Online Algorithms pt. 2

Experts problem

There are n experts $E_1, \dots, E_n$ and T days.
On each day $t = 1 \dots T$:

1. Each expert makes a prediction (e.g. Yes/No)
2. You make a prediction (e.g. YES/No)
3. The real outcome is revealed

Goal: At time T, (# of mistakes you make) $\leq$ (# of mistakes of best experts) + (small)

Last time:

Special case 1: $\exists$ one perfect (0 mistake) expert

We showed: alg which only makes $\log_2(n)$ mistakes

Special case 2: $\exists$ an expert who makes $\leq M$ mistakes

. . .

Special case 2

Promise: one of the n experts makes $\leq M$ mistakes

Alg: For day $t = 1 \ldots T$

1. Let m^t_i = # mistakes i-th expert has made so far
2. Set $\text{Plausible}_t \subseteq [n]$ to be the experts who have made $\leq M$ mistakes so far
3. Predict majority vote of Plausible_t experts

Thm: This Alg can make $\geq (M+1) \cdot \log(n)$ mistakes

Pf:

Experts / Days	$\underbrace{1 \quad 2 \quad \ldots}_{\text{half}+1}$	$\underbrace{\quad\quad}_{\text{half}+1}$	n	Alg	Real outcome
1	Y Y ... Y			Y	N
⋮				Y	N
M+1	Y Y ... Y			Y	N
M+2		Y Y ... Y		Y	N
⋮				Y	N
2M+2		Y Y ... Y	⋮	Y	N

Can continue for $(M+1)\log_2(N)$ rounds ∎

Goal: (# of mistakes) $\leq$ (best expert) + (small)

Weighted majority alg (given parameter $0<\varepsilon<1$)

Initialize weights $w_1^{(1)}, \dots, w_n^{(1)} = 1$

For $t = 1, \dots, T$:

1. Output weighted majority of experts according to $w^{(t)}$ weights.
2. Set $w_i^{(t+1)} = (1-\varepsilon)\cdot w_i^{(t)}$ if i-th expert erred.

   Set $w_i^{(t+1)} = w_i^{(t)}$ otherwise.

Example w/ $\varepsilon = 1/2$:

		Experts 1	2	3	4	Alg	Real
Day 1:	Weights	1	1	1	1		
	Predictions	Y	Y	Y	N	Y	N
Day 2:	W	1/2	1/2	1/2	1		
	P	Y	N	N	Y	Y	N
Day 3:	W	1/4	1/2	1/2	1/2		

Suppose we run WM alg w/ $\varepsilon = 1/2$.

Fact 1: Let $W^{(t)}$ = total weight of experts at time t.
Suppose WM made a mistake at time t.
Then $W^{(t+1)} \le \frac{3}{4} \cdot W^{(t)}$. (Halve $\geq$ 1/2 the weight)

Fact 2: Suppose expert i makes m_i mistakes.
Then $w_i^{(T+1)} \le (1/2)^{m_i}$.

Thm: Let M = # mistakes WM alg makes w/ $\varepsilon = 1/2$.
Then $M \le 2.4 \cdot (OPT + \log_2(n))$.

Pf: Let i^* = best expert.
Then $w_{i^*}^{(T+1)} = (1/2)^{OPT}$. So $W^{(t+1)} \geq (1/2)^{OPT}$.

Alg makes M mistakes $\Rightarrow W^{(t+1)} \le \left(\frac{3}{4}\right)^M \cdot W^{(1)} = \left(\frac{3}{4}\right)^M \cdot n$

$$\therefore \left(\frac{1}{2}\right)^{OPT} \le \left(\frac{3}{4}\right)^M \cdot n \Rightarrow \underbrace{\left(\frac{4}{3}\right)^M}_{2^{\log_2(4/3) \cdot M}} \le 2^{OPT} \cdot n \Rightarrow M \le 2.4 \cdot (OPT + \log_2(n)). \qquad \blacksquare$$

Can't beat 2

E1	E2	Alg = follow best expert so far (tie = E1)	Real = opposite of Alg
1	0	1	0
1	0	0	1
1	0	1	0
1	0	0	1
⋮			
1	0	1	0

\# of Alg mistakes = T
\# of mistakes of best expert = T/2

$\therefore$ # of Alg mistakes $\geq 2 \cdot OPT$

The fix: Let Alg be randomized

Alg's prediction is probability distribution (p_1, p_2).
Alg follows expert E_i w/prob p_i.

Randomized weighted majority (w/ parameter ε)

Initialize weights $w_1^{(1)}, \dots, w_n^{(1)} = 1$

For $t = 1, \dots, T$:

1. Set $W^{(t)} = w_1^{(t)} + \dots + w_n^{(t)}$
2. Guess $b \in \{0,1\}$ w/ prob $\frac{1}{W^{(t)}} \cdot \sum_{i=1}^{n} w_i^{(n)} \cdot \mathbb{1}[E_i \text{ guesses } b \text{ today}]$
3. Set $w_i^{(t+1)} = (1-\varepsilon) \cdot w_i^{(t)}$ if i-th expert erred.

