Lecture 20: Online Learning/Prediction Part 2: Randomization & Generalizations

Recall: Game: Each day $t = 1, 2, \ldots, T$:

- experts $1 \ldots N$ predict $e_1^t, \ldots, e_N^t \in \{0,1\}$
- you take action $a^t \in \{0,1\}$
- outcome o^t occurs

Mistake: $a^t \neq o^t$ (you)
$e_i^t \neq o^t$ (expert)

"Weighted Majority Alg":

- Initialize weights $w_1^1, w_2^1, \ldots, w_N^1 = 1$.
- a^t = weighted majo of $e_1^t, \ldots, e_N^t$ with w_i^t weights
- $w_i^{t+1} = \begin{cases} w_i^t & \text{if } e_i^t = o^t \\ (1-\varepsilon) w_i^t & \text{if } e_i^t \neq o^t. \end{cases}$

(our e.g. was $\varepsilon = .02 \Leftrightarrow \frac{\varepsilon}{2} = .01$)

Analysis: $\#\text{Your Mistakes} \le \underbrace{(2 + \tfrac{\varepsilon}{2})}_{\text{ish}}(\#\text{Best expert's mistakes}) + \underbrace{\tfrac{2}{\varepsilon}}_{\text{ish}} \ln N$

Sad :(

Fact: For any deterministic strat, can't beat factor 2.
Because if

e_1	e_2
0	1
1	0
0	1
1	0
...	...

, you do some fixed

a
a^1
a^2
a^3
a^4
...

, and o might be

o
$\neg a^1$
$\neg a^2$
$\neg a^3$
$\neg a^4$
...

rem: If you do W.M. alg, your a^t's alternate, o^t's alternate, experts $\approx$ "equally bad", $w_1^t \approx w_2^t \ \forall t$.

idea: Randomize! [If it's a "close call" on what to predict flip a coin!]

Random model: [[As usual, if the "outcome-revealing adversary" can see what action you play, it's just like the deterministic case.]]

[[It's like the same setup as the paging problem...]]

"Adversary" picking outcomes can: • know your randomized strategy
• see experts predictions

can't: • see algorithm's "coin flips"
• see algorithm's actions.

[[actually: it'll be okay if Adv. can see $a^1, \dots, a^{t-1}$ before deciding on o^t, but never mind for now]]

So can: • fix randomized alg
• then look at worst possible table

e_1	e_2	e_3	a	o
0	1	1		0
1	0	1		1
1	1	1		0

• then compute $\mathbb{E}[\#\text{ your mistakes}]$

[[What's most natural randomized alg?]]

Randomized Weighted Majority

At time t, if total weight is $\Phi^t = w_1^t + w_2^t + \cdots + w_N^t$,
say total weight of experts predicting 1 is $q \cdot \Phi^t$,
" " " " " 0 is $(1-q) \cdot \Phi^t$.

Then choose $a^t = \begin{cases} 1 & \text{with prob } q \\ 0 & \text{" " } 1-q. \end{cases}$ // old rule was 1 if $q \geq \frac{1}{2}$, 0 if $q < 1/2$.

Still penalize wrong experts' weight by $1-\varepsilon$ factor.

Analysis: Fix all expert predictions & outcomes [[for all time]] [[worst case]].

Play Randomized Weighted Maj, define random variables

$$M_t = \begin{cases} 1 & \text{if } a^t \neq o^t \quad \text{[[mistake by you at time } t\text{]]} \\ 0 & \text{if } a^t = o^t \end{cases}$$

$t = 1 \dots T$.

$$\mathbb{E}[\text{your mistakes}] = \mathbb{E}[M_1 + M_2 + \cdots + M_T]$$
$$= \mathbb{E}[M_1] + \cdots + \mathbb{E}[M_T] \quad [\text{linearity of expectation}]$$
$$= f_1 + \cdots + f_T, \quad f_t = \Pr[\text{you mistake at time } t]$$

[But f_t has another interpretation...]

[By def. of the alg.] f_t = fraction of ~~weight~~ Φ^t coming from experts making mistake on day t.

e.g.: $o^t = 0$

Φ^t goes down by $\varepsilon \cdot f_t \cdot \Phi^t$

their trust is down-weighted.

$$\therefore \boxed{\Phi^{t+1} \le (1 - \varepsilon \cdot f_t)\,\Phi^t}$$

[in det. case, just used $f_t \ge \frac{1}{2}$ whenever you make mistake]

$1 - \varepsilon \cdot f_t \approx e^{-\varepsilon \cdot f_t}$

$\because 1 + x \approx e^x$

In fact, $1 - x \le e^{-x} \; \forall x$

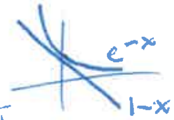

Still have

$$\Phi^{T+1} \ge w_{i^*}^T = (1-\varepsilon)^{\#BM} \leftarrow \#\text{ mistakes of best expert } i^*$$

$$\boxed{\Rightarrow \Phi^{t+1} \le e^{-\varepsilon \cdot f_t}\,\Phi^t} \quad (\Phi^1 = N)$$

$$\therefore \Phi^{(T+1)} \le N \cdot e^{-\varepsilon f_1} e^{-\varepsilon f_2} \cdot \cdots \cdot e^{-\varepsilon f_T}$$
$$= N e^{-\varepsilon \underbrace{(f_1 + \cdots + f_T)}_{\mathbb{E}[\text{your mistakes}]}} \; !$$

$$\therefore \; N \cdot e^{-\varepsilon \cdot \mathbb{E}[\#YM]} \ge (1-\varepsilon)^{\#BM}$$

[take ln]

$$\ln N - \varepsilon \cdot \mathbb{E}[\#YM] \ge (\#BM)(\ln(1-\varepsilon))$$
$$\approx -(\#BM)(\varepsilon + \tfrac{1}{2}\varepsilon^2)$$

calculus: $\ln(1-\varepsilon) = -\varepsilon - \frac{1}{2}\varepsilon^2 - \ldots$

$$\Rightarrow \quad \varepsilon\,\mathbb{E}[\#YM] \lesssim (\varepsilon + \tfrac{1}{2}\varepsilon^2) \cdot \#BM + \ln N$$

$$\Rightarrow \quad \mathbb{E}[\#YM] \le \left(1 + \frac{\varepsilon}{2}\right) \cdot (\#BM) + \frac{\ln N}{\varepsilon}$$

[Ratio is $1 + \frac{\varepsilon}{2}$!]

