ON THE IMPOSSIBILITY OF GLOBAL CONVERGENCE IN MULTI-LOSS OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Under mild regularity conditions, gradient-based methods converge globally to a critical point in the single-loss setting. This is known to break down for vanilla gradient descent when moving to *multi-loss* optimization, but can we hope to build some algorithm with global guarantees? We negatively resolve this open problem by proving that desirable convergence properties cannot simultaneously hold for *any* algorithm. Our result has more to do with the existence of games with no satisfactory outcomes, than with algorithms per se. More explicitly we construct a two-player game with zero-sum interactions whose losses are both coercive and analytic, but whose only simultaneous critical point is a strict maximum. Any 'reasonable' algorithm, defined to avoid strict maxima, will therefore fail to converge. This is fundamentally different from single losses, where coercivity implies existence of a global minimum. Moreover, we prove that a wide range of existing gradient-based methods almost surely have bounded but non-convergent iterates in a constructed *zero-sum* game for suitably small learning rates. It nonetheless remains an open question whether such behavior can arise in high-dimensional games of interest to ML practitioners, such as GANs or multi-agent RL.

1 INTRODUCTION

Problem Setting As multi-agent architectures proliferate in machine learning, it is becoming increasingly important to understand the dynamics of gradient-based methods when optimizing multiple interacting goals, otherwise known as *differentiable games*. This framework encompasses GANs (Goodfellow et al., 2014), intrinsic curiosity (Pathak et al., 2017), imaginative agents (Racanière et al., 2017), synthetic gradients (Jaderberg et al., 2017), hierarchical reinforcement learning (Wayne & Abbott, 2014; Vezhnevets et al., 2017) and multi-agent RL in general (Busoniu et al., 2008). The interactions between learning agents make for vastly more complex mechanics: naively applying gradient descent on each loss simultaneously is known to diverge even in simple bilinear games.

Related Work A large number of methods have recently been proposed to alleviate the failings of simultaneous gradient descent: adaptations of single-loss algorithms such as Extragradient (EG) (Azizian et al., 2019) and Optimistic Mirror Descent (OMD) (Daskalakis et al., 2018), Alternating Gradient Descent (AGD) for finite regret (Bailey et al., 2019), Consensus Optimization (CO) for GAN training (Mescheder et al., 2017), Competitive Gradient Descent (CGD) based on solving a bilinear approximation of the loss functions (Schaefer & Anandkumar, 2019), Symplectic Gradient Adjustment (SGA) based on a novel decomposition of game mechanics (Balduzzi et al., 2018; Letcher et al., 2019a), and opponent-shaping algorithms including Learning with Opponent-Learning Awareness (LOLA) (Foerster et al., 2018) and its convergent counterpart, Stable Opponent Shaping (SOS) (Letcher et al., 2019b). Let $\mathcal{A}$ be this set of algorithms.

Each has shown promising theoretical implications and empirical results, but none offers insight into global convergence in the *non-convex* setting, which includes the vast majority of machine learning applications. One of the main roadblocks compared with single-loss optimization has been noted by Schaefer & Anandkumar (2019): "a convergence proof in the nonconvex case analogue to Lee et al. (2016) is still out of reach in the competitive setting. A major obstacle to this end is the identification of a suitable measure of progress (which is given by the function value in the single agent setting), since norms of gradients can not be expected to decay monotonously for competitive dynamics in non-convex-concave games."

It has been established that Hamiltonian Gradient Descent converges in two-player zero-sum games under a "sufficiently bilinear" condition by Abernethy et al. (2019), but this algorithm is unsuitable for optimization as it cannot distinguish between minimization and maximization (Hsieh et al., 2020, Appendix C.4). Global convergence has also been established for some algorithms in a few special cases: potential and Hamiltonian games (Balduzzi et al., 2018), zero-sum games satisfying the two-sided Polyak-Łojasiewicz condition (Yang et al., 2020), zero-sum linear quadratic games (Zhang et al., 2019) and zero-sum games whose loss and first three derivatives are bounded (Mangoubi & Vishnoi, 2020). These are significant contributions with several applications of interest, but do not include any of the architectures mentioned above. Finally, Balduzzi et al. (2020) show that GD dynamics are bounded under a 'negative sentiment' assumption in smooth markets, which do include GANs – but this does not imply convergence, as we will show.

On the other hand, *failure* of global convergence has been shown for the Multiplicative Weights Update method by Palaiopanos et al. (2017), for policy-gradient algorithms by Mazumdar et al. (2020), and for simultaneous and alternating gradient descent (simGD and AGD) by Vlatakis-Gkaragkounis et al. (2019); Bailey et al. (2019), with interesting connections to Poincaré recurrence. Nonetheless, nothing is claimed about other optimization methods. Farnia & Ozdaglar (2020) show that GANs may have no Nash equilibria, but it does *not* follow that algorithms fail to converge since there may be locally-attracting but non-Nash critical points (Mazumdar et al., 2019, Example 2).