   Set $w_i^{(t+1)} = w_i^{(t)}$ otherwise.

$$\mathrm{OPT} + \varepsilon \cdot \underbrace{\cancel{\mathrm{OPT}}}_{}^{\le T} + \frac{1}{\varepsilon} \cdot \log_2(n)$$

Thm: Set $M = \#$ mistakes. Then $\overbrace{E[M] \le (1+\varepsilon) \cdot \mathrm{OPT} + \frac{1}{\varepsilon} \log_2(n)}$

Set $\varepsilon = \sqrt{\frac{\log_2(n)}{T}}$. Then $E[M] \le \mathrm{OPT} + 2\sqrt{T \log_2(n)}$

So $E[M]/T \le \mathrm{OPT}/T + \boxed{2\sqrt{\log_2(n)/T}} \xrightarrow{T \to \infty} 0$

"no regret algorithm"

General setting

There are n actions you can take each step (corresponding to n experts)

On day t, you play action $i_t \in [n]$,

action i has cost $c_i^{(t)} \in [0,1]$,

Goal: (Alg's cost) $\leq$ (OPT action) + small

$$\sum_{t=1}^{T} c_{i_t}^{(t)} \leq \min_{i^*} \left\{ \sum_{t=1}^{T} c_{i^*}^{(t)} \right\} + \boxed{\text{small}} \longrightarrow \text{Your "regret"}$$

Multiplicative Weights Algorithm

Initialize weights $w_1^{(1)}, \ldots, w_n^{(1)} = 1$

For $t = 1, \ldots, T$

1. Set $W^{(t)} = w_1^{(t)} + \cdots + w_n^{(t)}$
2. Play action i w/prob $w_i^{(t)} / W^{(t)}$
3. Observe costs $c_1^{(t)}, \ldots, c_n^{(t)}$
4. Set $w_i^{(t+1)} = w_i^{(t)} \cdot (1 - \varepsilon \cdot c_i^{(t)})$

Thm: Set $\varepsilon = \sqrt{\log_2(n)/T}$.

Then $E[\text{Regret}] \leq O(\sqrt{T \log_2(n)})$

$\therefore$ $E[\text{Regret}]/T \leq O(\sqrt{\log_2(n)/T})$ "no regret"

Zero sum games

Let $M = \begin{bmatrix} M_{11} & & M_{1n} \\ \vdots & \ddots & \vdots \\ M_{n1} & & M_{nn} \end{bmatrix}$ be an $n \times n$ payoff matrix, with $-1 \le M_{ij} \le 1$.

Given $p = \begin{bmatrix} p_1 \\ \vdots \\ p_n \end{bmatrix}$ and $q = \begin{bmatrix} q_1 \\ \vdots \\ q_n \end{bmatrix}$, $\text{Score}(p,q) = \sum_{i,j} p_i M_{ij} q_j = p^T M q$.

Minimax Thm: $\min_q \max_p p^T M q = \max_p \min_q p^T M q$

(col player 1st) (row player 1st)

Easy direction: $\geq$

Hard direction: $\leq$ (via strong duality of LPs)

Today: proof using multiplicative weights

Minimax proof

$$\begin{array}{ccc} \text{Col player's actions} & & \\ \text{Col 1} & \cdots & \text{Col } n \\ \downarrow & & \downarrow \end{array}$$
$$\begin{bmatrix} & & \\ & \text{Payoff Matrix } M & \\ & & \end{bmatrix}$$

For $t = 1, \ldots, T$ rounds:

1. Col player selects mixed strat $q^{(t)} = (q_1^{(t)}, \ldots, q_n^{(t)})$ using multiplicative weights
2. Row player "best responds" $p^{(t)} = \text{argmax}\{(p^{(t)})^T M q^{(t)}\}$
3. Cost vector $c^{(t)} = (c_1^{(t)}, \ldots, c_n^{(t)}) = p^{(t)} \cdot M$ for col player is revealed.
4. Col player's expected loss $= p^{(t)} M q^{(t)}$.

Def: $\bar{p} = \sum_{t=1}^{T} p^{(t)}$ $\quad \bar{q} = \sum_{t=1}^{T} q^{(t)}$

Goal: $\bar{p}$ and $\bar{q}$ are "approx minimax strats":

$$\max_p p M \bar{q} \leq \min_q \bar{p} M q + \text{small}$$

1. Col player picks $q^{(t)}$. 2. Row player sets $p^{(t)} = \max\{pMq^{(t)}\}$. 3. Cost $= p^{(t)}M$.

Fact: $$\max_p pM\bar{q} \le \min_q \bar{p}Mq + O\left(\sqrt{\frac{\log(n)}{T}}\right)$$

Pf: $$\begin{aligned} \max_p pM\bar{q} &= \max_p \frac{1}{T}\sum_{t=1}^{T} pMq^{(t)} \\ &\le \frac{1}{T}\sum_{t=1}^{T} \max_p pMq^{(t)} \\ &= \frac{1}{T}\sum_{t=1}^{T} \boxed{p^{(t)}Mq^{(t)}} \end{aligned}$$

= Alg's round t expected loss

$$\begin{aligned} &\le \min_{i\in[n]} \frac{1}{T}\sum_{t=1}^{T} (p^{(t)}M)_i + O\left(\sqrt{\frac{\log(n)}{T}}\right). \\ &= \min_{i\in[n]} (\bar{p}M)_i + O\left(\sqrt{\frac{\log(n)}{T}}\right) \\ &= \min_q \bar{p}Mq + O\left(\sqrt{\frac{\log(n)}{T}}\right). \quad \square \end{aligned}$$

Pf of minimax: $$\begin{aligned} \min_q \max_p pMq &\le \max_p pM\bar{q} \\ &\le \min_q \bar{p}Mq + O\left(\sqrt{\frac{\log(n)}{T}}\right) \\ &\le \max_p \min_q pMq + \boxed{O\left(\sqrt{\frac{\log(n)}{T}}\right)} \xrightarrow{T\to 0} 0 \quad \square \end{aligned}$$