[Now want to generalize the setup somewhat.
Once we get a good alg. in generalized setup, we'll see it can be used to...]

- solve flows, 2-player zero-sum games, LPs....

R.W.M.'s strategy is equivalent to...

- at time t, let $p_i^t = \frac{w_i^t}{\Phi^t}$ $\forall i = 1 \ldots N$

  these form a "prob. distribution" [they sum to 1]

- play expert i's prediction with prob. p_i^t

~~Could imagine Adversary's strat. is... pick an expert i, do their prediction or its opposite.~~
~~[a weird way to look at things, granted]~~
~~[could also let adversary~~

New game:

- Expert i is like a slot machine
- Alg. chooses $p_1^t, \ldots, p_N^t$ at time t: interpretation is "play slot machine i with prob. p_i".
- "Adversary" sets loss/reward $\ell_1^t, \ldots, \ell_N^t$ between $-1, +1$
- "Loss" at time t is $p^t \cdot \ell^t = p_1^t \ell_1^t + \cdots + p_N^t \ell_N^t$

Goal: small loss compared with best (least) value of $\sum_{t=1}^{T} \ell_{i^*}^t$ over $1 \le i^* \le N$

[What you'd suffer if you just did 100% on expert/slot machine i^* every time.]

"Hedge" Alg., AKA "Multiplicative Weights":

- Same as R.W.M., but ~~every weight~~ every weight changed each round:
  ~~"punished" [illegible] each round~~

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

$1-\varepsilon$ when i^{th} "loss" is $+1$, but $1+\varepsilon$ when "loss" is -1.

[No change if loss is 0.]

Analysis: [[Very similar to R.W.M.]] 5

$$\Phi^{t+1} = \sum_{i=1}^{N} w_i^{t+1} = \sum_{i=1}^{N} (1 - l_i^t \varepsilon) w_i^t$$

$$= \sum_i w_i^t - \varepsilon \sum_i l_i^t w_i^t$$

$$= \Phi^t - \varepsilon \sum_i l_i^t p_i^t \Phi^t$$

$$= \Big(1 - \varepsilon \underbrace{\sum_i l_i^t \cdot p_i^t}_{\text{your loss @ time } t: \text{ "} YL_t \text{"}}\Big) \Phi^t$$

$(1 - x \le e^{-x} \; \forall x)$

$$\le e^{-YL_t \cdot \varepsilon} \Phi^t$$

$$\therefore \; \Phi^{(T+1)} \le \underset{\hat{\Phi}^1}{N} \cdot e^{-\varepsilon YL_1} e^{-\varepsilon YL_2} \cdots e^{-\varepsilon YL_T} = N e^{-\varepsilon (YL)} \quad ①$$

total loss.

OTOH: ~~Potential loss~~, of "expert" i^*: $(1 - \varepsilon l_{i^*}^1)(1 - \varepsilon l_{i^*}^2) \cdots (1 - \varepsilon l_{i^*}^T)$. ②

$\Phi^{(T+1)} \geq$ final weight

$$② \le ① \Rightarrow \ln ② \le \ln ① \Rightarrow \ln(1 - \varepsilon l_{i^*}^1) + \cdots + \ln(1 - \varepsilon l_{i^*}^T) \le -\varepsilon \cdot (YL) + \ln N$$

$$\downarrow$$

$$\approx -\varepsilon l_{i^*}^1$$

$\ln(1-x) \approx -x - \frac{1}{2}x^2$

$\geq -x - x^2$

[[for small x]]

$$\geq -\varepsilon l_{i^*}^1 - \varepsilon^2 (l_{i^*}^1)^2$$

$$\geq -\varepsilon l_{i^*}^1 - \varepsilon^2.$$

$$\therefore \quad -\varepsilon \big(\underbrace{l_{i^*}^1 + \cdots + l_{i^*}^T}\big) - \varepsilon^2 T \le (-\varepsilon) \cdot (YL) + \ln N$$

$\div (-\varepsilon)$

"best expert loss"

$$\rightarrow (\text{total loss of } i^*) + \varepsilon T \geq YL - \frac{\ln N}{\varepsilon}$$

$$\Rightarrow \quad YL \le (\text{loss of } i^*) + \varepsilon T + \frac{\ln N}{\varepsilon}. \quad :)$$

$$\Rightarrow \quad \frac{1}{T}(YL) \le \frac{1}{T}(\text{loss of } i^*) + \varepsilon + \frac{\ln N}{\varepsilon T}$$

your avg. loss — avg. loss of best fixed slot machine — small — diminishes over time

E.g., can choose $\varepsilon = \sqrt{\frac{\ln N}{T}}$.

$\Rightarrow$ your avg. loss $\leq$ avg. loss of best fixed i^* $+ \underbrace{\frac{2\sqrt{\ln N}}{\sqrt{T}}}_{\text{decays to } 0 \text{ as } T \to \infty}$

"learning to play perfectly over time"