Anticipating the Future for Better Performance: Optimistic Gradient Methods for Learning in Games

Yu-Guan Hsieh

Université Grenoble Alpes, France

July 5th, 2022

Optimization as Minimization

$$\min_{x \in \mathcal{X}} \ell(x)$$

- Inverse problem (MRI, CT, . . .)
- Power system management
- Machine learning

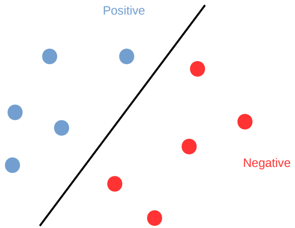

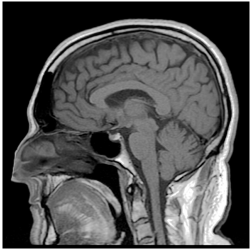

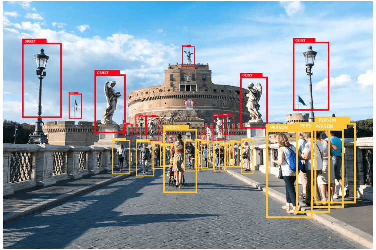

Optimization Beyond Minimization

Learning in an environment that is reactive

Probably due to the presence of multiple agents $\longrightarrow$ game theory

- Explicit: games, interaction of robots, autonomous vehicles
- Implicit: robust optimization, generative adversarial networks (GANs)

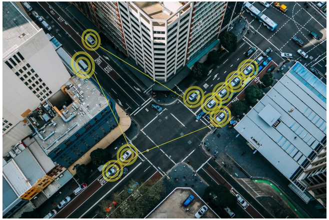

Motivating Example: Generative Adversarial Networks (GANs)

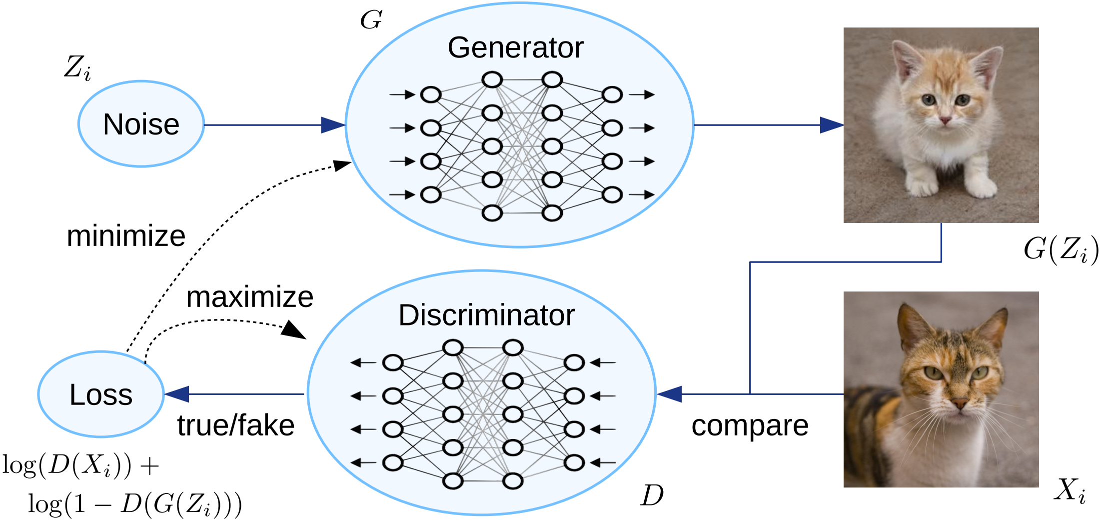

Motivating Example: Generative Adversarial Networks (GANs)

- GauGAN (Nvidia):

Motivating Example: Generative Adversarial Networks (GANs)

- GauGAN (Nvidia):  

- Distribution matching: Domain adaptation [Tzeng et al. 2017], Imitation learning [Ho and Ermon 2016]

Motivating Example: Generative Adversarial Networks (GANs)

- GauGAN (Nvidia):


- Distribution matching: Domain adaptation [Tzeng et al. 2017], Imitation learning [Ho and Ermon 2016]
- GANs are hard to train due to the interaction of multiple agents

Motivating Example: Generative Adversarial Networks (GANs)

- GauGAN (Nvidia):

- Distribution matching: Domain adaptation [Tzeng et al. 2017], Imitation learning [Ho and Ermon 2016]
- GANs are hard to train due to the interaction of multiple agents
- More recently:
  - DALL-E (Open AI), Imagen (Google): Diffusion model
  - Parti (Google): Autoregressive model + GAN

A small cactus wearing a straw hat and neon sunglasses in the Sahara desert.

Motivating Example: Multi-Agent Reinforcement Learning

- Self-interest agents coexist in a shared environment
- Collaboration, coordination, competition, etc.

How to learn a good policy that performs well in a multi-agent environment?

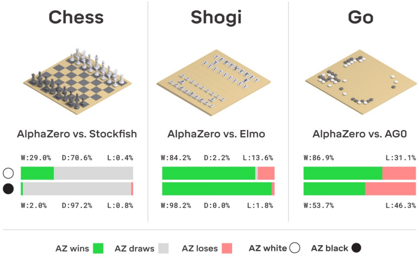

Outline

Disclaimer

- This presentation is about intuitions and theories
- We focus on normal-form monotone games

In this talk, we will not cover

- Experiments on real-world applications
- General-sum games
- Non-monotone landscapes
- Extensive-form games

Outline

Normal-Form Finite Games

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$
- Each player has a finite set of actions $\mathcal{A}^i$ and a payoff function $u^i \colon \prod_{i \in \mathcal{N}} \mathcal{A}^i \to \mathbb{R}$

Normal-Form Finite Games

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$
- Each player has a finite set of actions $\mathcal{A}^i$ and a payoff function $u^i \colon \prod_{i \in \mathcal{N}} \mathcal{A}^i \to \mathbb{R}$
- Examples: Rock-Paper-Scissors

	rock	paper	scissors
rock	0,0	-1,1	1,-1
paper	1,-1	0,0	-1,1
scissors	-1,1	1,-1	0,0

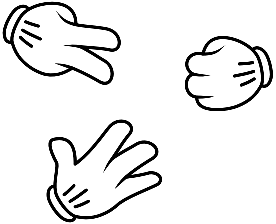

Normal-Form Finite Games

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$
- Each player has a finite set of actions $\mathcal{A}^i$ and a payoff function $u^i \colon \prod_{i \in \mathcal{N}} \mathcal{A}^i \to \mathbb{R}$
- Examples: Rock-Paper-Scissors

	rock	paper	scissors
rock	0,0	-1,1	1,-1
paper	1,-1	0,0	-1,1
scissors	-1,1	1,-1	0,0

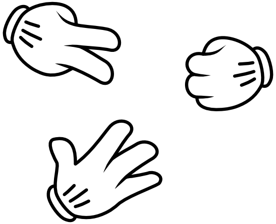

- A Nash equilibrium $\mathbf{a}_\star = (a^i, \mathbf{a}_\star^{-i})$ is a joint action profile from which no player has incentive to deviate unilaterally, i.e., for all $i \in \mathcal{N}$, $a^i \in \mathcal{A}^i$, $u^i(a_\star^i, \mathbf{a}_\star^{-i}) \geq u^i(a^i, \mathbf{a}_\star^{-i})$

Normal-Form Finite Games

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$
- Each player has a finite set of actions $\mathcal{A}^i$ and a payoff function $u^i \colon \prod_{i \in \mathcal{N}} \mathcal{A}^i \to \mathbb{R}$
- Examples: Rock-Paper-Scissors

	rock	paper	scissors
rock	0,0	-1,1	1,-1
paper	1,1	0,0	-1,1
scissors	-1,1	1,-1	0,0

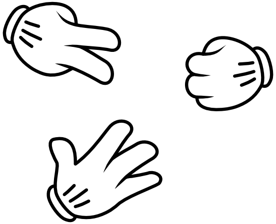

- A Nash equilibrium $\mathbf{a}_\star = (a^i, \mathbf{a}_\star^{-i})$ is a joint action profile from which no player has incentive to deviate unilaterally, i.e., for all $i \in \mathcal{N}$, $a^i \in \mathcal{A}^i$, $u^i(a_\star^i, \mathbf{a}_\star^{-i}) \geq u^i(a^i, \mathbf{a}_\star^{-i})$

Continuous Games

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$
- Each player has a convex closed action set $\mathcal{X}^i \subseteq \mathbb{R}^{d^i}$ and a loss function $\ell^i \colon \prod_{i \in \mathcal{N}} \mathcal{X}^i \to \mathbb{R}$

Continuous Games

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$
- Each player has a convex closed action set $\mathcal{X}^i \subseteq \mathbb{R}^{d^i}$ and a loss function $\ell^i \colon \prod_{i \in \mathcal{N}} \mathcal{X}^i \to \mathbb{R}$
- Examples: Auctions, Cournot Duopoly, GANs

Continuous Games

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$
- Each player has a convex closed action set $\mathcal{X}^i \subseteq \mathbb{R}^{d^i}$ and a loss function $\ell^i \colon \prod_{i \in \mathcal{N}} \mathcal{X}^i \to \mathbb{R}$
- Examples: Auctions, Cournot Duopoly, GANs
- A Nash equilibrium $\mathbf{x}_\star = (x^i_\star, \mathbf{x}^{-i}_\star)$ is a joint action profile from which no player has incentive to deviate unilaterally, i.e., for all $i \in \mathcal{N}$, $x^i \in \mathcal{X}^i$, $\ell^i(x^i_\star, \mathbf{x}^{-i}_\star) \le \ell^i(x^i, \mathbf{x}^{-i}_\star)$

