Last-Iterate Convergence: Zero-Sum Games and Constrained Min-Max Optimization

Ioannis Panageas
Singapore University of Technology and Design (SUTD)

Constantinos Daskalakis
MIT

Min-max Optimization

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

- Zero-sum: $f(\mathbf{x}, \mathbf{y})$ represents the payment of the $\mathcal{Y}$ player to the $\mathcal{X}$.
- Classic def: $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top A\mathbf{y}$ (bilinear), $\mathcal{X}, \mathcal{Y}$ are simplex.
- Applications: Generative Adversarial networks (GANs), Game Theory, etc.

Min-max Optimization

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

- Zero-
- Class
- Applic

1. GANs: Objective non-convex non-concave!
2. NE not guaranteed to exist!
3. Gradient Descent Oscillates for unconstrained case.

ne $\mathcal{X}$.

y, etc.

Min-max Optimization

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

- Zero-s … he $\mathcal{X}$.
- Class
- Applic … y, etc.

1. GANs: Objective non-convex non-concave!
2. NE not guaranteed to exist!
3. Gradient Descent Oscillates for unconstrained case.

[Daskalakis et al, ICLR18]: Show last iterate convergence of Optimistic Gradient Descent (OGD) for bilinear functions in the unconstrained case!

Min-max Optimization (cont.)

$$\min_{\mathbf{y} \in \Delta_m} \max_{\mathbf{x} \in \Delta_n} f(\mathbf{x}, \mathbf{y}).$$

- Captures linear programming.
- Daskalakis et al does not apply!
- The analogue of OGD is optimistic multiplicative weights update.

Last-Iterate vs Time Average

Last-Iterate convergence: $\lim_{T\to\infty} z^T$ exists!	
Time Average: $\lim_{T\to\infty} \frac{\sum_{j=1}^{T} z^j}{T}$ exists!	

Last-Iterate vs Time Average

Last-Iterate convergence: $\lim_{T\to\infty} z^T$ exists!

Time Average: $\lim_{T\to\infty} \frac{\sum_{j=1}^{T} z^j}{T}$ exists!

- Last-iterate implies time average (converse not true).

Last-Iterate vs Time Average

Last-Iterate convergence: $\lim_{T\to\infty} z^T$ exists!

Time Average: $\lim_{T\to\infty} \frac{\sum_{j=1}^{T} z^j}{T}$ exists!

- Last-iterate implies time average (converse not true).
- Convergence rate $\Omega\left(\frac{1}{T}\right)$ for time average if sequence is positive, last-iterate potentially faster!

Last-Iterate vs Time Average

Last-Iterate convergence: $\lim_{T\to\infty} z^T$ exists!

Time Average: $\lim_{T\to\infty} \frac{\sum_{j=1}^{T} z^j}{T}$ exists!

- Last-iterate implies time average (converse not true).
- Convergence rate $\Omega\left(\frac{1}{T}\right)$ for time average if sequence is positive, last-iterate potentially faster!
- Last-iterate gives better predictions!

Last-Iterate vs Time Average

Last-Iterate convergence: $\lim_{T\to\infty} z^T$ exists!

Time Average: $\lim_{T\to\infty} \frac{\sum_{j=1}^{T} z^j}{T}$ exists!

- Last-iterate implies time average (converse not true).
- Convergence rate $\Omega\left(\frac{1}{T}\right)$ for time average if sequence is positive, last-iterate potentially faster!
- Last-iterate gives better predictions!
- Example (2 slides later).

Multiplicative Weights Update

$$\text{MWU} \equiv \text{FTRL with entropy regularizer,}$$

boils down to:

$$x_i^{t+1} = x_i^t \frac{e^{\eta(A\mathbf{y}^t)_i}}{\sum_{j=1}^n x_j^t e^{\eta(A\mathbf{y}^t)_j}} \text{ for all } i \in [n],$$

$$y_i^{t+1} = y_i^t \frac{e^{-\eta(A^\top\mathbf{x}^t)_i}}{\sum_{j=1}^m y_j^t e^{-\eta(A^\top\mathbf{x}^t)_j}} \text{ for all } i \in [m].$$

Multiplicative Weights Update (cont.)

