Linear Last-iterate Convergence in Constrained Saddle-point Optimization

Chen-Yu Wei
University of Southern California

Collaborators

Chung-Wei Lee

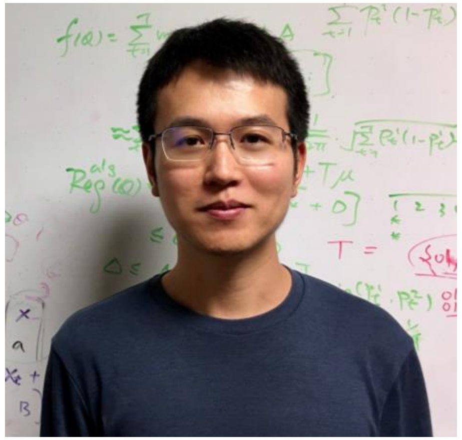

Haipeng Luo

Mengxiao Zhang

Background

Today's topic

Two-player zero-sum game / saddle-point problem: $$\min_x \max_y f(x, y)$$

The goal is to find an (approximate) solution of it. (a.k.a. equilibrium)

Applications:

- Generative Adversarial Networks
- Primal-dual methods (e.g., max flow, MDP, optimal transport)
- Multi-agent (reinforcement) learning problems (e.g., game of Go)

We will focus on a specific class of algorithm **"optimistic mirror descent".**

Formalization

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y)$$

Define $\mathcal{Z} \triangleq \mathcal{X} \times \mathcal{Y}$ and write $z = (x, y) \in \mathcal{Z}$

Assumption 1. Convex-concave game

$\mathcal{X} \subseteq \mathbb{R}^A, \mathcal{Y} \subseteq \mathbb{R}^B$ are convex sets

$f(z) = f(x, y)$ is convex in x and concave in y

Assumption 2. Smoothness

Let $F(z) = \begin{bmatrix} \nabla_x f(x, y) \\ -\nabla_y f(x, y) \end{bmatrix}$ Assume $\|F(z) - F(z')\| \leq L\|z - z'\|$

A simple first-order method

Gradient Descent (GD)