Finally, Hsieh et al. (2020) uploaded a preprint just after the completion of this work with a similar focus to ours. They prove that generalized Robbins-Monro schemes may converge with arbitrarily high probability to spurious attractors. This includes simGD, AGD, stochastic EG, optimistic gradient and Kiefer-Wolfowitz. However, Hsieh et al. (2020) focus on the possible occurrence of undesirable convergence phenomena for *stochastic* algorithms. We instead prove that desirable convergence properties cannot simultaneously hold for all algorithms (including *deterministic*). Moreover, their results apply only to decreasing step-sizes whereas ours include constant step-sizes. These distinctions are further highlighted by Hsieh et al. (2020) in the further related work section. Taken together, our works give a fuller picture of the failure of global convergence in multi-loss optimization.

Contribution We prove that global convergence is fundamentally incompatible with the 'reasonable' requirement that algorithms avoid strict maxima and converge only to critical points. We construct a two-player game with zero-sum interactions whose losses are coercive and analytic, but whose only critical point is a strict maximum. Any reasonable algorithm must either diverge to infinite losses or have bounded non-convergent iterates ('cycle'). Restricting the game class even further, we then construct a zero-sum game in which all algorithms in $\mathcal{A}$ are proven to cycle.

One might hope that global convergence could at least be guaranteed in games with strict *minima* and no other critical points. On the contrary we show that strict minima can have arbitrarily small regions of attraction, in the sense that reasonable algorithms will fail to converge there with arbitrarily high probability for fixed initial parameter distribution.

It may be that cycles do not arise in high-dimensional games of interest including GANs. This paper is meant only to prove that global convergence comes at the cost of reasonable requirements in a simple two-player game with zero-sum interactions.

2 Background

2.1 Single losses: global convergence of gradient descent

Given a continuously differentiable function $f : \mathbb{R}^d \to \mathbb{R}$, let

$$\theta_{k+1} = \theta_k - \alpha \nabla f(\theta_k)$$

be the iterates of gradient descent with learning rate α, initialised at θ_0. Under standard regularity conditions, gradient descent converges globally to critical points:

Proposition 1. *Assume $f \in C^2$ has compact sublevel sets and is either analytic or has isolated critical points. For any $\theta_0 \in \mathbb{R}^d$, define $U_0 = \{f(\theta) \leq f(\theta_0)\}$ and let $L < \infty$ be a Lipschitz constant for ∇f in U_0. Then for any $0 < \alpha < 2/L$ we have $\lim_k \theta_k = \bar{\theta}$ for some critical point $\bar{\theta}$.*

The requirements for convergence are relatively mild:

1. f has compact sublevel sets iff f is coercive, $\lim_{\|\theta\|\to\infty} f(\theta) = \infty$, which mostly holds in machine learning since f is a loss function.
2. f has isolated critical points if it is a Morse function (nondegenerate Hessian at critical points), which holds for almost all C^2 functions. More precisely, Morse functions form an open, dense subset of all functions $f \in C^2(\mathbb{R}^d, \mathbb{R})$ in the Whitney C^2-topology.
3. Lipschitz continuity of ∇f is *not* assumed, which would fail even for cubic polynomials.

The goal of this paper is to prove that similar (even weaker) guarantees *cannot* be obtained in the multi-loss setting – not only for GD, but for any reasonable algorithm. This has to do with the more complex nature of gradient vector fields arising from multiple losses.

2.2 Differentiable games

Following Balduzzi et al. (2018), we frame the problem of multi-loss optimization as a differentiable game among cooperating and competing agents/players. These may simply be different internal components of a single system, like the generator and discriminator in GANs.

Definition 1. *A* differentiable game *is a set of n agents with parameters $\theta = (\theta^1, \dots, \theta^n) \in \mathbb{R}^d$ and twice continuously differentiable losses $L^i : \mathbb{R}^d \to \mathbb{R}$, where $\theta^i \in \mathbb{R}^{d_i}$ for each i and $\sum_i d_i = d$.*

Losses are *not* assumed to be convex/concave in any of the parameters. In practice, losses need only be differentiable almost-everywhere: think of neural nets with rectified linear units.

