PRINCETON UNIVERSITY

On the *Global* Convergence of Gradient-Based Learning in Continuous Zero-Sum Games

"Tony" Runzhe Yang

Feb. 14, 2020

https://runzhe-yang.science

Minimax Problems in Machine Learning

$$\min_G \max_D V(D,G) = \mathbb{E}_{\boldsymbol{x}\sim p_{\text{data}}(x)}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z}\sim p_{\boldsymbol{z}}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z})))]$$

Generative Adversarial Networks (GANs)

[Goodfellow et al., 2014]

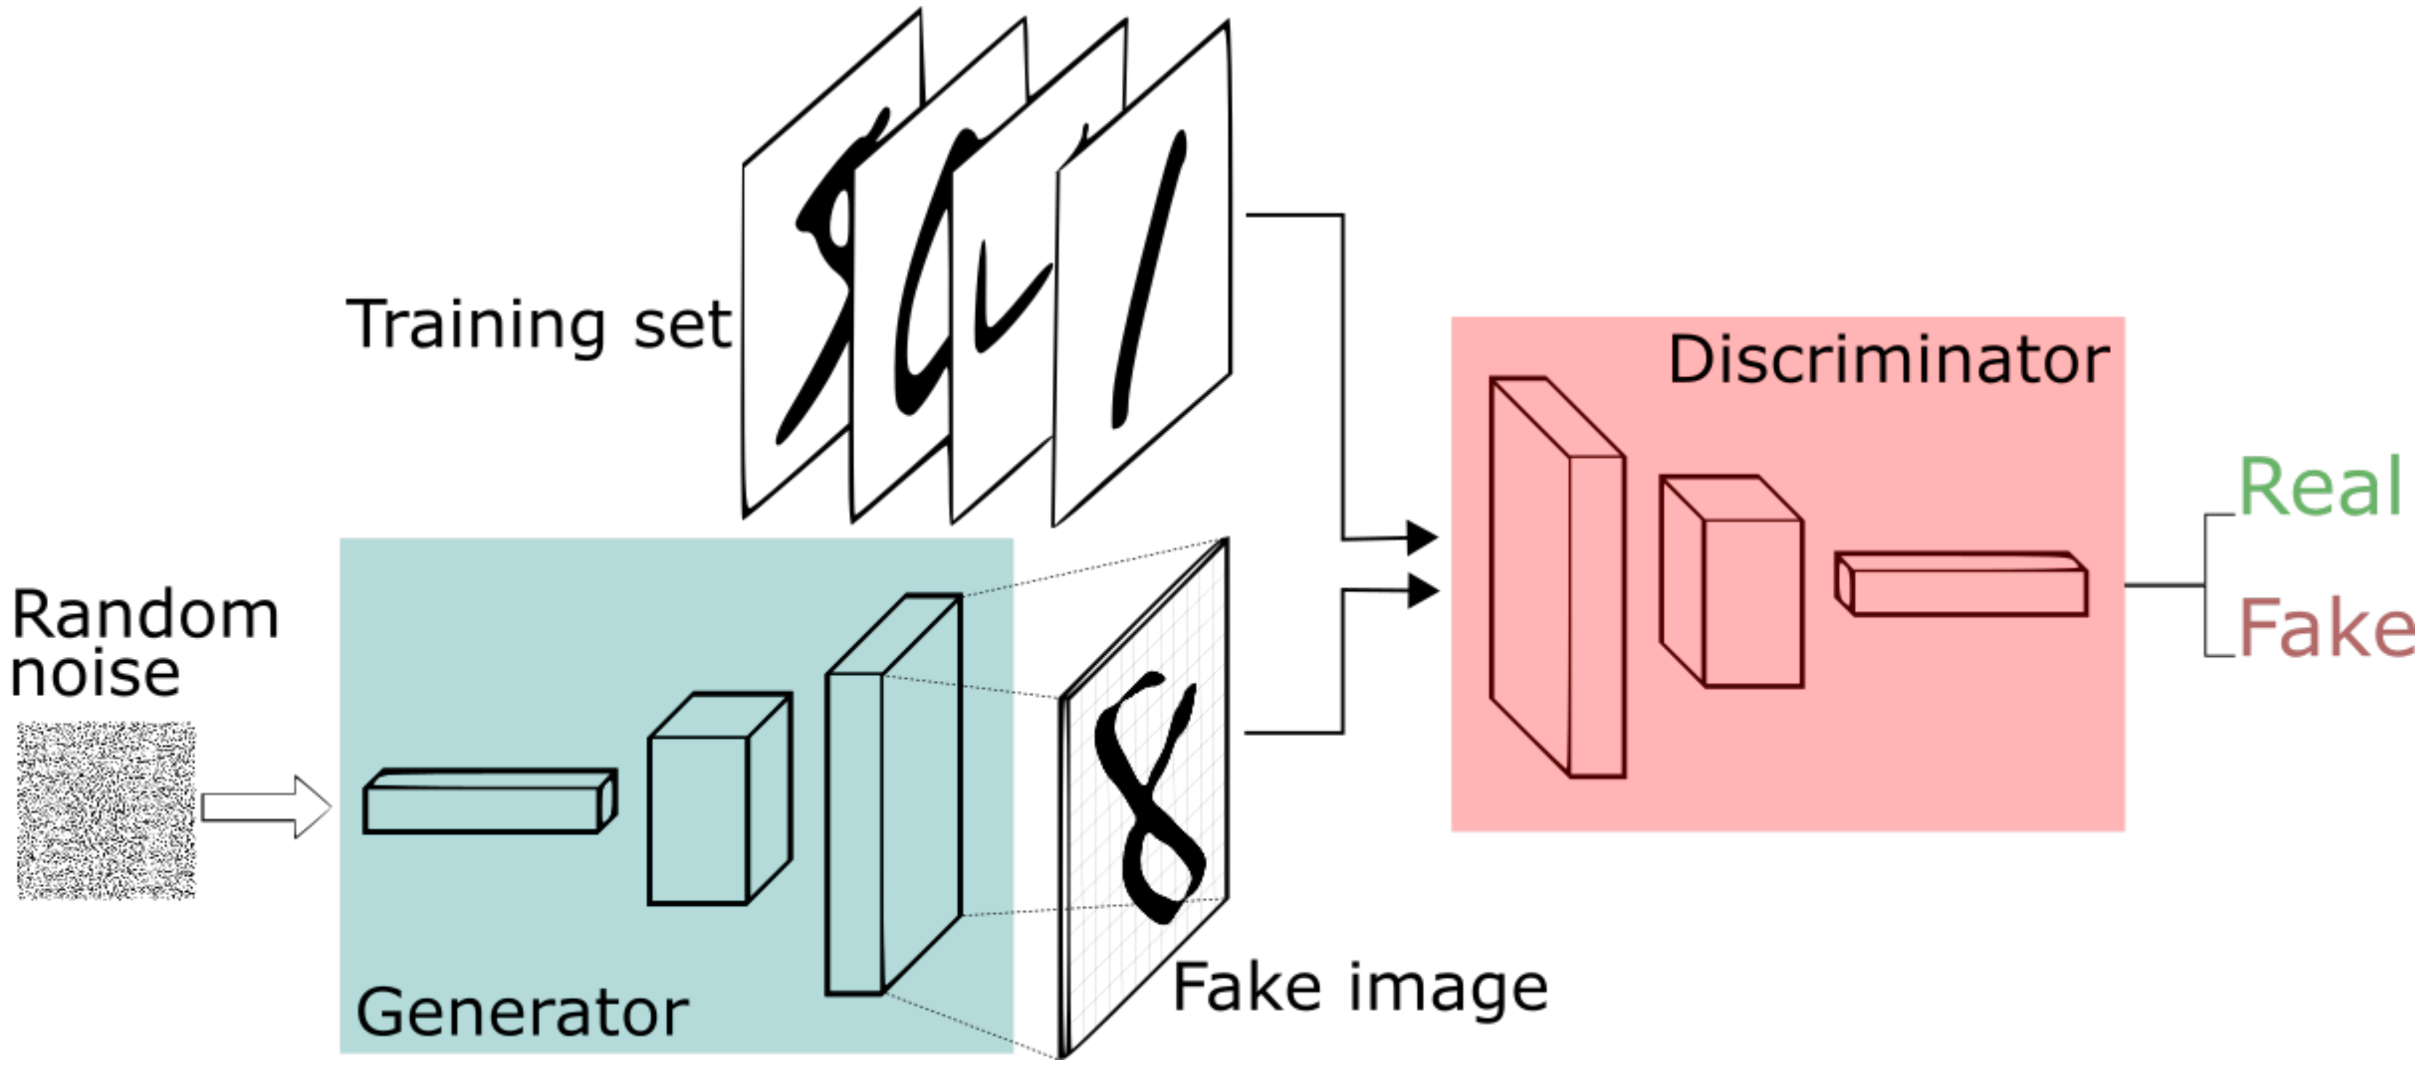

GANs Illustration: https://medium.freecodecamp.org/an-intuitive-introduction-to-generative-adversarial-networks-gans-7a2264a81394

