Algorithmic Game Theory

CoReLab (NTUA)

Lecture 2:
Nash Equilibria in Bimatrix Games

2-player 0-sum game: Rock-Paper-Scissors

- 2-players: “row” and “column”
- Payoff of row player = loss of column player
- Nash Equilibrium:

$$x = \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}, y = \begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}$$

			
	0,0	-1,1	1,-1
	1,-1	0,0	-1,1
	-1,1	1,-1	0,0

Nash Equilibrium: A set of strategies such that no player can increase their payoff by deviating unilaterally from their strategy.

2-player 0-sum game

- Row player has m (pure) strategies and column player has n.
- Strategy ⟶ probability vector over the pure strategies

Row player: $\boldsymbol{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{bmatrix} \in \Delta_m$ Column player: $\boldsymbol{y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} \in \Delta_n$

- A 2-player game is a tuple (R, C) of playoff matrices of size $m \times n$ for the row and the column player respectively.
  The game is zero-sum if $R = -C$.

2-player 0-sum game

$(\boldsymbol{x}, \boldsymbol{y})$ is a Nash Equilibrium $\Leftrightarrow$ $\begin{array}{l} \boldsymbol{x}^T R \boldsymbol{y} \geq \boldsymbol{x'}^T R \boldsymbol{y} \ \ \forall \boldsymbol{x'} \in \Delta_m \\ \boldsymbol{x}^T C \boldsymbol{y} \geq \boldsymbol{x}^T C \boldsymbol{y'} \ \ \forall \boldsymbol{y'} \in \Delta_n \end{array}$

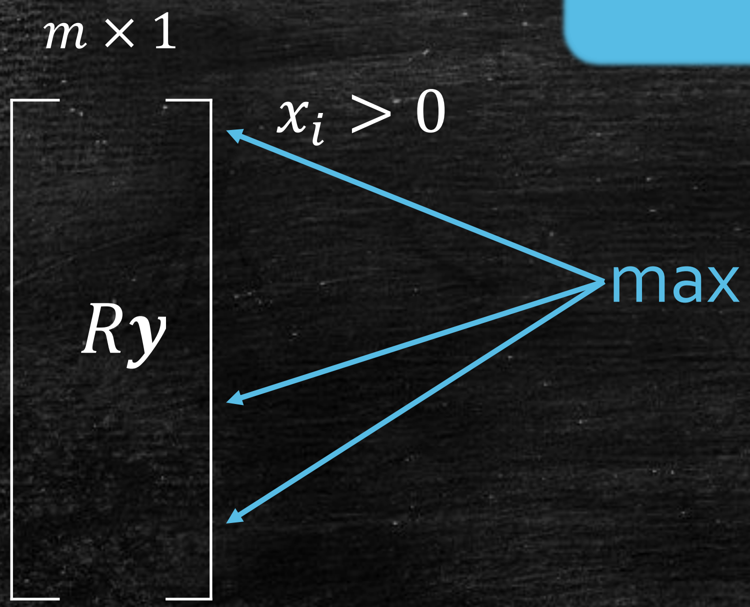

$m \times 1$

$x_i > 0$

$R\boldsymbol{y}$

max

- If row player knew $\boldsymbol{y}$, she would choose $\boldsymbol{x}$ so that its positive coordinates correspond to the max values of $R\boldsymbol{y}$.

$$\boldsymbol{x}^T R \boldsymbol{y} \geq \boldsymbol{e}_i^T R \boldsymbol{y} \ \forall i \in [m]$$

$$\boldsymbol{x}^T R \boldsymbol{y} \geq \boldsymbol{x}^T R \boldsymbol{e}_j \ \forall j \in [n]$$

2-player 0-sum game: Example

$$R = -C = \begin{bmatrix} 3 & -1 \\ -2 & 1 \end{bmatrix}$$

- If row player chooses $\boldsymbol{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$, then $\boldsymbol{x}^T C = \begin{bmatrix} -3x_1 + 2x_2 \\ x_1 - x_2 \end{bmatrix} = -\boldsymbol{x}^T R$.
- Then the column player responds optimally and gets $\max(-3x_1 + 2x_2, x_1 - x_2)$ so the row player gets $-\max(-3x_1 + 2x_2, x_1 - x_2) = \min(3x_1 - 2x_2, -x_1 + x_2)$.
- Row player chooses $\boldsymbol{x}$ to `maximize` $\min(3x_1 - 2x_2, -x_1 + x_2)$.