If $n = 1$, the 'game' is simply to minimise a given loss function. We write $\nabla_i L^k = \nabla_{\theta^i} L^k$ and $\nabla_{ij} L^k = \nabla_{\theta^j} \nabla_{\theta^i} L^k$ for any i, j, k, and define the simultaneous gradient of the game

$$\xi = \left(\nabla_1 L^1, \dots, \nabla_n L^n\right)^T \in \mathbb{R}^d$$

as the concatenation of each player's gradient. If each agent independently minimises their loss using GD with learning rate α, the parameter update for all agents is given by $\theta \leftarrow \theta - \alpha\xi(\theta)$. We call this simultaneous gradient descent (simGD), or GD for short. We call $\bar{\theta}$ a *critical point* if $\xi(\bar{\theta}) = 0$. Now introduce the 'Hessian' (or Jacobian) of the game as the block matrix

$$H = \nabla\xi = \begin{pmatrix} \nabla_{11} L^1 & \cdots & \nabla_{1n} L^1 \\ \vdots & \ddots & \vdots \\ \nabla_{n1} L^n & \cdots & \nabla_{nn} L^n \end{pmatrix} \in \mathbb{R}^{d\times d}.$$

Importantly note that H is not symmetric in general unless $n = 1$, in which case we recover the usual Hessian $H = \nabla^2 L$. However H can be decomposed into symmetric and anti-symmetric components as $H = S + A$ (Balduzzi et al., 2018). A second useful decomposition has appeared recently in (Letcher et al., 2019b) and (Schaefer & Anandkumar, 2019): $H = H_d + H_o$ where H_d and H_o are the matrices of diagonal and off-diagonal blocks; formally, $H_d = \bigoplus_i \nabla_{ii} L^i$. One solution concept for differentiable games, analogous to the single-loss case, is defined as follows.

Definition 2. *A critical point $\bar{\theta}$ is a (strict, local)* minimum *if $H(\bar{\theta}) \succ 0$.*[1]

These were named (strict) stable fixed points by Balduzzi et al. (2018), but the term is usually reserved in dynamical systems to the larger class defined by Hessian eigenvalues with positive real parts, which is implied but not equivalent to $H \succ 0$ for non-symmetric matrices.

In particular, strict minima are (differential) Nash equilibria as defined by Mazumdar et al. (2019), since diagonal blocks must also be positive definite: $\nabla_{ii} L^i(\bar{\theta}) \succ 0$. The converse does not hold.

Algorithm class. This paper is concerned with any algorithm whose iterates are obtained by initialising θ_0 and applying a function F to the previous iterates, namely $\theta_{k+1} = F(\theta_k, \dots, \theta_0)$. This holds for all gradient-based methods (deterministic or stochastic); most of them are only functions

[1] For non-symmetric matrices, positive definiteness is defined as $H \succ 0$ iff $u^T H u > 0$ for all non-zero $u \in \mathbb{R}^d$. This is equivalent to the symmetric part S of H being positive definite.

of the current iterate θ_k, so that $\theta_k = F^k(\theta_0)$. All probabilistic statements in this paper assume that θ_0 is initialised following any bounded and continuous measure ν on $\mathbb{R}^d$. Continuity is a weak requirement and widely holds across machine learning, while boundedness mostly holds in practice since the bounded region can be made large enough to accommodate required initial points.

For single-player games, the goal of such algorithms is for θ_k to converge to a local (perhaps global) minimum as $k \to \infty$. The goal is less clear for differentiable games, but is generally to reach a minimum or a Nash equilibrium. In the case of GANs the goal might be to reach parameters that produce realistic images, which is more challenging to define formally.

Throughout the text we use the term *(limit) cycle* to mean *bounded but non-convergent iterates*. This terminology is used because bounded iterates are non-convergent if and only if they have at least two accumulation points, between which they must 'cycle' infinitely often. This is not to be taken literally: the set of accumulation points may not even be connected. Hsieh et al. (2020) provide a more complete characterisation of these cycles.

Game class. Expecting global guarantees in *all* differentiable games is excessive, since every continuous dynamical system arises as simultaneous GD on the loss functions of a differentiable game (Balduzzi et al., 2020, Lemma 1). For this reason, the aforementioned authors have introduced a vastly more tractable class of games called *markets*.

Definition 3. *A (smooth) market is a differentiable game where interactions between players are pairwise zero-sum, namely,*

$$L^i(\theta) = L^i(\theta^i) + \sum_{j \neq i} g_{ij}(\theta^i, \theta^j)$$

with $g_{ij}(\theta^i, \theta^j) + g_{ji}(\theta^j, \theta^i) = 0$ *for all* i, j.

This generalises zero-sum games while remaining amenable to optimization and aggregation, meaning that "we can draw conclusions about the gradient-based dynamics of the collective by summing over properties of its members" (Balduzzi et al., 2020). Moreover, this class captures a large number of applications including GANs and related architectures, intrinsic curiosity modules, adversarial training, task-suites and population self-play. One would modestly hope for *some* reasonable algorithm to converge globally in markets. We will prove that even this is too much to ask.