Minimax Problems in Machine Learning

$$\min_{\pi} \max_{C} V(\pi, C) = \mathbb{E}_{\tau\sim\pi}[-C(\tau)] + \log\left(E_{\tau\sim\pi}[\pi(\tau)/\exp(-C(\tau))]\right)$$

Maximum Entropy Inverse Reinforcement Learning (MaxEnt IRL)

[Ziebart et al., 2008]&[Finn et al., 2016 - slightly different here]

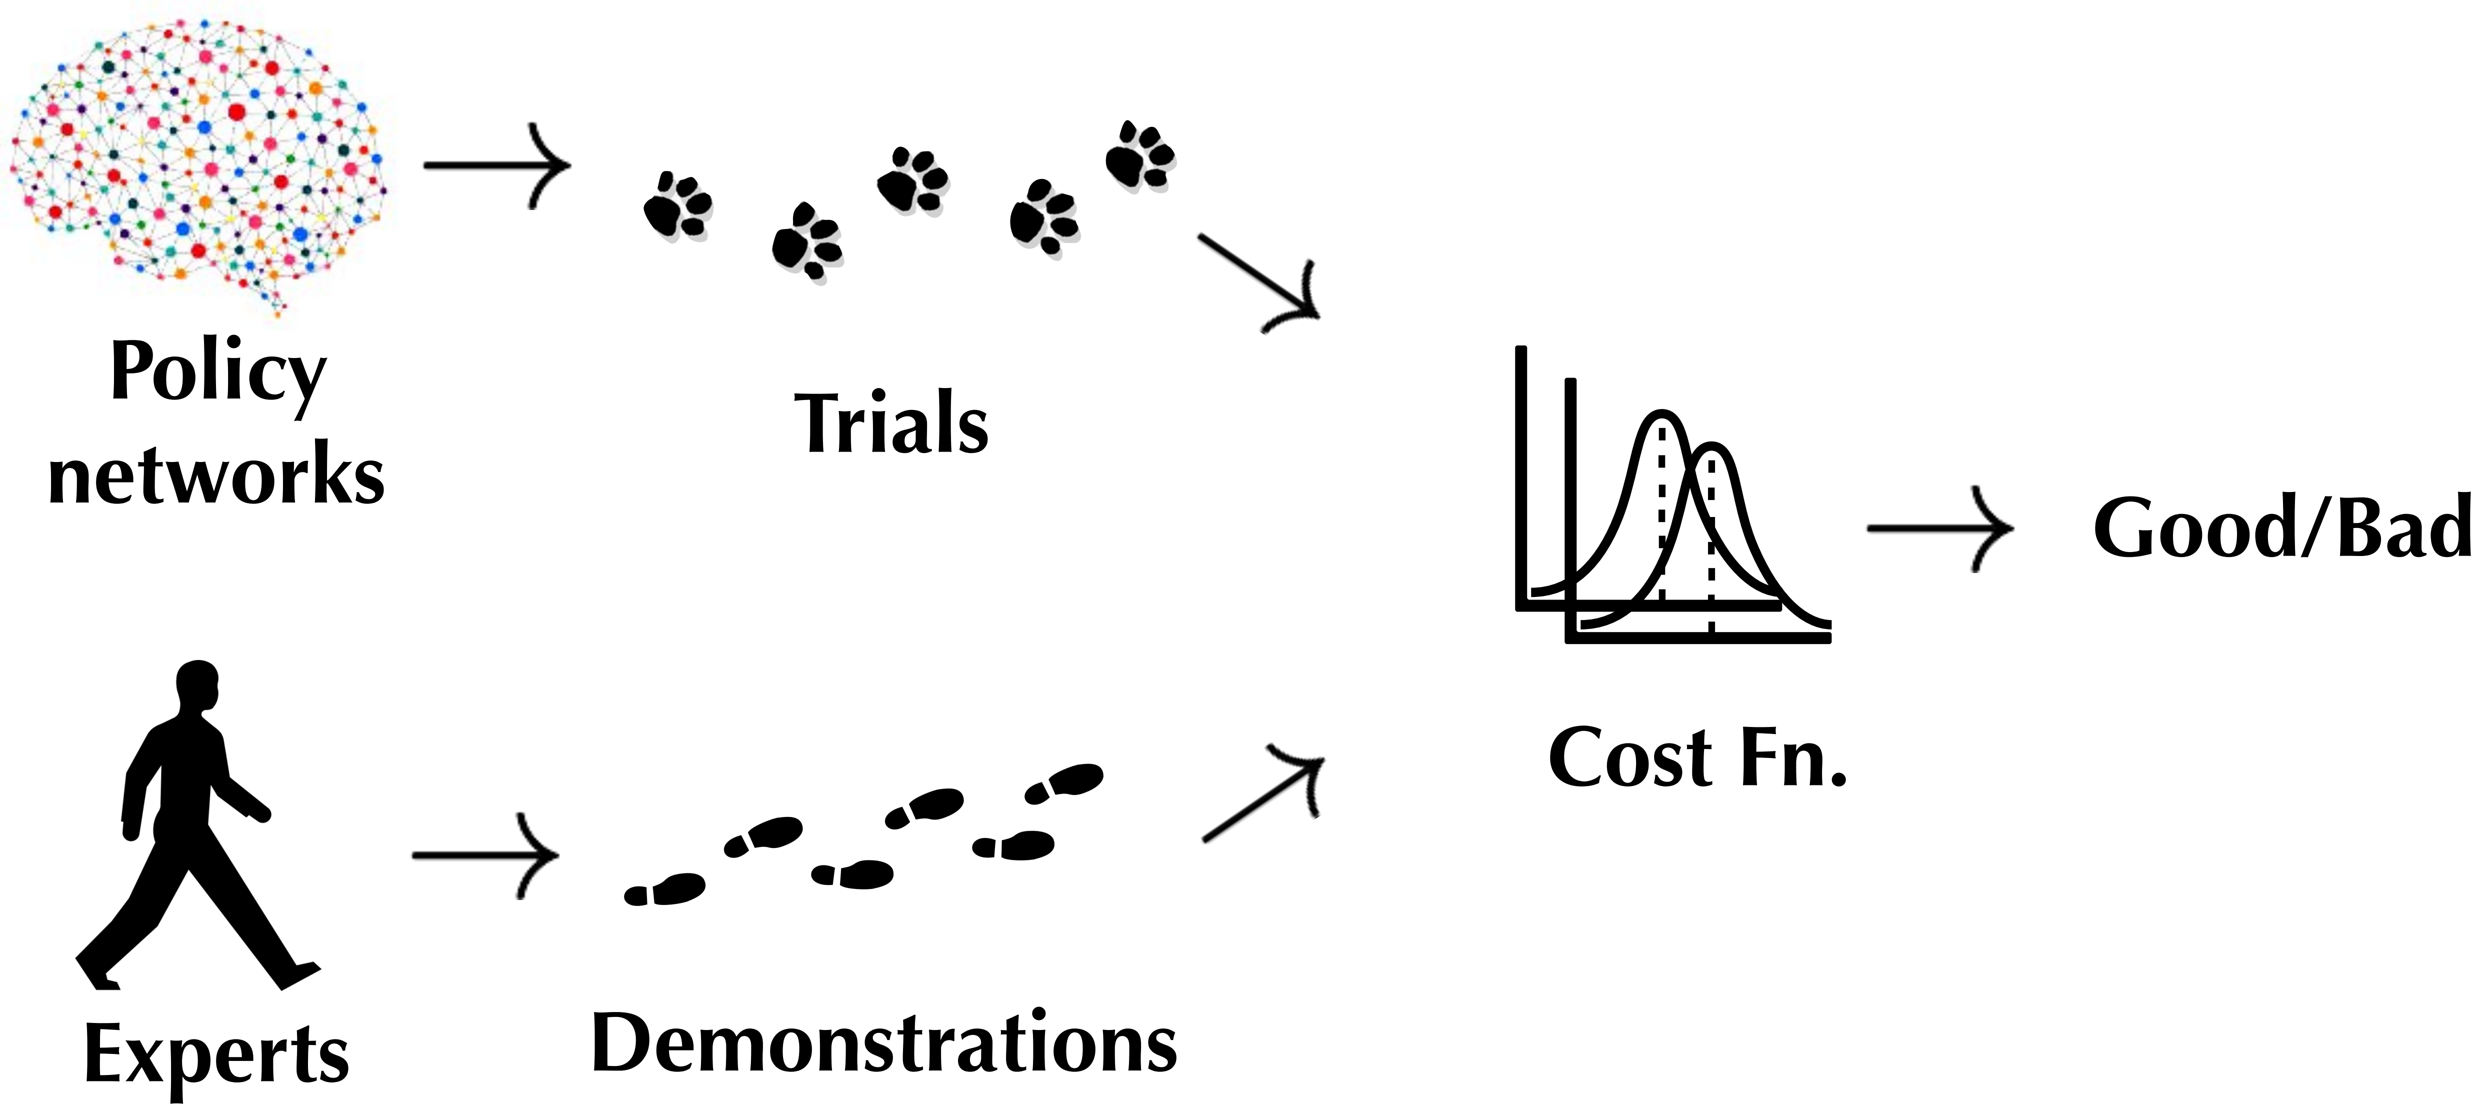

IRL Illustration: https://runzhe-yang.science/demo/Inverse_Reinforcement_Learning.pdf

Minimax Problems in Game Theory

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} S(\boldsymbol{x}, \boldsymbol{y})$$

	Agent 1	**Agent 2**
Action space:	$\boldsymbol{x} \in \mathbb{R}^n$	$\boldsymbol{y} \in \mathbb{R}^m$
Payoff function:	$-S(\boldsymbol{x}, \boldsymbol{y})$	$S(\boldsymbol{x}, \boldsymbol{y})$

sequential game - agent 1 moves first (slower)

Gradient Descent-Ascent Dynamics of GANs

$$\min_G \max_D V(D,G) = \mathbb{E}_{\boldsymbol{x}\sim p_{\text{data}}(x)}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z}\sim p_{\boldsymbol{z}}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z})))]$$