$$x_{t+1} = \Pi_{\mathcal{X}}\left[x_t - \eta \nabla_x f(x_t, y_t)\right]$$

$$y_{t+1} = \Pi_{\mathcal{Y}}\left[y_t + \eta \nabla_y f(x_t, y_t)\right]$$

which can be simplified as

$$z_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z_t)\right]$$

The dynamics of GD

$$f(x, y) = x^\top \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} y$$

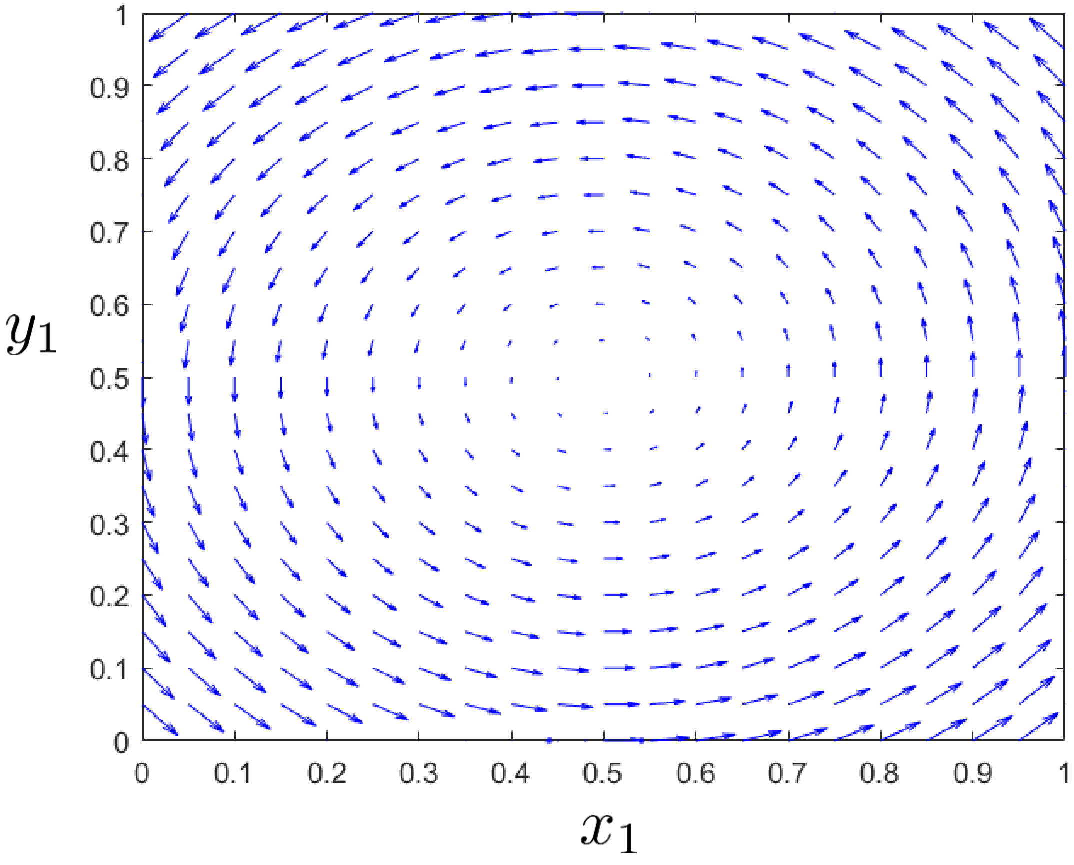

Update direction

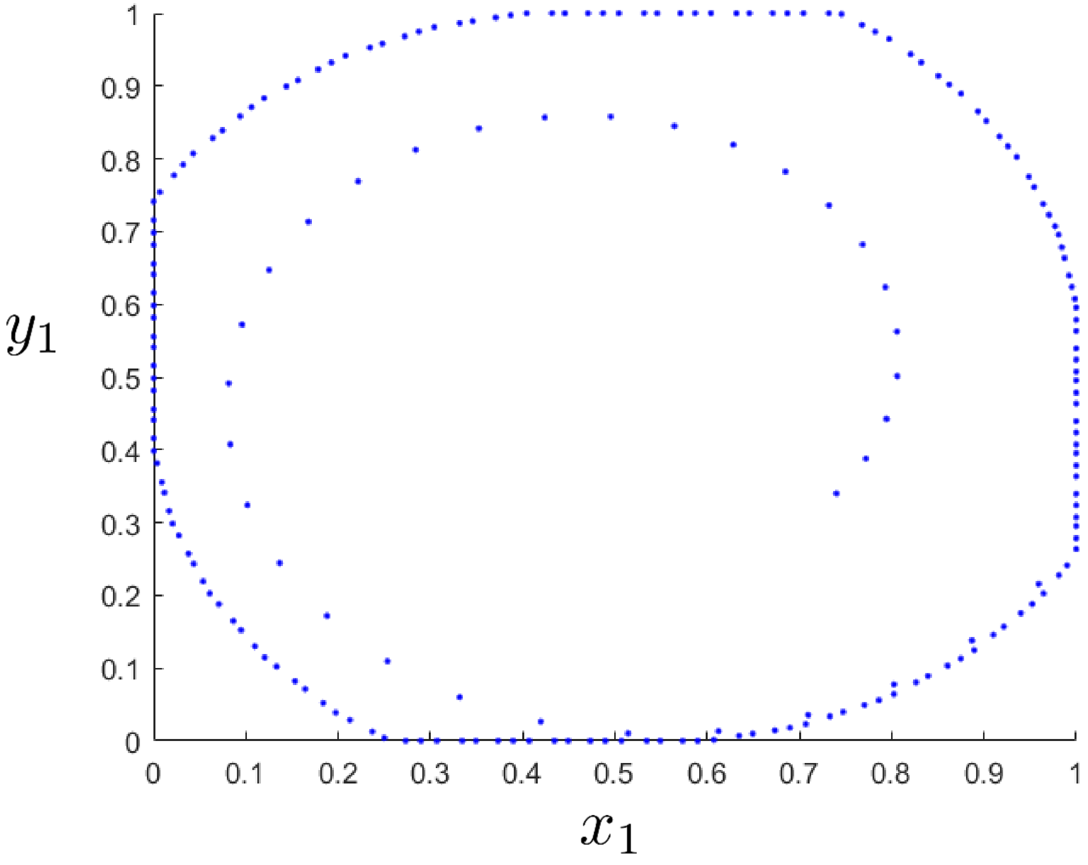

Trajectory ($\eta = 0.1$)