2.3 Reasonable algorithms

We wish to prove that global convergence is at odds with weak, 'reasonable' desiderata. The first requirement is that fixed points of an optimization algorithm F are critical points. Formally,

$$F(\theta) = \theta \implies \xi(\theta) = 0\,. \tag{R1}$$

If not, some agent i could strictly improve its losses by following the gradient $-\nabla_i L^i \neq 0$. There is no reason for a gradient-based algorithm to stop improving if its gradient is non-zero.

The second requirement is that algorithms avoid strict maxima. Analogous to strict minima, they are defined for single losses by a negative-definite Hessian $H \prec 0$. Converging to such a point $\bar{\theta}$ is the opposite goal of any meaningful algorithm since moving *anywhere* away from $\bar{\theta}$ decreases the loss. There are multiple ways of generalising this concept for multiple losses, but Proposition 2 below justifies that $H \prec 0$ is the weakest one.

Proposition 2. *Write* $\lambda(A) = \mathrm{Re}(\mathrm{Spec}(A))$ *for real parts of the eigenvalues of a matrix* A. *We have the following implications, and none of them are equivalences.*

$$\begin{tikzcd}
 & \max\lambda(H) < 0 \arrow[r, Rightarrow] \arrow[dr, Rightarrow] & \min\lambda(H) < 0 \arrow[dr, Rightarrow] & \\
H \prec 0 \arrow[ur, Rightarrow] \arrow[dr, Rightarrow] & & & \min\lambda(S) < 0 \\
 & \max\lambda(H_d) < 0 \arrow[r, Rightarrow] \arrow[ur, Rightarrow] & \min\lambda(H_d) < 0 \arrow[ur, Rightarrow] &
\end{tikzcd}$$

Definition 4. *A critical point* $\bar{\theta}$ *is a (strict, local)* maximum *if* $H(\bar{\theta}) \prec 0$.

Imposing that algorithms avoid strict maxima is therefore the weakest possible requirement of its kind. Note that the bottom-left implication Proposition 2 is equivalent to $\nabla_{ii} L^i \prec 0$ for all i, so strict maxima are also strict maxima of each player's individual loss function. Players can *all* decrease their losses by moving *anywhere* away from them. It is exceedingly reasonable to ask that optimization algorithms avoid these points almost surely. Formally, we require that for any strict maximum $\bar{\theta}$ and bounded region U there are hyperparameters such that

$$\mu\left(\{\theta_0 \in U \mid \lim_k \theta_k = \bar{\theta}\}\right) = 0\,. \tag{R2}$$

μ denotes Lebesgue measure. Hyperparameters may depend on the given game and the region U, as is typical for learning rates in gradient-based methods.

Definition 5 (Reason)**.** *An algorithm is* reasonable *if it satisfies R1 and R2.*

Reason is not equivalent to *rationality* or *self-interest*. Reason is much weaker, imposing only that agents are well-behaved regarding strict maxima even if their individual behavior is not self-interested. For instance, SGA agents do not behave out of self-interest (Balduzzi et al., 2018).

3 Global Convergence in Differentiable Games

3.1 Reasonable algorithms fail to converge globally

Our main contribution is to show that global guarantees do not exist for *any* reasonable algorithm. First recall that global convergence should not be expected in *all* games, since there may be a divergent direction with minimal loss (imagine minimising $L = e^x$). It should however be asked that algorithms have bounded iterates in *coercive* games, defined by coercive losses

$$\lim_{\|\theta\| \to \infty} L^i(\theta) = \infty$$

for all i. Indeed, unbounded iterates in coercive games would lead to infinite losses for *all* agents, the worst possible outcome. Given bounded iterates, convergence should hold if the Hessian is nondegenerate at critical points (which must therefore be isolated, recall Proposition 1). We call such a game *nondegenerate*. This condition can also be replaced by analyticity of the loss. In the spirit of weakest assumptions, we ask for convergence when *both* conditions hold.

Definition 6 (Globality)**.** *An algorithm is* global *if, in a coercive, analytic and nondegenerate game, for any fixed θ_0, iterates θ_k are bounded and converge for suitable hyperparameters.* (G1)

Note that GD is global for *single-player* games by Proposition 1. Unfortunately, reason and globality are fundamentally at odds as soon as we move to two-player markets.

