Lecture 21: Solving linear programs via online learning / multiplicative weights

Recap: Game:
- "Slot machines" ("experts") $i = 1 \ldots N$
- At time $t = 1, 2, \ldots, T$:

  Alg. chooses probabilities $p_1^t, p_2^t, \ldots, p_N^t$

  (interpretation: play machine i with prob. p_i^t)
- "Adversary" sets losses $\ell_1^t, \ldots, \ell_N^t$ between $-1, +1$
- Alg.'s loss at time t is $p^t \cdot \ell^t = p_1^t \ell_1^t + \cdots + p_N^t \ell_N^t$

Compare Alg's total loss (over T days) to best (smallest) value of $\sum_{t=1}^{T} \ell_i^t$ among $i = 1 \ldots N$

↑ total loss if Alg. just did $p_i^t = 1$ $\forall t$.

Hedge/Multiplicative Weights Alg:

Parameter: $0 < \varepsilon < 1$.

Initialize $w_1^1 = w_2^1 = \cdots = w_N^1 = 1$. $\qquad \Phi^t := w_1^t + \cdots + w_N^t$

At time t:
- Set $p_i^t = \dfrac{w_i^t}{\Phi^t}$ $\quad \forall i = 1 \ldots N$
- After losses revealed, set $w_i^{t+1} = (1 - \ell_i^t \varepsilon) w_i^t$ $\quad \forall i = 1 \ldots N$

Analysis: We showed:
1. $\Phi^{T+1} \le N \cdot e^{-\varepsilon (YL)}$, $\qquad YL :=$ your total loss
2. $\Phi^{T+1} \geq (1 - \varepsilon \ell_{i^*}^1) \cdots (1 - \varepsilon \ell_{i^*}^T)$ for any i^*

$(2) \le (1) \Rightarrow \ln(2) \le \ln(1)$

$$\Rightarrow \ln(1-\varepsilon l^1_{i*}) + \dots + \ln(1-\varepsilon l^T_{i*}) \le -\varepsilon(YL) + \ln N \quad (3)$$

Now use $\ln(1-x) \approx -x - \frac{1}{2}x^2 \geq -x - x^2$ for small x to get

$$\ln(1-\varepsilon l) \geq -\varepsilon l - \varepsilon^2 l^2 \geq -\varepsilon l - \varepsilon^2, \quad \text{since } -1 \le l \le 1 \Rightarrow l^2 \le 1.$$

$\therefore (3) \Rightarrow -\varepsilon l^1_{i*} - \varepsilon^2 - \varepsilon l^2_{i*} - \varepsilon^2 - \dots - \varepsilon l^T_{i*} - \varepsilon^2 \le -\varepsilon(YL) + \ln N.$

$$\Rightarrow \varepsilon \cdot (YL) \le \varepsilon(l^1_{i*} + \dots + l^T_{i*}) + \varepsilon^2 T + \ln N$$

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

[[or, divided by T...]]

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

- $\frac{1}{T}(YL)$: your avg. loss / day
- $\frac{1}{T}(\text{loss of always } i^*)$: avg loss/day of "always i^*"
- ε: small
- $\frac{\ln N}{\varepsilon T}$: diminishes over time

A good choice of ε: balance $\varepsilon = \frac{\ln N}{\varepsilon T} \Leftrightarrow \varepsilon^2 = \frac{\ln N}{T}$

$$\Leftrightarrow \varepsilon = \sqrt{\frac{\ln N}{T}}$$

$$\Rightarrow \frac{1}{T}(YL) \le \frac{1}{T}(BL) + \frac{2\sqrt{\ln N}}{\sqrt{T}}$$

$\leftarrow$ diminishes over time, small once $T \gg \ln N$.

Solving LPs with this.

$\rightarrow$ zero-sum games.

[[Zero-sum games are a kind of problem from Game Theory / Economics. They're a special case of LPs (like flows), though actually they're kind of "equiv."; you can prove every LP can be reduced to one. We won't, tho; will just be content to solve them.]]

Zero-Sum Games (e.g.: Rock-Paper-Scissors)

2 players, Alice & Bob.

Alice → N_1 "actions"; Bob → N_2 "actions"

Alice \ Bob	R	P	S
R	0	+1	-1
P	-1	0	+1
S	+1	-1	0

$N_1 = N_2 = \{R, P, S\}$

[N_1, N_2 don't have to be same]

"Payoff matrix" M: $N_1 \times N_2$

M_{ab} = how much Alice pays Bob if she plays a, he plays b

[loss to Alice, gain to Bob; these sum to zero]

WLOG, $|M_{a,b}| \le 1 \quad \forall a,b$.

Who plays first?

→ Play "simultaneously", and each may use a "mixed strategy" = probability distrib. on actions

If Alice uses $p_1, p_2, \ldots, p_{N_1}$
Bob uses $q_1, q_2, \ldots, q_{N_2}$, Alice's expected loss is $\sum_{a,b} p_a q_b M_{ab}$.
(= Bob's expected gain).

What are their "optimal strats"?
How to compute?! [Not an LP, seemingly...]

