Lecture 19: Online Learning (Part 1)

Learning/prediction in "Mistake-Bounded Model"

[[A bit different from the "competitive ratio" model...]]

still an "online" kind of problem, but...

Versus C.R.: - look at cost difference [[versus ratio]]

- don't compare to best poss. decisions but to best fixed policy "in hindsight".

[[Not so clear yet, but will be. Kind of like situation where you have several possible algs you could use, want to not regret just picking one & sticking with it at end of day.]]

Game: • Each day $t = 1, 2, \ldots, T$, you must predict 0 or 1.

"Bitcoin down" "Bitcoin up" [[and you buy/sell accordingly]]

a^t := your action [[prediction]] e.g.

• ~~Then the~~ Later each day, an outcome o^t occurs.

• "Mistake" if $a^t \neq o^t$. [[Your goal, of course, is to make few mistakes over the course of time.]]

[[Well, so far there's absolutely nothing for you to go on, so we need to have more...]]

• There are N quote-unquote "experts".

[[Don't assume anything about how much they "know". They're just people with predictions.]]

At start of each day [[before you predict]], experts make predictions $e_1^t, e_2^t, \ldots, e_N^t \in \{0,1\}$

[[As I said, these may or may not have anything to do with outcomes. They have to do with your...]]

Regret: After T days, how are you vs. Best Expert?

e.g. $N=4$:

$\times$ = mistake

	e_1	e_2	e_3	e_4	a	o
$t=1$	1	1	1	0 $\times$	1	1
2	1 $\times$	0	1 $\times$	1 $\times$	1 $\times$	0
3	0	1	1	0	1	1
4	1 $\times$	1 $\times$	1 $\times$	0	1 $\times$	0

Best expert: [retrospectively] 1 mistake

You: 2 mistakes

- Dream(?) # Mistakes ≤ 1.01 (# Best expert's Mistakes)
- OK: # Mistakes $\le 3(\text{——————}) + O(\log N)$

"additive term" [unavoidable]

[To see why additive term is unavoidable, suppose....]

Say I promise some mystery expert will make 0 mistakes.

[What should you do? Well, if anyone ever makes a mistake, henceforth ignore them. Among the rest... go with majority?]

"Majority Alg":
- Dump any expert making a mistake
- O/w predict via majority.

Analysis: Each time you make a mistake, so did $\geq \frac{1}{2}$ of (remaining) experts

$\Rightarrow$ you dump $\geq \frac{1}{2}$ of remaining experts.

$\therefore$ you make $\le \log_2 N + 1$ mistakes [Total, indep. of T].

[vs. 0 of Best Expert). [This is optimal, and shows necessity of additive $O(\log N)$.]

Say Best Expert will make 5 (e.g.) mistakes

Could still do "Majority" Alg.... "resetting" if all experts get dumped

[What's the worst scenario now?...]

You make up to $\log_2 N + 1$ mistakes till Best Expert's 1st,
+ " more " " " " 2nd,
+ " " " " " " 3rd
4th
5th
+ up to $\log_2 N + 1$ more.

$\therefore$ Your Mistakes $\leq (\log_2 N + 1)(\#\text{ Best Expert Mistakes}) + \log_2 N + 1$

sad :(OK

[Better alg: Don't completely discount an expert after they make a mistake, just trust them less.]

"Weighted Majority" Alg.: You store a "weight" ("trust amount") w_i^t for each expert i @ time t.

- Initially, $w_1^1 = w_2^1 = w_3^1 = \cdots = w_N^1 = 1$.
- On day t, predict according to weighted majority, with w_i^t weights.
- After seeing outcome, halve the weight of each mistaken expert:

$$w_i^{t+1} = \begin{cases} w_i^t & \text{if } e_i^t = o_i \\ w_i^t/2 & \text{if } e_i^t \neq o_i. \end{cases}$$

Analysis idea: . Let $\Phi^t = w_1^t + w_2^t + \cdots + w_N^t$ [total "trust" @ time t]

① If you make a lot of mistakes, Φ^t becomes small.