2-player 0-sum game: Example

We can formulate $\max_{x_1,x_2} \min(3x_1 - 2x_2, -x_1 + x_2)$ as a LP:

$$\begin{aligned} & \max z \\ \text{s.t.} \quad & 3x_1 - 2x_2 \geq z \\ & -x_1 + x_2 \geq z \\ & x_1 + x_2 = 1 \\ & x_1, x_2 \geq 0 \end{aligned}$$

- The solution of this LP $\boldsymbol{x} = \begin{bmatrix} 3/7 \\ 4/7 \end{bmatrix}$ and the solution of its dual LP $\boldsymbol{y} = \begin{bmatrix} 2/7 \\ 5/7 \end{bmatrix}$ is a Nash Equilibrium where the payoff of the row player is $1/7$.

Quick LP recap!

2-player 0-sum game

(Primal) LP for the row player:

$$\max z$$

s.t. $\boldsymbol{x}^T R \geq \underbrace{[z \quad \cdots \quad z]^T}_{n}$

$$\sum_{i=1}^{m} x_i = 1$$

$$x_i \geq 0 \ \forall i \in [m]$$

Solution: $\max_{\boldsymbol{x}} \min_{\boldsymbol{y}} \boldsymbol{x}^T R \boldsymbol{y}$

(Dual) LP for the column player:

$$\max \ z'$$

s.t. $C\boldsymbol{y} \geq \underbrace{[z' \quad \cdots \quad z']}_{m}$

$$\sum_{i=1}^{n} y_j = 1$$

$$\mathrm{y_j} \geq 0 \ \forall j \in [n]$$

Solution: $\max_{\boldsymbol{y}} \min_{\boldsymbol{x}} \boldsymbol{x}^T C \boldsymbol{y} = -\min_{\boldsymbol{y}} \max_{\boldsymbol{x}} \boldsymbol{x}^T R \boldsymbol{y}$

2-player 0-sum game

- If $(\boldsymbol{x}, z)$ and $(\boldsymbol{y}, z')$ are optimal solutions for the row and column player LPs respectively, then $(\boldsymbol{x}, \boldsymbol{y})$ is a N.E. with $z = -z'$ (and vise versa).

- $(\boldsymbol{x}, z)$ is feasible $\Rightarrow \boldsymbol{x}^T R \geq \underbrace{[z \quad \cdots \quad z]^T}_{n} \xRightarrow[\sum_{i=1}^{n} y_j = 1]{} \boldsymbol{x}^T R \boldsymbol{y} \geq z$

  Therefore if the row player plays $\boldsymbol{x}$ against $\boldsymbol{y}$ she gains at least z.

- $(\boldsymbol{y}, z')$ is feasible $\Rightarrow C\boldsymbol{y} \geq \underbrace{[z' \quad \cdots \quad z']}_{m} \xRightarrow[\sum_{i=1}^{n} x'_j = 1]{} \boldsymbol{x}'^T R \boldsymbol{y} \leq -z' \;\; \forall x'$

  Therefore the row player gains at most $-z' = z$ if the column player plays $\boldsymbol{y}$.

  Thus, $\boldsymbol{x}$ is best response and the same holds for $\boldsymbol{y}$.

2-player 0-sum game

- There is always a N.E.
- It doesn't matter who plays first (Minmax Theorem)
- A pair of any two solutions of the LPs is a N.E.
- No matter what strategies the players follow, the payoffs in a N.E. are the same (a.k.a. value of the game).
- The set of equilibria and the sets of optimal strategies are all convex.

