
Poincaré Recurrence, Cycles and Spurious Equilibria in Gradient-Descent-Ascent for Non-Convex Non-Concave Zero-Sum Games

Lampros Flokas*
Department of Computer Science
Columbia University
New York, NY 10025
lamflokas@cs.columbia.edu

Emmanouil V. Vlatakis-Gkaragkounis*
Department of Computer Science
Columbia University
New York, NY 10025
emvlatakis@cs.columbia.edu

Georgios Piliouras
Engineering Systems and Design
Singapore University of Technology and Design
Singapore
georgios@sutd.edu.sg

Abstract

We study a wide class of non-convex non-concave min-max games that generalizes over standard bilinear zero-sum games. In this class, players control the inputs of a smooth function whose output is being applied to a bilinear zero-sum game. This class of games is motivated by the indirect nature of the competition in Generative Adversarial Networks, where players control the parameters of a neural network while the actual competition happens between the distributions that the generator and discriminator capture. We establish theoretically, that depending on the specific instance of the problem gradient-descent-ascent dynamics can exhibit a variety of behaviors antithetical to convergence to the game theoretically meaningful min-max solution. Specifically, different forms of recurrent behavior (including periodicity and Poincaré recurrence) are possible as well as convergence to spurious (non-min-max) equilibria for a positive measure of initial conditions. At the technical level, our analysis combines tools from optimization theory, game theory and dynamical systems.

1 Introduction

Min-max optimization is a problem of interest in several communities including Optimization, Game Theory and Machine Learning. In its most general form, given an objective function $r : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ and we would like to solve the following problem

$$(\theta^*, \phi^*) = \arg \min_{\theta \in \mathbb{R}^n} \arg \max_{\phi \in \mathbb{R}^m} r(\theta, \phi). \quad (1)$$

This problem is much more complicated compared to classical minimization problems, as even understanding under which conditions such a solution is meaning-full is far from trivial [DP18, MPR⁺17, OSG⁺18, JNJ19]. What is even more demanding is understanding what kind of algorithms/dynamics are able to solve this problem when a solution is well defined.

*Equal contribution

Recently this problem has attracted renewed interest motivated by the advent of Generative Adversarial Networks (GANs) and their numerous applications [GPM⁺14, RMC16, IZZE17, GPM⁺14, ZXL17, ACB17, LTH⁺17, SGZ⁺16]. A classical GAN architecture mainly revolves around the competition between two players, the generator and the discriminator. On the one hand, the generator aims to train a neural network based generative model that can generate high fidelity samples from a target distribution. On the other hand, the discriminator’s goal is to train a neural network classifier than can distinguish between the samples of the target distribution and artificially generated samples. While one could consider each of the tasks in isolation, it is the competitive interaction between the generator and the discriminator that has lead to the resounding success of GANs. It is the "criticism" from a powerful discriminator that pushes the generator to capture the target distribution more accurately and it is the access to high fidelity artificial samples from a good generator that gives rise to better discriminators. Machine Learning researchers and practitioners have tried to formalize this competition using the min-max optimization framework mentioned above with great success [AGL⁺17, Ma18, GXC⁺18, YFW⁺19].

One of the main limitations of this framework however is that to this day efficiently training GANs can be a notoriously difficult task [SGZ⁺16, MPPS17, MPP18, KAHK17]. Addressing this limitation has been the object of interest for a long line work in the recent years [MGN18, MPPS17, PV16, RMC16, TGB⁺17, BSM17, GAA⁺17]. Despite the intensified study, very little is known about efficiently solving general min-max optimization problems. Even for the relatively simple case of bilinear games, the little results that are known have usually a negative flavour. For example, the continuous time analogue of standard game dynamics such as gradient-descent-ascent or multiplicative weights lead to cyclic or recurrent behavior [PS14, MPP18] whereas when they are actually run in discrete-time² they lead to divergence and chaos [BP18, CP19, BP19a]. While positive results for the case of bilinear games exist, like extra-gradient (optimistic) training ([DISZ18, MLZ⁺19a, DP19]) and other techniques [BRM⁺18, GHP⁺19a, GBV⁺19, ALW19], these results fail to generalize to complex non-convex non-concave settings [OSG⁺18, LLRY18, SRL18]. In fact, for the case of non-convex-concave optimization, game theoretic interpretations of equilibria might not even be meaningful [MR18, JNJ19, ADLH19].

In order to shed some light to this intellectually challenging problem, we propose a quite general class of min-max optimization problems that includes bilinear games as well as a wide range of non-convex non-concave games. In this class of problems, each player submits its own decision vector just like in general min-max optimization problems. Then each decision vector is processed separately by a (potentially different) smooth function. Each player finally gets rewarded by plugging in the processed decision vectors to a simple bilinear game. More concretely, there are functions $F : \mathbb{R}^n \rightarrow \mathbb{R}^N$ and $G : \mathbb{R}^m \rightarrow \mathbb{R}^M$ and a matrix $U_{N \times M}$ such that

$$r(\boldsymbol{\theta}, \boldsymbol{\phi}) = F(\boldsymbol{\theta})^\top U G(\boldsymbol{\phi}). \tag{2}$$

We call the resulting class of problems Hidden Bilinear Games.