Proximal-Point Method (Rockafellar'76)

$$z_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z_{t+1})\right]$$

$$\Updownarrow$$

$$(x_{t+1}, y_{t+1}) = \operatorname{argmin}_{x \in \mathcal{X}, y \in \mathcal{Y}} \left\{ f(x, y) + \frac{1}{2\eta}\|x - x_t\|^2 - \frac{1}{2\eta}\|y - y_t\|^2 \right\}$$

The dynamics of Proximal-Point Method

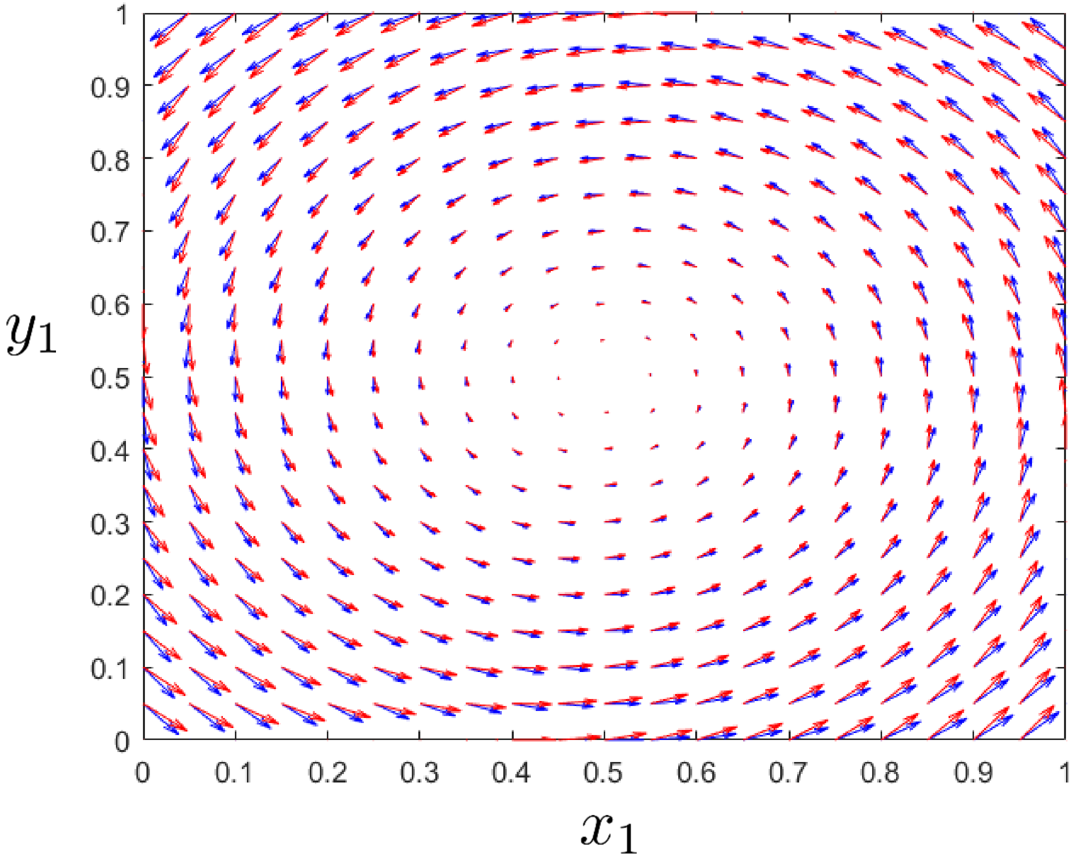

Update direction

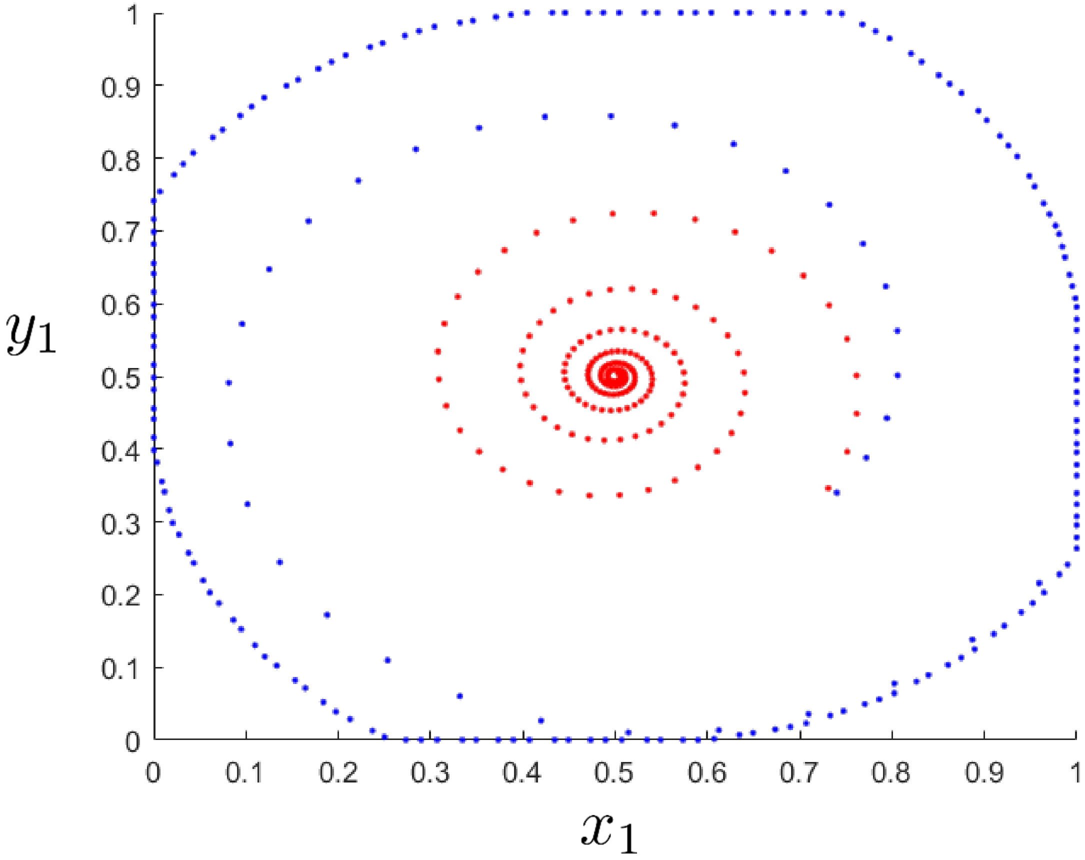

Trajectory ($\eta = 0.1$)