Mixed Extensions of Finite Games as Continuous Games

- A mixed strategy $x^i \in \Delta(\mathcal{A}^i) \subset \mathbb{R}^{\mathrm{card}(\mathcal{A}^i)}$ for player i is a probability distribution over $\mathcal{A}^i$ ($\Delta(\mathcal{A}^i)$ is the probability simplex on $\mathcal{A}^i$)

Mixed Extensions of Finite Games as Continuous Games

- A mixed strategy $x^i \in \Delta(\mathcal{A}^i) \subset \mathbb{R}^{\mathrm{card}(\mathcal{A}^i)}$ for player i is a probability distribution over $\mathcal{A}^i$ ($\Delta(\mathcal{A}^i)$ is the probability simplex on $\mathcal{A}^i$)
- For a joint mixed strategy $\mathbf{x}$, we can associate the payoff $u^i(\mathbf{x}) = \mathbb{E}_{\mathbf{a}\sim\mathbf{x}}\, u^i(\mathbf{a})$

Mixed Extensions of Finite Games as Continuous Games

- A mixed strategy $x^i \in \Delta(\mathcal{A}^i) \subset \mathbb{R}^{\mathrm{card}(\mathcal{A}^i)}$ for player i is a probability distribution over $\mathcal{A}^i$ ($\Delta(\mathcal{A}^i)$ is the probability simplex on $\mathcal{A}^i$)
- For a joint mixed strategy $\mathbf{x}$, we can associate the payoff $u^i(\mathbf{x}) = \mathbb{E}_{\mathbf{a}\sim\mathbf{x}}\, u^i(\mathbf{a})$
- Mixed strategy Nash equilibrium always exists

Mixed Extensions of Finite Games as Continuous Games

- A mixed strategy $x^i \in \Delta(\mathcal{A}^i) \subset \mathbb{R}^{\mathrm{card}(\mathcal{A}^i)}$ for player i is a probability distribution over $\mathcal{A}^i$ ($\Delta(\mathcal{A}^i)$ is the probability simplex on $\mathcal{A}^i$)
- For a joint mixed strategy $\mathbf{x}$, we can associate the payoff $u^i(\mathbf{x}) = \mathbb{E}_{\mathbf{a}\sim\mathbf{x}}\, u^i(\mathbf{a})$
- Mixed strategy Nash equilibrium always exists
- Take $\ell^i = -u^i$ to get the previous form

Notations and Assumptions

- Joint action $\mathbf{x} = (x^i)_{i\in\mathcal{N}} = (x^i, \mathbf{x}^{-i})$
- Important: we assume $\ell^i(\cdot, \mathbf{x}^{-i})$ to be convex

In previous slide, $\ell^i(\mathbf{x}) = -\mathbb{E}_{\mathbf{a}\sim\mathbf{x}}\, u^i(\mathbf{a}) = -\sum_{\mathbf{a}\in\prod_{i\in\mathcal{N}}\mathcal{A}^i}\left(\prod_{i\in\mathcal{N}} x^i(a^i)\right)u^i(\mathbf{a})$ is linear in x^i

Notations and Assumptions

- Joint action $\mathbf{x} = (x^i)_{i \in \mathcal{N}} = (x^i, \mathbf{x}^{-i})$
- Important: we assume $\ell^i(\cdot, \mathbf{x}^{-i})$ to be convex

  In previous slide, $\ell^i(\mathbf{x}) = -\mathbb{E}_{\mathbf{a} \sim \mathbf{x}}\, u^i(\mathbf{a}) = -\sum_{\mathbf{a} \in \prod_{i \in \mathcal{N}} \mathcal{A}^i} \left(\prod_{i \in \mathcal{N}} x^i(a^i) \right) u^i(\mathbf{a})$ is linear in x^i
- For simplicity: we consider unconstrained setup, i.e., $\mathcal{X}^i = \mathbb{R}^{d^i}$

Notations and Assumptions

- Joint action $\mathbf{x} = (x^i)_{i \in \mathcal{N}} = (x^i, \mathbf{x}^{-i})$
- Important: we assume $\ell^i(\cdot, \mathbf{x}^{-i})$ to be convex

  In previous slide, $\ell^i(\mathbf{x}) = -\mathbb{E}_{\mathbf{a} \sim \mathbf{x}}\, u^i(\mathbf{a}) = -\sum_{\mathbf{a} \in \prod_{i \in \mathcal{N}} \mathcal{A}^i} \left(\prod_{i \in \mathcal{N}} x^i(a^i) \right) u^i(\mathbf{a})$ is linear in x^i
- For simplicity: we consider unconstrained setup, i.e., $\mathcal{X}^i = \mathbb{R}^{d^i}$
- Joint vector field: $\mathbf{V}(\mathbf{X}) = (\nabla_i\, \ell^i(\mathbf{X}))_{i \in \mathcal{N}}$
- Nash equilibria: $\mathcal{X}_\star = \{\mathbf{x}_\star : \mathbf{V}(\mathbf{x}_\star) = 0\}$

Learning in Continuous Games

At each round $t = 1, 2, \ldots$, each player $i \in \mathcal{N}$

- Plays an action $x_t^i \in \mathcal{X}^i$
- Suffers loss $\ell^i(\mathbf{x}_t)$ and receives feedback g_t^i

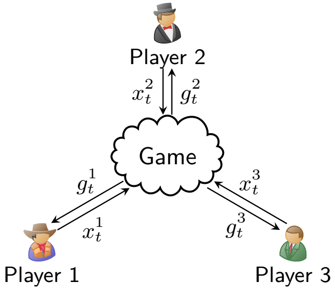

Learning in Continuous Games

At each round $t = 1, 2, \ldots$, each player $i \in \mathcal{N}$

- Plays an action $x_t^i \in \mathcal{X}^i$
- Suffers loss $\ell^i(\mathbf{x}_t)$ and receives feedback g_t^i

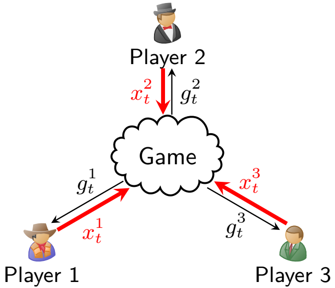

Learning in Continuous Games

At each round $t = 1, 2, \ldots$, each player $i \in \mathcal{N}$

- Plays an action $x_t^i \in \mathcal{X}^i$
- Suffers loss $\ell^i(\mathbf{x}_t)$ and receives feedback g_t^i

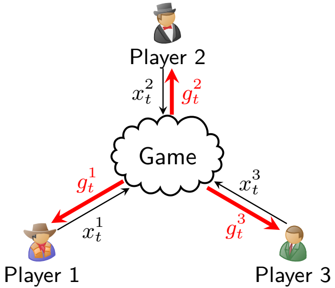

Learning in Continuous Games

At each round $t = 1, 2, \ldots$, each player $i \in \mathcal{N}$

- Plays an action $x_t^i \in \mathcal{X}^i$
- Suffers loss $\ell^i(\mathbf{x}_t)$ and receives feedback g_t^i

- First-order feedback so that $g_t^i \approx \nabla_i \ell^i(\mathbf{x}_t)$

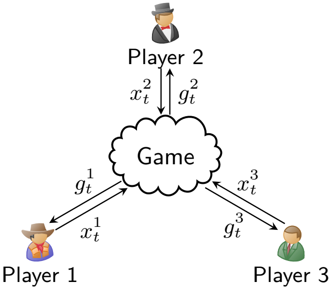

Learning in Continuous Games

At each round $t = 1, 2, \ldots$, each player $i \in \mathcal{N}$

- Plays an action $x_t^i \in \mathcal{X}^i$
- Suffers loss $\ell^i(\mathbf{x}_t)$ and receives feedback g_t^i

- First-order feedback so that $g_t^i \approx \nabla_i \ell^i(\mathbf{x}_t)$
- Regret of player i with respect to $p^i \in \mathcal{X}^i$ is

$$\mathrm{Reg}_T^i(p^i) = \sum_{t=1}^{T} \Big(\underbrace{\ell^i(x_t^i, \mathbf{x}_t^{-i}) - \ell^i(p^i, \mathbf{x}_t^{-i})}_{\text{cost of not playing } p^i \text{ in round } t} \Big)$$

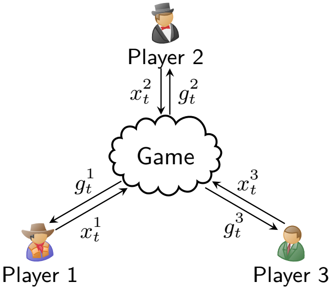

Learning in Continuous Games

At each round $t = 1, 2, \dots$, each player $i \in \mathcal{N}$

- Plays an action $x_t^i \in \mathcal{X}^i$
- Suffers loss $\ell^i(\mathbf{x}_t)$ and receives feedback g_t^i

- First-order feedback so that $g_t^i \approx \nabla_i \ell^i(\mathbf{x}_t)$
- Regret of player i with respect to $p^i \in \mathcal{X}^i$ is

$$\mathrm{Reg}_T^i(p^i) = \sum_{t=1}^{T} \Big(\underbrace{\ell^i(x_t^i, \mathbf{x}_t^{-i}) - \ell^i(p^i, \mathbf{x}_t^{-i})}_{\text{cost of not playing } p^i \text{ in round } t} \Big)$$

- No-regret if $\mathrm{Reg}_T^i(p^i) = o(T)$

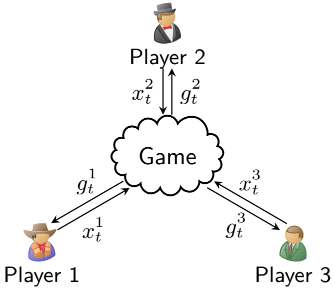

Learning in Continuous Games

At each round $t = 1, 2, \ldots$, each player $i \in \mathcal{N}$

- Plays an action $x_t^i \in \mathcal{X}^i$
- Suffers loss $\ell^i(\mathbf{x}_t)$ and receives feedback g_t^i