The motivation behind the proposed class of games is actually the setting of training GANs itself. During the training process of GANs, the discriminator and the generator "submit" the parameters of their corresponding neural network architectures, denoted as $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ in our problem formulation. However, deep networks introduce nonlinearities in mapping their parameters to their output space which we capture through the non-convex functions F, G . Thus, even though hidden bilinear games do not demonstrate the full complexity of modern GAN architectures and training, they manage to capture two of its most pervasive properties: *i) the indirect competition of the generator and the discriminator* and *ii) the non-convex non-concave nature of training GANs*. Both features are markedly missing from simple bilinear games.

Our results. We provide, the first to our own knowledge, global analysis of gradient-descent-ascent for a class of non-convex non-concave zero-sum games that by design includes both features of bilinear zero-sum games as well as of single-agent non-convex optimization. Our analysis focuses on the (smoother) continuous time dynamics (Section 4,5) but we also discuss the implications for discrete time (Section 7). The unified thread of our results is that gradient-descent-ascent can exhibit a variety of behaviors antithetical to convergence to the min-max solution. In fact, convergence to a set of parameters that implement the desired min-max solution (as e.g. GANs require), if it actually

²Interestingly, running alternating gradient-descent-ascent in discrete-time results once again in recurrent behavior [BGP19].

happens, is more of an accident due to fortuitous system initialization rather than an implication of the adversarial network architecture.

Informally, we prove that these dynamics exhibit conservation laws, akin to energy conservation in physics. Thus, in contrast to them making progress over time their natural tendencies is to "cycle" through their parameter space. If the hidden bilinear game U is 2×2 (e.g. Matching Pennies) with an interior Nash equilibrium, then the behavior is typically periodic (Theorem 3). If it is a higher dimensional game (e.g. akin to Rock-Paper-Scissors) then even more complex behavior is possible. Specifically, the system is formally analogous to Poincaré recurrent systems (e.g. many body problem in physics) (Theorems 6, 7). Due to the non-convexity of the operators F, G , the system can actually sometimes get stuck at equilibria, however, these fixed points may be merely artifacts of the nonlinearities of F, G instead of meaningful solutions to the underline minmax problem U . (Theorem 8).

In Section 7, we show that moving from continuous to discrete time, only enhances the disequilibrium properties of the dynamics. Specifically, instead of energy conservation now energy increases over time leading away from equilibrium (Theorem 9), whilst spurious (non-minmax) equilibria are still an issue (Theorem 10). Despite these negative results, there are some positive news, as at least in some cases we can show that time-averaging over these non-equilibrium trajectories (or equivalently choosing a distribution of parameters instead of a single set of parameters) can recover the min-max equilibrium (Theorem 4). Technically our results combine tools from dynamical systems (e.g. Poincaré recurrence theorem, Poincaré-Bendixson theorem, Liouville's theorem) along with tools from game theory and non-convex optimization.

Understanding the intricacies of GAN training requires broadening our vocabulary and horizons in terms of what type of long term behaviors are possible and developing new techniques that can hopefully counter them.

The structure of the rest of the paper is as follows. In Section 2 we will present key results from prior work on the problem of min-max optimization. In Section 3 we will present the main mathematical tools for our analysis. Sections 4 through 6 will be devoted to studying interesting special cases of hidden bilinear games. Section 8 will be the conclusion of our work.

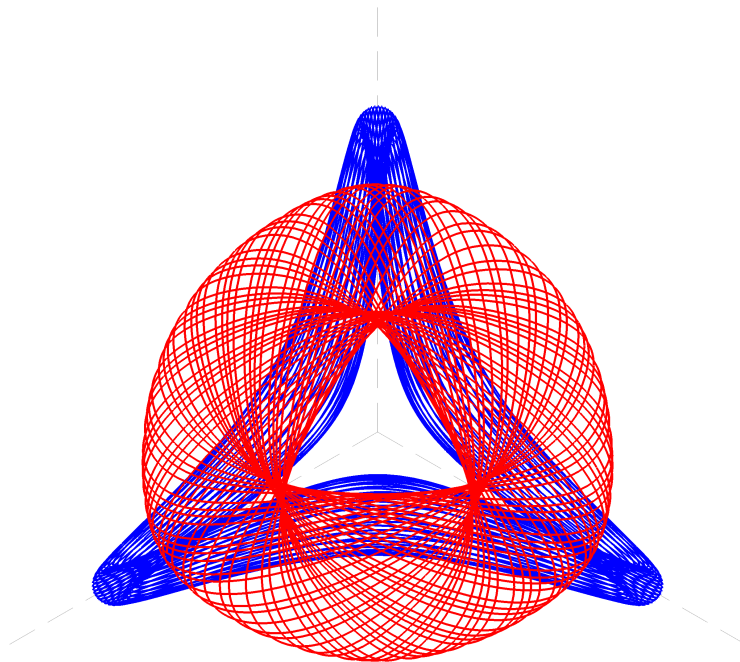


Figure 1: Trajectories of a single player using gradient-descent-ascent dynamics for a hidden Rock-Paper-Scissors game with sigmoid activations. The different colors correspond to different initializations of the dynamics. The trajectories exhibit Poincaré recurrence as expected by Theorem 7.