Extra Gradient Descent (EG) (Tseng'95)

$$z'_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z_t)\right]$$ Prediction Step

$$z_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z'_{t+1})\right]$$ Real Update

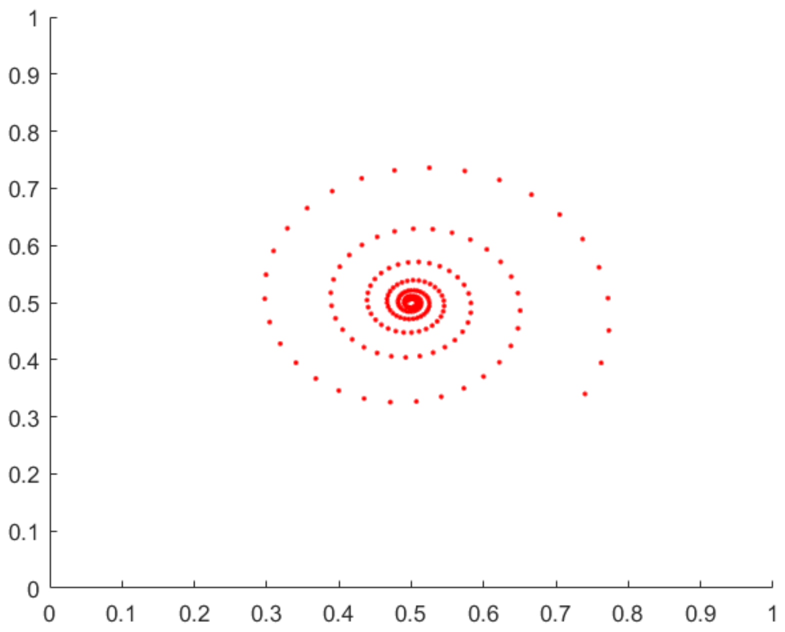

Optimistic Gradient Descent (OGD) (Popov'80)

$$z'_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z'_t)\right]$$ Prediction Step

$$z_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z'_{t+1})\right]$$ Real Update