Fictitious play

What if the players don't know the matrices R, C?

Example: $R = -C = \begin{bmatrix} 2 & 1 & 0 \\ 2 & 0 & 3 \\ -1 & 3 & -3 \end{bmatrix}$ (not known to the players).

Fictitious play

$$R = \begin{bmatrix} 2 & 1 & 0 \\ 2 & 0 & 3 \\ -1 & 3 & -3 \end{bmatrix}$$

☐ In the 1st round:

Row player plays arbitrary strategy $i_1 = 1$ and column player $j_1 = 3$.

- Row player observes $R\boldsymbol{e}_{j_1} = \begin{bmatrix} 0 \\ 3 \\ -3 \end{bmatrix} \begin{matrix} \\ \leftarrow \text{max} \\ \\ \end{matrix}$
- Column player observes $\boldsymbol{e}_{i_1}^T$C$= [-2 \quad -1 \quad \underset{\substack{\uparrow \\ \text{max}}}{0}]$

☐ In the 2nd round:

Row player plays $i_2 = 2$ and column player $j_2 = 3$.

- Row player observes $R\boldsymbol{e}_{j_2} = \begin{bmatrix} 0 \\ 3 \\ -3 \end{bmatrix}$
- Column player observes $\boldsymbol{e}_{i_2}^T$C$= [-2 \quad 0 \quad -3]$

Fictitious play

$$R = \begin{bmatrix} 2 & 1 & 0 \\ 2 & 0 & 3 \\ -1 & 3 & -3 \end{bmatrix}$$

- In the 3rd round:

Row player chooses the max row of $\frac{1}{2}R\left(\boldsymbol{e}_{j_1} + \boldsymbol{e}_{j_2}\right) = \begin{bmatrix} 0 \\ 3 \\ -3 \end{bmatrix} \longleftarrow i_3 = 2.$

Column player chooses the max column of $\frac{1}{2}\left(\boldsymbol{e}_{i_1}^T + \boldsymbol{e}_{i_2}^T\right)C = \begin{bmatrix} -2 & \underset{\substack{\uparrow \\ j_3 = 2}}{-\frac{1}{2}} & -\frac{3}{2} \end{bmatrix}$

- Row player observes $R\boldsymbol{e}_{j_3} = \begin{bmatrix} 1 \\ 0 \\ 3 \end{bmatrix}$
- Column player observes $\boldsymbol{e}_{i_3}^T C = \begin{bmatrix} -2 & 0 & -3 \end{bmatrix}$

Fictitious play

☐ In round t:

Row player chooses the max row of $\frac{1}{t-1} R \sum_{\tau \leq t-1} \boldsymbol{e}_{j_\tau}$ denoted as i_t.

Column player chooses the max column of $\frac{1}{t-1} \sum_{\tau \leq t-1} \boldsymbol{e}_{j_\tau}^T C$ denoted as j_t.

- Row player observes $R\boldsymbol{e}_{j_t}$.
- Column player observes $\boldsymbol{e}_{i_t}^T$C.

Each player plays a best response to her opponent's historical strategy:

$$\boldsymbol{x}_t = \frac{1}{t} \sum_{\tau \leq t} \boldsymbol{e}_{i_\tau} \ , \ \boldsymbol{y}_t = \frac{1}{t} \sum_{\tau \leq t} \boldsymbol{e}_{j_\tau}$$

Fictitious play

- $\forall t \geq 1 \ \max_i \boldsymbol{e}_i^T R \boldsymbol{y}_t \geq v \geq \min_j \boldsymbol{x}_t^T R \boldsymbol{e}_j$ and the game converges for $t \to \infty$.
  [J.Robinson '50]
- $\forall \epsilon > 0, \forall t \geq \left(\frac{R_{max}}{\epsilon}\right)^{\Omega(n+m)}$ it holds $\left|\max_i \boldsymbol{e}_i^T R \boldsymbol{y}_t - \min_j \boldsymbol{x}_t^T R \boldsymbol{e}_j\right| \leq \epsilon$.

