a general class of min-max optimization problems, where each agent selects its own vectors of parameters which are then processed separately by smooth functions. Each agent receives their respective payoff after entering the outputs of the processed decision vectors as inputs to a standard convex-concave game. Formally, there exist functions $\mathbf{F} : \mathbb{R}^N \to X \subset \mathbb{R}^n$ and $\mathbf{G} : \mathbb{R}^M \to Y \subset \mathbb{R}^m$ and a continuous convex-concave function $L : X \times Y \to \mathbb{R}$, such that the min-max game is

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^N} \max_{\boldsymbol{\phi} \in \mathbb{R}^M} L(\mathbf{F}(\boldsymbol{\theta}), \mathbf{G}(\boldsymbol{\phi})). \qquad \text{(Hidden Convex-Concave (HCC))}$$

We call this class of min-max problems Hidden Convex-Concave Games. It generalizes the recently defined hidden bilinear games of [65].

Our solution concept. Out of all the local Nash equilibria of HCC games, there exists a special subclass, the vectors $(\boldsymbol{\theta}^*, \boldsymbol{\phi}^*)$ that implement the von Neumann solution of the convex-concave game. This solution has a strong and intuitive game theoretic justification. Indeed, it is stable even if the agents could perform arbitrary deviations directly on the output spaces X, Y. These parameter combinations $(\boldsymbol{\theta}^*, \boldsymbol{\phi}^*)$ "solve" the "hidden" convex-concave L and thus we call them *von Neumann solutions*. Naturally, HCCs will typically have numerous local saddle/Nash equilibria/fixed points that do not satisfy this property. Instead, they correspond to stationary points of the $\mathbf{F}, \mathbf{G}$ where their output is stuck, e.g., due to an unfortunate initialization. At these points the agents may be receiving payoffs which can be arbitrarily smaller/larger than the game theoretic value of game L. Fortunately, we show that Gradient Descent Ascent (GDA) strongly favors von Neumann solutions over generic fixed points.

Our results. In this work, we study the behavior of continuous GDA dynamics for the class of HCC games where each coordinate of $\mathbf{F}, \mathbf{G}$ is controlled by disjoint sets of variables. In a nutshell, we show that GDA trajectories stabilize around or converge to the corresponding von Neumann solutions of the hidden game. Despite restricting our attention to a subset of HCC games, our analysis has to overcome unique hurdles not shared by standard convex concave games.

Challenges of HCC games. In convex-concave games, deriving the stability of the von Neumann solutions relies on the Euclidean distance from the equilibrium being a Lyapunov function. In contrast, in HCC games where optimization happens in the parameter space of $\boldsymbol{\theta}, \boldsymbol{\phi}$, the non-linear nature of $\mathbf{F}, \mathbf{G}$ distorts the convex-concave landscape in the output space. Thus, the Euclidean distance will not be in general a Lyapunov function. Moreover, the existence of *any* Lyapunov function for the trajectories in the output space of $\mathbf{F}, \mathbf{G}$ does not translate to a well-defined function in the parameter space (unless $\mathbf{F}, \mathbf{G}$ are trivial, invertible maps). Worse yet, even if L has a unique solution in the output space, this solution could be implemented by multiple equilibria in the parameter space and thus each of them can not be individually globally attracting. Clearly any transfer of stability or convergence properties from the output to the parameter space needs to be initialization dependent. It is worth mentioning that similar challenges like transfering results from the output to the input space was also faced in the simpler class of hidden bilinear games. However, [65] to sidestep this issue assume the restricitve requirement of $\mathbf{F}$,$\mathbf{G}$ to be invertible operators. Our results go beyond this simplified case requiring new proof techniques. Our results go beyond this simplified case requiring new proof techniques. Specifically, we show how to combine the powerful technologies of the the Center-Stable Manifold Theorem, typically used to argue convergence to equilibria in non-convex optimization settings [34, 52, 54, 53, 35], along with a novel Lyapunov function argument to prove that almost all initial conditions converge to the our game theoretic solution.

Lyapunov Stability. Our first step is to construct an initialization-dependent Lyapunov function that accounts for the curvature induced by the operators $\mathbf{F}$ and $\mathbf{G}$ (Lemma 2). Leveraging a potentially infinite number of initialization-dependent Lyapunov functions in Theorem 5 we prove that under mild assumptions the outputs of $\mathbf{F}, \mathbf{G}$ stabilize around the von Neumann solution of L.

Convergence. Mirroring convex concave games, we require strict convexity or concavity of L to provide convergence guarantees to von Neumann solutions (Theorem 6). Barring initializations where von Neumann solutions are not reachable due to the limitations imposed by $\mathbf{F}$ and $\mathbf{G}$, the set of von Neumann solutions are globally asymptotically stable (Corollary 1). Even in non-strict HCC games, we can add regularization terms to make L strictly convex concave. Small amounts of regularization allows for convergence without significantly perturbing the von Neumann solution (Theorem 7) while

increasing regularization enables exponentially faster convergence rates (Theorem 8). Similar to the aforementioned theoretical work, our model of HCC games provides a formal and theoretical tractable testbed for evaluating the performance of different training methods in GAN inspired architectures. As a concrete example, [36] recently proved the success of WGAN training for learning the parameters of non-linearly transformed Gaussian distributions, where for simplicity they replaced the typical Lipschitz constraint of the discriminator function with a quadratic regularizer. Interestingly, we can elucidate on why regularized learning is actually necessary by establishing a formal connection to HCC games. On top of other such ML applications, our game theoretic framework can furthermore capture and generalize evolutionary game theoretic models. [57] analyze a model of evolutionary competition between two species (host-parasite). The outcome of this competition depends on their respective phenotypes (informally their properties, e.g., agility, camouflage, etc.). These phenotypes are encoded via functions that map input vectors (here genotype/DNA sequences) to phenotypes. While [57] proved that learning in these games does not converge to equilibria and typically cycles for almost all initial conditions, we can explicitly construct initial conditions that do not satisfy our definition of safety and end up converging to artificial fixed points. Safety conditions aside, we show that a slight variation of the evolutionary/learning algorithm suffices to resolve the cycling issues and for the dynamics to equilibrate to the von Neumann solution. Hence, we provide the first instance of team zero-sum games [62], a notoriously hard generalization of zero-sum games with a large duality gap, that is solvable by decentralized dynamics.

Organization. In Section 2 we provide some preliminary notation, the definition of our model and some useful technical lemmas. Section 3 is devoted to the presentation of our the main results. Section 4 discusses applications of our framework to specific GAN formulations. Section 5 concludes our work with a discussion of future directions and challenges. We defer the full proofs of our results as well as further discussion on applications to the Appendix.