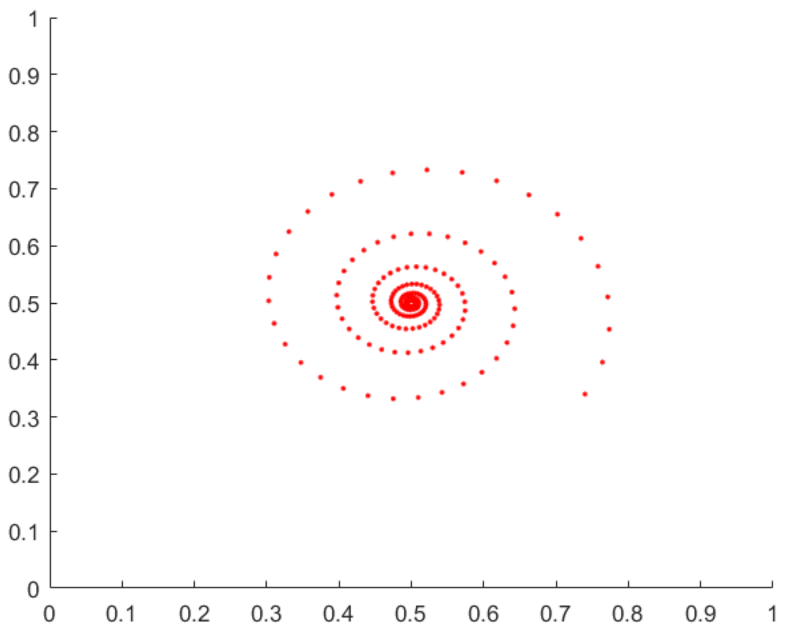

Comparing EG and OGD

EG

$$z'_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z_t)\right]$$
$$z_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z'_{t+1})\right]$$

OGD

$$z'_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z'_t)\right]$$
$$z_{t+1} = \Pi_{\mathcal{Z}}\left[z_t - \eta F(z'_{t+1})\right]$$

Algorithm	Converge under self-play?	No regret against adversary?
GD		√
EG	√	
OGD	√	√

Single gradient call per update

Run alone, it is a no-regret algorithm

Optimistic Multiplicative Weight Update

$$f(x, y) = x^\top G y$$

$$x'_{t+1,i} \propto x_{t,i} e^{-\eta (G y'_t)_i}$$ Prediction Step

$$x_{t+1,i} \propto x_{t,i} e^{-\eta (G y'_{t+1})_i}$$ Real Update

$$y'_{t+1,i} \propto y_{t,i} e^{\eta (G^\top x'_t)_i}$$ Prediction Step

$$y_{t+1,i} \propto y_{t,i} e^{\eta (G^\top x'_{t+1})_i}$$ Real Update

Related works and contributions (1/2)

Results for OGD in smooth convex-concave games

paper	function	iterate	constrained setting	rate
Popov'80		last	√	asymptotic
Rakhlin & Sridharan'13		average	√	$\Theta(1/t)$
Gidel et al'19	strongly convex/concave	last	√	$e^{-\lambda t}$
Liang & Stokes'19	bilinear	last		$e^{-\lambda t}$
Ours	SP-MS (β)	last	√	$(\lambda t)^{-\frac{1}{\beta}}$, $e^{-\lambda t}$
Ours (special case)	bilinear in polytope	last	√	$e^{-\lambda t}$
Golowich et al.'20	2^{nd}-order smoothness	last		$\Theta(1/\sqrt{t})$

We also show that in some **non-polytope bilinear** games (parabolic boundary), the rate can be $\Omega(1/t^2)$

Related works and contributions (2/2)