$(\boldsymbol{x}, \boldsymbol{y})$ is an ϵ-approximate N.E. $\Leftrightarrow$

$$\begin{aligned} \boldsymbol{x}^T R \boldsymbol{y} &\geq \boldsymbol{x}'^T R \boldsymbol{y} - \epsilon \quad \forall \boldsymbol{x}' \in \Delta_m \\ \boldsymbol{x}^T C \boldsymbol{y} &\geq \boldsymbol{x}^T C \boldsymbol{y}' - \epsilon \quad \forall \boldsymbol{y}' \in \Delta_n \end{aligned}$$

⬇

$\forall \epsilon > 0, \forall t \geq \left(\frac{R_{max}}{\epsilon}\right)^{\Omega(n+m)}$: $(\boldsymbol{x}_t, \boldsymbol{y}_t)$ is an ϵ-approximate N.E. of the game.

Multiplicative Weights Update Method

- The player has n strategies.
- At each round t, the player chooses a mixed strategy $\boldsymbol{p}^t$ and suffers loss $\boldsymbol{p}^t\boldsymbol{l}^t$ where $\boldsymbol{l}^t \in [0,1]^n$ is determined by nature or an adversary.
- We need to minimize the cumulative loss $L_T = \sum_{t=1}^{\mathrm{T}} \boldsymbol{p}^t\boldsymbol{l}^t$ compared to the cumulative loss of the best fixed strategy k with the $\min_{\mathrm{i}} \sum_{t=1}^{\mathrm{T}} l_i^t = L_T(k)$.

Idea: Maintain weights over the strategies to "remember" their performance so far.

Multiplicative Weights Update Method

- Choose $\delta \in [0,1]$. Initialize weight vector $\boldsymbol{w}^1 = \begin{bmatrix}\frac{1}{n} & \cdots & \frac{1}{n}\end{bmatrix}$.
- At round t:
  - Choose strategy i with probability $p_i^t = \frac{w_i^t}{\sum_{j=1}^{n} w_j^t}$
  - Given the loss vector update the weights:

$$w_i^{t+1} = w_i^t \cdot u_\delta\left(l_i^t\right)$$

where $u_\delta\colon [0,1] \to [0,1]$ satisfies $(1-\delta)^x \leq u_\delta(x) \leq 1 - \delta x \ \forall x \in [0,1]$

Idea: Maintain weights over the strategies to "remember" their performance so far.

Multiplicative Weights Update Method

- ❑ Choose $\delta \in [0,1]$. Initialize weight vector $\boldsymbol{w}^1 = \left[\frac{1}{n} \quad \cdots \quad \frac{1}{n}\right]$.
- ❑ At round t:
  - Choose strategy i with probability $p_i^t = \frac{w_i^t}{\sum_{j=1}^{n} w_j^t}$
  - Given the loss vector update the weights:

$$w_i^{t+1} = w_i^t \cdot u_\delta\left(l_i^t\right)$$

where $u_\delta : [0,1] \rightarrow [0,1]$ satisfies $(1-\delta)^x \leq u_\delta(x) \leq 1 - \delta x \; \forall x \in [0,1]$

For all T and any sequence of loss vectors:

$$L_T \leq \frac{\ln n}{\delta} + \frac{\ln\left(\frac{1}{1-\delta}\right)}{\delta} \min_i L_T(i)$$

Multiplicative Weights Update Method

For all T and any sequence of loss vectors:

$$L_T \leq \frac{\ln n}{\delta} + \frac{\ln\left(\frac{1}{1-\delta}\right)}{\delta} \min_i L_T(i)$$

For $\delta = \min(\frac{1}{2}, \sqrt{\frac{\ln n}{T}})$, the regret of the player $\frac{L_T}{T} - \min_i \frac{L_T(i)}{T} \leq \sqrt{\frac{2 \ln n}{T}}$.