- First-order feedback so that $g_t^i \approx \nabla_i \ell^i(\mathbf{x}_t)$
- Regret of player i with respect to $p^i \in \mathcal{X}^i$ is

$$\mathrm{Reg}_T^i(p^i) = \sum_{t=1}^{T} \Big(\underbrace{\ell^i(x_t^i, \mathbf{x}_t^{-i}) - \ell^i(p^i, \mathbf{x}_t^{-i})}_{\text{cost of not playing } p^i \text{ in round } t} \Big)$$

- No-regret if $\mathrm{Reg}_T^i(p^i) = o(T)$
- Players can be adversarial or optimizing their own benefit

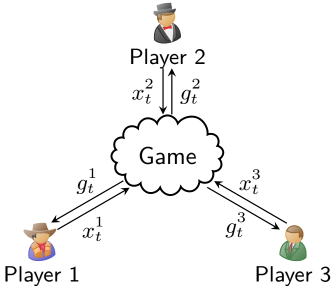

Online Learning

At each round $t = 1, 2, \ldots$, the learner

- Plays an action $x_t \in \mathcal{X}$
- Suffers loss $\ell_t(x_t)$ and receives feedback g_t [$\ell_t \colon \mathcal{X} \to \mathbb{R}$]

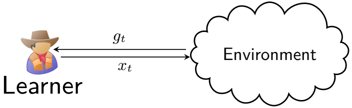

Online Learning

At each round $t = 1, 2, \ldots$, the learner

- Plays an action $x_t \in \mathcal{X}$
- Suffers loss $\ell_t(x_t)$ and receives feedback g_t [$\ell_t \colon \mathcal{X} \to \mathbb{R}$]

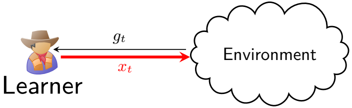

Online Learning

At each round $t = 1, 2, \ldots$, the learner

- Plays an action $x_t \in \mathcal{X}$
- Suffers loss $\ell_t(x_t)$ and receives feedback g_t $[\ell_t \colon \mathcal{X} \to \mathbb{R}]$

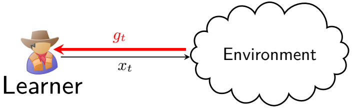

Online Learning

At each round $t = 1, 2, \ldots$, the learner

- Plays an action $x_t \in \mathcal{X}$
- Suffers loss $\ell_t(x_t)$ and receives feedback g_t [$\ell_t \colon \mathcal{X} \to \mathbb{R}$]

- First-order feedback $g_t \approx \nabla \ell_t(x_t)$

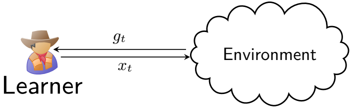

Online Learning

At each round $t = 1, 2, \ldots$, the learner

- Plays an action $x_t \in \mathcal{X}$
- Suffers loss $\ell_t(x_t)$ and receives feedback g_t $[\ell_t\colon \mathcal{X} \to \mathbb{R}]$

- First-order feedback $g_t \approx \nabla \ell_t(x_t)$
- Regret of the learner with respect to $p \in \mathcal{X}$ is

$$\mathrm{Reg}_T(p^i) = \sum_{t=1}^{T} \Big(\underbrace{\ell_t(x_t) - \ell_t(p)}_{\text{cost of not playing } p \text{ in round } t} \Big)$$

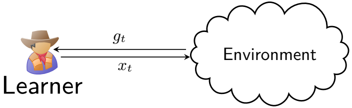

Online Learning

At each round $t = 1, 2, \ldots,$ the learner

- Plays an action $x_t \in \mathcal{X}$
- Suffers loss $\ell_t(x_t)$ and receives feedback g_t [$\ell_t : \mathcal{X} \to \mathbb{R}$]

- First-order feedback $g_t \approx \nabla \ell_t(x_t)$
- Regret of the learner with respect to $p \in \mathcal{X}$ is

$$\mathrm{Reg}_T(p^i) = \sum_{t=1}^{T} \Big(\underbrace{\ell_t(x_t) - \ell_t(p)}_{\text{cost of not playing } p \text{ in round } t} \Big)$$

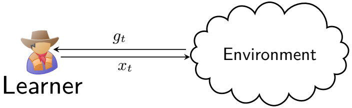

- The environment can be adversarial, stochastic, multi-agent $\ell_t = \ell^i(\cdot, \mathbf{x}_t^{-i})$, etc.

Outline

The Failure of Vanilla Gradient Method in Bilinear Games

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

Unique Nash equilibrium: $(0, 0)$

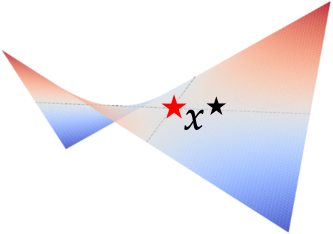

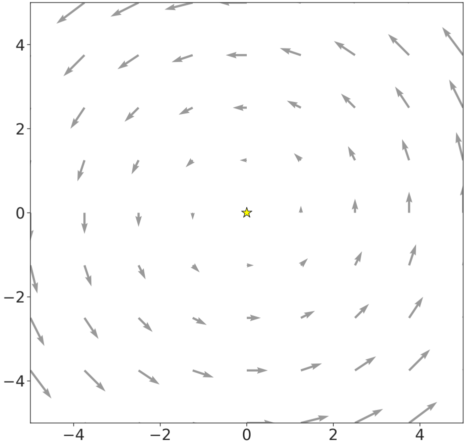

The Failure of Vanilla Gradient Method in Bilinear Games

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

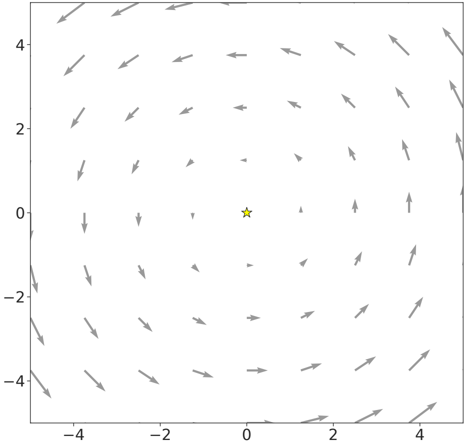

The Failure of Vanilla Gradient Method in Bilinear Games

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \, \ell^1(\mathbf{x}), \nabla_\phi \, \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- Gradient descent

$$\begin{aligned} \theta_{t+1} &= \theta_t - \eta_t \, \nabla_\theta \, \ell^1(\theta_t, \phi_t) \\ \phi_{t+1} &= \phi_t - \eta_t \, \nabla_\phi \, \ell^2(\theta_t, \phi_t) \end{aligned}$$

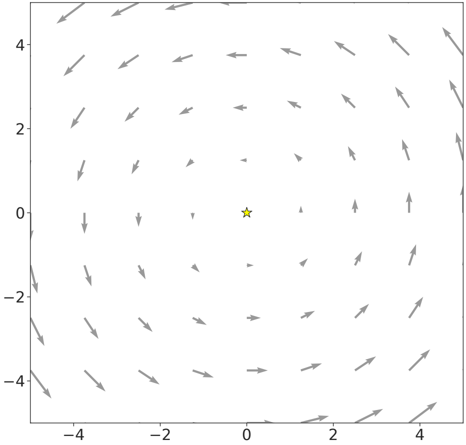

The Failure of Vanilla Gradient Method in Bilinear Games

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- Gradient descent

$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_t)$$

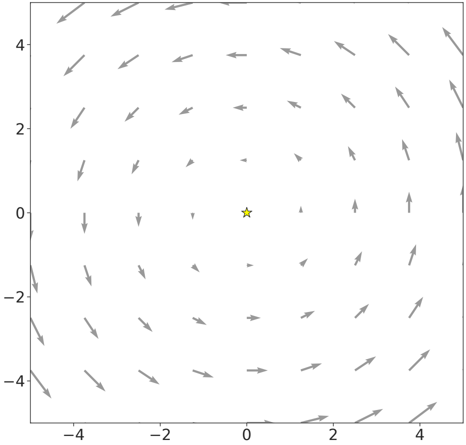

The Failure of Vanilla Gradient Method in Bilinear Games

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- Gradient descent

$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_t)$$

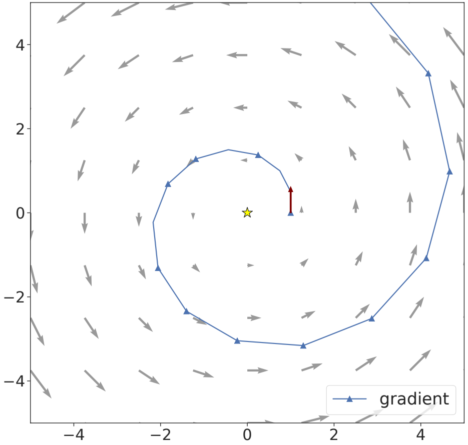

Extra-Gradient to the Rescue

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- Extra-gradient [Korpelevich 1976]

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_t)$$
$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

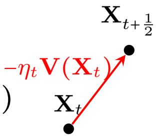

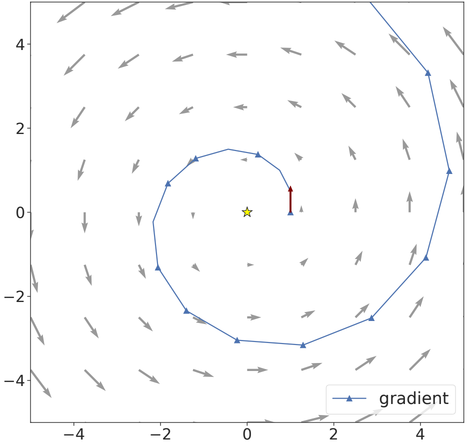

Extra-Gradient to the Rescue

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta,\phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- Extra-gradient [Korpelevich 1976]

