Regret minimization

Set of actions $X = \{1 \ldots N\}$, Number of steps X

In every step $t \in T$:

Agent A chooses prob. distr. $p^t = (p_1^t \ldots p_N^t)$ on X,
where $p_i^t = P_r[A \text{ select action } i]$

Adversary then gives a loss vector $\ell^t = (\ell_1^t \ldots \ell_N^t) \in [-1,1]^N$
wher ℓ_i^t = loss for selecting $i \in X$
A experiences expected loss $\ell_A^t = \sum_i^N p_i^t \cdot \ell_i^t$

Cumulative loss $L_A^T = \sum_{t=1}^T \ell_A^t$, $\left(i \in X : L_i^T = \sum_{t=1}^T \ell_i^t \right)$

Comparision class α = set of agents

$L_{\alpha,min}^T = \min_{B \in \alpha} L_B^T$ = cumulative loss of best agent from α

we want to minimize extern regret $R_{A,\alpha}^T = L_A^T - L_{\alpha,min}^T$

typically, we consider $\alpha_X = \{$ agents that always play the same action with prob $1 \}$

$$\left.\begin{aligned} R_A^T &= R_{A,\alpha_X}^T \\ L_{min}^T &= L_{\alpha_X,min}^T \end{aligned}\right\} \text{ for simplification}$$

$\{$agents that select actions with prob $1\}$
– the action may change each iteration !

Thm: $\forall$ agent A, $\forall T \in \mathbb{N}$: $\exists$ loss vectors $R_{A,\alpha_{ALL}}^T \geq T(1-\frac{1}{N})$

Bw: $\forall t$, by pigeonhole principle $\Rightarrow \exists i_t \in X : p_{i_t}^t \leq \frac{1}{N}$

Adversary selects $\ell_{i_t} = 0$, $\ell_i = 1 \quad i \neq i_t \Rightarrow L_A^T \geq T(1-\frac{1}{N})$
$\exists B \in \alpha_{ALL}$, who plays $i_1 \ldots i_t \Rightarrow L_B^T = 0$ $\boxtimes$

Greedy algorithm

– lets play something with the lowest cumulative loss.

Input: $T \in \mathbb{N}$, $X = \{1 - N\}$

Output: p^t, $t = 1 \_ T$

$p^1 = (1, 0, 0, \dots, 0)$ ⟶ random initial state

for $t = 2, \dots, T$:

$L^{t-1}_{MIN} = \min_{j \in X} L^{t-1}_j$ (best cumulative loss of action so far)

$S^{t-1} = \{ i \in X : L^{t-1}_i = L^{t-1}_{min} \}$ (best actions so far)

$h = \min S^{t-1}$

$p^t_h = 1$, $p^t_j = 0$ for $j \neq h$ ⟶ this is the bad part

– I experience loss $\leq N$.

Thm: $L^T_{greedy} \leq N \cdot L^T_{min} + N - 1$

everytime it increases

Dk: If loss experienced, then $|S^{t-1}|$ decreases.

This happens $\leq N$ times, before L^t_{min} increases.

$\Rightarrow L^T_{greedy} \leq N \cdot L^T_{min} + N - |S^T| \leq N \cdot L^T_{min} + N - 1$ $\boxtimes$

It is actually deterministic.

Prop: $\forall$ deterministic D, $T \in \mathbb{N}$: $\exists$ loss vectors : $L^T_D = T$, $L^T_{min} = \lfloor \frac{T}{N} \rfloor$

Dk: D chooses some $i_t \in X$ at step t with prob 1. Set $\ell^t_{i_t} = 1$, $\ell^t_i = 0$ $i \neq i_t$

$\Rightarrow L^T_D = T$. Pigeonhole principle $\Rightarrow \exists j \in X$ that equals i_t in $\leq \lfloor \frac{T}{N} \rfloor$ steps.