2 Preliminaries

2.1 Notation

Vectors are denoted in boldface $\mathbf{x}, \mathbf{y}$ unless otherwise indicated are considered as column vectors. We use $\|\cdot\|$ to denote the ℓ_2-norm. For a function $f : \mathbb{R}^d \to \mathbb{R}$ we use ∇f to denote its gradient. For functions of two vector arguments, $f(\mathbf{x}, \mathbf{y}) : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \to \mathbb{R}$, we use $\nabla_{\mathbf{x}} f, \nabla_{\mathbf{y}} f$ to denote its partial gradient. For the time derivative we will use the dot accent abbreviation, i.e., $\dot{\mathbf{x}} = \frac{d}{dt}[\mathbf{x}(t)]$. A function f will belong to C^r if it is r times continuously differentiable. Additionally, $f \circ g = f(g(\cdot))$ denotes the composition of f, g. Finally, the term "sigmoid" function refers to $\sigma : \mathbb{R} \to \mathbb{R}$ such that $\sigma(x) = (1 + e^{-x})^{-1}$.

2.2 Hidden Convex Concave Games

We will begin our discussion by defining the notion of convex concave functions as well as strictly convex concave functions. Note that our definition of strictly convex concave functions is a superset of strictly convex strictly concave functions that are usually studied in the literature.

Definition 1. *$L : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is convex concave if for every $\mathbf{y} \in \mathbb{R}^n$ $L(\cdot, \mathbf{y})$ is convex and for every $\mathbf{x} \in \mathbb{R}^m$ $L(\mathbf{x}, \cdot)$ is concave. Function L will be called strictly convex concave if it is convex concave and for every $\mathbf{x} \times \mathbf{y} \in \mathbb{R}^n \times \mathbb{R}^m$ either $L(\cdot, \mathbf{y})$ is strictly convex or $L(\mathbf{x}, \cdot)$ is strictly concave.*

At the center of our definition of HCC games is a convex concave utility function L. Additionally, each player of the game is equipped with a set of operator functions. The minimization player is equipped with n functions $f_i : \mathbb{R}^{n_i} \to \mathbb{R}$ while the maximization player is equipped with m functions $g_j : \mathbb{R}^{m_j} \to \mathbb{R}$. We will assume in the rest of our discussion that f_i, g_j, L are all C^2 functions. The inputs $\boldsymbol{\theta}_i \in \mathbb{R}^{n_i}$ and $\boldsymbol{\phi}_j \in \mathbb{R}^{m_j}$ are grouped in two vectors

$$\boldsymbol{\theta} = \begin{bmatrix}\boldsymbol{\theta}_1 & \cdots & \boldsymbol{\theta}_n\end{bmatrix}^\top \qquad \mathbf{F}(\boldsymbol{\theta}) = \begin{bmatrix}f_1(\boldsymbol{\theta}_1) & \cdots & f_n(\boldsymbol{\theta}_n)\end{bmatrix}^\top$$
$$\boldsymbol{\phi} = \begin{bmatrix}\boldsymbol{\phi}_1 & \cdots & \boldsymbol{\phi}_m\end{bmatrix}^\top \qquad \mathbf{G}(\boldsymbol{\phi}) = \begin{bmatrix}g_1(\boldsymbol{\phi}_1) & \cdots & g_m(\boldsymbol{\phi}_m)\end{bmatrix}^\top$$

We are ready to define the hidden convex concave game

$$(\boldsymbol{\theta}^*, \boldsymbol{\phi}^*) = \min_{\boldsymbol{\theta} \in \mathbb{R}^N} \max_{\boldsymbol{\phi} \in \mathbb{R}^M} L(\mathbf{F}(\boldsymbol{\theta}), \mathbf{G}(\boldsymbol{\phi})).$$

where $N = \sum_{i=1}^{n} n_i$ and $M = \sum_{j=1}^{m} m_j$. Given a convex concave function L, all stationary points of L are (global) Nash equilibria of the min-max game. We will call the set of all equilibria of L, von Neumann solutions of L and denote them by Solution(L). Unfortunately, Solution(L) can be empty for games defined over the entire $\mathbb{R}^n \times \mathbb{R}^m$. For games defined over convex compact sets, the existence of at least one solution is guaranteed by von Neumann's minimax theorem. Our definition of HCC games can capture games on restricted domains by choosing appropriately bounded functions f_i and g_j. In the following sections, we will just assume that Solution(L) is not empty. We note that our results hold for both bounded and unbounded f_i and g_j. We are now ready to write down the equations of the GDA dynamics for a HCC game:

$$\begin{aligned} \dot{\boldsymbol{\theta}}_i &= -\nabla_{\boldsymbol{\theta}_i} L(\mathbf{F}(\boldsymbol{\theta}), \mathbf{G}(\boldsymbol{\phi})) &&= -\nabla_{\boldsymbol{\theta}_i} f_i(\boldsymbol{\theta}_i) \frac{\partial L}{\partial f_i}(\mathbf{F}(\boldsymbol{\theta}), \mathbf{G}(\boldsymbol{\phi})) \\ \dot{\boldsymbol{\phi}}_j &= \nabla_{\boldsymbol{\phi}_j} L(\mathbf{F}(\boldsymbol{\theta}), \mathbf{G}(\boldsymbol{\phi})) &&= \nabla_{\boldsymbol{\phi}_j} g_j(\boldsymbol{\phi}_j) \frac{\partial L}{\partial g_j}(\mathbf{F}(\boldsymbol{\theta}), \mathbf{G}(\boldsymbol{\phi})) \end{aligned} \tag{1}$$

2.3 Reparametrization

The following lemma is useful in studying the dynamics of hidden games.