2 Related Work

Non-equilibrating dynamics in game theory. [KLPT11] established non-convergence for a continuous-time variant of Multiplicative Weights Update (MWU), known as the replicator dynamic, for a $2 \times 2 \times 2$ game and showed that as a result the system converges to states whose social welfare dominates that of all Nash equilibria. [PPP17] proved the existence of Li-Yorke chaos in MWU dynamics of 2×2 potential games. From the perspective of evolutionary game theory, which typically studies continuous time dynamics, numerous nonconvergence results are known but again typically for small games, e.g., [San10]. [PS14] shows that replicator dynamics exhibit a specific type of near periodic behavior in bilinear (network) zero-sum games, which is known as Poincaré recurrence. Recently, [MPP18] generalized these results to more general continuous time variants of FTRL dynamics (e.g. gradient-descent-ascent). Cycles arise also in evolutionary team competition [PS18] as well as in network competition [NMP18]. Technically, [PS18] is the closest paper to our own as it studies evolutionary competition between Boolean functions, however, the dynamics in the two models are different and that paper is strictly focused on periodic systems. The papers in the category of cyclic/recurrent dynamics combine delicate arguments such as volume preservation and the existence of constants of motions (“energy preservation”). In this paper we provide a wide generalization of these type of results by establishing cycles and recurrence type of behavior for a large class of non-convex non-concave games. In the case of discrete time dynamics, such as standard gradient-descent-ascent, the system trajectories are first order approximations of the above motion and these conservation arguments do not hold exactly. Instead, even in bilinear games, the “energy” slowly increases over time [BP18] implying chaotic divergence away from equilibrium [CP19]. We extend such energy increase results to non-linear settings.

Learning in zero-sum games and connections to GANs. Several recent papers have shown positive results about convergence to equilibria in (mostly bilinear) zero-sum games for suitable adapted variants of first-order methods and then apply these techniques to Generative Adversarial Networks (GANs) showing improved performance (e.g. [DISZ18, DP19]). [BRM⁺18] made use of conservation laws of learning dynamics in zero-sum games (e.g. [BP19b]) to develop new algorithms for training GANs that add a new component to the vector field that aims at minimizing this energy function. Different energy shrinking techniques for convergence in GANs (non-convex saddle point problems) exploit connections to variational inequalities and employ mirror descent techniques with an extra gradient step [GBVL18, MLZ⁺19a]. Moreover, adding negative momentum can help with stability in zero-sum games [GHP⁺19b]. Game theoretic inspired methods such as time-averaging work well in practice for a wide range of architectures [YFW⁺19].

3 Preliminaries

3.1 Notation

Vectors are denoted in boldface \mathbf{x}, \mathbf{y} unless otherwise indicated are considered as column vectors. We use $\|\cdot\|$ corresponds to denote the ℓ_2 -norm. For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ we use ∇f to denote its gradient. For functions of two vector arguments, $f(\mathbf{x}, \mathbf{y}) : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$, 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. The term “sigmoid” function refers to $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ such that $\sigma(x) = (1 + e^{-x})^{-1}$. Finally, we use $\mathcal{P}(\cdot)$, operating over a set, to denote its (Lebesgue) measure.

3.2 Definitions

Definition 1 (Hidden Bilinear Zero-Sum Game). *In a hidden bilinear zero-sum game there are two players, each one equipped with a smooth function $\mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^N$ and $\mathbf{G} : \mathbb{R}^m \rightarrow \mathbb{R}^M$ and a payoff matrix $U_{N \times M}$ such that each player inputs its own decision vector $\boldsymbol{\theta} \in \mathbb{R}^n$ and $\boldsymbol{\phi} \in \mathbb{R}^m$ and is trying to maximize or minimize $r(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbf{F}(\boldsymbol{\theta})^\top U \mathbf{G}(\boldsymbol{\phi})$ respectively.*

In this work we will mostly study continuous time dynamics of solutions for the problem of Equation 1 for hidden bilinear zero-sum games but we will also make some important connections to discrete time dynamics that are also prevalent in practice. In order to make this distinction clear, let us define the following terms.

Definition 2 (Continuous Time Dynamical System). *A system of ordinary differential equations $\dot{\mathbf{x}} = f(\mathbf{x})$ where $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ will be called a continuous time dynamical system. Solutions of the equation $f(\mathbf{x}) = 0$ are called the fixed points of the dynamical system.*

We will call f the *vector field* of the dynamical system. In order to understand the properties of continuous time dynamical systems, we will often need to study their behaviour given different initial conditions. This behaviour is captured by the flow of the dynamical system. More precisely,

Definition 3. *If f is Lipschitz-continuous, there exists a continuous map $\Phi(\mathbf{x}_0, t) : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ called flow of the dynamical system such that for all $\mathbf{x}_0 \in \mathbb{R}^d$ we have that $\Phi(\mathbf{x}_0, t)$ is the unique solution of the problem $\{\dot{\mathbf{x}} = f(\mathbf{x}), \mathbf{x}(0) = \mathbf{x}_0\}$. We will refer to $\Phi(\mathbf{x}_0, t)$ as a trajectory or orbit of the dynamical system.*

