Algorithmic game theory

Martin Balko

7th lecture

November 21st 2022

Regret minimization

Our notation

Our notation

- Agent A with actions $X = \{1, \dots, N\}$ selects a probability distribution $p^t = (p_1^t, \dots, p_N^t)$ at every step $t = 1, \dots, T$.

Our notation

- Agent A with actions $X = \{1, \dots, N\}$ selects a probability distribution $p^t = (p_1^t, \dots, p_N^t)$ at every step $t = 1, \dots, T$.
- After that, the adversary environment selects a loss vector $\ell^t = (\ell_1^t, \dots, \ell_N^t) \in [-1, 1]^N$ at every step $t = 1, \dots, T$.

Our notation

- Agent A with actions $X = \{1, \ldots, N\}$ selects a probability distribution $p^t = (p_1^t, \ldots, p_N^t)$ at every step $t = 1, \ldots, T$.
- After that, the adversary environment selects a loss vector $\ell^t = (\ell_1^t, \ldots, \ell_N^t) \in [-1, 1]^N$ at every step $t = 1, \ldots, T$.
- The agent A receives loss $\ell_A^t = \sum_{i=1}^N p_i^t \ell_i^t$ at step t.

Our notation

- Agent A with actions $X = \{1, \ldots, N\}$ selects a probability distribution $p^t = (p_1^t, \ldots, p_N^t)$ at every step $t = 1, \ldots, T$.
- After that, the adversary environment selects a loss vector $\ell^t = (\ell_1^t, \ldots, \ell_N^t) \in [-1, 1]^N$ at every step $t = 1, \ldots, T$.
- The agent A receives loss $\ell_A^t = \sum_{i=1}^N p_i^t \ell_i^t$ at step t. His cumulative loss after all steps is $L_A^T = \sum_{t=1}^T \ell_A^t$.

Our notation

- Agent A with actions $X = \{1, \dots, N\}$ selects a probability distribution $p^t = (p_1^t, \dots, p_N^t)$ at every step $t = 1, \dots, T$.
- After that, the adversary environment selects a loss vector $\ell^t = (\ell_1^t, \dots, \ell_N^t) \in [-1, 1]^N$ at every step $t = 1, \dots, T$.
- The agent A receives loss $\ell_A^t = \sum_{i=1}^N p_i^t \ell_i^t$ at step t. His cumulative loss after all steps is $L_A^T = \sum_{t=1}^T \ell_A^t$. The cumulative loss of action i is $L_i^T = \sum_{t=1}^T \ell_i^t$.

Our notation

- Agent A with actions $X = \{1, \ldots, N\}$ selects a probability distribution $p^t = (p_1^t, \ldots, p_N^t)$ at every step $t = 1, \ldots, T$.
- After that, the adversary environment selects a loss vector $\ell^t = (\ell_1^t, \ldots, \ell_N^t) \in [-1, 1]^N$ at every step $t = 1, \ldots, T$.
- The agent A receives loss $\ell_A^t = \sum_{i=1}^N p_i^t \ell_i^t$ at step t. His cumulative loss after all steps is $L_A^T = \sum_{t=1}^T \ell_A^t$. The cumulative loss of action i is $L_i^T = \sum_{t=1}^T \ell_i^t$.
- Given a comparison class $\mathcal{A}_X$ of agents A_i that select a single action i in all steps, we let $L_{min}^T = \min_{i \in X}\{L_{A_i}^T\}$ be the minimum cumulative loss of an agent from $\mathcal{A}_X$.

Our notation

- Agent A with actions $X = \{1, \dots, N\}$ selects a probability distribution $p^t = (p_1^t, \dots, p_N^t)$ at every step $t = 1, \dots, T$.
- After that, the adversary environment selects a loss vector $\ell^t = (\ell_1^t, \dots, \ell_N^t) \in [-1, 1]^N$ at every step $t = 1, \dots, T$.
- The agent A receives loss $\ell_A^t = \sum_{i=1}^N p_i^t \ell_i^t$ at step t. His cumulative loss after all steps is $L_A^T = \sum_{t=1}^T \ell_A^t$. The cumulative loss of action i is $L_i^T = \sum_{t=1}^T \ell_i^t$.
- Given a comparison class $\mathcal{A}_X$ of agents A_i that select a single action i in all steps, we let $L_{min}^T = \min_{i \in X} \{L_{A_i}^T\}$ be the minimum cumulative loss of an agent from $\mathcal{A}_X$.
- Our goal is to minimize the external regret $R_A^T = L_A^T - L_{min}^T$.