Lemma 1. *Let $k : \mathbb{R}^d \to \mathbb{R}$ be a C^2 function. Let $h : \mathbb{R} \to \mathbb{R}$ be a C^1 function and $\mathbf{x}(t)$ denote the unique solution of the dynamical system Σ_1. Then the unique solution for dynamical system Σ_2 is $\mathbf{z}(t) = \mathbf{x}(\int_0^t h(s)\mathrm{d}s)$*

$$\left\{ \begin{array}{ccc} \dot{\mathbf{x}} & = & \nabla k(\mathbf{x}) \\ \mathbf{x}(0) & = & \mathbf{x}_{init} \end{array} \right\} : \Sigma_1 \quad \left\{ \begin{array}{ccc} \dot{\mathbf{z}} & = & h(t)\nabla k(\mathbf{z}) \\ \mathbf{z}(0) & = & \mathbf{x}_{init} \end{array} \right\} : \Sigma_2 \tag{2}$$

By choosing $h(t) = -\partial L(\mathbf{F}(t), \mathbf{G}(t))/\partial f_i$ and $h(t) = \partial L(\mathbf{F}(t), \mathbf{G}(t))/\partial g_j$ respectively, we can connect the dynamics of each $\boldsymbol{\theta}_i$ and $\boldsymbol{\phi}_j$ under Equation (1) to gradient ascent on f_i and g_j. Applying Lemma 1, we get that trajectories of $\boldsymbol{\theta}_i$ and $\boldsymbol{\phi}_j$ under Equation (1) are restricted to be subsets of the corresponding gradient ascent trajectories with the same initializations. For example, in Figure 1 $\theta_i(t)$ can not escape the purple section if it is initialized at (a) neither the orange section if it is initialiazed at (f). This limits the attainable values that $f_i(t)$ and $g_j(t)$ can take for a specific initialization. Let us thus define the following:

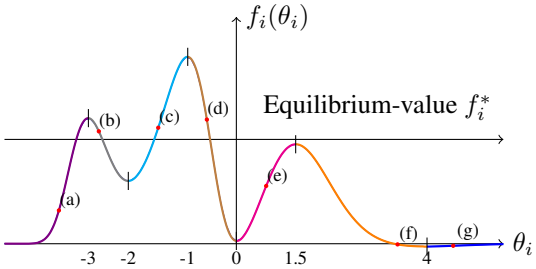

Figure 1: Neither Gradient Descent nor Ascent can traverse stationary points. An immediate consequence of Lemma 1 is that if we initialize in the above example $\theta_i(0)$ at (a), $f_i(\theta_i(t))$ can not escape the purple section. This extends to cases where $\boldsymbol{\theta}_i$ is vector of variables.

Definition 2. *For each initialization $\mathbf{x}(0)$ of Σ_1, $\mathrm{Im}_k(\mathbf{x}(0))$ is the image of $k \circ \mathbf{x} : \mathbb{R} \to \mathbb{R}$.*

Applying Definition 2 in the above example, $\mathrm{Im}_{f_i}(\theta_i(0)) = (f_i(-2), f_i(-1))$ if θ_i is initialized at (c). Additionally, observe that in each colored section $f_i(\theta_i(t))$ uniquely identifies $\theta_i(t)$. Generally, even in the case that $\boldsymbol{\theta}_i$ are vectors, Lemma 1 implies that for a given $\boldsymbol{\theta}_i(0)$, $f_i(\boldsymbol{\theta}_i(t))$ uniquely identifies $\boldsymbol{\theta}_i(t)$. As a result we get that a new dynamical system involving only f_i and g_j

Theorem 1. *For each initialization $(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0))$ of Equation (1), there are C^1 functions $X_{\boldsymbol{\theta}_i(0)}$, $X_{\boldsymbol{\phi}_j(0)}$ such that $\boldsymbol{\theta}_i(t) = X_{\boldsymbol{\theta}_i(0)}(f_i(t))$ and $\boldsymbol{\phi}_j(t) = X_{\boldsymbol{\phi}_j(0)}(g_j(t))$. If $(\boldsymbol{\theta}(t), \boldsymbol{\phi}(t))$ satisfy Equation (1) then $f_i(t) = f_i(\boldsymbol{\theta}_i(t))$ and $g_j(t) = g_j(\boldsymbol{\phi}_j(t))$ satisfy*

$$\begin{aligned} \dot{f}_i &= -\|\nabla_{\boldsymbol{\theta}_i} f_i(X_{\boldsymbol{\theta}_i(0)}(f_i))\|^2 \frac{\partial L}{\partial f_i}(\mathbf{F}, \mathbf{G}) \\ \dot{g}_j &= \|\nabla_{\boldsymbol{\phi}_j} g_j(X_{\boldsymbol{\phi}_j(0)}(g_j))\|^2 \frac{\partial L}{\partial g_j}(\mathbf{F}, \mathbf{G}) \end{aligned} \tag{3}$$

By determining the ranges of f_i and g_j, an initialization clearly dictates if a von Neumann solution is attainable. In Figure 1 for example, any point of the pink, orange or blue colored section like (e), (f) or (g) can not converge to a von Neumann solution with $f_i(\theta_i) = f_i^*$. The notion of *safety* captures which initializations can converge to a given element of Solution(L).

Definition 3. . *We will call the initialization $(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0))$ safe for a $(\mathbf{p}, \mathbf{q}) \in$ Solution(L) if $\boldsymbol{\phi}_i(0)$ and $\boldsymbol{\theta}_j(0)$ are not stationary points of f_i and g_j respectively and $p_i \in \mathrm{Im}_{f_i}(\boldsymbol{\theta}_i(0))$ and $q_j \in \mathrm{Im}_{g_j}(\boldsymbol{\phi}_j(0))$.*

Leveraging the Center-Stable Manifold Theorem [55], the following observation shows that under mild assumptions almost all initializations are safe:

Theorem 2. *If f_i and g_j have isolated stationary points, strict saddle points, compact sublevel-sets, both $p_i \in (\max \mathrm{LocalMin}(f_i), \min \mathrm{LocalMax}(f_i))$ and $q_j \in (\max \mathrm{LocalMin}(g_j), \min \mathrm{LocalMax}(g_j))$, then almost all initializations are safe for a $(\mathbf{p}, \mathbf{q}) \in$ Solution(L).*

Finally, in the following sections we use some fundamental notions of stability. We call an equilibrium $\mathbf{x}^*$ of an autonomous dynamical system $\dot{\mathbf{x}} = \mathcal{D}(\mathbf{x}(t))$ *stable* if for every neighborhood U of $\mathbf{x}^*$ there is a neighborhood V of $\mathbf{x}^*$ such that if $\mathbf{x}(0) \in V$ then $\mathbf{x}(t) \in U$ for all $t \geq 0$. We call a set S *asymptotically stable* if there exists a neighborhood $\mathcal{R}$ such that for any initialization $\mathbf{x}(0) \in \mathcal{R}$, $\mathbf{x}(t)$ approaches S as $t \to +\infty$. If $\mathcal{R}$ is the whole space the set *globally asymptotically stable*.