In this work we will be mainly study the gradient-descent-ascent dynamics for the problem of Equation 1. The continuous (discrete) time version of the dynamics (with learning rate α) are based on the following equations:

$$(\text{CGDA}) : \begin{cases} \dot{\boldsymbol{\theta}} &= -\nabla_{\boldsymbol{\theta}} r(\boldsymbol{\theta}, \boldsymbol{\phi}) \\ \dot{\boldsymbol{\phi}} &= \nabla_{\boldsymbol{\phi}} r(\boldsymbol{\theta}, \boldsymbol{\phi}) \end{cases} \quad (\text{DGDA}) : \begin{cases} \boldsymbol{\theta}_{k+1} &= \boldsymbol{\theta}_k - \alpha \nabla_{\boldsymbol{\theta}} r(\boldsymbol{\theta}_k, \boldsymbol{\phi}_k) \\ \boldsymbol{\phi}_{k+1} &= \boldsymbol{\phi}_k + \alpha \nabla_{\boldsymbol{\phi}} r(\boldsymbol{\theta}_k, \boldsymbol{\phi}_k) \end{cases}$$

A key notion in our analysis is that of (Poincaré) recurrence. Intuitively, a dynamical system is recurrent if, after a sufficiently long (but finite) time, almost every state returns arbitrarily close to the system's initial state.

Definition 4. *A point $\mathbf{x} \in \mathbb{R}^d$ is said to be recurrent under the flow Φ , if for every neighborhood $U \subseteq \mathbb{R}^d$ of \mathbf{x} , there exists an increasing sequence of times t_n such that $\lim_{n \rightarrow \infty} t_n = \infty$ and $\Phi(\mathbf{x}, t_n) \in U$ for all n . Moreover, the flow Φ is called Poincaré recurrent in non-zero measure set $A \subseteq \mathbb{R}^d$ if the set of the non-recurrent points in A has zero measure.*

4 Cycles in hidden bilinear games with two strategies

In this section we will focus on a particular case of hidden biinear games where both the generator and the discriminator play only two strategies. Let U be our zero-sum game and without loss of generality we can assume that there are functions $f : \mathbb{R}^n \rightarrow [0, 1]$ and $g : \mathbb{R}^m \rightarrow [0, 1]$ such that

$$\mathbf{F}(\boldsymbol{\theta}) = \begin{pmatrix} f(\boldsymbol{\theta}) \\ 1 - f(\boldsymbol{\theta}) \end{pmatrix} \quad U = \begin{pmatrix} u_{0,0} & u_{0,1} \\ u_{1,0} & u_{1,1} \end{pmatrix} \quad \mathbf{G}(\boldsymbol{\phi}) = \begin{pmatrix} g(\boldsymbol{\phi}) \\ 1 - g(\boldsymbol{\phi}) \end{pmatrix}$$

Let us assume that the hidden bi-linear game has a unique mixed Nash equilibrium (p, q) :

$$v = u_{0,0} - u_{0,1} - u_{1,0} + u_{1,1} \neq 0, \quad q = -\frac{u_{0,1} - u_{1,1}}{v} \in (0, 1), \quad p = -\frac{u_{1,0} - u_{1,1}}{v} \in (0, 1)$$

Then we can write down the equations of gradient-descent-ascent :
$$\begin{cases} \dot{\boldsymbol{\theta}} &= -v \nabla f(\boldsymbol{\theta})(g(\boldsymbol{\phi}) - q) \\ \dot{\boldsymbol{\phi}} &= v \nabla g(\boldsymbol{\phi})(f(\boldsymbol{\theta}) - p) \end{cases} \quad (3)$$

In order to analyze the behavior of this system, we would like to understand the topology of the trajectories of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, at least individually. The following lemma makes a connection between the trajectories of each variable in the min-max optimization system of Equation 3 and simple gradient ascent dynamics.

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

$$\left\{ \begin{array}{l} \dot{\mathbf{x}} = \nabla k(\mathbf{x}) \\ \mathbf{x}(0) = \mathbf{x}_0 \end{array} \right\} : \Sigma_1 \quad \left\{ \begin{array}{l} \dot{\mathbf{z}} = h(t) \nabla k(\mathbf{z}) \\ \mathbf{z}(0) = \mathbf{x}_0 \end{array} \right\} : \Sigma_2$$

By applying the previous result for $\boldsymbol{\theta}$ with $k = f$ and $h(t) = -v(g(\boldsymbol{\phi}(t)) - q)$, we get that even under the dynamics of Equation 3, $\boldsymbol{\theta}$ remains on a trajectory of the simple gradient ascent dynamics with initial condition $\boldsymbol{\theta}(0)$. This necessarily affects the possible values of f and g given the initial conditions. Let us define the sets of values attainable for each initialization.

Definition 5. For each $\theta(0)$, $f_{\theta(0)}$ is the set of possible values of $f(\theta(t))$ can attain under gradient ascent dynamics. Similarly, we define $g_{\phi(0)}$ the corresponding set for g .

What is special about the trajectories of gradient ascent is that along this curve f is strictly increasing (For a detailed explanation, reader could check the proof of Theorem 1 in the Appendix) and therefore each point $\theta(t)$ in the trajectory has a unique value for f . Therefore even in the system of Equation 3, $f(\theta(t))$ uniquely identifies $\theta(t)$. This can be formalized in the next theorem.