Example: [Bailey et. al., EC18]

Matching pennies game, classic MWU.

$$A = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

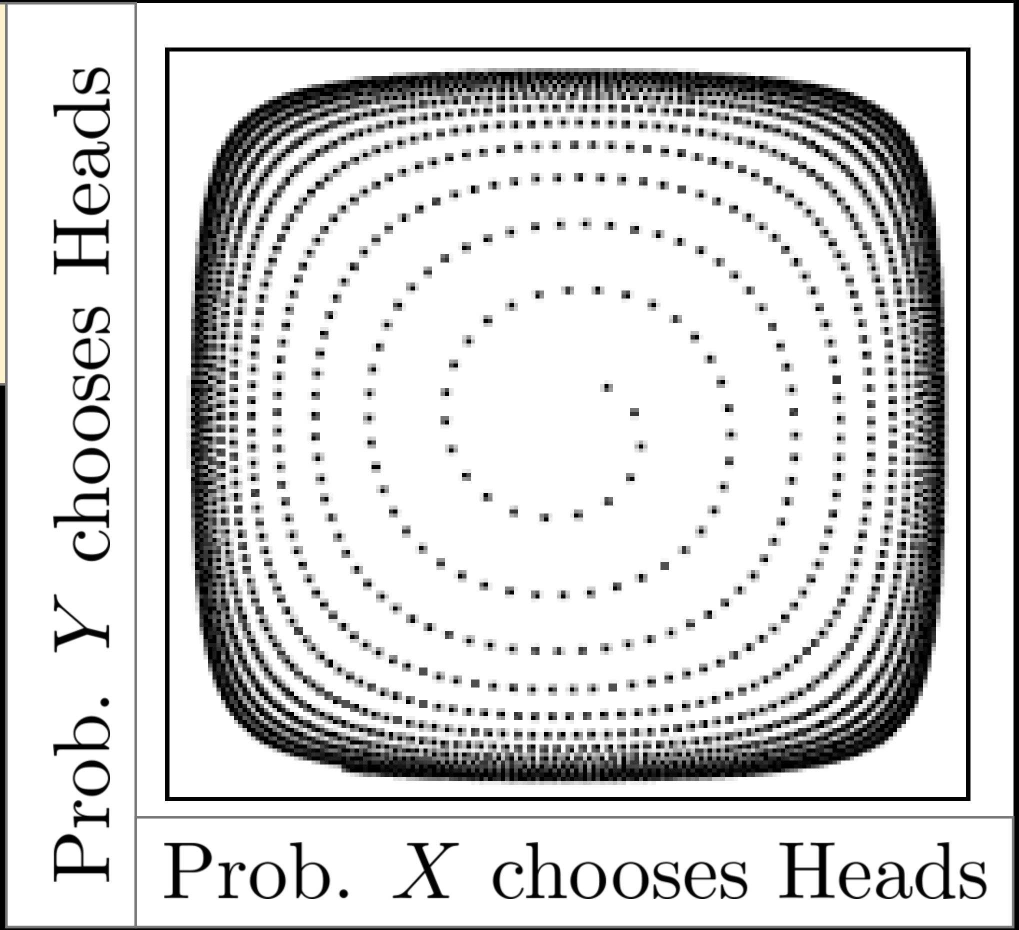

Multiplicative Weights Update (cont.)

Example: [Bailey et. al., EC18]

Matching pennies game, classic MWU.

$$A = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

- Oscillates towards the boundary!
- Time average converges to Nash eq. (MWU is no-regret):

$$\lim_{T\to\infty} \frac{\sum_{j=1}^{T} x^{j}_{\text{Heads}}}{T} = \frac{1}{2}$$

$$\lim_{T\to\infty} \frac{\sum_{j=1}^{T} y^{j}_{\text{Heads}}}{T} = \frac{1}{2}$$