3 Learning in Hidden Convex Concave Games

3.1 General Case

Our main results are based on designing a Lyapunov function for the dynamics of Equation (3):

Lemma 2. *If L is convex concave and $(\boldsymbol{\phi}(0), \boldsymbol{\theta}(0))$ is a safe for $(\mathbf{p}, \mathbf{q}) \in$ Solution(L), then the following quantity is non-increasing under the dynamics of Equation* (3)*:*

$$H(\mathbf{F}, \mathbf{G}) = \sum_{i=1}^{N} \int_{p_i}^{f_i} \frac{z - p_i}{\|\nabla f_i(X_{\boldsymbol{\theta}_i(0)}(z))\|^2} \mathrm{d}z + \sum_{j=1}^{M} \int_{q_j}^{g_j} \frac{z - q_j}{\|\nabla g_j(X_{\boldsymbol{\phi}_j(0)}(z))\|^2} \mathrm{d}z \tag{4}$$

Observe that our Lyapunov function here is not the distance to $(\mathbf{p}, \mathbf{q})$ as in a classical convex concave game. The gradient terms account for the non constant multiplicative terms in Equation (3). Indeed if the game was not hidden and f_i and g_j were the identity functions then H would coincide with the Euclidean distance to $(\mathbf{p}, \mathbf{q})$. Our first theorem employs the above Lyapunov function to show that $(\mathbf{p}, \mathbf{q})$ is stable for Equation (3).

Theorem 3. *If L is convex concave and $(\boldsymbol{\phi}(0), \boldsymbol{\theta}(0))$ is a safe for $(\mathbf{p}, \mathbf{q}) \in$ Solution(L), then $(\mathbf{p}, \mathbf{q})$ is stable for Equation* (3)*.*

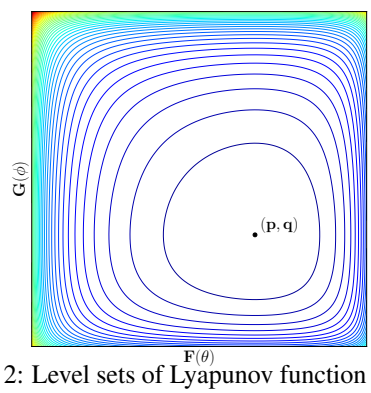

Figure 2: Level sets of Lyapunov function of Equation (4) for both $\mathbf{F}$ and $\mathbf{G}$ being one dimensional sigmoid functions.

Clearly, for the special case of globally invertible functions $\mathbf{F}, \mathbf{G}$ we could come up with an equivalent Lyapunov function in the $\boldsymbol{\theta}, \boldsymbol{\phi}$-space. In this case it is straightforward to transfer the stability results from the induced dynamical system of $\mathbf{F}, \mathbf{G}$ (Equation (3)) to the initial dynamical system of $\boldsymbol{\theta}, \boldsymbol{\phi}$ (Equation (1)). For example we can prove the following result:

Theorem 4. *If f_i and g_j are sigmoid functions and L is convex concave and there is a $(\boldsymbol{\phi}(0), \boldsymbol{\theta}(0))$ that is safe for $(\mathbf{p}, \mathbf{q}) \in$ Solution(L), then $(\mathbf{F}^{-1}(\mathbf{p}), \mathbf{G}^{-1}(\mathbf{q}))$ is stable for Equation* (1)*.*

In the general case though, stability may not be guaranteed in the parameter space of Equation (1). We will instead prove a weaker notion of stability, which we call hidden stability. Hidden stability captures that if $(\mathbf{F}(\boldsymbol{\theta}(0)), \mathbf{G}(\boldsymbol{\phi}(0)))$ is close to a von Neumann solution, then $(\mathbf{F}(\boldsymbol{\theta}(t)), \mathbf{G}(\boldsymbol{\phi}(t)))$ will remain close to that solution. Even though hidden stability is weaker, it is essentially what we are interested in, as the output space determines the utility that each player gets. Here we provide sufficient conditions for hidden stability.

Theorem 5 (Hidden Stability)**.** *Let $(\mathbf{p}, \mathbf{q}) \in$ Solution(L). Let R_{f_i} and R_{g_j} be the set of regular values*[1] *of f_i and g_j respectively. Assume that there is a $\xi > 0$ such that $[p_i - \xi, p_i + \xi] \subseteq R_{f_i}$ and $[q_j - \xi, q_j + \xi] \subseteq R_{g_j}$. Define*

$$r(t) = \|\mathbf{F}(\boldsymbol{\theta}(t)) - \mathbf{p}\|^2 + \|\mathbf{G}(\boldsymbol{\phi}(t)) - \mathbf{q}\|^2.$$

[1] A value $a \in \mathrm{Im}\, f$ is called a regular value of f if $\forall q \in \mathrm{dom}\, f : f(q) = a$, it holds $\nabla f(q) \neq \mathbf{0}$.

If f_i and g_j are proper functions[2]*, then for every $\epsilon > 0$, there is an $\delta > 0$ such that*

$$r(0) < \delta \implies \forall t \geq 0 : r(t) < \epsilon.$$

Unfortunately hidden stability still does not imply convergence to von Neumann solutions. [65] studied hidden bilinear games and proved that $\dot{H} = 0$ for this special class of HCC games. Hence, a trajectory is restricted to be a subset of a level set of H which is bounded away from the equilibrium as shown in Figure 2. To sidestep this, we will require in the next subsection the hidden game to be strictly convex concave.

3.2 Hidden strictly convex concave games

In this subsection we focus on the case where L is a strictly convex concave function. Based on Definition 1, a strictly convex concave game is not necessarily strictly convex strictly concave and thus it may have a continum of von Neumann solutions. Despite this, LaSalle's invariance principle, combined with the strict convexity concavity, allows us to prove that if $(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0))$ is safe for $Z \subseteq \text{Solution}(L)$ then Z is locally asymptotically stable for Equation (3).

Lemma 3. *Let L be strictly convex concave and $Z \subset$ Solution(L) is the non empty set of equilbria of L for which $(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0))$ is safe. Then Z is locally asymptotically stable for Equation* (3)*.*

The above lemma however does not suffice to prove that for an arbitrary initialization $(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0))$, $(\mathbf{F}(t), \mathbf{G}(t))$ approaches Z as $t \to +\infty$. In other words, *a-priori* it is unclear if $(\mathbf{F}(\boldsymbol{\theta}(0)), \mathbf{G}(\boldsymbol{\phi}(0)))$ is necessarily inside the region of attraction (ROA) of Z. To get a refined estimate of the ROA of Z, we analyze the behavior of H as f_i and g_j approach the boundaries of $\mathrm{Im}_{f_i}(\boldsymbol{\theta}_i(0))$ and $\mathrm{Im}_{g_j}(\boldsymbol{\phi}_j(0))$ and more precisely we show that the level sets of H are bounded. Once again the corresponding analysis is trivial for convex concave games, since the level sets are spheres around the equilibria.