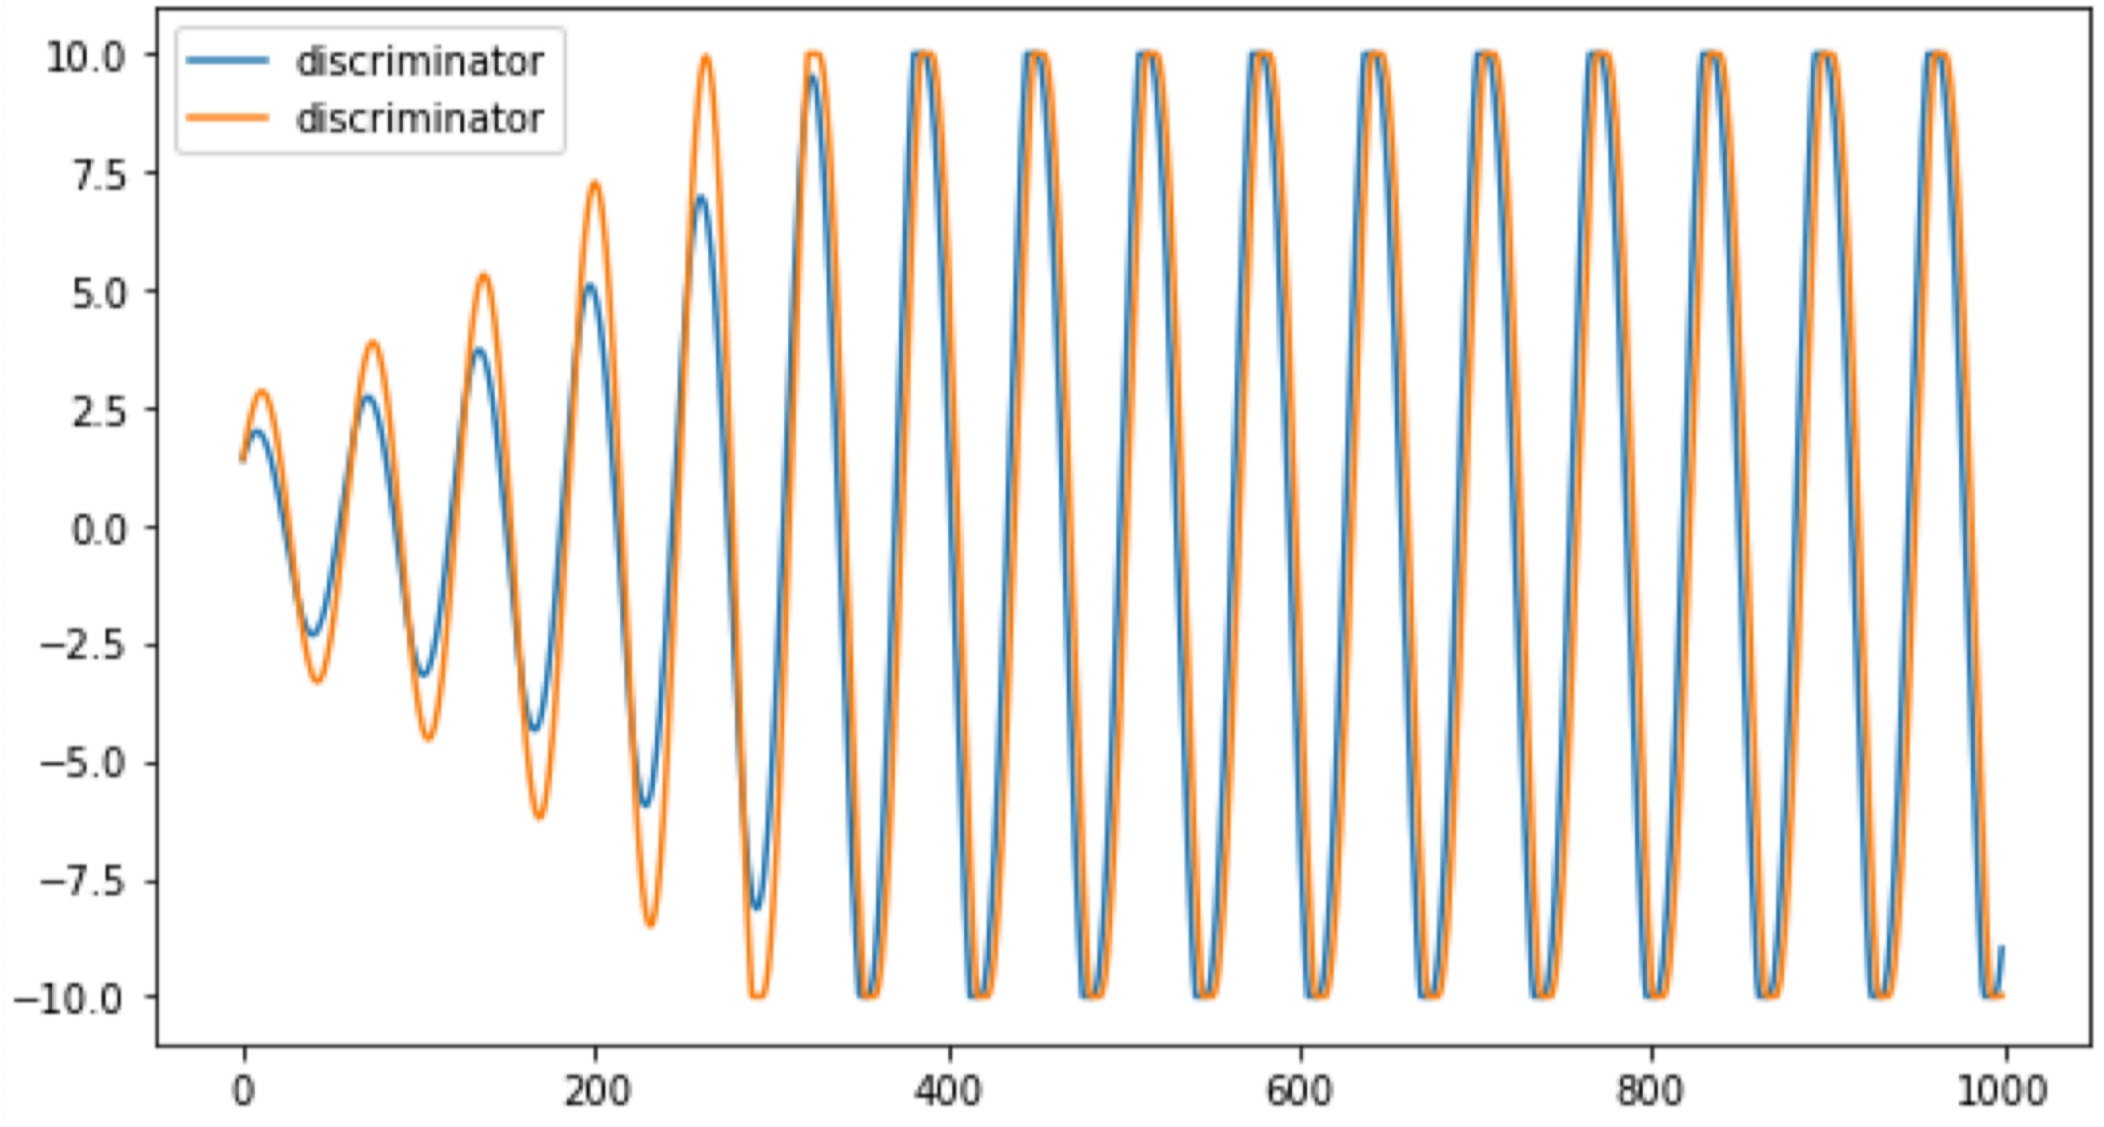

$$\theta_D = \theta_D + \eta \nabla_{\theta_G} V(D,G)$$

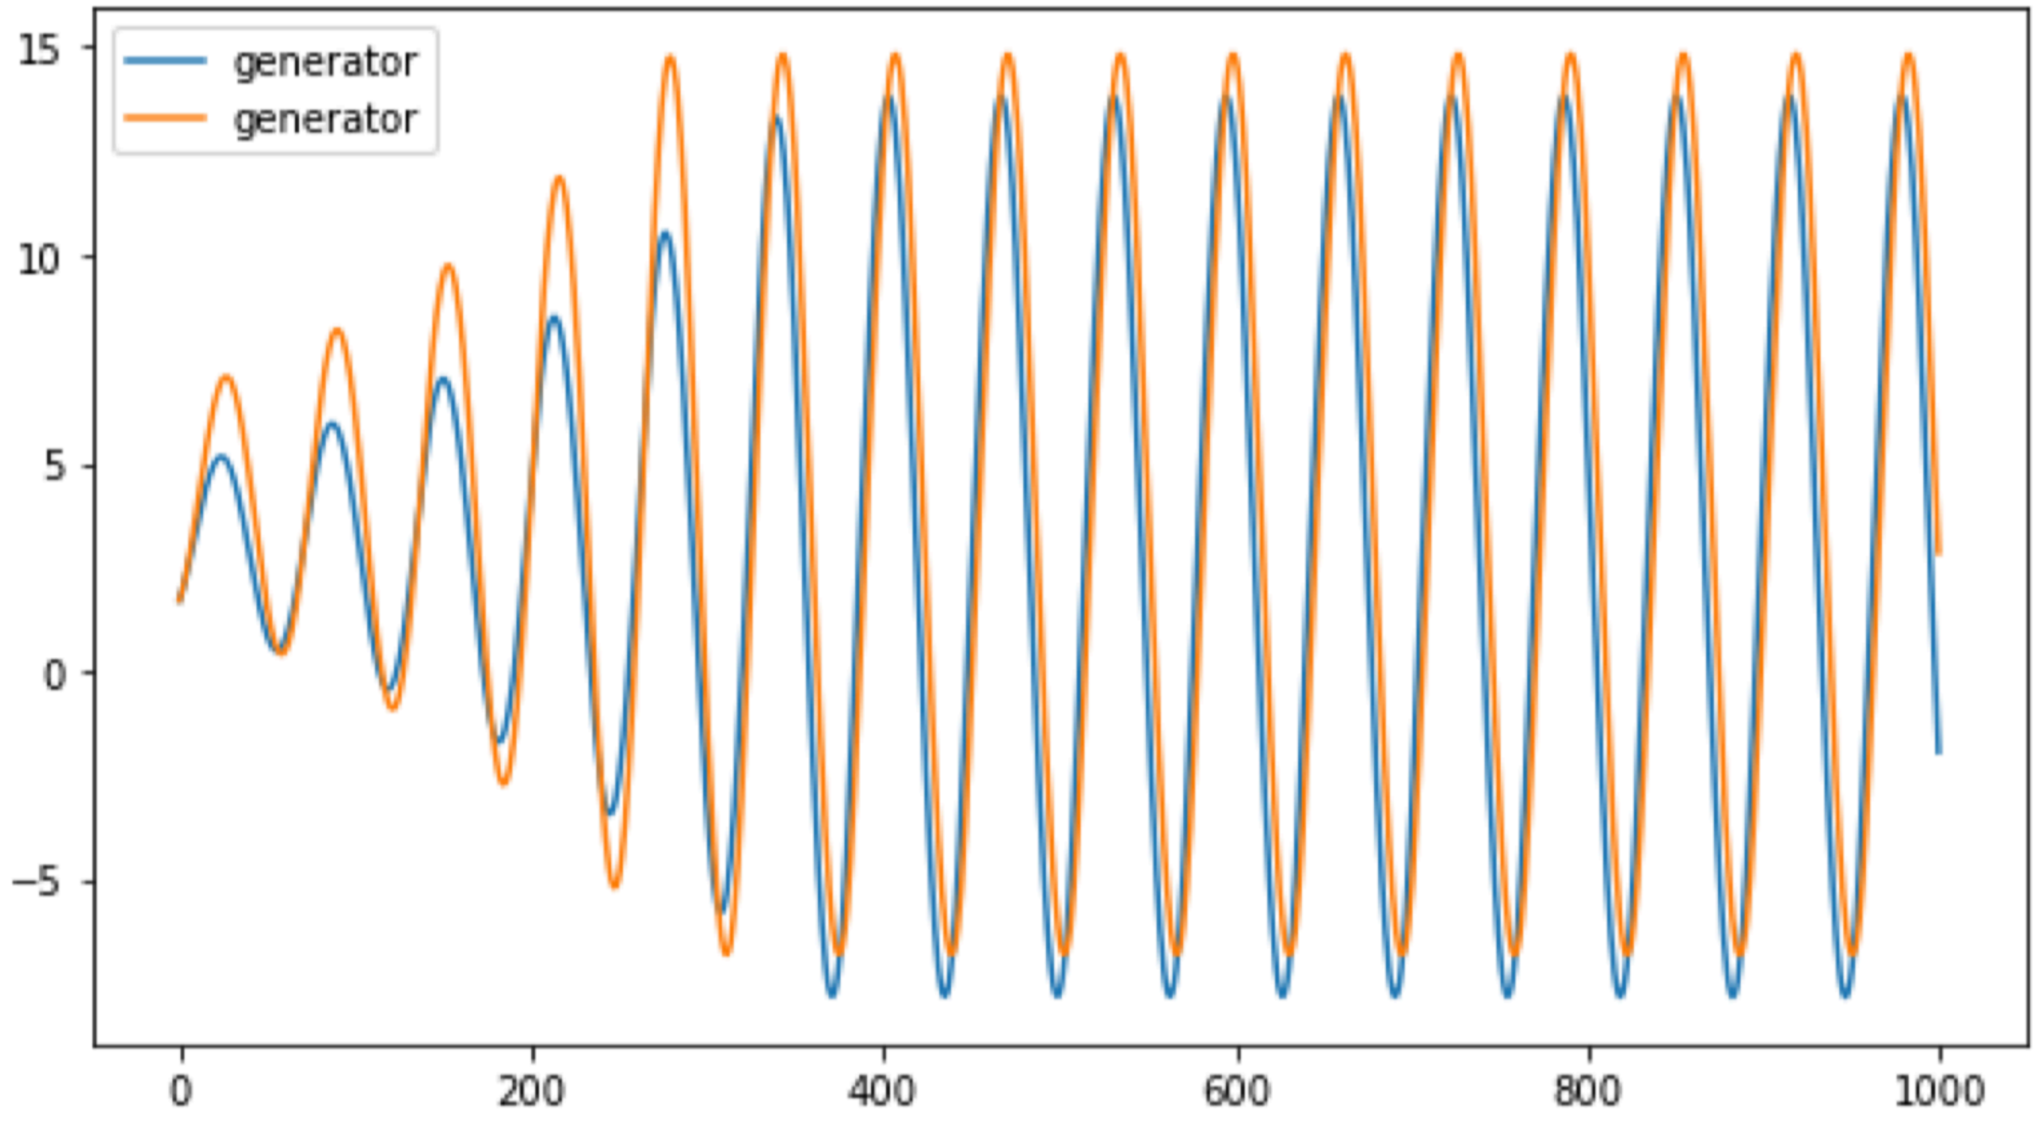

$$\theta_G = \theta_G - \eta \nabla_{\theta_G} V(D,G)$$

Daskalakis et al., Training GANs with Optimism, ICLR 2018

Gradient Descent-Ascent Dynamics

$$\boxed{\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} S(\boldsymbol{x}, \boldsymbol{y})}$$

Continuous-time system:
$$\begin{cases} \tau_x \dot{\boldsymbol{x}} = -\nabla_{\boldsymbol{x}} S(\boldsymbol{x}, \boldsymbol{y}) \\ \tau_y \dot{\boldsymbol{y}} = \nabla_{\boldsymbol{y}} S(\boldsymbol{x}, \boldsymbol{y}) \end{cases}$$

Two questions:

1. **When does the system converge?**
2. **Where does the system converge, if it converges?**