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_t)$$
$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

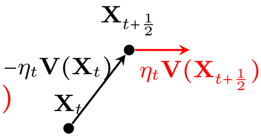

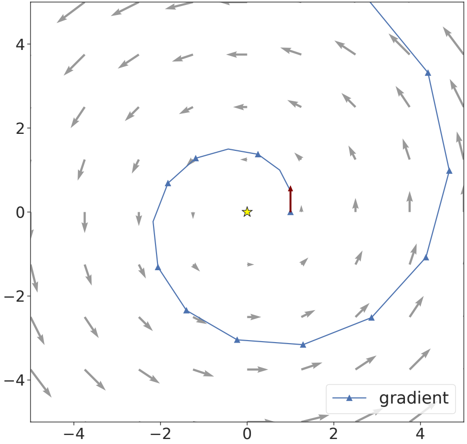

Extra-Gradient to the Rescue

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- Extra-gradient [Korpelevich 1976]

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_t)$$
$$\mathbf{X}_{t+1} = \mathbf{X}_t - {\color{red}\eta_t \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})}$$

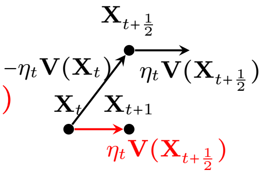

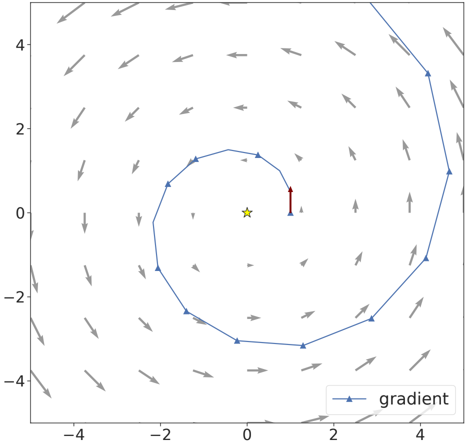

Extra-Gradient to the Rescue

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- Extra-gradient [Korpelevich 1976]

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_t)$$
$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

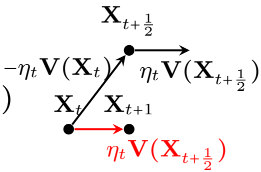

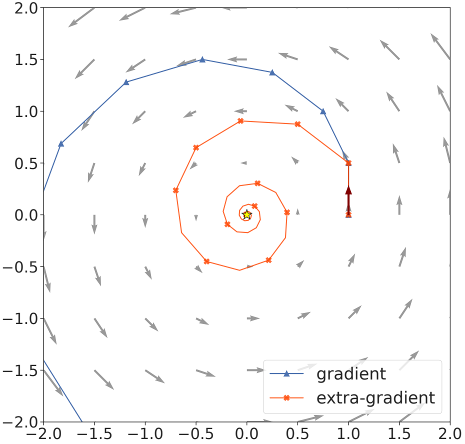

Extra-Gradient to the Rescue

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Jo

Extra-gradient is not an online algorithm

- Extra-gradient [Korpelevich 1976]

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_t)$$
$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

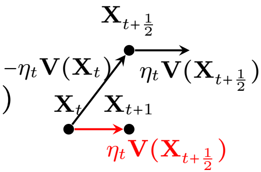

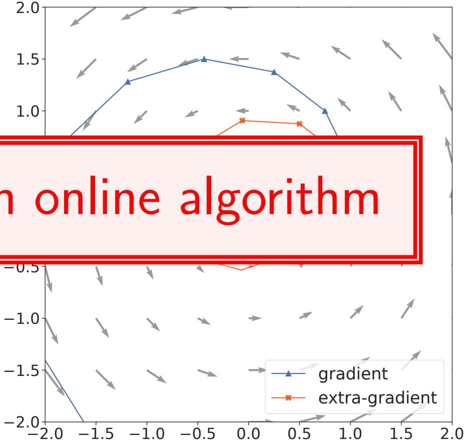

Optimistic Gradient: Online Variant of Extra-Gradient

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- OG – Optimistic gradient [Popov 1980]

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \boxed{\mathbf{V}(\mathbf{X}_{t-\frac{1}{2}})}$$

$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1} \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

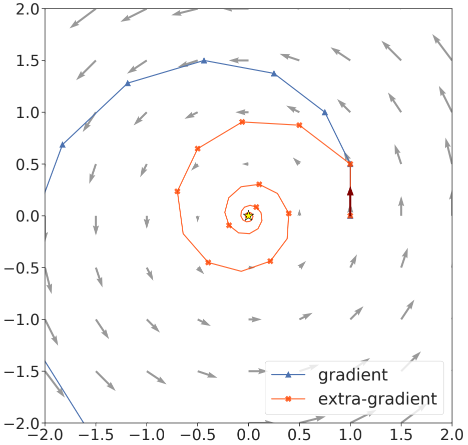

Optimistic Gradient: Online Variant of Extra-Gradient

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- OG – Optimistic gradient [Popov 1980]

$$X^i_{t+\frac{1}{2}} = X^i_{t-\frac{1}{2}} - 2\eta^i_t g^i_t + \eta^i_{t-1} g^i_{t-1} \quad (x^i_t = X^i_{t+\frac{1}{2}})$$

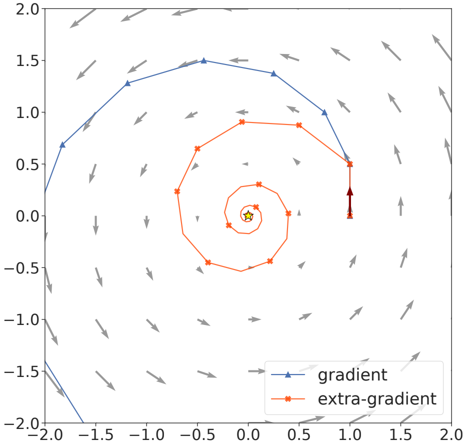

Optimistic Gradient: Online Variant of Extra-Gradient

- Two-player planar bilinear zero-sum game

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Joint vector field

$$\mathbf{V}(\mathbf{x}) \coloneqq (\nabla_\theta \ell^1(\mathbf{x}), \nabla_\phi \ell^2(\mathbf{x})) = (\phi, -\theta)$$

- OG – Optimistic gradient [Popov 1980]

$$X^i_{t+\frac{1}{2}} = X^i_{t-\frac{1}{2}} - 2\eta^i_t g^i_t + \eta^i_{t-1} g^i_{t-1} \quad \left(x^i_t = X^i_{t+\frac{1}{2}} \right)$$

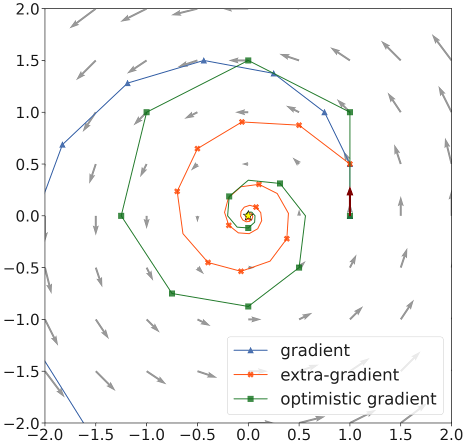

Intuitions Behind Extra-Gradient / Optimistic Gradient

- Better discretization of the continuous flow [Lu 21]

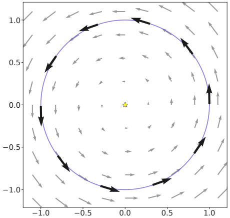

Intuitions Behind Extra-Gradient / Optimistic Gradient

- Better discretization of the continuous flow [Lu 21]
- Look into the future, anticipating the landscape

Intuitions Behind Extra-Gradient / Optimistic Gradient

- Better discretization of the continuous flow [Lu 21]
- Look into the future, anticipating the landscape
  - Learning with recency bias [Rakhlin and Sridharan 13, Syrgkanis et al. 15]
    Compare with methods that 'knows the future'

Intuitions Behind Extra-Gradient / Optimistic Gradient

- Better discretization of the continuous flow [Lu 21]
- Look into the future, anticipating the landscape
  - Learning with recency bias [Rakhlin and Sridharan 13, Syrgkanis et al. 15]
    Compare with methods that 'knows the future'
  - Approximation of proximal point (PP) methods [Mokhtari et al. 20]

Intuitions Behind Extra-Gradient / Optimistic Gradient

- Better discretization of the continuous flow [Lu 21]
- Look into the future, anticipating the landscape
  - Learning with recency bias [Rakhlin and Sridharan 13, Syrgkanis et al. 15]
    Compare with methods that 'knows the future'
  - Approximation of proximal point (PP) methods [Mokhtari et al. 20]
- Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

Intuitions Behind Extra-Gradient / Optimistic Gradient

$$\boxed{\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}_t, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1} \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})}$$

Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

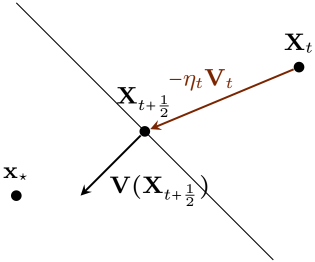

Intuitions Behind Extra-Gradient / Optimistic Gradient

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}_t, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1} \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

- Consider the hyperplan

$$\mathcal{H} \coloneqq \{\mathbf{x} : \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x} \rangle = 0\}$$

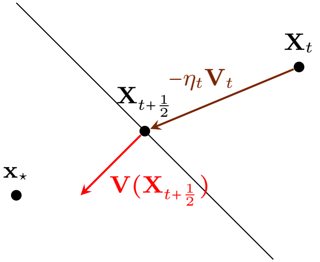

Intuitions Behind Extra-Gradient / Optimistic Gradient

$$\boxed{\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}_t, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1}\mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})}$$

Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

- Consider the hyperplan

$$\mathcal{H} \coloneqq \{\mathbf{x} : \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x}\rangle = 0\}$$

- Assumption: $\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x}_\star\rangle \geq 0$

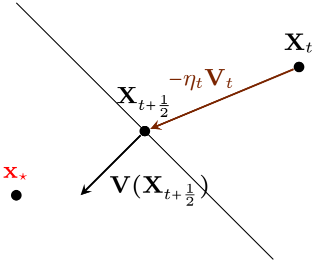

Intuitions Behind Extra-Gradient / Optimistic Gradient

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}_t, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1} \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

- Consider the hyperplan

$$\mathcal{H} := \{\mathbf{x} : \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x} \rangle = 0\}$$

- Assumption: $\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x}_\star \rangle \geq 0$

  Monotone: $\langle \mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x}), \mathbf{x}' - \mathbf{x} \rangle \geq 0$

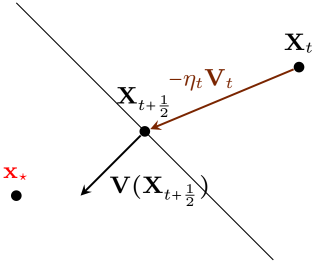

Intuitions Behind Extra-Gradient / Optimistic Gradient

$$\boxed{\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}_t, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1} \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})}$$

Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

- Consider the hyperplan

$$\mathcal{H} \coloneqq \{\mathbf{x} : \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x} \rangle = 0\}$$

- Assumption: $\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x}_\star \rangle \geq 0$
- If $\mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}) \approx \mathbf{V}_t$ then

$$\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{X}_t \rangle = -\eta_{t+1} \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{V}_t \rangle \leq 0$$

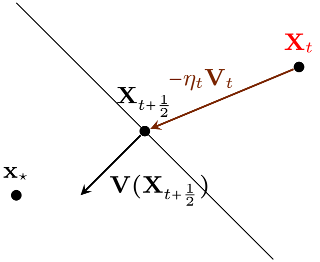

Intuitions Behind Extra-Gradient / Optimistic Gradient

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}_t, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1} \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

- Consider the hyperplan

$$\mathcal{H} \coloneqq \{\mathbf{x} : \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x} \rangle = 0\}$$

- Assumption: $\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x}_\star \rangle \geq 0$
- If $\mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}) \approx \mathbf{V}_t$ then

$$\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{X}_t \rangle = -\eta_{t+1} \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{V}_t \rangle \leq 0$$

Lipschiz: $\|\mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x})\| \leq \beta \|\mathbf{x}' - \mathbf{x}\|$

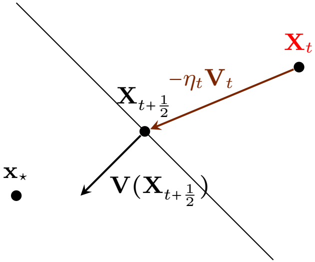

Intuitions Behind Extra-Gradient / Optimistic Gradient

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \mathbf{V}_t, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1}\mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$$

Relaxed projection onto a separating hyperplan [Tseng 00, Facchinei and Pang 03]

- Consider the hyperplan

$$\mathcal{H} \coloneqq \{\mathbf{x} : \langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x}\rangle = 0\}$$

- Assumption: $\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{x}_\star\rangle \geq 0$
- If $\mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}) \approx \mathbf{V}_t$ then

$$\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{X}_{t+\frac{1}{2}} - \mathbf{X}_t\rangle = -\eta_{t+1}\langle \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}}), \mathbf{V}_t\rangle \leq 0$$

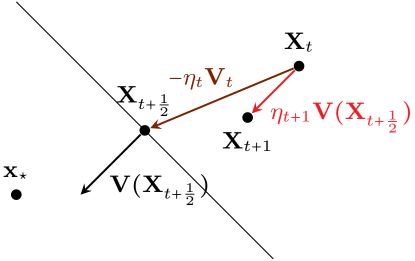

- The update step moves the iterate closer to the solutions

Theoretical Guarantees

- Bounded gradient feedback: $\exists\, G > 0,\ \forall\, t,\ g_t^i \leq G$
- Learning rate $\eta_t = \Theta(1/\sqrt{t})$
- $\mathcal{O}\left(\sqrt{t}\right)$ minimax-optimal regret in the adversarial regime

	Adversarial	Same algorithm + Lipschitz operator + M			
	Bounded feedback	-	-	Strongly M	Error bound
	$\mathrm{Reg}_t\,/t$	$\mathrm{Reg}_t\,/t$	$\\|\mathbf{V}(\mathbf{x}_t)\\|$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
GD	$1/\sqrt{t}$	✗	✗	$e^{-\rho_1 t}$	✗
OG	$1/\sqrt{t}$ Chiang et al. 12	$1/t$ H. et al. 19	$1/\sqrt{t}$ Cai et al. 22	$e^{-\rho_2 t}\ (\rho_2 \geq \rho_1)$ Mokhtari et al. 20	$e^{-\rho t}$ Wei et al. 21

Theoretical Guarantees

- Joint vector field: $\mathbf{V}(\mathbf{X}) = (\nabla_i \ell^i(\mathbf{X}))_{i \in \mathcal{N}}$; Nash equilibria: $\mathcal{X}_\star$
- Lipschitz operator: $\|\mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x})\| \leq \beta \|\mathbf{x}' - \mathbf{x}\|$
- Monotone (M): $\langle \mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x}), \mathbf{x}' - \mathbf{x} \rangle \geq 0$

	Adversarial	Same algorithm + Lipschitz operator + M			
	Bounded feedback	-	-	Strongly M	Error bound
	Reg_t / t	Reg_t / t	$\|\mathbf{V}(\mathbf{x}_t)\|$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
GD	$1/\sqrt{t}$	✗	✗	$e^{-\rho_1 t}$	✗
OG	$1/\sqrt{t}$ Chiang et al. 12	$1/t$ H. et al. 19	$1/\sqrt{t}$ Cai et al. 22	$e^{-\rho_2 t}$ $(\rho_2 \geq \rho_1)$ Mokhtari et al. 20	$e^{-\rho t}$ Wei et al. 21

Theoretical Guarantees

- Joint vector field: $\mathbf{V}(\mathbf{X}) = (\nabla_i \ell^i(\mathbf{X}))_{i\in\mathcal{N}}$; Nash equilibria: $\mathcal{X}_\star$
- Lipschitz operator: $\|\mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x})\| \le \beta\|\mathbf{x}' - \mathbf{x}\|$
- OG achieves constant regret

	Adversarial	Same algorithm + Lipschitz operator + M			
	Bounded feedback	-	-	Strongly M	Error bound
	Reg_t/t	Reg_t/t	$\|\mathbf{V}(\mathbf{x}_t)\|$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
GD	$1/\sqrt{t}$	✗	✗	$e^{-\rho_1 t}$	✗
OG	$1/\sqrt{t}$ Chiang et al. 12	$1/t$ H. et al. 19	$1/\sqrt{t}$ Cai et al. 22	$e^{-\rho_2 t}$ $(\rho_2 \ge \rho_1)$ Mokhtari et al. 20	$e^{-\rho t}$ Wei et al. 21

Theoretical Guarantees

- Joint vector field: $\mathbf{V}(\mathbf{X}) = (\nabla_i \ell^i(\mathbf{X}))_{i\in\mathcal{N}}$; Nash equilibria: $\mathcal{X}_\star$
- Lipschitz operator: $\|\mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x})\| \le \beta\|\mathbf{x}' - \mathbf{x}\|$
- OG achieves last-iterate convergence

	Adversarial	Same algorithm + Lipschitz operator + M			
	Bounded feedback Reg_t/t	- Reg_t/t	- $\|\mathbf{V}(\mathbf{x}_t)\|$	Strongly M $\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	Error bound $\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
GD	$1/\sqrt{t}$	✗	✗	$e^{-\rho_1 t}$	✗
OG	$1/\sqrt{t}$ Chiang et al. 12	$1/t$ H. et al. 19	$1/\sqrt{t}$ Cai et al. 22	$e^{-\rho_2 t}$ $(\rho_2 \ge \rho_1)$ Mokhtari et al. 20	$e^{-\rho t}$ Wei et al. 21

Theoretical Guarantees

- Joint vector field: $\mathbf{V}(\mathbf{X}) = (\nabla_i \ell^i(\mathbf{X}))_{i\in\mathcal{N}}$; Nash equilibria: $\mathcal{X}_\star$
- Lipschitz operator: $\|\mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x})\| \leq \beta\|\mathbf{x}' - \mathbf{x}\|$
- Strongly monotone: $\exists\, \alpha > 0,\ \langle \mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x}), \mathbf{x}' - \mathbf{x}\rangle \geq \alpha\|\mathbf{x}' - \mathbf{x}\|^2$

	Adversarial	Same algorithm + Lipschitz operator + M			
	Bounded feedback	-	-	Strongly M	Error bound
	Reg_t/t	Reg_t/t	$\|\mathbf{V}(\mathbf{x}_t)\|$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
GD	$1/\sqrt{t}$	✗	✗	$e^{-\rho_1 t}$	✗
OG	$1/\sqrt{t}$ Chiang et al. 12	$1/t$ H. et al. 19	$1/\sqrt{t}$ Cai et al. 22	$e^{-\rho_2 t}$ $(\rho_2 \geq \rho_1)$ Mokhtari et al. 20	$e^{-\rho t}$ Wei et al. 21

Theoretical Guarantees

- Joint vector field: $\mathbf{V}(\mathbf{X}) = (\nabla_i \ell^i(\mathbf{X}))_{i\in\mathcal{N}}$; Nash equilibria: $\mathcal{X}_\star$
- Lipschitz operator: $\|\mathbf{V}(\mathbf{x}') - \mathbf{V}(\mathbf{x})\| \leq \beta\|\mathbf{x}' - \mathbf{x}\|$
- Error bound / Metric (sub-)regularity: $\exists\, \tau > 0, \|\mathbf{V}(\mathbf{x})\| \geq \tau\, \mathrm{dist}(\mathbf{x}, \mathcal{X}_\star)$