Theorem 1. For each $\theta(0), \phi(0)$, under the dynamics of Equation 3, there are C^1 functions $(X_{\theta(0)}, X_{\phi(0)})$ such that $X_{\theta(0)} : f_{\theta(0)} \rightarrow \mathbb{R}^n, X_{\phi(0)} : g_{\phi(0)} \rightarrow \mathbb{R}^n$ and $\theta(t) = X_{\theta(0)}(f(t)), \phi(t) = X_{\phi(0)}(g(t))$.

Equipped with these results, we are able to reduce this complicated dynamical system of θ and ϕ to a planar dynamical system involving f and g alone.

Lemma 2. If $\theta(t)$ and $\phi(t)$ are solutions to Equation 3 with initial conditions $(\theta(0), \phi(0))$, then we have that $f(t) = f(\theta(t))$ and $g(t) = g(\phi(t))$ satisfy the following equations

$$\begin{aligned}\dot{f} &= -v \|\nabla f(X_{\theta(0)}(f))\|^2 (g - q) \\ \dot{g} &= v \|\nabla g(X_{\phi(0)}(g))\|^2 (f - p)\end{aligned}\tag{4}$$

As one can observe both from Equation 3 and Equation 4, fixed points of the gradient-descent-ascent dynamics correspond to either solutions of $f(\theta) = p$ and $g(\phi) = q$ or stationary points of f and g or even some combinations of the aforementioned conditions. Although, all of them are fixed points of the dynamical system, only the former equilibria are game theoretically meaningful. We will therefore define a subset of initial conditions for Equation 3 such that convergence to game theoretically meaningful fixed points may actually be feasible:

Definition 6. We will call the initialization $(\theta(0), \phi(0))$ safe for Equation 3 if $\theta(0)$ and $\phi(0)$ are not stationary points of f and g respectively and $p \in f_{\theta(0)}$ and $q \in g_{\phi(0)}$.

For safe initial conditions we can show that gradient-descent-ascent dynamics applied in the class of the hidden bilinear zero-sum game mimic properties and behaviors of conservative/Hamiltonian physical systems [BP19b], like an ideal pendulum or an ideal spring-mass system. In such systems, there is a notion of energy that remains constant over time and hence the system trajectories lie on level sets of these functions. To motivate further this intuition, it is easy to check that for the simplified case where $\|\nabla f\| = \|\nabla g\| = 1$ the level sets correspond to cycles centered at the Nash equilibrium and the system as a whole captures gradient-descent-ascent for a bilinear 2×2 zero-sum game (e.g. Matching Pennies).

Theorem 2. Let $\theta(0)$ and $\phi(0)$ be safe initial conditions. Then for the system of Equation 3, the following quantity is time-invariant

$$H(f, g) = \int_p^f \frac{z - p}{\|\nabla f(X_{\theta(0)}(z))\|^2} dz + \int_q^g \frac{z - q}{\|\nabla g(X_{\phi(0)}(z))\|^2} dz$$

The existence of this invariant immediately guarantees that Nash Equilibrium (p, q) cannot be reached if the dynamical system is not initialized there. Taking advantage of the planarity of the induced system - a necessary condition of Poincaré-Bendixson Theorem - we can prove that:

Theorem 3. Let $\theta(0)$ and $\phi(0)$ be safe initial conditions. Then for the system of Equation 3, the orbit $(\theta(t), \phi(t))$ is periodic.

On a positive note, we can prove that the time averages of f and g as well as the time averages of expected utilities of both players converge to their Nash equilibrium values.

Theorem 4. Let $\theta(0)$ and $\phi(0)$ be safe initial conditions and $(\mathbf{P}, \mathbf{Q}) = \left(\begin{pmatrix} p \\ 1-p \end{pmatrix}, \begin{pmatrix} q \\ 1-q \end{pmatrix} \right)$, then for the system of Equation 3

$$\lim_{T \rightarrow \infty} \frac{\int_0^T f(\theta(t)) dt}{T} = p, \quad \lim_{T \rightarrow \infty} \frac{\int_0^T r(\theta(t), \phi(t)) dt}{T} = \mathbf{P}^\top U \mathbf{Q}, \quad \lim_{T \rightarrow \infty} \frac{\int_0^T g(\phi(t)) dt}{T} = q$$

5 Poincaré recurrence in hidden bilinear games with more strategies

In this section we will extend our results by allowing both the generator and the discriminator to play hidden bilinear games with more than two strategies. We will specifically study the case of hidden bilinear games where each coordinate of the vector valued functions F and G is controlled by disjoint subsets of the variables θ and ϕ , i.e.