Application to zero-sum game

If $\boldsymbol{p}^t$ and $\boldsymbol{q}^t$ are the mixed strategies played by the players of the game who use the MWU algorithm, then $(\frac{1}{T}\sum_{t\leq T}\boldsymbol{p}^t, \frac{1}{T}\sum_{t\leq T}\boldsymbol{q}^t)$ is a $O\left(\frac{1}{\sqrt{T}}\right)$-approximate N.E.

Nash’s Theorem

- Let there be n players, each player p with a set of strategies S_p (set of mixed strategies denoted by $\Delta_{\mathrm{S_p}}$).
- There is a utility function for every player $u_p: \Delta_{\mathrm{S_1}} \times \cdots \times \Delta_{\mathrm{S_n}} \to \mathbb{R}$.
- A mixed strategy profile $\vec{\boldsymbol{x}} = (\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_n)$ is a N.E. if for every player p:

$$u_p(\vec{\boldsymbol{x}}) \geq u_p(\boldsymbol{x}'_p; \vec{\boldsymbol{x}}_{-p})$$

Nash's Theorem

Nash's Theorem (1950)	Every (finite) game has a Nash Equilibrium.
	↑
Brouwer's Theorem (1911)	Every continuous function from a closed compact convex set to itself has a fixed point.
	↑
Sperner's Lemma (1950)	Every proper coloring of a triangulation has a panchromatic triangle.
	↑
Parity Argument (1990)	If a directed graph has an unbalanced node, then it must have another.

Nash's Theorem

Nash's Theorem (1950) Every (finite) game has a Nash Equilibrium.

↑

Brouwer's Theorem (1911) Every continuous function from a closed compact convex set to itself has a fixed point.

- The set of mixed strategy profiles $\Delta_{S_1} \times \cdots \times \Delta_{S_n} = \Delta$ is closed, compact and convex.
- We need to find a continuous $f: \Delta \to \Delta$.

Take $\vec{y} = f(\vec{x})$ where $\boldsymbol{y}_p(s_p) = \frac{\boldsymbol{x}_p(s_p) + G_{p,s_p}(\vec{\boldsymbol{x}})}{1 + \sum_{s'_p \in S_p} G_{p,s'_p}(\vec{\boldsymbol{x}})}$

and $G_{p,s_p}(\vec{\boldsymbol{x}}) = u_p(s_p; \vec{\boldsymbol{x}}_{-p}) \dot{-} u_p(\vec{\boldsymbol{x}})$ is the gain of player p if she switched to strategy s_p.

Nash's Theorem

Brouwer's Theorem (1911)	Every continuous function from a closed compact convex set to itself has a fixed point.

↑

Sperner's Lemma (1950)	Every proper coloring of a triangulation has a panchromatic triangle.