	Adversarial	Same algorithm + Lipschitz operator + M			
	Bounded feedback	-	-	Strongly M	Error bound
	Reg_t/t	Reg_t/t	$\|\mathbf{V}(\mathbf{x}_t)\|$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
GD	$1/\sqrt{t}$	✗	✗	$e^{-\rho_1 t}$	✗
OG	$1/\sqrt{t}$ Chiang et al. 12	$1/t$ H. et al. 19	$1/\sqrt{t}$ Cai et al. 22	$e^{-\rho_2 t}$ $(\rho_2 \geq \rho_1)$ Mokhtari et al. 20	$e^{-\rho t}$ Wei et al. 21

Outline

Jérôme Malick

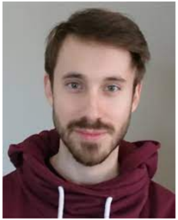

Franck Iutzeler

Panaytois Mertikopoulos

Kimon Antonakopoulos

Volkan Cevher

1. Y-G H., K. Antonakopoulos., V. Cevher, and P. Mertikopoulos. *No-Regret Learning in Games with Noisy Feedback: Faster Rates and Adaptivity via Learning Rate Separation.* arXiv preprint arXiv:2206.06015, 2022.
2. Y-G H., K. Antonakopoulos., and P. Mertikopoulos. *Adaptive Learning in Continuous Games: Optimal Regret Bounds and Convergence to Nash Equilibrium.* In **COLT**, 2021.
3. Y-G H., F. Iutzeler, J. Malick, and P. Mertikopoulos. *Explore Aggressively, Update Conservatively: Stochastic Extragradient Methods with Variable Stepsize Scaling.* In **NeurIPS**, 2020.
4. Y-G H., F. Iutzeler, J. Malick, and P. Mertikopoulos. *On the Convergence of Single-Call Stochastic Extra-Gradient Methods.* In **NeurIPS**, 2019.

The Burden of Learning Rate Tuning

Consider the same bilinear problem

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

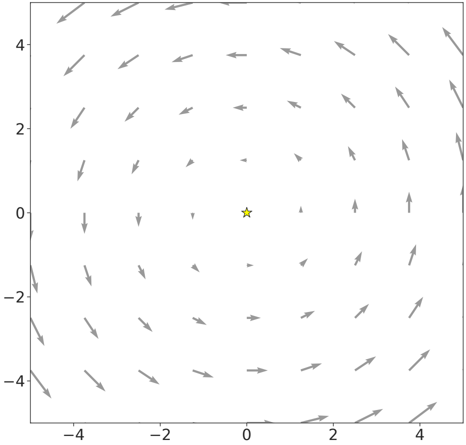

The Burden of Learning Rate Tuning

Consider the same bilinear problem

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Fast convergence is guaranteed with suitably tune learning rate

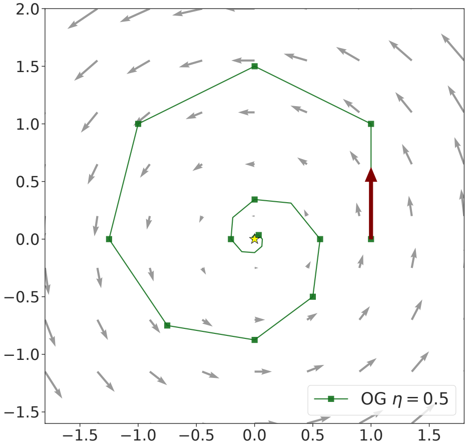

The Burden of Learning Rate Tuning

Consider the same bilinear problem

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Fast convergence is guaranteed with suitably tune learning rate
- But small perturbation can result in divergence

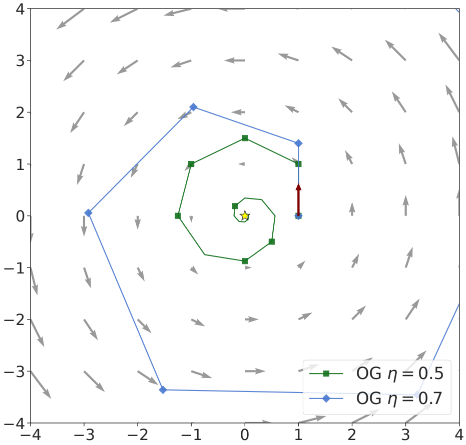

The Burden of Learning Rate Tuning

Consider the same bilinear problem

$$\ell^1(\mathbf{x}) = -\ell^2(\mathbf{x}) = \theta\phi \quad [\mathbf{x} = (\theta, \phi) \in \mathbb{R}^2]$$

- Fast convergence is guaranteed with suitably tune learning rate
- But small perturbation can result in divergence
- Robustness to adversarial requires vanishing learning rate that causes slow convergence

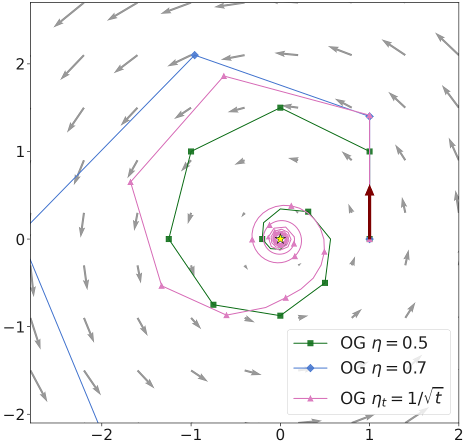

Adaptive Optimistic Dual Averaging

- OptDA – Optimistic dual averaging

$$X^i_{t+\frac{1}{2}} = X^i_t - \eta^i_t g^i_{t-1}$$

$$X^i_{t+1} = X^i_1 - \boxed{\eta^i_{t+1} \sum_{s=1}^{t} g^i_s}$$

Standard

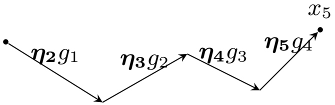

DA-type

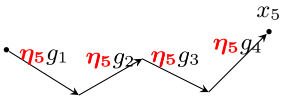

(More robust against adversaries)

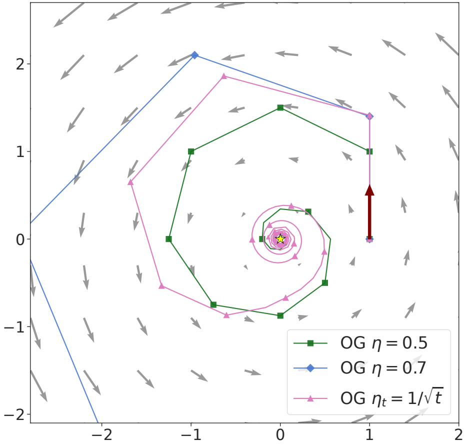

Adaptive Optimistic Dual Averaging

- OptDA – Optimistic dual averaging

$$X^i_{t+\frac{1}{2}} = X^i_t - \eta^i_t g^i_{t-1}$$

$$X^i_{t+1} = X^i_1 - \eta^i_{t+1} \sum_{s=1}^{t} g^i_s$$

- AdaOptDA uses learning rate

$$\eta^i_t = \frac{1}{\sqrt{1 + \boxed{\sum_{s=1}^{t-1} \|g^i_s - g^i_{s-1}\|^2}}}$$

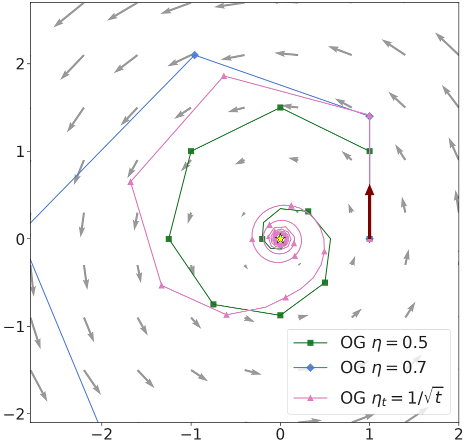

Adaptive Optimistic Dual Averaging

- OptDA – Optimistic dual averaging

$$X^i_{t+\frac{1}{2}} = X^i_t - \eta^i_t g^i_{t-1}$$

$$X^i_{t+1} = X^i_1 - \eta^i_{t+1} \sum_{s=1}^{t} g^i_s$$

- AdaOptDA uses learning rate

$$\eta^i_t = \frac{1}{\sqrt{1 + \sum_{s=1}^{t-1} \|g^i_s - g^i_{s-1}\|^2}}$$

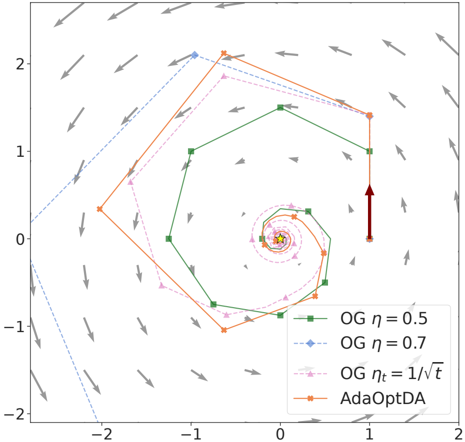

Second Issue: Stochasticity Breaks Optimistic Gradient

- Draw $\mathcal{L}_1(\mathbf{x}) = 3\theta\phi$ or $\mathcal{L}_2(\mathbf{x}) = -\theta\phi$ with equal probability so

$$\ell^1 = -\ell^2 = (\mathcal{L}_1 + \mathcal{L}_2)/2 = \theta\phi$$

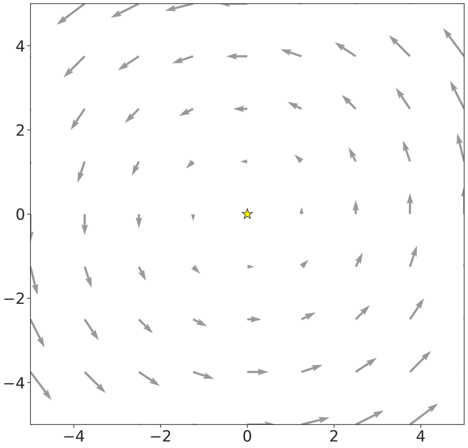

Second Issue: Stochasticity Breaks Optimistic Gradient

- Draw $\mathcal{L}_1(\mathbf{x}) = 3\theta\phi$ or $\mathcal{L}_2(\mathbf{x}) = -\theta\phi$ with equal probability so

$$\ell^1 = -\ell^2 = (\mathcal{L}_1 + \mathcal{L}_2)/2 = \theta\phi$$

- Stochastic estimate $\mathbb{E}[\hat{\mathbf{V}}_{t+\frac{1}{2}}] = \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$

$$\hat{\mathbf{V}}_{t+\frac{1}{2}} = \begin{cases} (3\phi_{t+\frac{1}{2}}, -3\theta_{t+\frac{1}{2}}) & \text{with prob. } 1/2 \\ (-\phi_{t+\frac{1}{2}}, \theta_{t+\frac{1}{2}}) & \text{with prob. } 1/2 \end{cases}$$