$$\theta = \begin{bmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_N \end{bmatrix} \quad F(\theta) = \begin{bmatrix} f_1(\theta_1) \\ f_2(\theta_2) \\ \vdots \\ f_N(\theta_N) \end{bmatrix} \quad \phi = \begin{bmatrix} \phi_1 \\ \phi_2 \\ \vdots \\ \phi_M \end{bmatrix} \quad G(\phi) = \begin{bmatrix} g_1(\phi_1) \\ g_2(\phi_2) \\ \vdots \\ g_M(\phi_M) \end{bmatrix} \quad (5)$$

where each function f_i and g_i takes an appropriately sized vector and returns a non-negative number. To account for possible constraints (e.g. that probabilities of each distribution must sum to one), we will incorporate this restriction using Lagrange Multipliers. The resulting problem becomes

$$\min_{\theta \in \mathbb{R}^n, \mu \in \mathbb{R}} \max_{\phi \in \mathbb{R}^m, \lambda \in \mathbb{R}} F(\theta)^\top U G(\phi) + \lambda \left(\sum_{i=1}^N f_i(\theta_i) - 1 \right) + \mu \left(\sum_{j=1}^M g_j(\phi_j) - 1 \right) \quad (6)$$

Writing down the equations of gradient-ascent-descent we get

$$\begin{aligned} \dot{\theta}_i &= -\nabla f_i(\theta_i) \left(\sum_{j=1}^M u_{i,j} g_j(\phi_j) + \lambda \right) & \dot{\phi}_j &= \nabla g_j(\phi_j) \left(\sum_{i=1}^N u_{i,j} f_i(\theta_i) + \mu \right) \\ \dot{\mu} &= -\left(\sum_{j=1}^M g_j(\phi_j) - 1 \right) & \dot{\lambda} &= \left(\sum_{i=1}^N f_i(\theta_i) - 1 \right) \end{aligned} \quad (7)$$

Once again we can show that along the trajectories of the system of Equation 7, θ_i can be uniquely identified by $f_i(\theta_i)$ given $\theta_i(0)$ and the same holds for the discriminator. This allows us to construct functions $X_{\theta_i(0)}$ and $X_{\phi_j(0)}$ just like in Theorem 1. We can now write down a dynamical system involving only f_i and g_j .

Lemma 3. *If $\theta(t)$ and $\phi(t)$ are solutions to Equation 7 with initial conditions $(\theta(0), \phi(0), \lambda(0), \mu(0))$, then we have that $f_i(t) = f_i(\theta_i(t))$ and $g_j(t) = g_j(\phi_j(t))$ satisfy the following equations*

$$\begin{aligned} \dot{f}_i &= -\|\nabla f_i(X_{\theta_i(0)}(f_i))\|^2 \left(\sum_{j=1}^M u_{i,j} g_j + \lambda \right) \\ \dot{g}_j &= \|\nabla g_j(X_{\phi_j(0)}(g_j))\|^2 \left(\sum_{i=1}^N u_{i,j} f_i + \mu \right) \end{aligned} \quad (8)$$

Similarly to the previous section, we can define a notion of safety for Equation 7. Let us assume that the hidden Game has a fully mixed Nash equilibrium (p, q) . Then we can define

Definition 7. *We will call the initialization $(\theta(0), \phi(0), \lambda(0), \mu(0))$ safe for Equation 7 if $\theta_i(0)$ and $\phi_j(0)$ are not stationary points of f_i and g_j respectively and $p_i \in f_{i\theta_i(0)}$ and $q_j \in g_{j\phi_j(0)}$.*

Theorem 5. *Assume that $(\theta(0), \phi(0), \lambda(0), \mu(0))$ is a safe initialization. Then there exist λ_* and μ_* such that the following quantity is time invariant:*

$$\begin{aligned} H(F, G, \lambda, \mu) &= \sum_{i=1}^N \int_{p_i}^{f_i} \frac{z - p_i}{\|\nabla f_i(X_{\theta_i(0)}(z))\|^2} dz + \sum_{j=1}^M \int_{q_j}^{g_j} \frac{z - q_j}{\|\nabla g_j(X_{\phi_j(0)}(z))\|^2} dz + \\ &\quad \int_{\lambda_*}^{\lambda} (z - \lambda_*) dz + \int_{\mu_*}^{\mu} (z - \mu_*) dz \end{aligned}$$

Given that even our reduced dynamical system has more than two state variables we cannot apply the Poincaré-Bendixson Theorem. Instead we can prove that there exists a one to one differentiable

transformation of our dynamical system so that the resulting system becomes divergence free. Applying Liouville’s formula, the flow of the transformed system is volume preserving. Combined with the invariant of Theorem 5, we can prove that the variables of the transformed system remain bounded. This gives us the following guarantees

Theorem 6. *Assume that $(\theta(0), \phi(0), \lambda(0), \mu(0))$ is a safe initialization. Then the trajectory under the dynamics of Equation 7 is diffeomorphic to one trajectory of a Poincaré recurrent flow.*

This result implies that if the corresponding trajectory of the Poincaré recurrent flow is itself recurrent, which almost all of them are, then the trajectory of the dynamics of Equation 7 is also recurrent. This is however not enough to reason about how often any of the trajectories of the dynamics of Equation 7 is recurrent. In order to prove that the flow of Equation 7 is Poincaré recurrent we will make some additional assumptions

Theorem 7. *Let f_i and g_j be sigmoid functions. Then the flow of Equation 7 is Poincaré recurrent. The same holds for all functions f_i and g_j that are one to one functions and for which all initializations are safe.*

It is worth noting that for the unconstrained version of the previous min-max problem we arrive at the same conclusions/theorems by repeating the above analysis without using the Lagrange multipliers.

