COLUMBIA UNIVERSITY

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

Lampros Flokas*, Emmanouil V. Vlatakis-Gkaragkounis*, Georgios Piliouras

Columbia University & Singapore University of Technology and Design

SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN

Introduction

⇒ Min-max optimization is a problem of interest in several communities including Optimization, Game Theory and Machine Learning.

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

⇒ Unfortunately for non-convex non-concave r little is known, not even when the problem has meaningful solutions.

⇒ Min-max optimization has attracted renewed interest thanks to the application of training GANs.

Bilinear Games vs GANs

⇒ Bilinear games

$$\arg\min_{\theta\in\Delta_n} \arg\max_{\phi\in\Delta_m} \theta^T U \phi$$

→ Convex Concave

→ Players control probabilities directly

⇒ GANs

$$\arg\min_{\theta\in\mathbb{R}^n} \arg\max_{\phi\in\mathbb{R}^m} r(\theta, \phi)$$

→ Non-Convex Non-Concave

→ Players control weights of Generators and Discriminators that compete with each other.

How can we bridge the gap in theoretical understanding between bilinear games and GANs?

Hidden Bilinear Games

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

⇒ Each player has a smooth function not necessarily concave or convex F and G

⇒ They get rewarded based on a hidden bilinear game whose Nash equilibria map to equilibria of this game

But: Can players use gradient descent to find them?

Gradient-Descent-Ascent

$$\underbrace{\begin{cases} \dot{\theta} = -\nabla_\theta r(\theta, \phi) \\ \dot{\phi} = \nabla_\phi r(\theta, \phi) \end{cases}}_{\text{Continuous time:CGDA}} \quad \text{VS} \quad \underbrace{\begin{cases} \theta_{k+1} = \theta_k - \alpha \nabla_\theta r(\theta_k, \phi_k) \\ \phi_{k+1} = \phi_k + \alpha \nabla_\phi r(\theta_k, \phi_k) \end{cases}}_{\text{Discrete time:DGDA}}$$

Hidden Bilinear Games with two strategies

⇒ Hidden Bilinear Game setup

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

⇒ Nash equilibrium is p, q and there is a scalar v

$$\dot{\theta} = -v \nabla f(\theta)(g(\phi) - q) \;\;\&\;\; \dot{\phi} = v \nabla g(\phi)(f(\theta) - p)$$

Safety of an initialization

⇒ Observe that $\theta(t)$ cannot traverse stationary points of f.

⇒ This traps $f(\theta(t))$ in an interval $f_{\theta(0)}$. Same thing applies to g.

⇒ We call $(\theta(0), \phi(0))$ safe if $\nabla f(\theta(0)) \neq \mathbf{0}$ and $\nabla g(\phi(0)) \neq \mathbf{0}$ and $p \in f_{\theta(0)}$ and $q \in g_{\phi(0)}$

Possible behaviours of $f(\theta(t))$ and $g(\phi(t))$

⇒ Periodic orbits around the (p, q) equilibrium

Theorem (Periodical Orbits)

Let $\theta(0)$ and $\phi(0)$ be safe initial conditions. Then the orbit $(\theta(t), \phi(t))$ is periodic. Additionally

$$\lim_{T\to\infty} \frac{\int_0^T f(\theta(t))\mathrm{d}t}{T} = p, \qquad \lim_{T\to\infty} \frac{\int_0^T g(\phi(t))\mathrm{d}t}{T} = q$$

⇒ Or even convergence to stationary points of f and g

Theorem (Spurious equilibria)

One can construct functions f and g 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.

Full version at `arxiv:TBA`

Main ideas behind the proof

⇒ Even if f, g are non-invertible, there exist smooth functions such that $\theta(t) = X_{\theta(0)}(f(t))$, $\phi(t) = X_{\phi(0)}(g(t))$.

⇒ The hidden game of f, g mimics properties and behaviors of conservative systems. The following quantity is time invariant:

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

⇒ Apply the Poincaré-Bendixon theorem on the planar dynamical system of (f, g) to get periodicity.