$\Rightarrow L^T_{min} \leq L^t_j \leq \lfloor \frac{T}{N} \rfloor$

Random algorithm

Input: $T \in \mathbb{N}$, $X = \{1 \ldots N\}$

Output: p^t, $t = 1 \ldots T$

$p^1 = \cancel{(1, 0, 0, \ldots, 0)} = (1/N, 1/N, \ldots, 1/N)$

for $t = 2, \ldots, T$:

$L_{MIN}^{t-1} = \min_{j \in X} L_j^{t-1}$ (best cumulative loss of action so far)

$S^{t-1} = \{ i \in X : L_i^{t-1} = L_{min}^{t-1} \}$ (best actions so far)

$\cancel{h = \min S^{t-1}}$

$\cancel{p_h^t = 1,\ p_j^t = 0 \text{ for } j \neq h}$

$$p_i^t = \begin{cases} 1/|S^{t-1}| & \text{if } i \in S^{t-1} \\ 0 & \text{else} \end{cases}$$

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

Pf: for $j \in \mathbb{N}$ let t_j = first step when L_{min}^t reaches j. Consider $t \in (t_j, t_{j+1}]$.

If $|S^t|$ shrinks from n' to $n' - k$, then we experience loss $\frac{k}{n'} \leq \frac{1}{n'} + \frac{1}{n'-1} + \ldots + \frac{1}{n'-k+1}$

$\Rightarrow$ before increase of L_{min}^t, we get loss of $\leq \frac{1}{1} + \frac{1}{2} + \ldots + \frac{1}{N} \leq 1 + \log N$

$\Rightarrow L_{RG}^T \leq (1 + \log N) \cdot L_{min}^T + \underbrace{\frac{1}{N} + \ldots + \frac{1}{|S^t| + 1}}_{\leq \log N}$

Polynomial weights alg.

Input: $T \in \mathbb{N}$, $X = \{1, \dots, N\}$ $\eta \in (0, \frac{1}{2}]$ → very simple, yet very powerful

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

$w_i^1 = 1$ for $i \in X$

$p^1 = (1/N, \dots, 1/N)$

for $t = 2 \dots T$:

$$w_i^t = w_i^{t-1}(1 - \eta \ell_i^{t-1})$$ → the bigger loss, the smaller number

$$W^t = \sum_i^N w_i^t$$

$$p_i^t = \frac{w_i^t}{W^t} \text{ for } \forall i \in X$$

Thm: $\forall \eta \in (0, \frac{1}{2}]$ $\forall$ loss vectors from $[-1, 1]^N$, $\forall k \in X$:

$$L_{pw}^T \le L_k^T + \eta Q_k^T + \log N / \eta \qquad \text{where } Q_k^T = \sum_{t=1}^T (\ell_k^t)^2$$

If $T \ge 4 \log N$, then choosing $\eta = \sqrt{\log N / T}$ gives

$$L_{pw}^T \le L_{MIN}^T + 2 \cdot \sqrt{T \cdot \log N}$$

→ the average R_A^T goes to 0.

Dk:

at t: $\ell_{pw}^t = \sum_{i=1}^N \frac{w_i^t}{W^t} \ell_i^t$ $\qquad W^1 = N$

$$W^{t+1} = W^t - \sum_{i=1}^N \eta w_i^t \ell_i^t = W^t(1 - \eta \ell_{pw}^t)$$

$w_i^{t+1} = w_i^t \cdot (1 - \eta \ell_i^t)$

$1 - z \le e^{-z}$

$$\Rightarrow W^{T+1} = N \cdot \prod_{t=1}^T (1 - \eta \cdot \ell_{pw}^t) \le N \cdot \prod_{t=1}^T e^{-\eta \ell_{pw}^t} = N \cdot e^{-\eta \sum_t \ell_{pw}^t}$$

$$\Rightarrow \log W^{T+1} \le \log N - \eta \sum_{t=1}^T \ell_{pw}^t = \log N - L_{pw}^T$$ → upperbound done.

$$W^{T+1} \ge w_k^{T+1} = \prod_{t=1}^T (1 - \eta \ell_k^t) \Rightarrow \log W^{T+1} \ge \sum_{t=1}^T \log(1 - \eta \ell_k^t) \ge$$

$\log(1-z) \ge -z - z^2$ if $z \le \frac{1}{2}$

$$\ge -\underbrace{\sum_{t=1}^T \eta \ell_k^T}_{-\eta L_k^T} - \underbrace{\sum_{t=1}^T \eta^2 (\ell_k^t)^2}_{-\eta^2 Q_k^T} \Rightarrow \eta L_k^T - \eta^2 Q_k^T \le \log W^{T+1} \le \log N - \eta L_{pw}^T \Rightarrow$$

$$L_{pw}^T \le L_k^T + \eta Q_k^T + \frac{\log N}{\eta}$$

$\boxtimes$