Theorem 6. *Let L be strictly convex concave and $Z \subset$ Solution(L) is the non empty set of equilbria of L for which $(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0))$ is safe. Under the dynamics of Equation* (1) *$(\mathbf{F}(\boldsymbol{\theta}(t)), \mathbf{G}(\boldsymbol{\theta}(t)))$ converges to a point in Z as $t \to \infty$.*

The theorem above guarantees convergence to a von Neumann solution for all initializations that are safe for at least one element of Solution(L). However, this is not the same as global asymptotic stability. To get even stronger guarantees, we can assume that all initializations are safe. In this case it is straightforward to get a global asymptotic stability result:

Corollary 1. *Let L be strictly convex concave and assume that all intitializations are safe for at least one element of Solution(L). The following set is globally asymptotically stable for continuous GDA dynamics.*

$$\{(\boldsymbol{\theta}^*, \boldsymbol{\phi}^*) \in \mathbb{R}^n \times \mathbb{R}^m : (F(\boldsymbol{\theta}^*), G(\boldsymbol{\phi}^*)) \in \textit{Solution}(L)\}$$

Notice that the above approach on global asymptotic convergence using Lyapunov arguments can be extended to other popular alternative gradient-based heuristics like variations of Hamiltonian Gradient Descent. For concision, we defer the exact statements, proofs in the supplement.

3.3 Convergence via regularization

Regularization is a key technique that works both in the practice of GANs [47, 33] and in the theory of convex concave games [56, 59, 60]. Our settings of hidden convex concave games allows for provable guarantees for regularization in a wide class of settings, bringing closer practical and theoretical guarantees. Let us have a utility $L(\mathbf{x}, \mathbf{y})$ that is convex concave but not strictly. Here we will propose a modified utility L' that is strictly convex strictly concave. Specifically we will choose

$$L'(\mathbf{x}, \mathbf{y}) = L(\mathbf{x}, \mathbf{y}) + \frac{\lambda}{2}\|\mathbf{x}\|^2 - \frac{\lambda}{2}\|\mathbf{y}\|^2$$

The choice of the parameter λ captures the trade-off between convergence to the original equilibrium of L and convergence speed. On the one hand, invoking the implicit function theorem, we get that for small λ the equilibria of L are not significantly perturbed.

Theorem 7. *If L is a convex concave function with invertible Hessians at all its equilibria, then for each $\epsilon > 0$ there is a $\lambda > 0$ such that L' has equilibria that are ϵ-close to the ones of L.*

[2] A function is proper if inverse images of compact subsets are compact.

Note that invertibility of the Hessian means that L must have a unique equilibrium. On the other hand increasing λ increases the rate of convergence of safe initializations to the perturbed equilibrium.

Theorem 8. *Let $(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0))$ be a safe initialization for the unique equilibrium of L' $(\mathbf{p}, \mathbf{q})$. If*

$$r(t) = \|\mathbf{F}(\boldsymbol{\theta}(t)) - \mathbf{p}\|^2 + \|\mathbf{G}(\boldsymbol{\phi}(t)) - \mathbf{q}\|^2$$

then there are initialization dependent constants $c_0, c_1 > 0$ such that $r(t) \leq c_0 \exp(-\lambda c_1 t)$.

4 Applications

In this section, we discuss how HCC framework can be used to give new insights in a variety of application areas including min-max training for GANs and Evolutionary Game Theory. We also describe applications of regularization to normal form zero sum games in Appendix D.3.

Hidden strictly convex-concave games. We will start our discussion with the fundamental generative architecture of [25]'s GAN. In the *vanilla GAN* architecture, as it is commonly referred, our goal is to find a generator distribution p_{G} that is close to an input data distribution p_{data}. To find such a generator function, we can use a discriminator D that "criticizes" the deviations of the generator from the input data distribution. For the case of a discrete p_{data} over a set $\mathcal{N}$, the minimax problem of [25] is the following:

$$\min_{\substack{p_{\text{G}}(x) \geq 0, \\ \sum_{x \in \mathcal{N}} p_{\text{G}}(x) = 1}} \max_{D \in (0,1)^{|\mathcal{N}|}} V(G, D)$$

where $V(G, D) = \sum_{x \in \mathcal{N}} p_{\text{data}}(x) \log(D(x)) + \sum_{x \in \mathcal{N}} p_{\text{G}}(x) \log(1 - D(x))$. The problem above can be formulated as a constrained strictly convex-concave hidden game. On the one hand, for a fixed discriminator D^*, the $V(G, D^*)$ is linear over the $p_{\text{G}}(x)$. On the other hand, for a fixed generator G^*, $V(G^*, D)$ is strongly-concave. We can implement the inequality constraints on both the generator probabilities and discriminator using sigmoid activations. For the equality constraint $\sum_{x \in \mathcal{N}} p_{\text{G}}(x) = 1$ we can introduce a Langrange multiplier. Having effectively removed the constraints, we can see in Figure 3, the dynamics of Equation (1) converge to the unique equilibrium of the game, an outcome consistent with our results in Corollary 1. While the Euclidean distance to the equilibrium is not monotonically decreasing, $H(t)$ is.

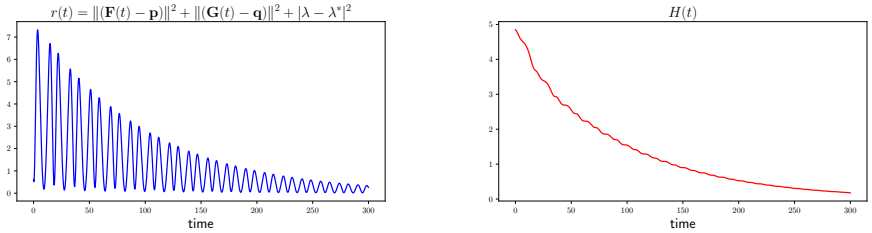

Figure 3: Both the ℓ_2 distance from the equilibrium and $H(t)$ converge to zero but only the latter does so monotonically For p_{data} we choose a fully mixed distribution of dimension $d = 4$. Given the sigmoid activations all the initializations are safe. We defer the detailed proof of convergence in Appendix