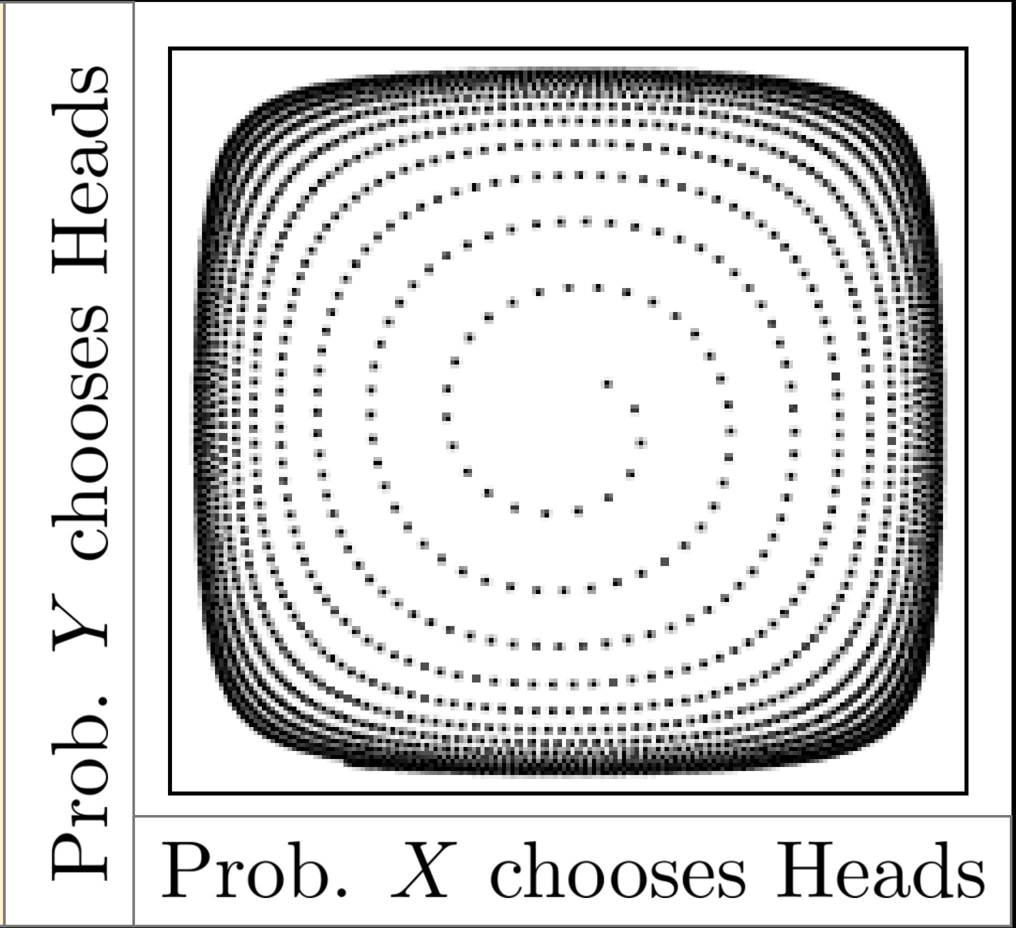

Optimistic Multiplicative Weights Update

Optimistic FTRL with entropy regularizer boils down to:

$$x_i^{t+1} = x_i^t \frac{e^{\mathbf{2}\eta(A\mathbf{y}^t)_i - \eta(A\mathbf{y}^{t-1})_i}}{\sum_{j=1}^n x_j^t e^{\mathbf{2}\eta(A\mathbf{y}^t)_j - \eta(A\mathbf{y}^{t-1})_j}} \text{ for all } i \in [n],$$

$$y_i^{t+1} = y_i^t \frac{e^{-\mathbf{2}\eta(A^\top\mathbf{x}^t)_i + \eta(A^\top\mathbf{x}^{t-1})_i}}{\sum_{j=1}^m y_j^t e^{-\mathbf{2}\eta(A^\top\mathbf{x}^t)_j + \eta(A^\top\mathbf{x}^{t-1})_j}} \text{ for all } i \in [m].$$