Our notation

- Agent A with actions $X = \{1, \dots, N\}$ selects a probability distribution $p^t = (p_1^t, \dots, p_N^t)$ at every step $t = 1, \dots, T$.
- After that, the adversary environment selects a loss vector $\ell^t = (\ell_1^t, \dots, \ell_N^t) \in [-1, 1]^N$ at every step $t = 1, \dots, T$.
- The agent A receives loss $\ell_A^t = \sum_{i=1}^N p_i^t \ell_i^t$ at step t. His cumulative loss after all steps is $L_A^T = \sum_{t=1}^T \ell_A^t$. The cumulative loss of action i is $L_i^T = \sum_{t=1}^T \ell_i^t$.
- Given a comparison class $\mathcal{A}_X$ of agents A_i that select a single action i in all steps, we let $L_{min}^T = \min_{i \in X}\{L_{A_i}^T\}$ be the minimum cumulative loss of an agent from $\mathcal{A}_X$.
- Our goal is to minimize the external regret $R_A^T = L_A^T - L_{min}^T$.
- We now restrict ourselves to losses from $\{0, 1\}$.

Example

Example

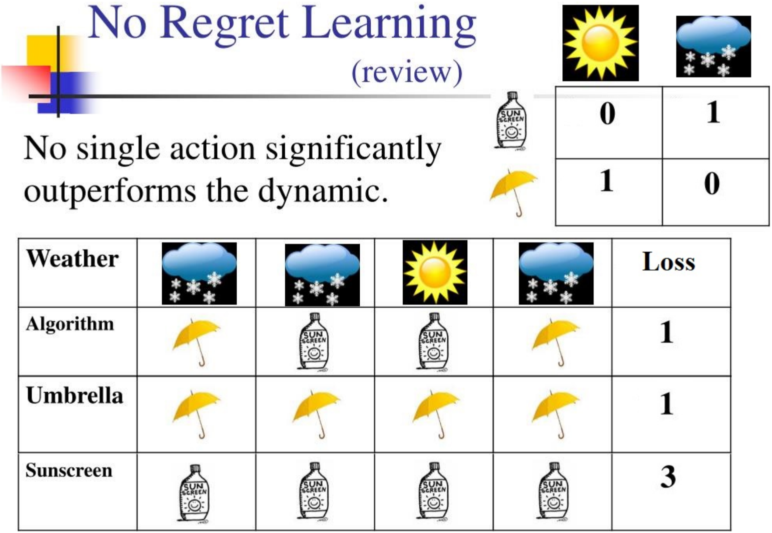

Source: No regret algorithms in games (Georgios Piliouras)

The Randomized greedy algorithm

The Randomized greedy algorithm

Algorithm 0.2: RANDOMIZED GREEDY ALGORITHM(X, T)

Input : A set of actions $X = \{1, \dots, N\}$ and a number of steps T.
Output : A probability distribution p^t for every $t \in \{1, \dots, T\}$.
$p^1 \leftarrow (1/N, \dots, 1/N)$,
for $t = 2, \dots, T$

do $\begin{cases} L_{min}^{t-1} \leftarrow \min_{j \in X}\{L_j^{t-1}\}, \\ S^{t-1} \leftarrow \{i \in X : L_i^{t-1} = L_{min}^{t-1}\}, \\ p_i^t \leftarrow 1/|S^{t-1}| \text{ for every } i \in S^{t-1} \text{ and } p_i^t \leftarrow 0 \text{ otherwise.} \end{cases}$

Output $\{p^t : t \in \{1, \dots, T\}\}$.

The Randomized greedy algorithm

Algorithm 0.3: RANDOMIZED GREEDY ALGORITHM(X, T)

Input : A set of actions $X = \{1, \dots, N\}$ and a number of steps T.
Output : A probability distribution p^t for every $t \in \{1, \dots, T\}$.
$p^1 \leftarrow (1/N, \dots, 1/N)$,
for $t = 2, \dots, T$

do $\begin{cases} L_{min}^{t-1} \leftarrow \min_{j \in X}\{L_j^{t-1}\}, \\ S^{t-1} \leftarrow \{i \in X \colon L_i^{t-1} = L_{min}^{t-1}\}, \\ p_i^t \leftarrow 1/|S^{t-1}| \text{ for every } i \in S^{t-1} \text{ and } p_i^t \leftarrow 0 \text{ otherwise.} \end{cases}$

Output $\{p^t \colon t \in \{1, \dots, T\}\}$.

Proposition 2.48

For any sequence of $\{0, 1\}$-valued loss vectors, we have

$$L_{\mathrm{RG}}^T \leq (1 + \ln N) \cdot L_{min}^T + \ln N.$$


WHEN YOU FIND THE NASH EQUILIBRIUM
NO RAGRETS
memegenerator.net

Thank you for your attention.