Optimality in Minimax Optimization

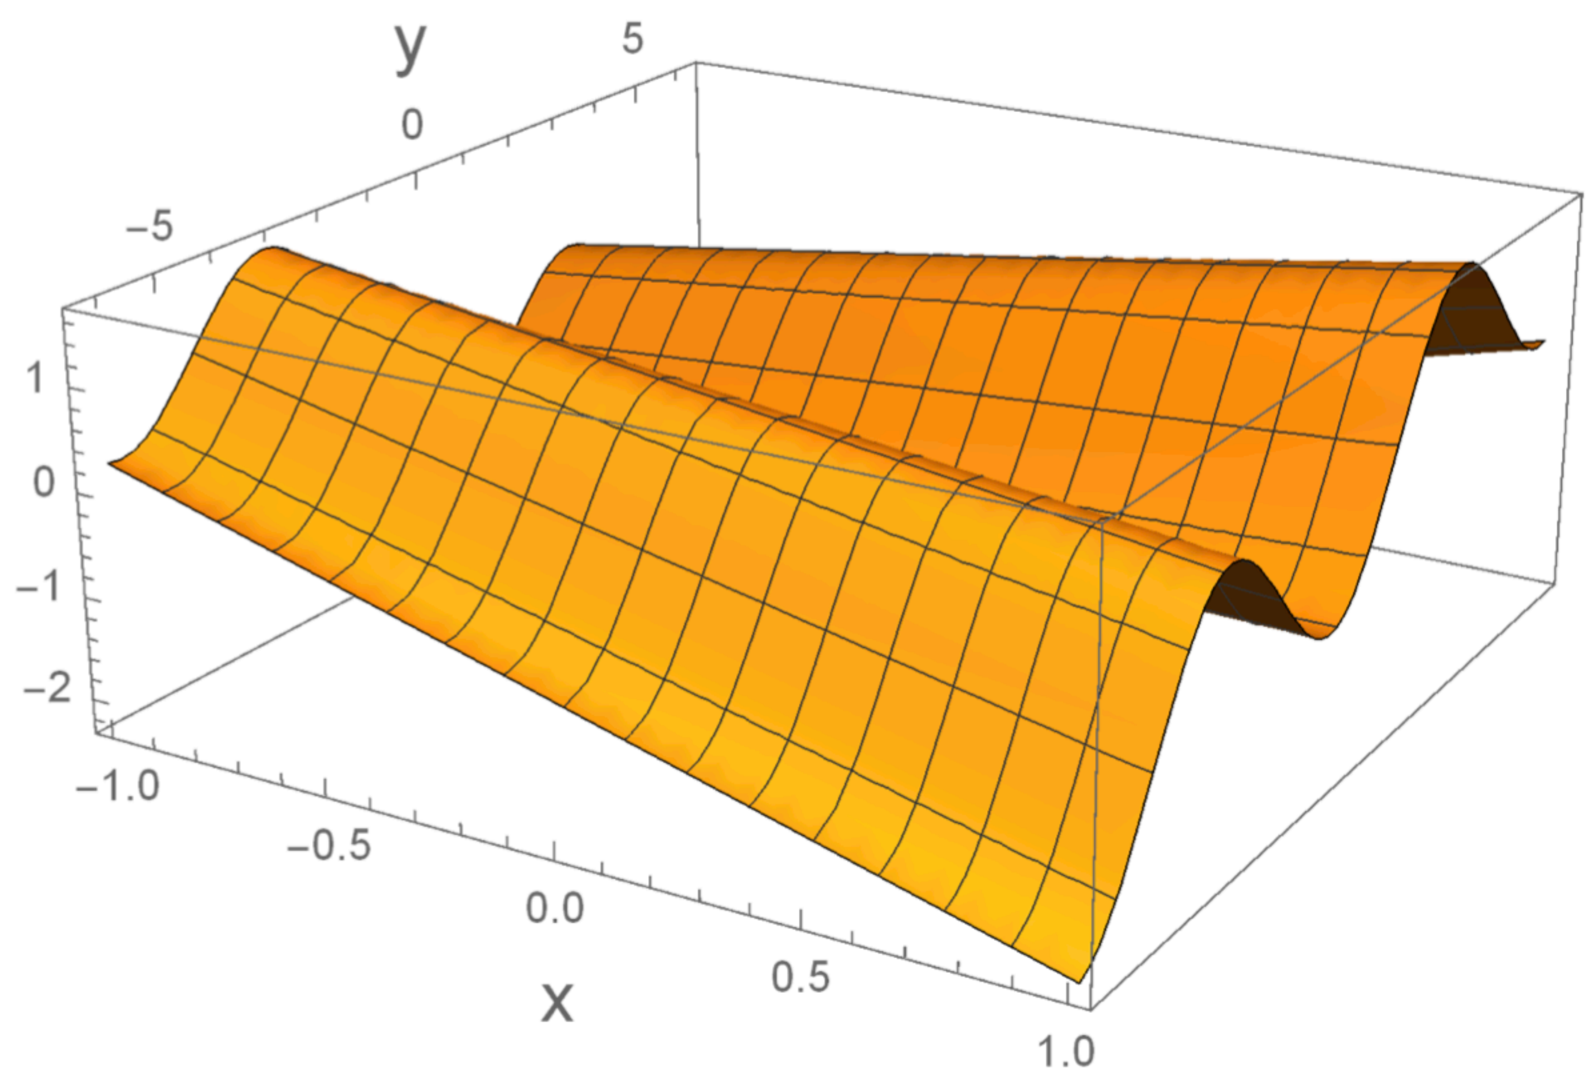

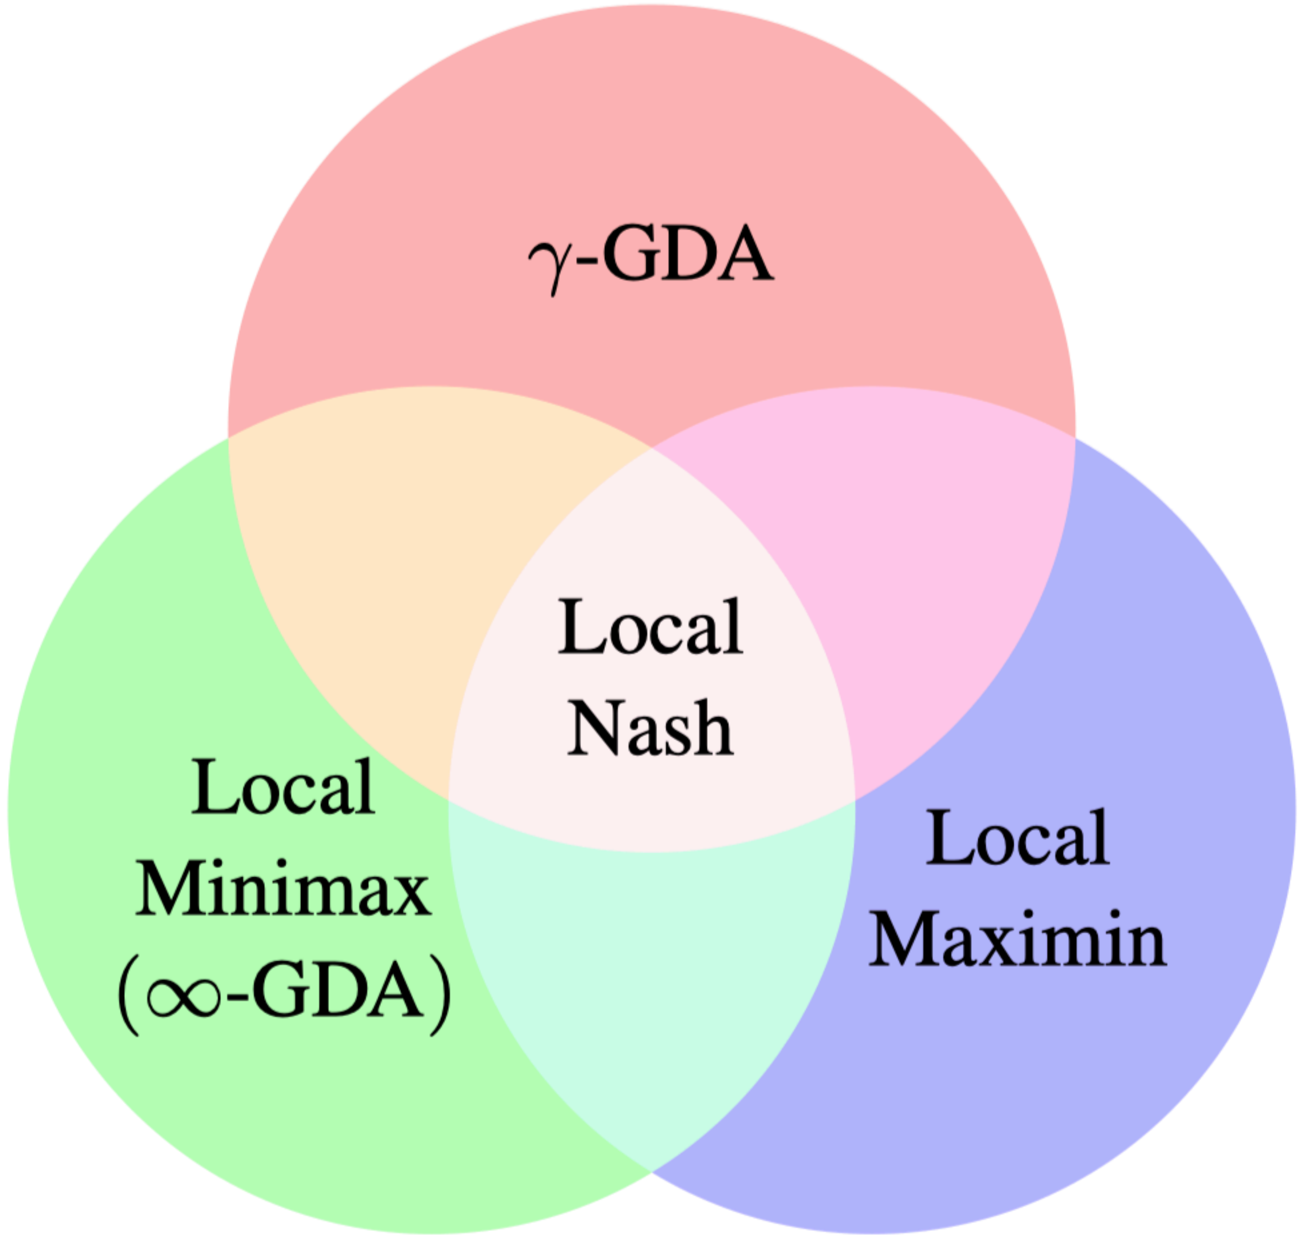

Where does the system converge, if it converges?

$$S(\boldsymbol{x}^*, \boldsymbol{y}) \leq S(\boldsymbol{x}^*, \boldsymbol{y}^*) \leq \max_{\boldsymbol{y}'} S(\boldsymbol{x}, \boldsymbol{y}')$$

Jin et al., What is Local Optimality in Nonconvex-Nonconcave Minimax Optimization? ArXiv, 2019

Kose-Uzawa Theorem: convex-concave payoff fn.

Continuous-time system: $$\begin{cases} \tau_x \dot{\boldsymbol{x}} = -\nabla_{\boldsymbol{x}} S(\boldsymbol{x}, \boldsymbol{y}) \\ \tau_y \dot{\boldsymbol{y}} = \nabla_{\boldsymbol{y}} S(\boldsymbol{x}, \boldsymbol{y}) \end{cases}$$

"If *S(x, y)* is strictly <u>convex</u> in *x* and strictly <u>concave</u> in *y*, then gradient descent-ascent converges a saddle point."

E.g., $S(x, y) = (x^2 - y^2)/2$

T. Kose. Solutions of Saddle Value Problems by Differential Equations. Econometrica, 24:59–70, 1956.
H. Uzawa. Gradient method for concave programming, II global stability in the strictly concave case. Stanford Univ., 1958.