Results for Optimistic Multiplicative Weight in bilinear games on simplexes

paper	Iterate	step size	assumption	rate
Rakhlin & Sridharan’13	average	constant	-	$1/t$
Daskalakis & Panageas’19	last	$< (AB)^{-1/z_i^*}$	unique equilibrium	asymptotic
Ours	last	constant	unique equilibrium	$e^{-\lambda t}$
Ours	last	constant	-	asymptotic

Convergence Analysis

Convergence Analysis (1/3)

Let's focus on the unconstrained setting

$$\boxed{\begin{aligned} z'_{t+1} &= z_t - \eta M_{t+1} \\ z_{t+1} &= z_t - \eta F(z'_{t+1}) \end{aligned}}$$

$$M_{t+1} = \begin{cases} 0 & \text{GD} \\ F(z_{t+1}) & \text{Proximal} \\ F(z_t) & \text{EG} \\ F(z'_t) & \text{OGD} \end{cases}$$

Equilibrium: $z_\star = (x_\star, y_\star)$ $\quad f(x_\star, y) \le f(x_\star, y_\star) \le f(x, y_\star)$

$$\begin{aligned} \|z_t - z_\star\|^2 &= \|z_{t+1} - z_\star + \eta F(z'_{t+1})\|^2 \\ &= \|z_{t+1} - z_\star\|^2 + 2\eta(z_{t+1} - z_\star) \cdot F(z'_{t+1}) + \eta^2 \|F(z'_{t+1})\|^2 \end{aligned}$$

$$\boxed{\begin{aligned} z'_{t+1} &= z_t - \eta M_{t+1} \\ z_{t+1} &= z_t - \eta F(z'_{t+1}) \end{aligned}}$$

$$M_{t+1} = \begin{cases} 0 & \text{GD} \\ F(z_{t+1}) & \text{Proximal} \\ F(z_t) & \text{EG} \\ F(z'_t) & \text{OGD} \end{cases}$$

$\forall\ z \in \mathcal{Z},$
$(z - z_\star) \cdot F(z)$
$= (x - x_\star) \cdot \nabla f_x(x, y) - (y - y_\star) \cdot \nabla f_y(x, y))$
$\geq f(x, y) - f(x_\star, y) - f(x, y) + f(x, y_\star) \geq 0$

$$\|z_{t+1} - z_\star\|^2 \leq \|z_t - z_\star\|^2 \underbrace{- 2\eta(z_{t+1} - z_\star) \cdot F(z'_{t+1}) - \eta^2 \|F(z'_{t+1})\|^2}$$

Proximal $\quad -2\eta(z_{t+1} - z_\star) \cdot F(z_{t+1}) - \eta^2 \|F(z'_{t+1})\|^2 \leq 0$

Other cases

$$- 2\eta(z_{t+1} - z_\star) \cdot F(z'_{t+1}) - \eta^2 \|F(z'_{t+1})\|^2$$

$$= \underbrace{-2\eta(z'_{t+1} - z_\star) \cdot F(z'_{t+1})}_{\leq 0} + 2\eta \underbrace{(z'_{t+1} - z_{t+1})}_{= \eta(F(z'_{t+1}) - M_{t+1})} \cdot F(z'_{t+1}) - \eta^2 \|F(z'_{t+1})\|^2$$

$$\leq 2\eta^2 \left(F(z'_{t+1}) - M_{t+1}\right) \cdot F(z'_{t+1}) - \eta^2 \|F(z'_{t+1})\|^2$$

$$= \eta^2 \|F(z'_{t+1}) - M_{t+1}\|^2 - \eta^2 \|M_{t+1}\|^2$$

For GD, this ≥ 0

$$\|z_{t+1} - z_\star\|^2 \le \|z_t - z_\star\|^2 + \eta^2\left(\|F(z'_{t+1}) - M_{t+1}\|^2 - \|M_{t+1}\|^2\right)$$

EG

$$z'_{t+1} = z_t - \eta F(z_t)$$
$$z_{t+1} = z_t - \eta F(z'_{t+1})$$

$$\begin{aligned}
&\|F(z'_{t+1}) - M_{t+1}\|^2 - \|M_{t+1}\|^2 \\
&= \|F(z'_{t+1}) - F(z_t)\|^2 - \|F(z_t)\|^2 \\
&\le L^2\|z'_{t+1} - z_t\|^2 - \|F(z_t)\|^2 \\
&\le \eta^2 L^2\|F(z_t)\|^2 - \|F(z_t)\|^2 \\
&\le -\frac{1}{2}\|F(z_t)\|^2
\end{aligned}$$