Alternate version 1: Hard on Alice:

- Alice must first announce her randomized strat. $p_1, \ldots, p_{N_1}$
- Bob may now choose his randomized strat. $q_1, \ldots, q_N$

But Bob's expected gain is $q_1 \left(\sum_a p_a M_{a,1} \right) + \cdots + q_{N_2} \left(\sum_a p_a M_{a,N_2} \right)$

Bob should just put 100% prob. on whichever of these is largest.

∴ Bob may as well be deterministic

Alice's goal: minimize $\max \left\{ \sum_a p_a M_{a,1}, \ldots, \sum_a p_a M_{a,N_2} \right\} \to v$

s.t. $p_1, \ldots, p_{N_1} \ge 0$

$p_1 + \cdots + p_{N_1} = 1$

$v \ge \sum_a p_a M_{a,1}$ ← consts

$\vdots$

$v \ge \sum_a p_a M_{a,N_2}$

An LP! Say its opt. value is L_{hard}: least expected loss in hard ver. for Alice.

Alt ver 2: <u>Easy on Alice</u>:

- Bob must first announce his randomized strat $q_1, \ldots, q_{N_2}$
- Alice may now choose hers
- May as well be deterministic
- Bob's optimal strategy is given by an LP; [[a maximization]]
  call its opt. value L_{easy}: least expected loss in this easy ver.

$$L_{easy} \le L \le L_{hard}$$

- L_{easy}: Alice goes second
- L: expected Alice-loss when they play (optimal strats) simultaneously
- L_{hard}: [[expected loss when Alice goes first]]

[[von Neumann]] Minimax Theorem: $L_{easy} = L_{hard}$ $(= L)$.

Proof 1 (hmwk): The two LPs, for L_{easy}, L_{hard}, are <u>duals</u>.

Proof 2 ... [[We'll show it, & we'll in fact show an algorithm to find the optimal strategies achieving L.]]

Need to show <u>$L_{hard} \le L_{easy}$.</u>

~~Consider your/Alice's actions $a = 1 \ldots N_1$ the slots/experts.~~
~~Play Hedge for a while...~~
~~On day 1, $p_1 = \cdots = p_{N_1} = \frac{1}{N_1}$ Let Adversary "think/play"~~

Consider playing "Hard on Alice" version T days in a row.
Treat Alice's/your options $a = 1 \ldots N_1$ as slots/experts
You/Alice will play $p_1^t, \ldots, p_{N_1}^t$ according to Hedge strategy on round t.
Adversary/Bob will play "best response" in Zero Sum Game, b^t, on round t.
↳ yields a "loss vector" for Hedge, $\ell_i^t = M_{i,b^t}$.
Alice updates with this loss vector.

e.g. Z.S.G.:

		Bob 1	Bob 2	Bob 3
Alice	1	+.3	−.2	−.4
	2	−.6	+.9	+.8

e.g.: $\varepsilon = .2$

day	Alice mixed strat	Bob's response	loss vector	Alice's expected loss
1	$(\frac{1}{2}, \frac{1}{2})$	2	$(-.2, +.9)$	$\frac{1}{2}(-.2) + \frac{1}{2}(.9) = +.35$
2	$w_1 = 1 + .2\varepsilon$, $w_2 = 1 - .9\varepsilon$ (.56, .44)	2	$(-.2, +.9)$	$.56(-.2) + .44(.9) = .284$
3	(.62, .38)	2		
⋮		2		
		2		
6	(.77, .23)	1	$(+.3, -.6)$	

$\frac{1}{T}$(Your/Alice loss) $\geq L_{hard}$, since You/Alice had to go first each time.

In Hedge ver, what – in hindsight – would be the best single play for You/Alice?

Define $q_1 = \frac{\text{frac times Bob's resp. 1}}{T}$, $q_2 = \frac{\text{frac times Bob resp. 2}}{T}$, ... $q_{N_2} = \frac{\_\_ N_2}{T}$.

$\frac{1}{T}$("Best Loss") is avg. value of best response by Alice to Bob playing mixed strat $q_1, \ldots, q_{N_2}$!

$\therefore$ ↳ $\leq L_{easy}$.

$\therefore$ after T rounds, we conclude

$$L_{hard} \leq \frac{1}{T}(\text{Your Loss}) \leq \frac{1}{T}(\text{Best Loss}) + \varepsilon + \frac{\ln N}{\varepsilon T}$$
$$\leq L_{easy} + \varepsilon + \frac{\ln N}{\varepsilon T}.$$

$\therefore$ must have $L_{hard} \leq L_{easy}$,
because if $L_{hard} > L_{easy} + \delta$, we could make $\varepsilon < \frac{\delta}{2}$, then T large enough so $\frac{\ln N}{\varepsilon T} < \frac{\delta}{2}$,
get $L_{hard} \leq L_{easy} + \delta$, $\Rightarrow\Leftarrow$.

$\therefore$ $L_{hard} \leq L_{easy}$. ✓

Moreover: the <u>average</u> $\bar{p}_1, \dots, \bar{p}_{N_1}$ of Alice's plays,
& the $q_1, \dots, q_{N_2}$ of Bob's plays

are near-optimal strategies,
only $\varepsilon + \frac{\ln N}{\varepsilon T}$ off from the optimal $L_{easy} = L = L_{hard}$ value.

$\therefore$ Can algorithmically find them:
Simulate for $T = \frac{\ln N}{\varepsilon^2}$ steps, and then $\bar{p}, q$ are within 2ε of optimal!