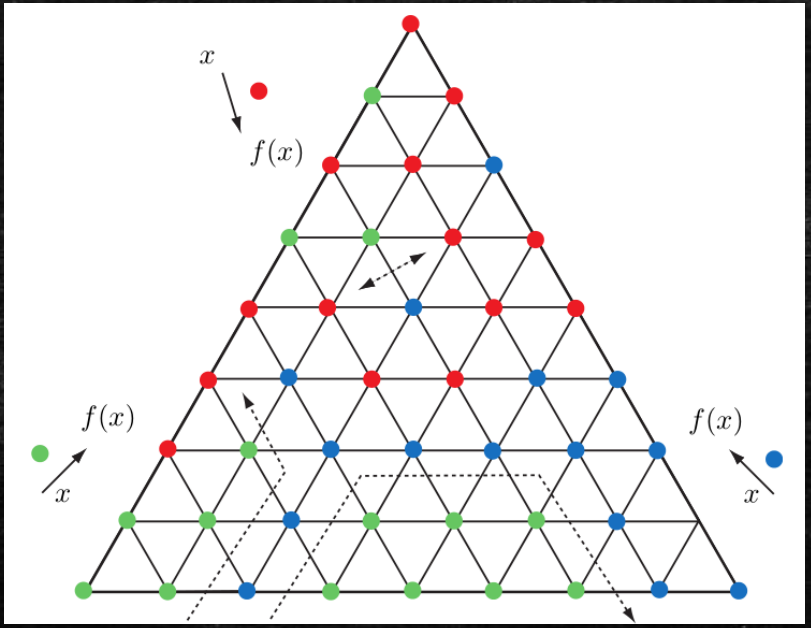

Nash's Theorem

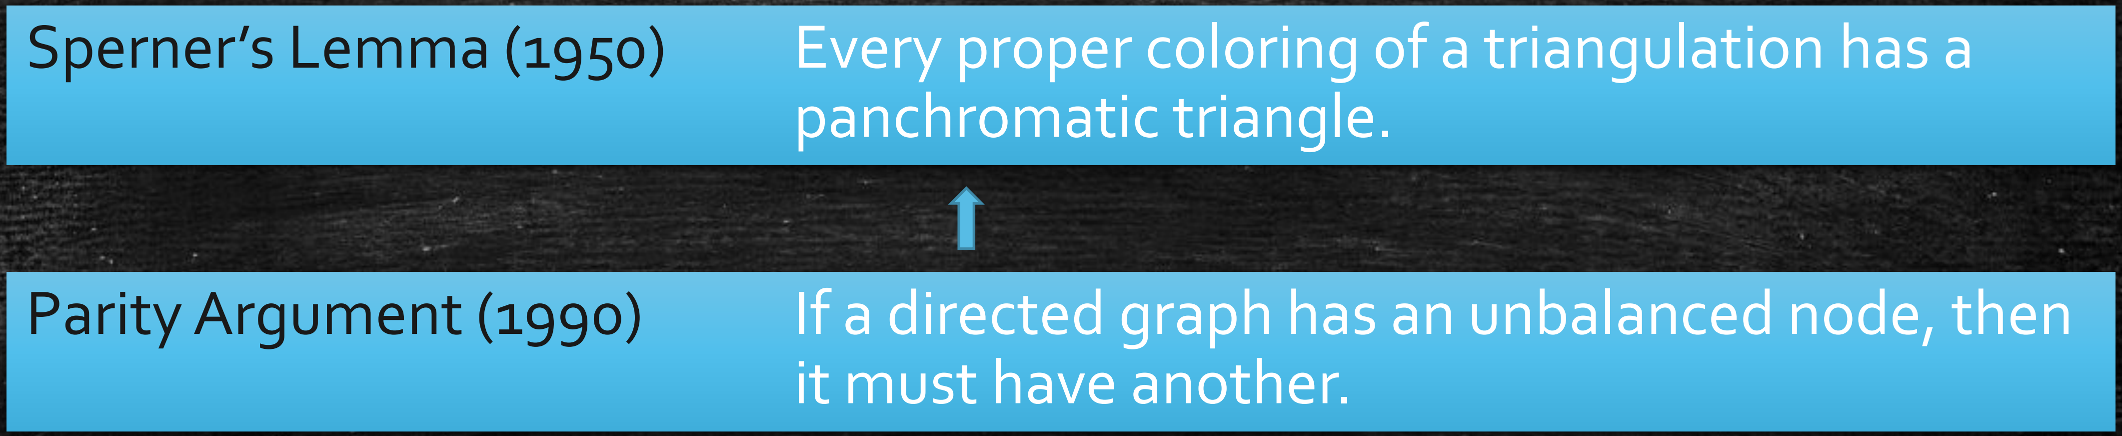

Sperner's Lemma (1950) — Every proper coloring of a triangulation has a panchromatic triangle.

↑

Parity Argument (1990) — If a directed graph has an unbalanced node, then it must have another.