Extensions

More than two strategies:

⇒ Let vectors $F(\theta)$, $G(\phi)$ have size (N, M) whose each coordinate is controlled by disjoint subsets of the vectors θ, ϕ, i.e.

$$\boldsymbol{F}(\theta) = \left[f_1(\theta_1) \cdots f_N(\theta_N)\right] \qquad \boldsymbol{G}(\phi) = \left[g_1(\phi_1) \cdots g_M(\phi_M)\right]$$

⇒ CGDA is bounded and diffeomorphic to a volume-preserving system.

⇒ **Poincaré recurrence** theorem states that volume-preserving bounded systems will, after a sufficiently long but finite time, recurturn to a state very close to their initial state.

Question: Discretization is critical? Answer: *No*

Theorem (Spurious equilibria)

One can choose a learning rate α and functions f, g (i.e sigmoids) for DGDA so that its energy H is always non-decreasing and for a positive measure set of initial conditions the trajectories converge to fixed points that do not correspond to equilibria of the hidden game.

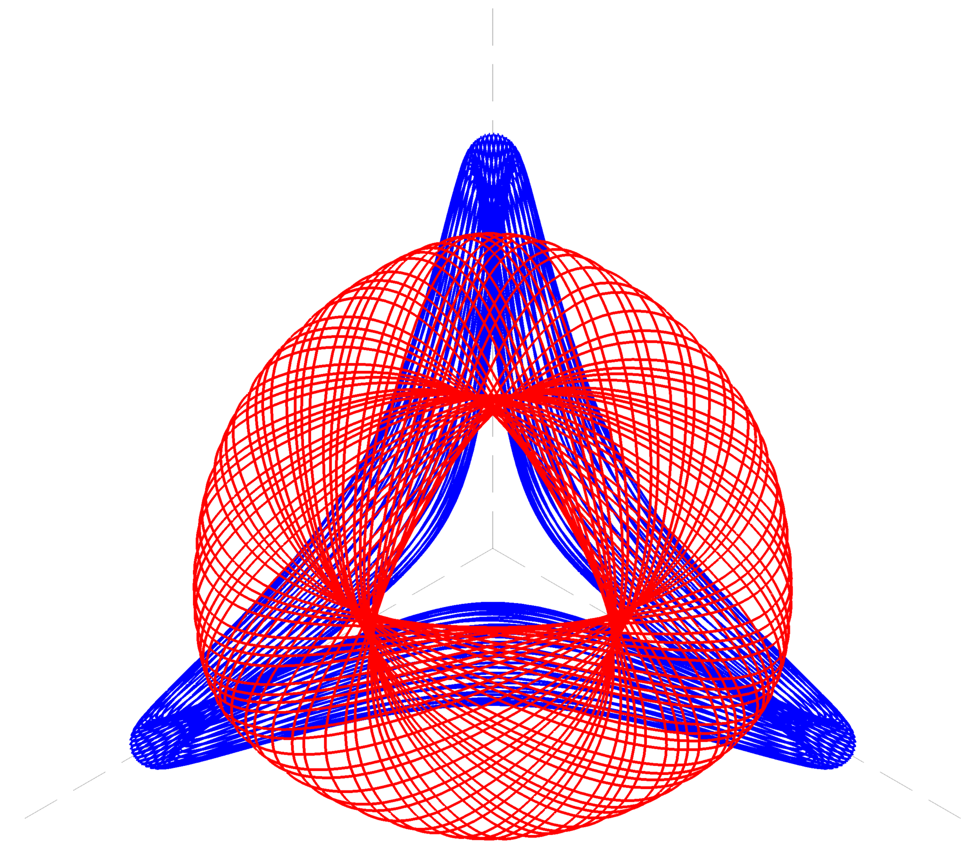

Figure: Rock-Paper-Scissor dynamics with sigmoid functions.