6 Spurious equilibria

In the previous sections we have analyzed the behavior of safe initializations and we have proved that they lead to either periodic or recurrent trajectories. For initializations that are not safe for some equilibrium of the hidden game, game theoretically interesting fixed points are not even realizable solutions. In fact we can prove something stronger:

Theorem 8. *One can construct functions f and g for the system of Equation 3 so that for a positive measure set of initial conditions the trajectories converge to fixed points that do not correspond to equilibria of the hidden game.*

The main idea behind our theorem is that we can construct functions f and g that have local optima that break the safety assumption. For a careful choice of the value of the local optima we can make these fixed points stable and then the Stable Manifold Theorem guarantees that a non zero measure set of points in the vicinity of the fixed point converges to it. Of course the idea of these constructions can be extended to our analysis of hidden games with more strategies.

7 Discrete Time Gradient-Ascent-Descent

In this section we will discuss the implications of our analysis of continuous time gradient-ascent-descent dynamics on the properties of their discrete time counterparts. In general, the behavior of discrete time dynamical systems can be significantly different [LY75, BP18, PPP17] so it is critical to perform this non-trivial analysis. We are able to show that the picture of non-equilibration persists for an interesting class of hidden bilinear games.

Theorem 9. *Let f_i and g_j be sigmoid functions. Then for the discretized version of the system of Equation 7 and for safe initializations, function H of Theorem 5 is non-decreasing.*

An immediate consequence of the above theorem is that the discretized system cannot converge to the equilibrium (\mathbf{p}, \mathbf{q}) if its not initialized there. For the case of non-safe initializations, the conclusions of Theorem 8 persist in this case as well.

Theorem 10. *One can choose a learning rate α and functions f and g for the discretized version of the system of Equation 3 so that for a positive measure set of initial conditions the trajectories converge to fixed points that do not correspond to equilibria of the hidden game.*

8 Conclusion

In this work, inspired broadly by the structure of the complex competition between generators and discriminators in GANs, we defined a broad class of non-convex non-concave min max optimization

games, which we call hidden bilinear zero-sum games. In this setting, we showed that gradient-descent-ascent behavior is considerably more complex than a straightforward convergence to the min-max solution that one might at first suspect. We showed that the trajectories even for the simplest but evocative 2x2 game exhibits cycles. In higher dimensional games, the induced dynamical system could exhibit even more complex behavior like Poincare recurrence. On the other hand, we explored safety conditions whose violation may result in convergence to spurious game-theoretically meaningless equilibria. Finally, we show that even for a simple but widespread family of functions like sigmoids discretizing gradient-descent-ascent can further intensify the disequilibrium phenomena resulting in divergence away from equilibrium.

As a consequence of this work numerous open problems emerge; Firstly, extending such recurrence results to more general families of functions, as well as examining possible generalizations to multi-player network zero-sum games are fascinating questions. Recently, there has been some progress in resolving cyclic behavior in simpler settings by employing different training algorithms/dynamics (e.g., [DISZ18, MLZ⁺19b, GHP⁺19b]). It would be interesting to examine if these algorithms could enhance equilibration in our setting as well. Additionally, the proposed safety conditions shows that a major source of spurious equilibria in GANs could be the bad local optima of the individual neural networks of the discriminator and the generator. Lessons learned from overparametrized neural network architectures that converge to global optima [DLL⁺18] could lead to improved efficiency in training GANs. Finally, analyzing different simplification/models of GANs where provable convergence is possible could lead to interesting comparisons as well as to the emergence of theoretically tractable hybrid models that capture both the hardness of GAN training (e.g. non-convergence, cycling, spurious equilibria, mode collapse, etc) as well as their power.

Acknowledgements

Georgios Piliouras acknowledges MOE AcRF Tier 2 Grant 2016-T2-1-170, grant PIE-SGP-AI-2018-01 and NRF 2018 Fellowship NRF-NRFF2018-07. Emmanouil-Vasileios Vlatakis-Gkaragkounis was supported by NSF CCF-1563155, NSF CCF-1814873, NSF CCF-1703925, NSF CCF-1763970. Finally this work was supported by the Onassis Foundation - Scholarship ID: F ZN 010-1/2017-2018.

References

- [ACB17] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017.
- [ADLH19] 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.
- [AGL⁺17] Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, and Yi Zhang. Generalization and equilibrium in generative adversarial nets (gans). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 224–232, 2017.
- [ALW19] Jacob Abernethy, Kevin A. Lai, and Andre Wibisono. Last-iterate convergence rates for min-max optimization. *CoRR*, abs/1906.02027, 2019.
- [BGP19] James P. Bailey, Gauthier Gidel, and Georgios Piliouras. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. *CoRR*, abs/1907.04392, 2019.
- [BP18] 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.
- [BP19a] James P Bailey and Georgios Piliouras. Fast and furious learning in zero-sum games: Vanishing regret with non-vanishing step sizes. In *NeurIPS*, 2019.
- [BP19b] James P. Bailey and Georgios Piliouras. Multi-agent learning in network zero-sum games is a hamiltonian system. In *18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2019.