Theorem 1. *There is a coercive, nondegenerate, analytic two-player market $\mathcal{M}$ whose only critical point is a strict maximum. In particular, algorithms only have four possible outcomes in $\mathcal{M}$:*

1. *Iterates are unbounded, and all players diverge to infinite loss. [Not global]*
2. *Iterates are bounded and converge to the strict maximum. [Not reasonable]*
3. *Iterates are bounded and converge to a non-critical point. [Not reasonable]*
4. *Iterates are bounded but do not converge (cycle). [Not global]*

Proof. Consider the analytic market $\mathcal{M}$ given by

$$L^1(x,y) = x^6/6 - x^2/2 + xy + \frac{1}{4}\left(\frac{y^4}{1+x^2} - \frac{x^4}{1+y^2}\right)$$
$$L^2(x,y) = y^6/6 - y^2/2 - xy - \frac{1}{4}\left(\frac{y^4}{1+x^2} - \frac{x^4}{1+y^2}\right)\,.$$

We prove in Appendix D that $\mathcal{M}$ is coercive, nondegenerate, and has a unique critical point at the origin, which is a strict maximum. □

Constructing an algorithm with global guarantees is therefore doomed to be unreasonable in that it will converge to strict maxima or non-critical points in $\mathcal{M}$.

None of the outcomes of $\mathcal{M}$ are satisfactory. The first three are highly objectionable, as already discussed. The fourth is less obvious, and may even have game-theoretic significance (Papadimitriou & Piliouras, 2019), but is counter-intuitive from an optimization standpoint. Terminating the iteration would lead to a non-critical point, much like the third outcome. Even if we let agents update parameters continuously as they play a game or solve a task, they will have oscillatory behavior and fail to produce consistent outcomes (e.g. when generating an image or playing Starcraft).

The hope for machine learning is that such predicaments do not arise in applications we care about, such as GANs or intrinsic curiosity. This may well be the case, but proving or disproving global convergence in these specific settings is beyond the scope of this paper.

Remark. Why can this approach not be used to *disprove* global convergence for single losses? One reason is that we cannot construct a coercive loss with no critical points other than strict maxima: coercive losses, unlike games, always have a global minimum.

3.2 What if there are strict minima?

One might wonder if it is purely the absence of strict minima that causes non-convergence, since strict minima are locally attracting under gradient dynamics. Can we guarantee global convergence if we impose existence of a minimum, and more, the absence of any other critical points?

Unfortunately, strict minima may have an arbitrarily small region of attraction. Assuming parameters are initialised following any bounded continuous measure ν on $\mathbb{R}^d$, we can always modify $\mathcal{M}$ by deforming a correspondingly small region around the origin, turning it into a minimum while leaving the dynamics unchanged outside of this region.

For a fixed initial distribution, any reasonable algorithm can therefore enter a limit cycle or diverge to infinite losses with arbitrarily high probability.

Theorem 2. *Given a reasonable algorithm with bounded continuous distribution on θ_0 and a real number $\epsilon > 0$, there exists a coercive, nondegenerate, almost-everywhere analytic two-player market $\mathcal{M}_\sigma$ with a strict minimum and no other critical points, such that θ_k either cycles or diverges to infinite losses for both players with probability at least $1 - \epsilon$.*

Proof. Let $0 < \sigma < 0.1$ and define

$$f_\sigma(\theta) = \begin{cases} (x^2 + y^2 - \sigma^2)/2 & \text{if } \|\theta\| \geq \sigma \\ (y^2 - 3x^2)(x^2 + y^2 - \sigma^2)/(2\sigma^2) & \text{otherwise,} \end{cases}$$

where $\theta = (x, y)$ and $\|\theta\| = \sqrt{x^2 + y^2}$ is the standard $L2$-norm. Note that f_σ is continuous since

$$\lim_{\|\theta\| \to \sigma^+} f_\sigma(x, y) = 0 = \lim_{\|\theta\| \to \sigma^-} f_\sigma(x)\,.$$

Now consider the two-player market $\mathcal{M}_\sigma$ given by

$$L^1(x, y) = x^6/6 - x^2 + f_\sigma(x, y) + xy + \frac{1}{4}\left(\frac{y^4}{1 + x^2} - \frac{x^4}{1 + y^2}\right)$$
$$L^2(x, y) = y^6/6 - f_\sigma(x, y) - xy - \frac{1}{4}\left(\frac{y^4}{1 + x^2} - \frac{x^4}{1 + y^2}\right)\,.$$

We prove in Appendix E that $\mathcal{M}_\sigma$ is a coercive, nondegenerate, almost-everywhere analytic game whose only critical point is a strict minimum at the origin. We then prove that θ_k cycles or diverges with probability at least $1 - \epsilon$, and plot iterates for each algorithm in $\mathcal{A}$ □