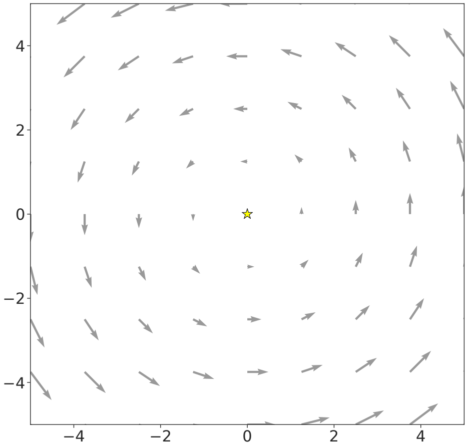

Second Issue: Stochasticity Breaks Optimistic Gradient

- Draw $\mathcal{L}_1(\mathbf{x}) = 3\theta\phi$ or $\mathcal{L}_2(\mathbf{x}) = -\theta\phi$ with equal probability so

$$\ell^1 = -\ell^2 = (\mathcal{L}_1 + \mathcal{L}_2)/2 = \theta\phi$$

- Stochastic estimate $\mathbb{E}[\hat{\mathbf{V}}_{t+\frac{1}{2}}] = \mathbf{V}(\mathbf{X}_{t+\frac{1}{2}})$

$$\hat{\mathbf{V}}_{t+\frac{1}{2}} = \begin{cases} (3\phi_{t+\frac{1}{2}}, -3\theta_{t+\frac{1}{2}}) & \text{with prob. } 1/2 \\ (-\phi_{t+\frac{1}{2}}, \theta_{t+\frac{1}{2}}) & \text{with prob. } 1/2 \end{cases}$$

- Optimistic gradient $[\mathbf{x}_t = \mathbf{X}_{t+\frac{1}{2}}]$

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \eta_t \hat{\mathbf{V}}_{t-\frac{1}{2}}, \quad \mathbf{X}_{t+1} = \mathbf{X}_t - \eta_{t+1} \hat{\mathbf{V}}_{t+\frac{1}{2}}$$

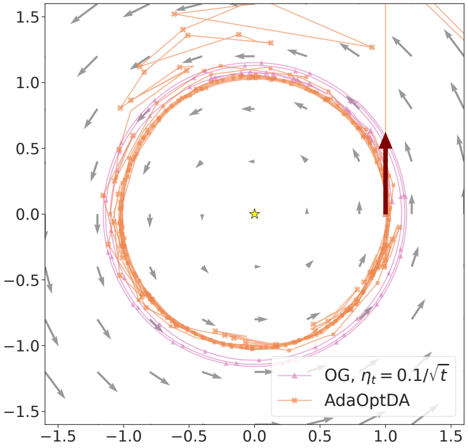

Scale Separation as a Remedy

- Draw $\mathcal{L}_1(\mathbf{x}) = 3\theta\phi$ or $\mathcal{L}_2(\mathbf{x}) = -\theta\phi$ with equal probability so

$$\ell^1 = -\ell^2 = (\mathcal{L}_1 + \mathcal{L}_2)/2$$

- OG+ $[\mathbf{x}_t = \mathbf{X}_{t+\frac{1}{2}}]$

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \boxed{\gamma_t} \hat{\mathbf{V}}_{t-\frac{1}{2}}$$

$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \hat{\mathbf{V}}_{t+\frac{1}{2}}$$

With $\gamma_t \geq \eta_t$

- This makes the noise an order smaller than the negative shift in the analysis

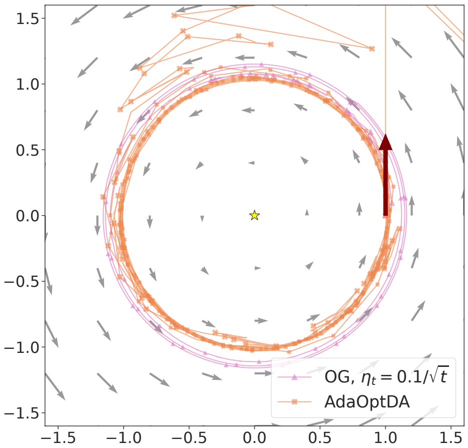

Scale Separation as a Remedy

- Draw $\mathcal{L}_1(\mathbf{x}) = 3\theta\phi$ or $\mathcal{L}_2(\mathbf{x}) = -\theta\phi$ with equal probability so

$$\ell^1 = -\ell^2 = (\mathcal{L}_1 + \mathcal{L}_2)/2$$

- OG+ $[\mathbf{x}_t = \mathbf{X}_{t+\frac{1}{2}}]$

$$\mathbf{X}_{t+\frac{1}{2}} = \mathbf{X}_t - \gamma_t \hat{\mathbf{V}}_{t-\frac{1}{2}}$$
$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta_t \hat{\mathbf{V}}_{t+\frac{1}{2}}$$

With $\gamma_t \geq \eta_t$

- This makes the noise an order smaller than the negative shift in the analysis

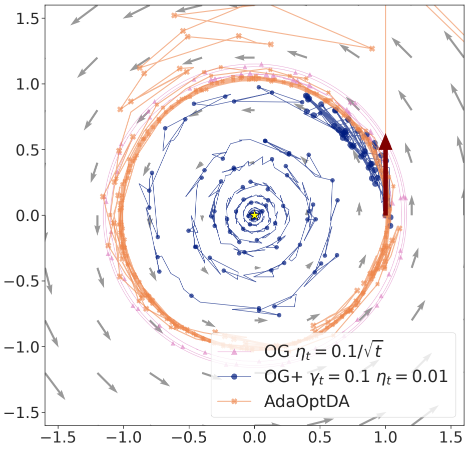

Adaptive Optimistic Dual Averaging with Scale Separation

- OptDA+ $[\gamma_t^i \geq \eta_t^i]$

$$X_{t+\frac{1}{2}}^i = X_t^i - \boxed{\gamma_t^i}\, g_{t-1}^i \quad X_{t+1}^i = X_1^i - \eta_{t+1}^i \sum_{s=1}^{t} g_s^i$$

- AdaOptDA+ uses learning rate

$$\gamma_t^i = \frac{1}{\left(1 + \sum_{s=1}^{t-2} \|g_s^i\|^2\right)^{\boxed{\frac{1}{4}}}}$$

$$\eta_t^i = \frac{1}{\sqrt{1 + \sum_{s=1}^{t-2} \left(\|g_s^i\|^2 + \|X_s^i - X_{s+1}^i\|^2\right)}}$$

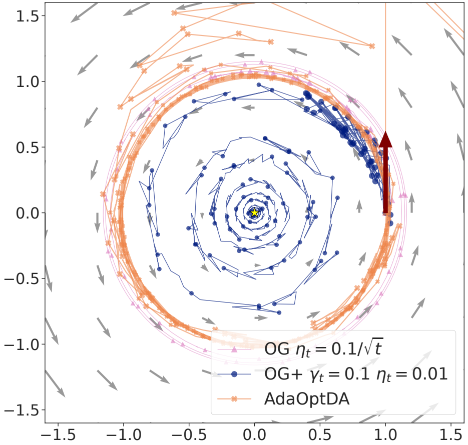

Adaptive Optimistic Dual Averaging with Scale Separation

- OptDA+ $[\gamma_t^i \geq \eta_t^i]$

$$X_{t+\frac{1}{2}}^i = X_t^i - \gamma_t^i g_{t-1}^i \quad X_{t+1}^i = X_1^i - \eta_{t+1}^i \sum_{s=1}^{t} g_s^i$$

- AdaOptDA+ uses learning rate

$$\gamma_t^i = \frac{1}{\left(1 + \sum_{s=1}^{t-2} \|g_s^i\|^2\right)^{\frac{1}{4}}}$$

$$\eta_t^i = \frac{1}{\sqrt{1 + \sum_{s=1}^{t-2} \left(\|g_s^i\|^2 + \|X_s^i - X_{s+1}^i\|^2\right)}}$$

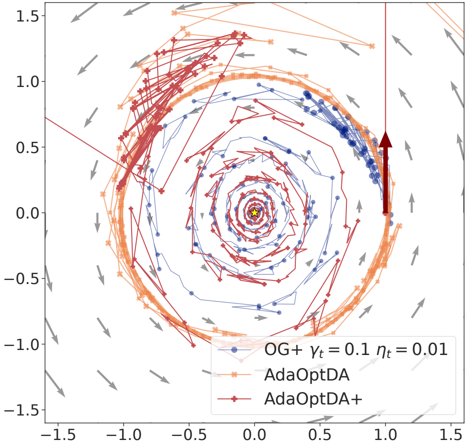

Adaptive Optimistic Dual Averaging with Scale Separation

- OptDA+ $[\gamma_t^i \geq \eta_t^i]$

$$X_{t+\frac{1}{2}}^i = X_t^i - \boxed{\gamma_t^i} g_{t-1}^i \quad X_{t+1}^i = X_1^i - \eta_{t+1}^i \sum_{s=1}^{t} g_s^i$$