Optimistic Multiplicative Weights Update

Optimistic FTRL with entropy regularizer boils down to:

$$x_i^{t+1} = x_i^t \frac{e^{\mathbf{2}\eta(A\mathbf{y}^t)_i - \eta(A\mathbf{y}^{t-1})_i}}{\sum_{j=1}^n x_j^t e^{\mathbf{2}\eta(A\mathbf{y}^t)_j - \eta(A\mathbf{y}^{t-1})_j}} \text{ for all } i \in [n],$$

$$y_i^{t+1} = y_i^t \frac{e^{-\mathbf{2}\eta(A^\top\mathbf{x}^t)_i + \eta(A^\top\mathbf{x}^{t-1})_i}}{\sum_{j=1}^m y_j^t e^{-\mathbf{2}\eta(A^\top\mathbf{x}^t)_j + \eta(A^\top\mathbf{x}^{t-1})_j}} \text{ for all } i \in [m].$$

Recall!
$$x_i^{t+1} = x_i^t \frac{e^{\eta(A\mathbf{y}^t)_i}}{\sum_{j=1}^n x_j^t e^{\eta(A\mathbf{y}^t)_j}},$$
$$y_i^{t+1} = y_i^t \frac{e^{-\eta(A^\top\mathbf{x}^t)_i}}{\sum_{j=1}^m y_j^t e^{-\eta(A^\top\mathbf{x}^t)_j}}.$$

In this talk

Theorem. *Let A be a $n \times m$ matrix and assume that*

$$\min_{\boldsymbol{y} \in \Delta_m} \max_{\boldsymbol{x} \in \Delta_n} \boldsymbol{x}^\top A \boldsymbol{y}$$

has a unique solution $(\boldsymbol{x}^*, \boldsymbol{y}^*)$. For stepsize η sufficiently small (depends on n, m, A), starting from a point $(\boldsymbol{x}^0, \boldsymbol{y}^0)$ in the interior of $\Delta_m \times \Delta_n$, it holds

$$\lim_{t \to \infty} (\boldsymbol{x}^t, \boldsymbol{y}^t) = (\boldsymbol{x}^*, \boldsymbol{y}^*),$$

under OMWU dynamics. The stepsize η is constant, i.e., does not scale with time.

About the proof

STEP 1: OMWU dynamics reaches a small neighborhood U around the NE $(\mathbf{x}^*, \mathbf{y}^*)$.

We show:
KL divergence from $(\mathbf{x}^t, \mathbf{y}^t)$ to $(\mathbf{x}^*, \mathbf{y}^*)$ decreases with time *unless*

$$(\mathbf{x}^t, \mathbf{y}^t) \in U.$$

About the proof

STEP 1: OMWU dynamics reaches a small neighborhood U around the NE $(\mathbf{x}^*, \mathbf{y}^*)$.

We show:

$$\sum_{i=1}^{n} x_i^* \ln(x_i^*/x_i^t) + \sum_{i=1}^{m} y_i^* \ln(y_i^*/y_i^t)$$

decreases with time *unless*

$$(\mathbf{x}^t, \mathbf{y}^t) \in U.$$

About the proof

STEP 2: OMWU dynamics is a contraction mapping inside U.

OMWU is a dynamical system of the form

$$(\mathbf{x}^{t+1}, \mathbf{y}^{t+1}, \mathbf{x}^{t}, \mathbf{y}^{t}) = \mathbf{G}(\mathbf{x}^{t}, \mathbf{y}^{t}, \mathbf{x}^{t-1}, \mathbf{y}^{t-1}).$$

We show:

The Jacobian of the map $\mathbf{G}$ has spectral radius less than one in U

Take away messages

- Optimism helps the trajectories of learning dynamics stabilize.
- Proof of convergence for the last iterate needs fundamentally different techniques! Dynamical systems to the rescue!

Take away messages – Questions

- Optimism helps the trajectories of learning dynamics stabilize.
- Proof of convergence for the last iterate needs fundamentally different techniques! Dynamical systems to the rescue!

- Rate of convergence (is it polynomial in n, m)?
- Generalize the result for convex-concave objectives.

Take away messages – Questions

- Optimism helps the trajectories of learning dynamics stabilize.
- Proof of convergence for the last iterate needs fundamentally different techniques! Dynamical systems to the rescue!

- Rate of convergence (is it polynomial in n, m)?
- Generalize the result for convex-concave objectives.

Thank you! ☺