3.3 How do existing algorithms behave?

All algorithms will either fail to be reasonable or global in $\mathcal{M}$. Nonetheless, it would be interesting to determine the specific failure that each algorithm in $\mathcal{A}$ exhibits. Each of them is defined formally

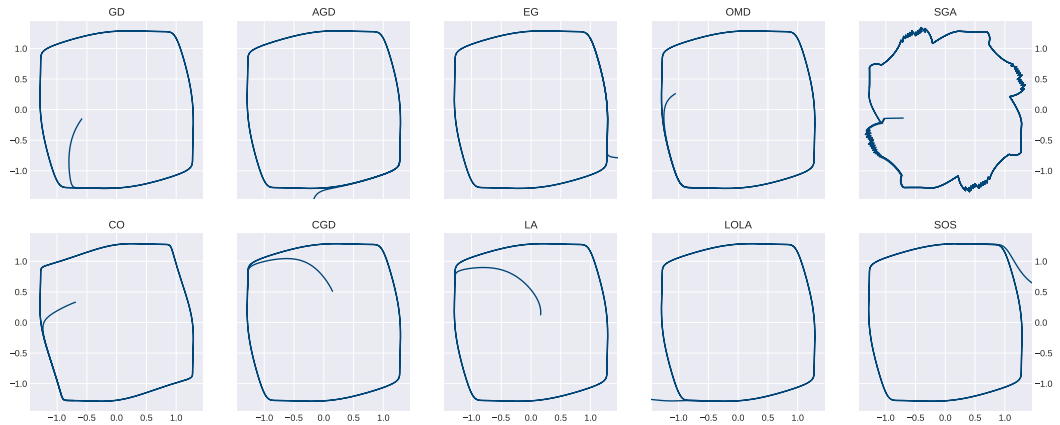

Figure 1: Algorithms in $\mathcal{A}$ fail to converge in $\mathcal{M}$ with $\alpha = \gamma = 0.01$. Single run with standard normal initialisation, 3000 iterations. The behavior of SGA is slightly different, explained by the presence of a non-continuous parameter λ jumping between ± 1 according to an alignment criterion.

in Appendix A. For suitably small hyperparameters, we expect each algorithm to be reasonable and moreover to have bounded iterates in $\mathcal{M}$. If this holds, they must cycle by Theorem 1.

This was witnessed experimentally across 1000 runs for $\alpha = \gamma = 0.01$, with every run resulting in cycles. A single such run is illustrated in Figure 1. Algorithms may follow one of the three other outcomes for other hyperparameters, for instance diverging to infinite loss if α is too large or converging to the strict maximum for CO if γ is too large. The point here is to characterise the 'regular' behavior which can be seen as that occurring for sufficiently small hyperparameters.

Instead of *proving* that algorithms must cycle in $\mathcal{M}$, we construct a *zero-sum* game $\mathcal{N}$ with similar properties as $\mathcal{M}$ and prove below that algorithms in $\mathcal{A}$ almost surely fail to converge there for small α, γ. This is stronger than proving the analogous result for $\mathcal{M}$, since $\mathcal{N}$ belongs to the even smaller class of zero-sum games which one might have hoped was well-behaved.

In this light, one might wish to extend Theorem 1 to zero-sum games. However, zero-sum games cannot be coercive since $L^1 \to \infty$ implies $L^2 \to -\infty$. It is therefore unclear whether global guarantees should be expected. Note however that $\mathcal{N}$ will be *weakly-coercive* in the sense that

$$\lim_{\|\theta^i\| \to \infty} L^i(\theta^i, \theta^{-i}) = \infty$$

for all i and fixed θ^{-i}.

Theorem 3. *There is a weakly-coercive, nondegenerate, analytic two-player zero-sum game $\mathcal{N}$ whose only critical point is a strict maximum. Algorithms in $\mathcal{A}$ almost surely have bounded non-convergent iterates in $\mathcal{N}$ for α, γ sufficiently small.*

Proof. Consider the analytic zero-sum game $\mathcal{N}$ given by

$$L^1 = xy - x^2/2 + y^2/2 + x^4/4 - y^4/4 = -L^2 \,.$$

We prove in Appendix F that $\mathcal{N}$ is weakly-coercive, nondegenerate, and has a unique critical point at the origin which is a strict maximum. We prove that algorithms in $\mathcal{A}$ have the origin as unique fixed points, with negative-definite Jacobian for α, γ small, hence failing to converge almost surely. We moreover prove that algorithms have bounded non-convergent iterates in $\mathcal{N}$ for α, γ sufficiently small. Iterates are plotted for a single run of each algorithm in Figure 3 with $\alpha = \gamma = 0.01$. □

As in $\mathcal{M}$, the behavior of each algorithm may differ for larger hyperparameters. EG and OMD may converge to a non-critical point for large α (see proof to see why this may occur), and all algorithms may converge to the strict maximum or have unbounded iterates for large α. All such outcomes are as unsatisfactory as for $\mathcal{M}$, except unbounded iteration which will not result in positive infinite losses for *both* players since $L^1 = -L^2$.