New Theorem: globally *less convex* in *y*

Continuous-time system: $\begin{cases} \tau_x \dot{\boldsymbol{x}} = -\nabla_{\boldsymbol{x}} S(\boldsymbol{x}, \boldsymbol{y}) \\ \tau_y \dot{\boldsymbol{y}} = \nabla_{\boldsymbol{y}} S(\boldsymbol{x}, \boldsymbol{y}) \end{cases}$

***S* is twice differentiable**

$$\lambda_{\inf}(S_{\boldsymbol{xx}}) := \inf_{\boldsymbol{x},\boldsymbol{y}} \lambda_{\min}(S_{\boldsymbol{xx}}(\boldsymbol{x}, \boldsymbol{y}))$$

$$\lambda_{\sup}(S_{\boldsymbol{yy}}) := \sup_{\boldsymbol{x},\boldsymbol{y}} \lambda_{\max}(S_{\boldsymbol{yy}}(\boldsymbol{x}, \boldsymbol{y}))$$

If $\tau_x^{-1} \lambda_{\inf}(S_{\boldsymbol{xx}}) > \tau_y^{-1} \lambda_{\sup}(S_{\boldsymbol{yy}})$ **and steady states exists, then all bounded trajectories converge.**

LaSalle's Invariance Principle

System: $\dot{\boldsymbol{z}} = f(\boldsymbol{z}), \boldsymbol{z} \in \mathbb{R}^k, f(\boldsymbol{z}^*) = \boldsymbol{0}$

$L : \mathbb{R}^k \to \mathbb{R}$

continuously differentiable.