$\eta \le \frac{1}{\sqrt{2}L}$

OGD

$$z'_{t+1} = z_t - \eta F(z'_t)$$
$$z_{t+1} = z_t - \eta F(z'_{t+1})$$

$$\begin{aligned}
&\|F(z'_{t+1}) - M_{t+1}\|^2 - \|M_{t+1}\|^2 \\
&= \|F(z'_{t+1}) - F(z'_t)\|^2 - \|F(z'_t)\|^2 \\
&\le L^2\|z'_{t+1} - z'_t\|^2 - \|F(z'_t)\|^2 \\
&\le \eta^2 L^2\|2F(z'_t) - F(z'_{t-1})\|^2 - \|F(z'_t)\|^2 \\
&\le 8\eta^2 L^2\|F(z'_t)\|^2 + 2\eta^2 L^2\|F(z'_{t-1})\|^2 - \|F(z'_t)\|^2 \\
&\le -\frac{1}{2}\|F(z'_t)\|^2 + \frac{1}{8}\|F(z'_{t-1})\|^2
\end{aligned}$$

$\eta \le \frac{1}{4L}$

Convergence Analysis

EG $$\|z_{t+1} - z_\star\|^2 \le \|z_t - z_\star\|^2 - \frac{1}{2}\eta^2\|F(z_t)\|^2$$

OGD $$\|z_{t+1} - z_\star\|^2 \le \|z_t - z_\star\|^2 - \frac{1}{2}\eta^2\|F(z'_t)\|^2 + \frac{1}{8}\eta^2\|F(z'_{t-1})\|^2$$

$$\Rightarrow \underbrace{\|z_{t+1} - z_\star\|^2 + \frac{1}{8}\eta^2\|F(z'_t)\|^2}_{\Theta_{t+1}} \le \underbrace{\|z_t - z_\star\|^2 + \frac{1}{8}\eta^2\|F(z'_{t-1})\|^2}_{\Theta_t} - \frac{3}{8}\eta^2\|F(z'_t)\|^2$$

Formal Result for the Constrained Setting

Define $\Theta_t = \|z_t - \Pi_{\mathcal{Z}_\star}(z_t)\|^2 + \frac{1}{16}\|z_t - z_t'\|^2$

$$\Theta_{t+1} \le \Theta_t - c\Theta_t \Rightarrow \Theta_t = \mathcal{O}\left(e^{-ct}\right)$$

$$\Theta_{t+1} \le \Theta_t - c\Theta_t^2 \Rightarrow \Theta_t = \mathcal{O}\left(1/t\right)$$

$$\Theta_{t+1} \le \Theta_t - c\Theta_t^3 \Rightarrow \Theta_t = \mathcal{O}\left(1/\sqrt{t}\right)$$

Then OGD with $\eta \le \frac{1}{8L}$ guarantees

$$\Theta_{t+1} \le \Theta_t - \frac{1}{2}\eta^2 \max_{z' \in \mathcal{Z}} \frac{[(z_{t+1} - z')^\top F(z_{t+1})]_+^2}{\|z_{t+1} - z'\|^2}$$

Convergence rate?

Relate this part to Θ_t (or Θ_{t+1})

Convergence Rate

Saddle-point Metric-subregularity (**SP-MS**) Condition:

$$\exists C > 0, \beta \geq 0 \qquad \forall z \in \mathcal{Z} \qquad \max_{z' \in \mathcal{Z}} \frac{(z - z')^\top F(z)}{\|z - z'\|} \geq C \|z - \Pi_{\mathcal{Z}_\star}(z)\|^{1+\beta}$$

$\beta = 0$: $\quad \Theta_{t+1} \leq e^{-\lambda t}$

bilinear games in polytopes
strongly-convex-strongly-concave games

$\beta > 0$: $\quad \Theta_{t+1} \leq (\lambda t)^{-1/\beta}$

Result for Optimistic Multiplicative Weight Update

Define $\Theta_t = \mathrm{KL}(z_\star, z_t) + \frac{1}{16}\mathrm{KL}(z'_t, z_t)$