Hidden convex-concave games & Regularizaiton. An even more interesting case is Wassertein GANs–WGANs [4]. One of the contributions of [36] is to show that WGANs trained with Stochastic GDA can learn the parameters of Gaussian distributions whose samples are transformed by non-linear activation functions. It is worth mentioning that the original WGAN formulation has a Lipschitz constraint in the discriminator function. For simplicity, [36] replaced this constraint with a quadratic regularizer. The min-max problem for the case of one-dimensional Gaussian $\mathcal{N}(0, \alpha_*^2)$ and linear discriminator $D_v(x) = v^\top x$ with x^2 activation is:

$$\begin{aligned}
\min_{\alpha \in \mathbb{R}} \max_{v \in \mathbb{R}} V_{\text{WGAN}}(G_\alpha, D_v) \qquad &= \mathbb{E}_{\mathbf{X} \sim p_{\text{data}}}[D(\mathbf{X})] - \mathbb{E}_{\mathbf{X} \sim p_{\text{G}}}[D(\mathbf{X})] - v^2/2 \\
&= \mathbb{E}_{x \sim \mathcal{N}(0, \alpha_*^2)^2}[vx] - \mathbb{E}_{x \sim \mathcal{N}(0, \alpha)^2}[vx] - v^2/2 \\
&= (\alpha_*^2 - \alpha^2)v - v^2/2
\end{aligned}$$

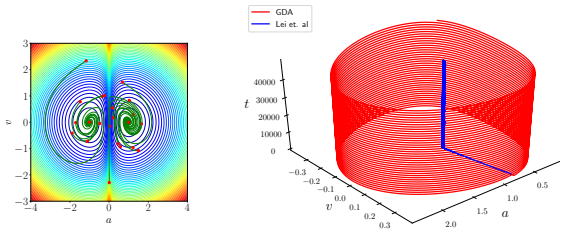

Figure 4: On the left, we show the trajectories of regularized GDA for $\alpha_*^2 = 1$ as well as the level sets of Equation (4). All trajectories (green curves-initialized at the red points) converge to one of the two equilibria $(0, 1)$ and $(0, -1)$ whereas without regularization, GDA would cycle on the level sets. In the right figure, we replace the exact expectations in V_{WGAN} with approximations via sampling and continuous time updates on α and v with discrete ones. For small learning rates and large sample sizes, unregularized GDA continues to cycle. In contrast, the regularization approach of [36] converges to the $(0, 1)$ equilibrium.

Observe that V_{WGAN} is not convex-concave but it can posed as a hidden strictly convex-concave game with $\mathbf{G}(\alpha) = (\alpha_*^2 - \alpha^2)$ and $\mathbf{F}(v) = v$. When computing expectations analytically without sampling, Theorem 6 guarantees convergence. In contrast, without the regularizer V_{WGAN} can be modeled as a hidden bilinear game and thus GDA dynamics cycle. Empirically, these results are robust to discrete and stochastic updates using sampling as shown in Figure 4. Therefore regularization in the work of [36] was a vital ingredient in their proof strategy and not just an implementation detail.

The two applications of HCC games in GANs are not isolated findings but instances of a broader pattern that connects HCC games and standard GAN formulations. As noted by [26], if updates in GAN applications were directly performed in the "functional space", i.e. the generator and discriminator outputs, then standard arguments from convex concave optimization would imply convergence to global Nash equilibria. Indeed, standard GAN formulations like the vanilla GAN [25], f-GAN [50] and WGAN [4] can all be thought of as convex concave games in the space of generator and discriminator outputs. Given that the connections between convex concave games and standard GAN objectives in the output space is missing from recent literature, in Appendix D.1 we show how one can apply Von Neumann's minimax theorem to derive the optimal generators and discriminators even in the non-realizable case. In practice, the updates happen in the parameter space and thus convexity arguments no longer apply. Our study of HCC games is a stepping stone towards bridging the gap in convergence guarantees between the case of direct updates in the output space and the parameter space.

Evolutionary Game Theory & Biology. The study of learning dynamics in games has always been strongly and inherently connected with mathematical models of biology and evolution. Typically, this line of research is studied under the name of *Evolutionary Game Theory* [28, 67]. Zero-sum games and variants thereof are of particular interest for this line of work as they encode settings of direct competition between species (e.g., prey-predator or host-parasite/virus). Even in the simplest such setting of matrix zero-sum games, used to capture competition between asexually reproducing species, it is well known that the emerging dynamics can be non-equilibrating and even chaotic [61, 58].

Studying the effects of evolutionary competition between sexually evolving species results in significantly more intricate models, as it does not suffice to merely keep track of the fractions of the different types of individuals that self-replicate. Instead it is necessary to keep a much more detailed accounting of the evolution of the frequencies of different genes that get reshuffled and recombined to create new individuals, whilst giving evolutionary preference to the most fit individuals given the current environment. Recent work on intersection of learning theory and game theory has provided concrete such game theoretic models [37, 13, 44, 43]. Due to the intricate nature of their dynamics, deciding even the simplest questions in regards to them (e.g. does genetic diversity survive or not?) is typically computationally hard [42].

A notable exception, where the dynamics of sexual evolution and, in fact, sexual competition have been relatively thoroughly understood, is the work of [57], on two species (host-parasite) antagonism. The outcome of this competition depends on their respective phenotypes (informally their properties, e.g., large wings versus small wings.) of the two species. The crucial assumption that makes this

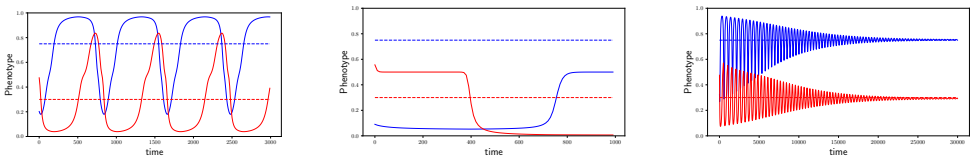

Figure 5: The above figures describes the evolution of the expected phenotype for two species A,B. The left one corresponds to a safe initialization leading to periodic trajectories. The middle one corresponds to an unsafe initialization where $\theta_1(0) = \theta_2(0)$. The dynamics converge albeit to a spurious equilibrium, that is different from the hidden game equilibrium (dash lines). Finally, the right one corresponds to a safe initialization of the regularized game, which converges to a slightly perturbed equilibrium (Theorem 7).

model theoretically tractable is that the phenotype for each species is a Boolean attribute (this assumption is also used [38]). Despite these simplifications, the dynamics are still *not equilibrating* and are, in fact, cyclic for almost all initial conditions. Two natural questions emerge: 1) *Is the almost everywhere condition necessary? I.e. Do there exist initial conditions which are not cyclic?* 2) *More importantly, can a slightly perturbed dynamic stabilize these systems and converge to a meaningful equilibrium?* Next, we will see how our framework addresses both of these questions.