- AdaOptDA+/v2 uses learning rate

$$\gamma_t^i = \frac{1}{\left(1 + \sum_{s=1}^{t-1} \|g_s^i - g_{s-1}^i\|^2\right)^{\frac{1}{4}}}$$

$$\eta_t^i = \frac{1}{\sqrt{1 + \sum_{s=1}^{t-1} \|g_s^i - g_{s-1}^i\|^2}}$$

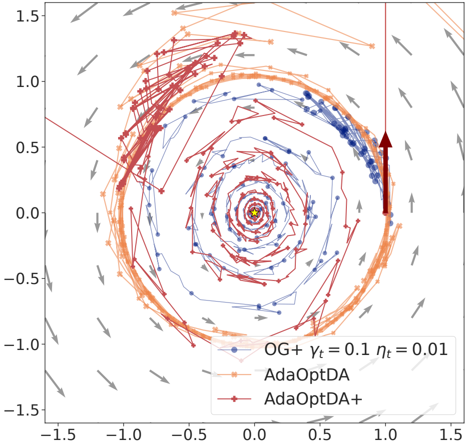

Adaptive Optimistic Dual Averaging with Scale Separation

- OptDA+ $[\gamma_t^i \geq \eta_t^i]$

$$X_{t+\frac{1}{2}}^i = X_t^i - \boxed{\gamma_t^i} g_{t-1}^i \quad X_{t+1}^i = X_1^i - \eta_{t+1}^i \sum_{s=1}^{t} g_s^i$$

- AdaOptDA+/v2 uses learning rate

$$\gamma_t^i = \frac{1}{\left(1 + \sum_{s=1}^{t-1} \|g_s^i - g_{s-1}^i\|^2\right)^{\frac{1}{4}}}$$

$$\eta_t^i = \frac{1}{\sqrt{1 + \sum_{s=1}^{t-1} \|g_s^i - g_{s-1}^i\|^2}}$$

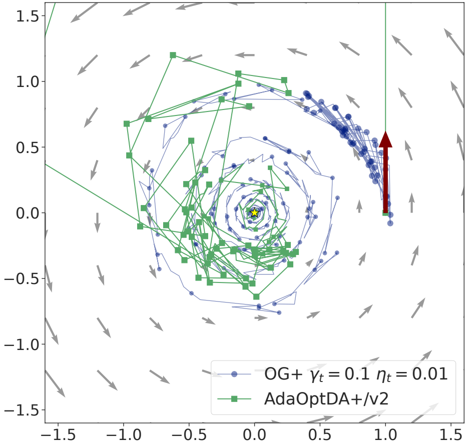

Convergence to Solution Under Multiplicative Noise

- $g_t^i = \nabla_i \ell^i(\mathbf{x}_t)(1 + \xi_t^i)$ where ξ_t^i is unbiased and has finite variance

Convergence to Solution Under Multiplicative Noise

- $g_t^i = \nabla_i \ell^i(\mathbf{x}_t)(1+\xi_t^i)$ where ξ_t^i is unbiased and has finite variance
- E.g., $\hat{\mathbf{V}}_{t+\frac{1}{2}}$ is $(3\phi_{t+\frac{1}{2}}, -3\theta_{t+\frac{1}{2}})$ or $(-\phi_{t+\frac{1}{2}}, \theta_{t+\frac{1}{2}})$ with probability one half for each

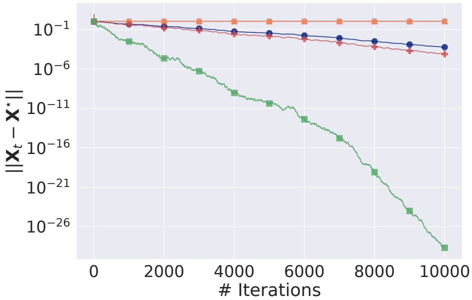

Base state $\mathbf{X}_t$

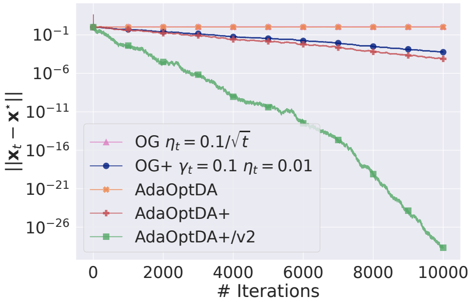

Played action $\mathbf{x}_t = \mathbf{X}_{t+\frac{1}{2}}$

Convergence to Solution Under Additive Noise

- $g_t^i = \boxed{\nabla_i \ell^i(\mathbf{x}_t) + \xi_t^i}$ where ξ_t^i is unbiased and has finite variance

Convergence to Solution Under Additive Noise

- $g_t^i = \nabla_i \ell^i(\mathbf{x}_t) + \xi_t^i$ where ξ_t^i is unbiased and has finite variance
- $\hat{\mathbf{V}}_{t+\frac{1}{2}} = (\phi_{t+\frac{1}{2}} + \xi_t^1, -\theta_{t+\frac{1}{2}} + \xi_t^2)$ where $\xi_t^1, \xi_t^2 \sim \mathcal{N}(0,1)$

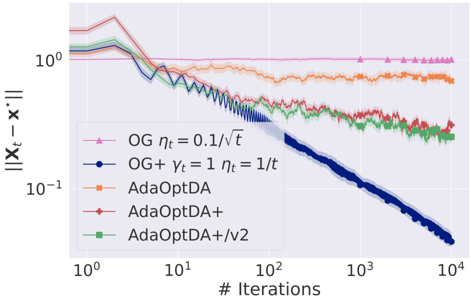

Base state $\mathbf{X}_t$

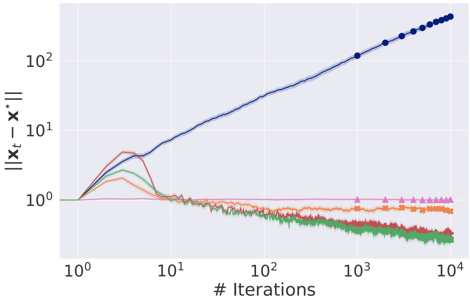

Played action $\mathbf{x}_t = \mathbf{X}_{t+\frac{1}{2}}$

Theoretical Guarantees Under Uncertainty (in Expectation)

		Adversarial	Same algorithm + Lipschitz operator + M			
		Bounded feedback	-	-	Strongly M	Error bound
		Reg_t/t	Reg_t/t	Cvg?	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
AdaOptDA	Det.	$1/\sqrt{t}$	$1/t$	✓	-	-
OG+	Mul.	$1/\sqrt{t}$	$1/t$	✓	$e^{-\rho t}$	$e^{-\rho t}$
	Add.		$1/\sqrt{t}$	✓	$1/\sqrt{t}$	$1/t^{1/6}$
AdaOptDA+	Mul.	$1/t^{1/4}$	$1/t$	✓	-	-
	Add.		$1/\sqrt{t}$	-	-	-

Theoretical Guarantees Under Uncertainty (in Expectation)

		Adversarial	Same algorithm + Lipschitz operator + M			
		Bounded feedback	-	-	Strongly M	Error bound
		$\mathrm{Reg}_t\,/t$	$\mathrm{Reg}_t\,/t$	Cvg?	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
AdaOptDA	Det.	$1/\sqrt{t}$	$1/t$	✓	-	-
OG+	Mul.	$1/\sqrt{t}$	$1/t$	✓	$e^{-\rho t}$	$e^{-\rho t}$
	Add.		$1/\sqrt{t}$	✓	$1/\sqrt{t}$	$1/t^{1/6}$
AdaOptDA+	Mul.	$1/t^{1/4}$	$1/t$	✓	-	-
	Add.		$1/\sqrt{t}$	-	-	-

Theoretical Guarantees Under Uncertainty (in Expectation)

		Adversarial	Same algorithm + Lipschitz operator + M			
		Bounded feedback	-	-	Strongly M	Error bound
		Reg_t/t	Reg_t/t	Cvg?	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$	$\mathrm{dist}(\mathbf{x}_t, \mathcal{X}_\star)$
AdaOptDA	Det.	$1/\sqrt{t}$	$1/t$	✓	-	-
OG+	Mul.	$1/\sqrt{t}$	$1/t$	✓	$e^{-\rho t}$	$e^{-\rho t}$
	Add.		$1/\sqrt{t}$	✓	$1/\sqrt{t}$	$1/t^{1/6}$
AdaOptDA+	Mul.	$1/t^{1/4}$	$1/t$	✓	-	-
	Add.		$1/\sqrt{t}$	-	-	-

Conclusion and Perspectives

- Iterative learning in a multi-agent environment
- Game theoretic model, algorithms with recency bias
- Automatic learning rate tuning, learning under uncertainty

Conclusion and Perspectives

- Iterative learning in a multi-agent environment
- Game theoretic model, algorithms with recency bias
- Automatic learning rate tuning, learning under uncertainty
- Perspectives:
  - Technical challenges (e.g., relaxed assumptions, convergence rates)
  - What about other algorithms?
  - Beyond no-regret and convergence

Conclusion and Perspectives

- Iterative learning in a multi-agent environment
- Game theoretic model, algorithms with recency bias
- Automatic learning rate tuning, learning under uncertainty
- Perspectives:
  - Technical challenges (e.g., relaxed assumptions, convergence rates)
  - What about other algorithms?
  - Beyond no-regret and convergence

Thank you for your attention