Then OMWU with $\eta \le \frac{1}{8L}$ guarantees

$$\Theta_{t+1} \le \Theta_t - \underbrace{\frac{15}{16}\left(\mathrm{KL}(z_{t+1}, z'_{t+1}) + \mathrm{KL}(z'_{t+1}, z_t)\right)}_{\zeta_t}$$

Phase 1: $\zeta_t \ge \eta^2 C_1 \mathrm{KL}(z_\star, z_{t+1})^2 \quad \Rightarrow \Theta_t \le 1/(\lambda_1 t)$

Phase 2 (when $\|z_t - z_\star\| \le \eta\xi$) : $\zeta_t \ge \eta^2 C_2 \mathrm{KL}(z_\star, z_{t+1}) \Rightarrow \Theta_t \le e^{-\lambda_2 t}$

Extension to Markov Games

Last-iterate Convergence of Decentralized Optimistic Gradient Descent/Ascent in Infinite-horizon Competitive Markov Games

Markov Games

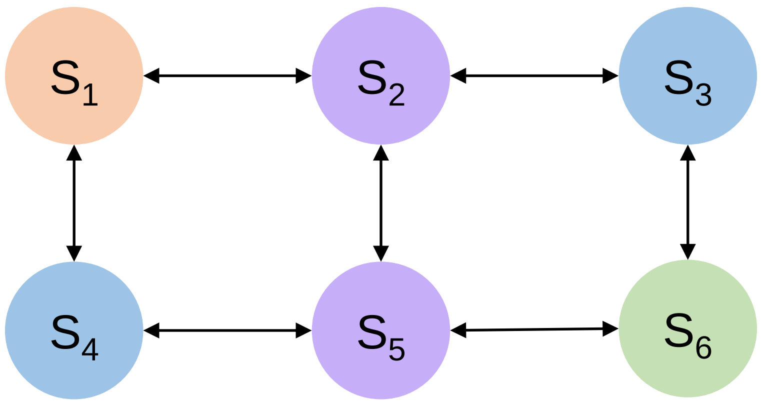

State: s

Player 1's action a, Player 2's action b

Loss/payoff function: $\sigma(s, a, b)$

Transition function: $s' \sim p(\cdot | s, a, b)$

Markov Games

Existing provably convergent algorithms usually belong to one of the following:

- Centralized algorithm
- Two-timescale algorithm (Golowich et al'20)
- Algorithm that only handles multi-stage games (Perolat et al.'18)

Is there an algorithm in general Markov games that

- Converges to Nash equilibria when both players use the same algorithm
- Converges to the best response when the other player converges to a fixed policy

A rational and convergent algorithm (Bowling & Veloso'01)

Markov Games

We propose an actor-critic algorithm where

The policy learner (actor) runs OGD on each state
The value learner (critic) slows down the variation of the game matrix

critic

$$Q_t(s,a,b) \triangleq \sigma(s,a,b) + \gamma \mathbb{E}_{s' \sim p(\cdot|s,a,b)} \left[V_{t-1}(s') \right]$$

$$V_t(s) = (1 - \alpha_t) V_{t-1}(s) + \alpha_t \sum_{a,b} x'_t(a|s) y'_t(b|s) Q_t(s,a,b)$$

actor

$$x'_{t+1}(\cdot|s) = \Pi_{\Delta_A} \left(x_t(\cdot|s) - \eta \sum_{b \in B} Q_t(s,\cdot,b) y'_t(b|s) \right)$$

$$x_{t+1}(\cdot|s) = \Pi_{\Delta_A} \left(x_t(\cdot|s) - \eta \sum_{b \in B} Q_{t+1}(s,\cdot,b) y'_{t+1}(b|s) \right)$$

Theorem for Markov Games

$$\sqrt{\frac{1}{S}\sum_{s}\|z_t(\cdot|s) - \Pi_{\mathcal{Z}^*(\cdot|s)}(z_t(\cdot|s))\|^2} = \mathcal{O}\left(\frac{S}{\eta^2 C^2 (1-\gamma)^2}\sqrt{\frac{1}{t}}\right)$$