② If some expert i^* makes few mistakes, $\Phi^t \geq w_{i^*}^t \geq$ large.

[In more detail:]

1. Suppose you make a mistake one day.
At least half the total weight Φ^t gets halved.
that contributing to wrong majority vote
we chose halving factor

$\Rightarrow \Phi^{t+1} \leq \frac{3}{4}\Phi^t$.

$\therefore \Phi^{(T+1)} \leq (3/4)^{\#\text{your mistakes (YM)}} \cdot \Phi^{(1)}, N$.
final weight

2. If expert i^* makes only BM ("Best Mistakes") mistakes, $w_{i^*}^{(T+1)} \geq \left(\frac{1}{2}\right)^{\#BM}$.
$\Phi^{(T+1)} \geq$

$\therefore \left(\frac{1}{2}\right)^{\#BM} \leq \left(\frac{3}{4}\right)^{\#YM} \cdot N$

$\Leftrightarrow$ (reciprocal) $2^{\#BM} \cdot N \geq \left(\frac{4}{3}\right)^{\#YM} \Rightarrow \left(\log_2 \frac{4}{3}\right) \cdot (\#YM) \leq \#BM + \log_2 N$.

$\Rightarrow$ Your Mistakes $\leq \frac{1}{\log_2 4/3} \cdot \left[(\#\text{Best expert Mistakes}) + \log_2 N\right]$

$\frac{1}{\log_2 4/3} \approx 2.41$

☺ Pretty good. [Better on multiplier, worse on additive term.]

Why be so harsh? Say we multiply wrong expert's weight by 0.98

Now: 1. $\geq$ half of Φ^t gets .98-timesed.
$\Rightarrow \Phi^{t+1} \leq \frac{1}{2}\Phi^t + \frac{1}{2} \cdot .98\,\Phi^t = .99\,\Phi^t$

2. $w_{i^*}^{(T+1)} \geq (.98)^{\#BM}$

$\Rightarrow (.98)^{\#BM} \leq (.99)^{\#YM} \cdot N$

(?)

(take logs) $(\#BM)\ln(.98) \le (\#YM)\ln(.99) + \ln N$. ⊛

What's $\ln(.99)$? [[Recall one of the greatest facts ever!!]]

$$e^x \approx 1+x \quad \text{for small } x$$

$$\Rightarrow \quad x \approx \ln(1+x) \qquad [[\text{took } \ln]]$$

Subs $x = -.01 \rightarrow -.01 \approx \ln(.99)$

$\therefore$ ⊛ $\approx$ $(-.02)(\#BM) \le (-.01)(\#YM) + \ln N$

$$(.01)(\#YM) \le (.02)(\#BM) + \ln N$$

$$\#YM \le \underset{\uparrow\ \ddot{\frown}}{2}(\#BM) + \underset{\uparrow\ \ddot{\smile}\ O(\log N)\ \ddot{\smile}}{100\ln N}$$

More precisely, using $\ln(1-\varepsilon) \approx -\varepsilon - \frac{1}{2}\varepsilon^2 - \dots$

$\rightsquigarrow$ If you penalize weight by factor $1-\varepsilon$,

get $\#YM \le 2(1+\varepsilon)(\#BM) + \frac{\ln N}{\varepsilon}$. $\ddot{\smile}$

Fact: No deterministic alg. can beat multiplicative factor 2.

Proof: Say just 2 experts:

e_1	e_2	a	o
0	1	a^1	
1	0	a^2	
0	1	a^3	
1	0	a^4	
0	1	⋮	
1	0	⋮	
0	1	⋮	
⋮	⋮		

Det. alg. $\Rightarrow a^1, a^2, a^3, \dots$ is fixed. An "adversary" designing outputs can always have $o^t \ne a^t$.

$\therefore$ Your Mistakes $= T$. But...

at least one expert must make $\le T/2$ mistakes!

Next time: Beating factor-2 with a randomized alg!