$E = \{\boldsymbol{z} \in \mathbb{R}^k : \dot{L}(\boldsymbol{z}) = \boldsymbol{0}\}$

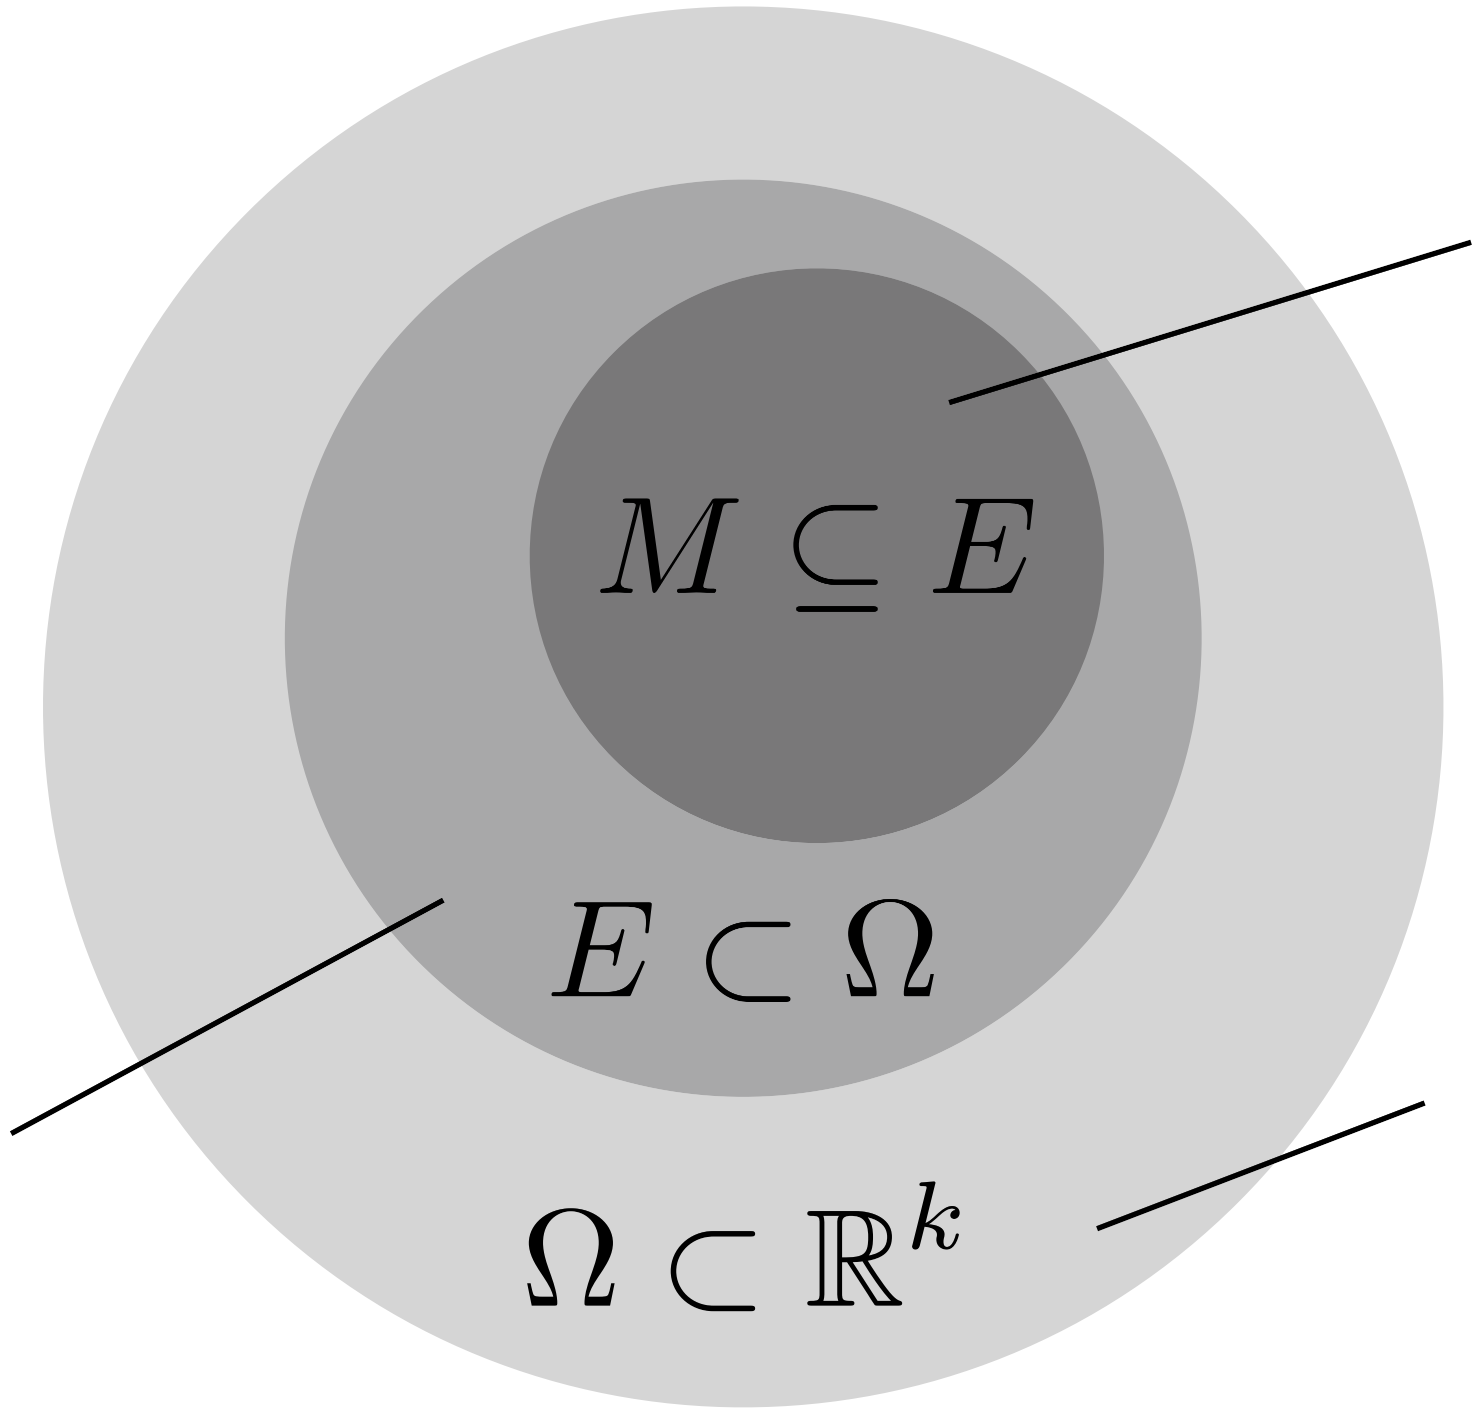

the <u>largest invariant set</u> in E

$\boldsymbol{z}(0) \in M \Rightarrow \boldsymbol{z}(t) \in M, \forall t \in \mathbb{R}$

compact <u>positive invariant set</u>

w.r.t. the system dynamics

$\boldsymbol{z}(0) \in \Omega \Rightarrow \boldsymbol{z}(t) \in \Omega, \forall t \geq 0$

$\dot{L}(\boldsymbol{z}(t)) \leq \boldsymbol{0}$ **in** Ω

$$\lim_{t \to \infty} \left(\inf_{\boldsymbol{p} \in M} \|\boldsymbol{z}(t) - p\| \right) = 0$$

Construction of a Lyapunov Function

Continuous-time system: $\begin{cases} \tau_x \dot{\boldsymbol{x}} = -S_{\boldsymbol{x}} \\ \tau_y \dot{\boldsymbol{y}} = S_{\boldsymbol{y}} \end{cases}$

$$\Rightarrow \begin{cases} \tau_x \ddot{\boldsymbol{x}} = -S_{\boldsymbol{xx}} \dot{\boldsymbol{x}} - S_{\boldsymbol{xy}} \dot{\boldsymbol{y}} \\ \tau_y \ddot{\boldsymbol{y}} = S_{\boldsymbol{yx}} \dot{\boldsymbol{x}} + S_{\boldsymbol{yy}} \dot{\boldsymbol{y}} \end{cases}$$

