Algorithmic game theory

Martin Balko

7th lecture

November 19th 2020

Regret minimization

Our notation

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 $t = 1, \ldots, T$.

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 $t = 1, \ldots, T$.

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 $X = \{1, \dots, N\}$ 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 $L_A^T = \sum_{t=1}^T \ell_A^t$.

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 $L_i^T = \sum_{t=1}^T \ell_i^t$.

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 $X = \{1, \ldots, N\}$ 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 goal is to minimize the external regret $R_A^T = L_A^T - L_{min}^T$.

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 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

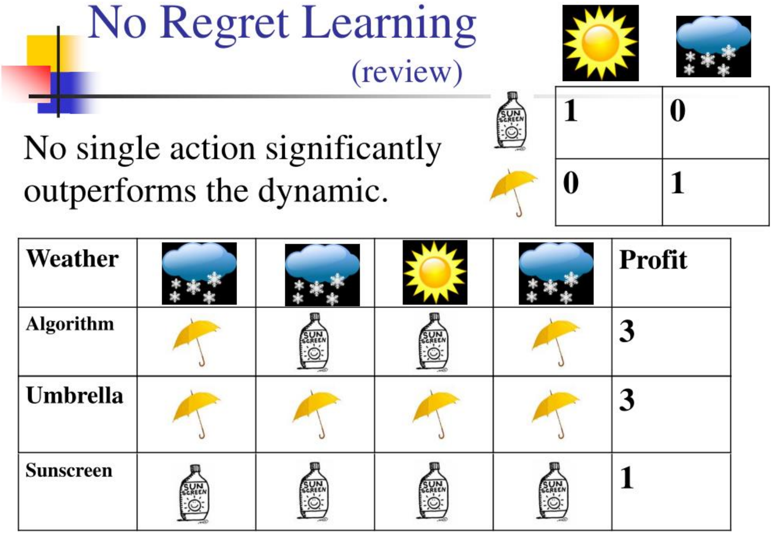

Source: No regret algorithms in games (Georgios Piliouras)

- Besides game theory, the "multiplicative weight update method" has many applications in various fields of science, for example in optimization, theoretical computer science, and machine learning.

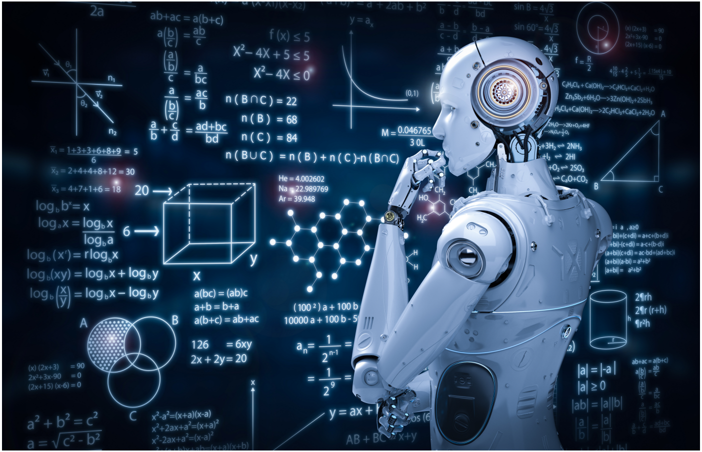

Sources: https://clubitc.ro

- Besides game theory, the “multiplicative weight update method” has many applications in various fields of science, for example in optimization, theoretical computer science, and machine learning.

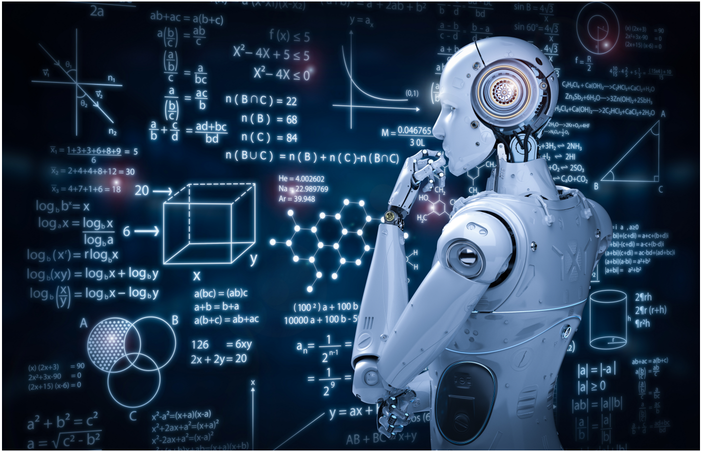

Sources: https://clubitc.ro

- See https://en.wikipedia.org/wiki/Multiplicative_weight_update_method#Applications

- Besides game theory, the “multiplicative weight update method” has many applications in various fields of science, for example in optimization, theoretical computer science, and machine learning.

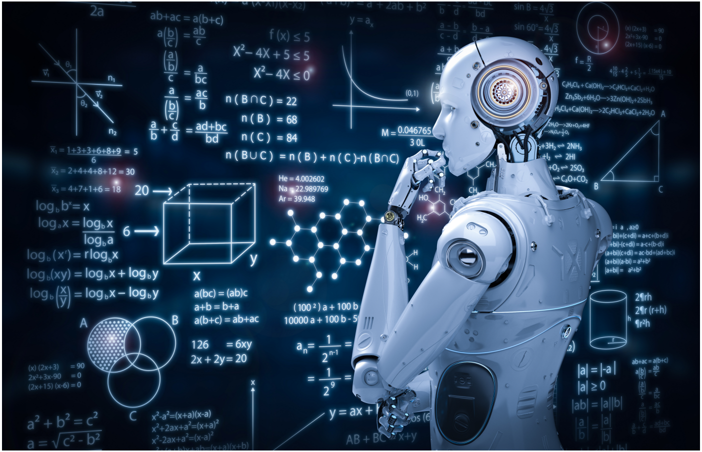

Sources: https://clubitc.ro

- See https://en.wikipedia.org/wiki/Multiplicative_weight_update_method#Applications

Thank you for your attention.