3.4 COROLLARY: THERE ARE NO SUITABLE MEASURES OF PROGRESS

A crucial step in proving global convergence of GD on single losses is showing that the set of accumulation points is a subset of critical points, using the function value as a 'measure of progress'. The fact that this fails for differentiable games implies that there can be no suitable measures of progress for reasonable algorithms with bounded iterates. We formalise this below, answering the question of Schaefer & Anandkumar (2019) quoted in the introduction.

Definition 7. *A measure of progress for an algorithm given by $\theta_{k+1} = F(\theta_k)$ is a continuous map $M : \mathbb{R}^d \to \mathbb{R}$, bounded below, such that $M(F(\theta)) \leq M(\theta)$ and $M(F(\theta)) = M(\theta)$ iff $F(\theta) = \theta$.*

Measures of progress are very similar to *descent functions*, as defined by Luenberger & Ye (1984), and somewhat akin to Lyapunov functions. The function value f is a measure of progress for single-loss GD under the usual regularity conditions, while the gradient norm $\|\xi\|$ is a measure of progress for GD in *strictly convex* differentiable games:

$$\|\xi(\theta - \alpha\xi)\|^2 = \|\xi\|^2 - \alpha\xi^T H^t \xi + o(\alpha) \leq \|\xi\|^2$$

for small α. Unfortunately, games like $\mathcal{M}$ prevent the existence of such measures in general.

Corollary 1. *There are no measures of progress for reasonable algorithms which produce bounded iterates in $\mathcal{M}$ or $\mathcal{N}$.*

Assuming the algorithm to be reasonable is necessary: any map is a measure of progress for the unreasonable algorithm $F(\theta) = \theta$. Assuming the algorithm to have bounded iterates in $\mathcal{M}$ or $\mathcal{N}$ is necessary: $M(\theta) = \exp(-\theta \cdot \mathbf{1})$ is a measure of progress for the reasonable but always-divergent algorithm $F(\theta) = \theta + \mathbf{1}$, where $\mathbf{1}$ is the constant vector of ones.

4 CONCLUSION

We have proven that global convergence is fundamentally at odds with weak, desirable requirements in multi-loss optimization. Any reasonable algorithm can cycle or diverge to infinite losses, even in two-player markets. This arises because coercive games, unlike losses, may have no critical points other than strict maxima. However, this is not the only point of failure: strict minima may have arbitrarily small regions of attraction, making convergence arbitrarily unlikely.

Limit cycles are not necessarily bad: they may even have game-theoretic significance (Papadimitriou & Piliouras, 2019). This paper nonetheless shows that some games have no satisfactory outcome in the usual sense, even in the class of two-player markets. Players should neither escape to infinite losses, nor converge to strict maxima or non-critical points, so cycling may be the lesser evil. The community is accustomed to optimization problems whose solutions are single points, but cycles may have to be accepted as solutions in themselves.

The hope for machine learning practitioners is that local minima with large regions of attraction prevent limit cycles from arising in applications of interest, including GANs. Proving or disproving this is an interesting and important avenue for further research, with real implications on what to expect when agents learn while interacting with others. Cycles may for instance be unacceptable in self-driving cars, where oscillatory predictions may have life-threatening implications.

REFERENCES

Jacob Abernethy, Kevin A. Lai, and Andre Wibisono. Last-iterate convergence rates for min-max optimization. *ArXiv e-prints*, 2019.

PA Absil, Robert Mahony, and Ben Andrews. Convergence of the iterates of descent methods for analytic cost functions. *SIAM Journal on Optimization*, 16, 01 2005.

Waïss Azizian, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. A tight and unified analysis of extragradient for a whole spectrum of differentiable games. *ArXiv e-prints*, 2019.

James P. Bailey, Gauthier Gidel, and Georgios Piliouras. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. *ArXiv e-prints*, 2019.

D. Balduzzi, S. Racaniere, J. Martens, J. Foerster, K. Tuyls, and T. Graepel. The Mechanics of n-Player Differentiable Games. *ICML*, 2018.

David Balduzzi, Wojciech M. Czarnecki, Tom Anthony, Ian Gemp, Edward Hughes, Joel Leibo, Georgios Piliouras, and Thore Graepel. Smooth markets: A basic mechanism for organizing gradient-based learners. In *International Conference on Learning Representations*, 2020.

Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. Springer-Verlag Berlin Heidelberg, 2006.

L. Busoniu, R. Babuska, and B. De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, March 2008.

C. Daskalakis, A. Ilyas, V. Syrgkanis, and H. Zeng. Training GANs with Optimism. In *International Conference on Learning Representations*, 2018.

Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. *NIPS*, 2018.

W. Decker, G.-M. Greuel, G. Pfister, and H. Schönemann. Singular 4-1-2 — A computer algebra system for polynomial computations. http://www.singular.uni-kl.de, 2019.

Farzan Farnia and Asuman Ozdaglar. Do GANs always have Nash equilibria? *ICML*, 2020.

J. N. Foerster, R. Y. Chen, M. Al-Shedivat, S. Whiteson, P. Abbeel, and I. Mordatch. Learning with Opponent-Learning Awareness. *AAMAS*, 2018.

I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Networks. *NIPS*, 2014.

Ya-Ping Hsieh, Panayotis Mertikopoulos, and Volkan Cevher. The limits of min-max optimization algorithms: convergence to spurious non-critical sets. *ArXiv e-prints*, 2020.

M. Jaderberg, W. M. Czarnecki, S. Osindero, O. Vinyals, A. Graves, D. Silver, and K. Kavukcuoglu. Decoupled Neural Interfaces using Synthetic Gradients. *ICML*, 2017.

K. Lange. *Optimization*. Springer Texts in Statistics. Springer New York, 2013.

J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. Gradient Descent Only Converges to Minimizers. In *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 1246–1257, 2016.

Alistair Letcher, David Balduzzi, Sébastien Racanière, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. Differentiable game mechanics. *Journal of Machine Learning Research*, 2019a.

Alistair Letcher, Jakob Foerster, David Balduzzi, Tim Rocktäschel, and Shimon Whiteson. Stable opponent shaping in differentiable games. In *International Conference on Learning Representations*, 2019b.

David Luenberger and Yinyu Ye. *Linear and Nonlinear Programming*, volume 67. 01 1984.

Oren Mangoubi and Nisheeth K. Vishnoi. A second-order equilibrium in nonconvex-nonconcave min-max optimization: Existence and algorithm. *ArXiv e-prints*, 2020.

Eric Mazumdar, Lillian J. Ratliff, Michael I. Jordan, and S. Shankar Sastry. Policy-gradient algorithms have no guarantees of convergence in linear quadratic games. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20, pp. 860–868. International Foundation for Autonomous Agents and Multiagent Systems, 2020.

Eric V. Mazumdar, Michael I. Jordan, and S. Shankar Sastry. On finding local nash equilibria (and only local nash equilibria) in zero-sum games. *ArXiv e-prints*, 2019.

L. Mescheder, S. Nowozin, and A. Geiger. The Numerics of GANs. *NIPS*, 2017.

Gerasimos Palaiopanos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. *NIPS*, 2017.

I. Panageas and G. Piliouras. Gradient Descent Only Converges to Minimizers: Non-Isolated Critical Points and Invariant Regions. In *ITCS 2017*, volume 67 of *Leibniz International Proceedings in Informatics*, pp. 2:1–2:12, 2017.

Christos Papadimitriou and Georgios Piliouras. Game dynamics as the meaning of a game. *SIGecom Exch.*, 16(2):53–63, May 2019.

D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven Exploration by Self-supervised Prediction. *ICML*, 2017.

S. Racanière, T. Weber, D. P. Reichert, L. Buesing, A. Guez, D. Jimenez Rezende, A. Puigdomènech Badia, O. Vinyals, N. Heess, Y. Li, R. Pascanu, P. Battaglia, D. Hassabis, D. Silver, and D. Wierstra. Imagination-Augmented Agents for Deep Reinforcement Learning. *NIPS*, 2017.

Florian Schaefer and Anima Anandkumar. Competitive gradient descent. *NIPS*, 2019.

M. Spivak. *Calculus On Manifolds: A Modern Approach To Classical Theorems Of Advanced Calculus*. Avalon Publishing, 1971.

A. S. Vezhnevets, S. Osindero, T. Schaul, N. Heess, M. Jaderberg, D. Silver, and K. Kavukcuoglu. FeUdal Networks for Hierarchical Reinforcement Learning. *ICML*, 2017.

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. *NIPS*, 2019.

G. Wayne and L. F. Abbott. Hierarchical control using networks trained with higher-level forward models. *Neural Computation*, 26(10):2163–2193, 2014.

Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance-reduced optimization for a class of nonconvex-nonconcave minimax problems. *ArXiv e-prints*, 2020.

Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. Policy optimization provably converges to nash equilibria in zero-sum linear quadratic games. *NIPS*, 2019.