To understand the connection these we will examine the model of [57] in more detail. Concretely, the phenotype of species A, B can be described as a Boolean function over the species genome which is encoded by a binary string (this acts as a simplified version of a DNA string). While the *phenotype* plays the dominant role for the survival of the species, sexual reproduction modifies only the *genotype* of an organism. As a result the species are actually involved in a hidden zero-sum game. More formally, each species is game-theoritically represented as a team of agents where each agent controls one bit of the genotype:

$$\begin{aligned}
&\mathsf{G}_A = (\mathsf{g}_1^A, \cdots, \mathsf{g}_n^A), \mathsf{G}_B = (\mathsf{g}_1^B, \cdots, \mathsf{g}_m^B)\\
&\mathsf{u}_A = L[\mathsf{Phenotype}_A(\mathsf{G}_A), \mathsf{Phenotype}_B(\mathsf{G}_B)]\\
&\mathsf{u}_B = -\mathsf{u}_A
\end{aligned}$$

Where $\mathsf{g}_i^A, \mathsf{g}_j^B \in \{0, 1\}$, $\mathsf{Phenotype}_A, \mathsf{Phenotype}_B$ is a Boolean function (e.g., $\mathsf{AND}, \mathsf{XOR}$) and L is a 2×2 matrix encoding a zero-sum game (e.g., Matching Pennies). Naturally, one can allow agents to use randomized/mixed strategies in which case the expected utilities of all agents/genes are defined using the standard multi-linear extension of utilities. Thus, these models of evolutionary sexual competition share the same basic structure as hidden linear-linear games, which explains their recurrent, non-equilibrating nature [65].

In Figure 5, each gene/agent g_i^A tunes one real variable θ_i such that $\Pr[\mathsf{g}_i^A = 1] = \sigma(\theta_i)$ and gene/agent g_j^B tunes one real variable ϕ_j correspondingly. Choosing as Boolean phenotype to be the XOR of two genes, almost all initializations are safe for any bilinear game with a mixed equilibrium. Actually, only the case $\theta_1(0) = \theta_2(0)$ or $\phi_1(0) = \phi_2(0)$ can be problematic, since for XOR the expected phenotype is bounded in $[0, 0.5]$ and a mixed equilibrium out of this range would be infeasible. Finally, leveraging Theorem 7, we can design a regularized version of the game such that the dynamics converge arbitrarily close to the true von Neumann solution of these games, which is encoded by the min-max strategies of the hidden bi-linear zero-sum game.

5 Discussion & Future Work

While this work is a promising first step towards understanding GAN training, significant challenges remain. Neural network architectures do not use disjoint set of parameters for each of the outputs. Additionally, the hidden competition of GANs can take place in an output space of probability distributions and classifiers whose vector spaces' dimension is typically infinite. On the bright side, we achieve to estabish last iterate convergence results which are, to the best of our knowledge, the first result of their kind for a wide class of non-convex non-concave games. Beyond ML applications, we believe that our framework could provide even further insights for evolutionary game theory, mathematical biology as well as team-zero-sum games. For example an interesting hybrid class of games could be network generalizations of team-zero-sums games, e.g. by combining [11] and [57].

References

[1] Jacob Abernethy, Kevin A Lai, Kfir Y Levy, and Jun-Kun Wang. Faster rates for convex-concave games. In *COLT*, 2018.

[2] Jacob Abernethy, Kevin A Lai, and Andre Wibisono. Last-iterate convergence rates for min-max optimization. *arXiv preprint arXiv:1906.02027*, 2019.

[3] Leonard Adolphs, Hadi Daneshmand, Aurélien Lucchi, and Thomas Hofmann. Local saddle point optimization: A curvature exploitation approach. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, pages 486–495, 2019.

[4] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017.

[5] James Bailey and Georgios Piliouras. Fast and furious learning in zero-sum games: vanishing regret with non-vanishing step sizes. In *Advances in Neural Information Processing Systems*, pages 12977–12987, 2019.

[6] James P Bailey, Gauthier Gidel, and Georgios Piliouras. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *COLT*, 2020.

[7] James P. Bailey and Georgios Piliouras. Multiplicative weights update in zero-sum games. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 321–338, 2018.

[8] David Balduzzi, Sébastien Racanière, James Martens, Jakob N. Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 363–372. PMLR, 2018.

[9] Sanjay P. Bhat and Dennis S. Bernstein. Nontangency-based lyapunov tests for convergence and stability in systems having a continuum of equilibria. *SIAM J. Control and Optimization*, 42(5):1745–1775, 2003.

[10] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, USA, 2004.

[11] Yang Cai and Costantinos Daskalakis. On minmax theorems for multiplayer games. In *ACM-SIAM Symposium on Discrete Algorithms*, SODA, pages 217–234, 2011.

[12] Nikolo Cesa-Bianchi and Gabor Lugoisi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[13] Erick Chastain, Adi Livnat, Christos Papadimitriou, and Umesh Vazirani. Algorithms, games, and evolution. *Proceedings of the National Academy of Sciences (PNAS)*, 111(29):10620–10623, 2014.

[14] Yun Kuen Cheung and Georgios Piliouras. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In *COLT*, 2019.

[15] Yun Kuen Cheung and Georgios Piliouras. Chaos, extremism and optimism: Volume analysis of learning in games. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

[16] Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training GANs with optimism. In *ICLR*, 2018.

[17] Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 9256–9266, 2018.

[18] Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, pages 27:1–27:18, 2019.

[19] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. *CoRR*, abs/2009.09623, 2020.

[20] Lawrence C Evans. Partial differential equations and monge-kantorovich mass transfer. *Current developments in mathematics*, 1997(1):65–126, 1997.

[21] Ky Fan. Minimax theorems. *Proceedings of the National Academy of Sciences of the United States of America*, 39(1):42, 1953.

[22] Farzan Farnia and Asuman E. Ozdaglar. Do gans always have nash equilibria? In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 3029–3039. PMLR, 2020.

[23] Tanner Fiez, Benjamin Chasnov, and Lillian Ratliff. Implicit learning dynamics in stackelberg games: Equilibria characterization, convergence analysis, and empirical study. In *ICML*, 2020.

[24] Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *ICLR*, 2019.

[25] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[26] Ian J. Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *CoRR*, abs/1701.00160, 2017.

[27] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2672–2680, 2014.

[28] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, 1998.

[29] 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*, pages 6938–6948, 2019.