- [BRM⁺18] David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In *International Conference on Machine Learning*, pages 363–372, 2018.
- [BSM17] David Berthelot, Tom Schumm, and Luke Metz. BEGAN: boundary equilibrium generative adversarial networks. *CoRR*, abs/1703.10717, 2017.
- [CP19] 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.
- [DISZ18] Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- [DLL⁺18] Simon S. Du, Jason D. Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. *CoRR*, abs/1811.03804, 2018.
- [DP18] 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.
- [DP19] 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.
- [GAA⁺17] Ishaan Gulrajani, Faruk Ahmed, Martín Arjovsky, Vincent Dumoulin, and Aaron C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5767–5777, 2017.
- [GBV⁺19] Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *ICLR*, 2019.
- [GBVL18] Gauthier Gidel, Hugo Berard, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial nets. *CoRR*, abs/1802.10551, 2018.
- [GHP⁺19a] Gauthier Gidel, Reyhane Askari Hemmat, Mohammad Pezeshki, Gabriel Huang, Rémi Lepriol, Simon Lacoste-Julien, and Ioannis Mitliagkas. Negative momentum for improved game dynamics. In *AISTATS*, 2019.
- [GHP⁺19b] Gauthier Gidel, Reyhane Askari Hemmat, Mohammad Pezeshki, Rémi Le Priol, Gabriel Huang, Simon Lacoste-Julien, and Ioannis Mitliagkas. Negative momentum for improved game dynamics. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, pages 1802–1811, 2019.
- [GPM⁺14] 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.
- [GXC⁺18] Hao Ge, Yin Xia, Xu Chen, Randall Berry, and Ying Wu. Fictitious GAN: training gans with historical models. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part I*, pages 122–137, 2018.

- [IZZE17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5967–5976, 2017.
- [JNJ19] 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.
- [KAHK17] Naveen Kodali, Jacob D. Abernethy, James Hays, and Zsolt Kira. On convergence and stability of gans. *CoRR*, abs/1705.07215, 2017.
- [KLPT11] Robert D. Kleinberg, Katrina Ligett, Georgios Piliouras, and Éva Tardos. Beyond the nash equilibrium barrier. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 7-9, 2011. Proceedings*, pages 125–140, 2011.
- [LLRY18] Qihang Lin, Mingrui Liu, Hassan Rafique, and Tianbao Yang. Solving weakly-convex-weakly-concave saddle-point problems as weakly-monotone variational inequality. *CoRR*, abs/1810.10207, 2018.
- [LTH⁺17] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenze Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 105–114, 2017.
- [LY75] Tien-Yien Li and James A. Yorke. Period three implies chaos. *The American Mathematical Monthly*, 82(10):985–992, 1975.
- [Ma18] Tengyu Ma. Generalization and equilibrium in generative adversarial nets (gans) (invited talk). In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, page 2, 2018.
- [MGN18] Lars M. Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 3478–3487, 2018.
- [MLZ⁺19a] 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 *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- [MLZ⁺19b] 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.
- [MPP18] 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.
- [MPPS17] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- [MPR⁺17] Tung Mai, Ioannis Panageas, Will Ratcliff, Vijay V. Vazirani, and Peter Yunker. Rock-paper-scissors, differential games and biological diversity. *CoRR*, abs/1710.11249, 2017.
- [MR18] Eric Mazumdar and Lillian J. Ratliff. On the convergence of gradient-based learning in continuous games. *CoRR*, abs/1804.05464, 2018.

- [NMP18] Sai Ganesh Nagarajan, Sameh Mohamed, and Georgios Piliouras. Three body problems in evolutionary game dynamics: Convergence, periodicity and limit cycles. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, pages 685–693, 2018.
- [OSG⁺18] Frans A. Oliehoek, Rahul Savani, Jose Gallego-Posada, Elise van der Pol, and Roderich Groß. Beyond local nash equilibria for adversarial networks. *CoRR*, abs/1806.07268, 2018.
- [PPP17] Gerasimos Palaioupanos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5872–5882, 2017.
- [PS14] 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.
- [PS18] Georgios Piliouras and Leonard J. Schulman. Learning dynamics and the co-evolution of competing sexual species. In *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, pages 59:1–59:3, 2018.
- [PV16] David Pfau and Oriol Vinyals. Connecting generative adversarial networks and actor-critic methods. *CoRR*, abs/1610.01945, 2016.
- [RMC16] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [San10] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.
- [SGZ⁺16] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2226–2234, 2016.
- [SRL18] Maziar Sanjabi, Meisam Razaviyayn, and Jason D. Lee. Solving non-convex non-concave min-max games under polyak-łojasiewicz condition. *CoRR*, abs/1812.02878, 2018.
- [TGB⁺17] Ilya O. Tolstikhin, Sylvain Gelly, Olivier Bousquet, Carl-Johann Simon-Gabriel, and Bernhard Schölkopf. Adagan: Boosting generative models. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5424–5433, 2017.
- [YFW⁺19] Yasin Yazıcı, Chuan-Sheng Foo, Stefan Winkler, Kim-Hui Yap, Georgios Piliouras, and Vijay Chandrasekhar. The unusual effectiveness of averaging in gan training. In *ICLR*, 2019.
- [ZXL17] Han Zhang, Tao Xu, and Hongsheng Li. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 5908–5916, 2017.