Scale transformation: $\boldsymbol{a} = \sqrt{\tau_x} \boldsymbol{x}$ **,** $\boldsymbol{b} = \sqrt{\tau_y} \boldsymbol{y}$

$$\begin{pmatrix} \ddot{\boldsymbol{a}} \\ \ddot{\boldsymbol{b}} \end{pmatrix} = - \begin{pmatrix} S_{\boldsymbol{xx}}/\tau_x & S_{\boldsymbol{xy}}/\sqrt{\tau_x \tau_y} \\ -S_{\boldsymbol{yx}}/\sqrt{\tau_x \tau_y} & -S_{\boldsymbol{yy}}/\tau_y \end{pmatrix} \begin{pmatrix} \dot{\boldsymbol{a}} \\ \dot{\boldsymbol{b}} \end{pmatrix}$$

Construction of a Lyapunov Function

$$\begin{pmatrix} \ddot{\boldsymbol{a}} \\ \ddot{\boldsymbol{b}} \end{pmatrix} = -\begin{pmatrix} S_{\boldsymbol{xx}}/\tau_x & S_{\boldsymbol{xy}}/\sqrt{\tau_x\tau_y} \\ -S_{\boldsymbol{yx}}/\sqrt{\tau_x\tau_y} & -S_{\boldsymbol{yy}}/\tau_y \end{pmatrix} \begin{pmatrix} \dot{\boldsymbol{a}} \\ \dot{\boldsymbol{b}} \end{pmatrix}$$

$$\parallel$$

$$K_r = \begin{pmatrix} rI_n & \mathbf{0} \\ \mathbf{0} & -rI_m \end{pmatrix}$$ **"electric"**

$$+$$

$$K_A = \begin{pmatrix} \mathbf{0} & S_{\boldsymbol{xy}}/\sqrt{\tau_x\tau_y} \\ -S_{\boldsymbol{yx}}/\sqrt{\tau_x\tau_y} & \mathbf{0} \end{pmatrix}$$ **"magnetic"**

$$+$$

$$K_S = \begin{pmatrix} S_{\boldsymbol{xx}}/\tau_x - rI_n & \mathbf{0} \\ \mathbf{0} & -(S_{\boldsymbol{yy}}/\tau_y - rI_m) \end{pmatrix}$$ **"frictional"**

Construction of a Lyapunov Function

$$\begin{pmatrix} \ddot{\boldsymbol{a}} \\ \ddot{\boldsymbol{b}} \end{pmatrix} = -\begin{pmatrix} S_{\boldsymbol{xx}}/\tau_x & S_{\boldsymbol{xy}}/\sqrt{\tau_x\tau_y} \\ -S_{\boldsymbol{yx}}/\sqrt{\tau_x\tau_y} & -S_{\boldsymbol{yy}}/\tau_y \end{pmatrix} \begin{pmatrix} \dot{\boldsymbol{a}} \\ \dot{\boldsymbol{b}} \end{pmatrix}$$

$$K_r = \begin{pmatrix} rI_n & \mathbf{0} \\ \mathbf{0} & -rI_m \end{pmatrix}$$

"electric"

Let $\boldsymbol{z} = \begin{pmatrix} \boldsymbol{a} \\ \boldsymbol{b} \end{pmatrix}$

"potential function"

$$\Phi(\boldsymbol{x}, \boldsymbol{y}) = -rS(\boldsymbol{x}, \boldsymbol{y})$$

$$-K_r\dot{\boldsymbol{z}} = -\Phi_{\boldsymbol{z}}$$

Construction of a Lyapunov Function

$$\begin{pmatrix} \ddot{\boldsymbol{a}} \\ \ddot{\boldsymbol{b}} \end{pmatrix} = - \begin{pmatrix} S_{\boldsymbol{xx}}/\tau_x & S_{\boldsymbol{xy}}/\sqrt{\tau_x \tau_y} \\ -S_{\boldsymbol{yx}}/\sqrt{\tau_x \tau_y} & -S_{\boldsymbol{yy}}/\tau_y \end{pmatrix} \begin{pmatrix} \dot{\boldsymbol{a}} \\ \dot{\boldsymbol{b}} \end{pmatrix}$$

$$K_A = \begin{pmatrix} \mathbf{0} & S_{\boldsymbol{xy}}/\sqrt{\tau_x \tau_y} \\ -S_{\boldsymbol{yx}}/\sqrt{\tau_x \tau_y} & \mathbf{0} \end{pmatrix}$$ **"magnetic"**

anti-symmetric $-K_A = K_A^\intercal$

$$\dot{\boldsymbol{z}}^\intercal K_A \dot{\boldsymbol{z}} = -\dot{\boldsymbol{z}}^\intercal K_A^\intercal \dot{\boldsymbol{z}} = 0$$

always perpendicular to the velocity $K_A \dot{\boldsymbol{z}} \perp \dot{\boldsymbol{z}}$

Construction of a Lyapunov Function

$$\begin{pmatrix} \ddot{\boldsymbol{a}} \\ \ddot{\boldsymbol{b}} \end{pmatrix} = -\begin{pmatrix} S_{\boldsymbol{xx}}/\tau_x & S_{\boldsymbol{xy}}/\sqrt{\tau_x\tau_y} \\ -S_{\boldsymbol{yx}}/\sqrt{\tau_x\tau_y} & -S_{\boldsymbol{yy}}/\tau_y \end{pmatrix} \begin{pmatrix} \dot{\boldsymbol{a}} \\ \dot{\boldsymbol{b}} \end{pmatrix}$$

$$\ddot{\boldsymbol{z}} = -\underbrace{\Phi_{\boldsymbol{z}}}_{\text{electric}} - \underbrace{K_A\dot{\boldsymbol{z}}}_{\text{magnetic}} - \underbrace{K_S\dot{\boldsymbol{z}}}_{\text{frictional}}$$

Energy conservation as a Lyapunov function:

$$L(\boldsymbol{z}) = \frac{1}{2}|\dot{\boldsymbol{z}}|^2 + \Phi(\boldsymbol{z}) \qquad \dot{L} = -\dot{\boldsymbol{z}}^\intercal K_S \dot{\boldsymbol{z}}$$

Construction of a Lyapunov Function

Energy conservation as a Lyapunov function:

$$L(\boldsymbol{z}) = \frac{1}{2}|\dot{\boldsymbol{z}}|^2 + \Phi(\boldsymbol{z}) \qquad \dot{L} = -\dot{\boldsymbol{z}}^\intercal K_S \dot{\boldsymbol{z}}$$

$$K_S = \begin{pmatrix} S_{\boldsymbol{xx}}/\tau_x - rI_n & \mathbf{0} \\ \mathbf{0} & -(S_{\boldsymbol{yy}}/\tau_y - rI_m) \end{pmatrix}$$ **"frictional"**

$$K_S \succ 0$$

If $\tau_x^{-1}\lambda_{\inf}(S_{\boldsymbol{xx}}) > \tau_y^{-1}\lambda_{\sup}(S_{\boldsymbol{yy}})$ **and steady states exists, then all bounded trajectories converge.**

Summary

$$\boxed{\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} S(\boldsymbol{x}, \boldsymbol{y})}$$

Continuous-time system: $\begin{cases} \tau_x \dot{\boldsymbol{x}} = -\nabla_{\boldsymbol{x}} S(\boldsymbol{x}, \boldsymbol{y}) \\ \tau_y \dot{\boldsymbol{y}} = \nabla_{\boldsymbol{y}} S(\boldsymbol{x}, \boldsymbol{y}) \end{cases}$

1. **System converges when** $\tau_x^{-1} \lambda_{\inf}(S_{\boldsymbol{xx}}) > \tau_y^{-1} \lambda_{\sup}(S_{\boldsymbol{yy}})$
2. **Boundedness, Discrete time approximation**
3. **Gradient projection, Rectified gradient descent-ascent**