[30] Yu-Guan Hsieh, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. Explore aggressively, update conservatively: Stochastic extragradient methods with variable stepsize scaling. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

[31] Chi Jin, Praneeth Netrapalli, and Michael I. Jordan. Minmax optimization: Stable limit points of gradient descent ascent are locally optimal. *CoRR*, abs/1902.00618, 2019.

[32] Hassan K Khalil. *Nonlinear systems; 3rd ed.* Prentice-Hall, Upper Saddle River, NJ, 2002.

[33] Karol Kurach, Mario Lucic, Xiaohua Zhai, Marcin Michalski, and Sylvain Gelly. A large-scale study on regularization and normalization in gans. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 3581–3590. PMLR, 2019.

[34] Jason D. Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I. Jordan, and Benjamin Recht. First-order methods almost always avoid saddle points. *CoRR*, abs/1710.07406, 2017.

[35] Jason D. Lee, Max Simchowitz, Michael I. Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, volume 49 of *JMLR Workshop and Conference Proceedings*, pages 1246–1257. JMLR.org, 2016.

[36] Qi Lei, Jason Lee, Alex Dimakis, and Constantinos Daskalakis. Sgd learns one-layer networks in wgans. In *International Conference on Machine Learning*, pages 5799–5808. PMLR, 2020.

[37] Adi Livnat, Christos Papadimitriou, Jonathan Dushoff, and Marcus W. Feldman. A mixability theory for the role of sex in evolution. *Proceedings of the National Academy of Sciences*, 105(50):19803–19808, 2008.

[38] Adi Livnat, Christos Papadimitriou, Aviad Rubinstein, Andrew Wan, and Gregory Valiant. Satisfiability and evolution. In *FOCS*, 2014.

[39] Chris J. Maddison, Daniel Paulin, Yee Whye Teh, Brendan O'Donoghue, and Arnaud Doucet. Hamiltonian descent methods. *CoRR*, abs/1809.05042, 2018.

[40] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.

[41] Eric Mazumdar and Lillian J Ratliff. Local nash equilibria are isolated, strict local nash equilibria in 'almost all' zero-sum continuous games. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 6899–6904. IEEE, 2019.

[42] R. Mehta, I. Panageas, G. Piliouras, and S. Yazdanbod. The Computational Complexity of Genetic Diversity. *European Symposium on Algorithms (ESA)*, 2016.

[43] Ruta Mehta, Ioannis Panageas, and Georgios Piliouras. Natural selection as an inhibitor of genetic diversity: Multiplicative weights updates algorithm and a conjecture of haploid genetics. In *Innovations in Theoretical Computer Science*, 2015.

[44] Reshef Meir and David Parke. A note on sex, evolution, and the multiplicative updates algorithm. In *AAMAS*, 2015.

[45] Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra(-gradient) mile. In *ICLR*, 2019.

[46] Panayotis Mertikopoulos, Christos H. Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2703–2717, 2018.

[47] Lars M. Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 3478–3487. PMLR, 2018.

[48] Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *International Conference on Artificial Intelligence and Statistics*, pages 1497–1507. PMLR, 2020.

[49] Yurii E. Nesterov. *Introductory Lectures on Convex Optimization - A Basic Course*, volume 87 of *Applied Optimization*. Springer, 2004.

[50] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 271–279, 2016.

[51] Brendan O'Donoghue and Chris J. Maddison. Hamiltonian descent for composite objectives. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 14443–14453, 2019.

[52] Ioannis Panageas and Georgios Piliouras. Gradient descent only converges to minimizers: Non-isolated critical points and invariant regions. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, pages 2:1–2:12, 2017.

[53] Ioannis Panageas and Georgios Piliouras. Gradient descent only converges to minimizers: Non-isolated critical points and invariant regions. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPIcs*, pages 2:1–2:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[54] Ioannis Panageas, Georgios Piliouras, and Xiao Wang. First-order methods almost always avoid saddle points: The case of vanishing step-sizes. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 6471–6480, 2019.

[55] Lawrence Perko. *Differential Equations and Dynamical Systems*. Springer, 3nd. edition, 1991.

[56] Julien Pérolat, Rémi Munos, Jean-Baptiste Lespiau, Shayegan Omidshafiei, Mark Rowland, Pedro A. Ortega, Neil Burch, Thomas W. Anthony, David Balduzzi, Bart De Vylder, Georgios Piliouras, Marc Lanctot, and Karl Tuyls. From poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization. *CoRR*, abs/2002.08456, 2020.

[57] Georgios Piliouras and Leonard J Schulman. Learning dynamics and the co-evolution of competing sexual species. In *ITCS*, 2018.

[58] Georgios Piliouras and Jeff S. Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via poincaré recurrence. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 861–873, 2014.

[59] Kevin Roth, Aurélien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 2018–2028, 2017.

[60] Maziar Sanjabi, Jimmy Ba, Meisam Razaviyayn, and Jason D. Lee. On the convergence and robustness of training gans with regularized optimal transport. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 7091–7101, 2018.

[61] Yuzuru Sato, Eizo Akiyama, and J. Doyne Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99(7):4748–4751, 2002.

[62] Leonard Schulman and Umesh V Vazirani. The duality gap for two-team zero-sum games. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[63] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.

[64] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

[65] Emmanouil-Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, and Georgios Piliouras. Poincaré recurrence, cycles and spurious equilibria in gradient-descent-ascent for non-convex non-concave zero-sum games. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 10450–10461, 2019.

[66] Jun-Kun Wang and Jacob D Abernethy. Acceleration through optimistic no-regret dynamics. In *Advances in Neural Information Processing Systems*, pages 3824–3834, 2018.

[67] J. W. Weibull. *Evolutionary Game Theory*. MIT Press; Cambridge, MA: Cambridge University Press., 1995.

[68] Guojun Zhang, Pascal Poupart, and Yaoliang Yu. Optimality and Stability in Non-Convex-Non-Concave Min-Max Optimization. *arXiv e-prints*, page arXiv:2002.11875, February 2020.

[69] Guojun Zhang and Yaoliang Yu. Convergence of gradient methods on bilinear zero-sum games. In *International Conference on Learning Representations*, 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]
   (c) Did you discuss any potential negative societal impacts of your work? [No] ***We do not see any ethical or future societal consequences of this work, since its motivation is strongly theoretical.***
   (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]
   (b) Did you include complete proofs of all theoretical results?[Yes]
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]
   (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)? [No] ***The presented experiments are used for illustrative purposes and only to validate the theoretical findings which are the core results of this work.***
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]