Experts II:

Last time we saw: that Hedge, which takes loss vectors $\ell^{(t)} \in [-1,1]^N$ and does the update $\omega_j^{(t+1)} = \omega_j^{(t)} \cdot e^{-\varepsilon \ell_j^{(t)}}$ has

$$\sum_t \langle \bar{p}^t, \bar{\ell}^t \rangle \leq \sum_t \langle \bar{p}^t, \bar{e}_i \rangle + \frac{\ln N}{\varepsilon} + \varepsilon T. \qquad \Rightarrow \text{Avg regret} \leq 2\varepsilon \text{ after } T \geq \frac{\ln N}{\varepsilon^2} \text{ steps.}$$

- And some extension to $\ell^{(t)} \in [-\rho, \rho]^N$, then $T \geq \frac{4\rho^2 \ln N}{\varepsilon^2}$ ensures avg. regret $\leq \varepsilon$.
- Same for gains, just set $g^{(t)} = -\ell^{(t)}$.

Fact: by setting $\varepsilon = \sqrt{\frac{\ln N}{T}}$ we get that regret $\leq 2\sqrt{T \ln N}$ after T steps.

Observation: this behaviour is asymptotically tight.

Suppose there are ~~2 coins (unbiased)~~ 2 experts, one is 1, other -1, and we're trying to predict the outcome of an unbiased coin.

Then whp we will get $1/2 \pm \Omega(\sqrt{T})$ fraction of flips being of one type.

And then better expert is $T/2 + \Omega(\sqrt{T})$ whereas we cannot be better than $T/2$ in expectation.

If we have an N-sided coin then one of the outcomes is likely to be $\sqrt{T \log N}$ more than the average.

Last time we alluded to the bandit setting, but not today, maybe later in the course.

Application 1: Solving LPs approximately.

Application 2: Minimax theorem for 2 player 0-sum games. (Let's do it first).

Let M be the payoff matrix for the row player. $M = [-1,1]^{m \times n}$ by scaling

- So if $x, y \in \Delta_m, \Delta_n$ respectively then $E[\text{payoff}] = x^T M y$
- Given x, the best response for the column player is

$$j^*(x) = \underset{j}{\operatorname{argmin}}\ x^T M e_j \quad \text{and let } C(x) \text{ be the payoff.}$$

Similarly for $y \in \Delta_n$

$$i^*(y) = \underset{i}{\operatorname{argmax}}\ e_i^T M y \qquad R(y) \text{ be the payoff.}$$

Thm [von Neumann's Minimax Thm].

$\forall$ finite 0sum game (2player). $M = [-1,1]^{m\times n}$ say.

$$\max_{x\in\Delta_m} C(x) = \min_{y\in\Delta_n} R(y). = V = \text{value of the game.}$$

Pf: Suppose not.

[Clearly BTW. $\max_{x\in\Delta_m} C(x) \le \min_{y\in\Delta_n} R(y)$

row player fixes strategy first $\Rightarrow$ has disadvantage

column player fixes strategy first $\Rightarrow$ row player has advantage.

Suppose $\exists$ gap between $\max_x C(x)$ and $\min_y R(y)$ say $\delta > 0$ gap. Get contradiction.

say achieved by y^*.

Treat each row of M as expert.

@ time t, row player produces $p^{(t)}$. $p^{(1)} = (1/m \cdots 1/m)$.

then column player plays best response $j^*(p^t). = j_t.$ to give $C(p^t)$ payoff.

$\Rightarrow g^t = M e_{j_t}$

Play for T time-steps. $T \ge \frac{4 \ln m}{\varepsilon^2}$

~~(by Hedge) Avg. Total payoff~~ $\cancel{= \frac{1}{T}\sum_{t=1}^{T} \langle p^t, g^t\rangle \ge \frac{1}{T}\sum_{t=1}^{T} \langle p^t, e_j\rangle - \epsilon}$ ~~$\forall$ possible strategies e_j for column player.~~

~~ugh!~~

~~If we define~~ $\cancel{\hat{x} = \frac{1}{T}\sum p^t \in \Delta_m}$

~~then this means that (our payoff) avg~~ $\cancel{\ge \frac{1}{T}\sum \langle \hat{x}, e_j\rangle - \epsilon \quad \forall j \quad \ge C(\hat{x}) - \epsilon}$

~~But at each time column player could have played~~ $\cancel{y^*}$

---

① by Hedge. $\frac{1}{T}\sum_t \langle p^t, g^t\rangle \ge \frac{1}{T}\sum_t \langle e_i, g^t\rangle - \varepsilon.$ after T steps.

$= R(\hat{y}) - \varepsilon$ if $\hat{y} = \frac{1}{T}\sum_t e_{j_t}$

② Our avg payoff $= \frac{1}{T}\sum \langle p^t, g^t\rangle = \frac{1}{T}\sum C(p^t) \le C(\hat{x})$. $\Rightarrow C(\hat{x}) \ge R(\hat{y}) - \varepsilon$

C is concave
(can use the minimizer on RHS for all p^t on the left,)

Make T large so that $\varepsilon < \delta$.

Solving LPs using experts:

Say want to solve $\max c^T x$ st $Ax \leq b$, $x \geq 0$.

will get an apx-soln. st $\langle a_i, x \rangle \leq b_i + \varepsilon \quad \forall i = 1..m$.

Assume we know that $\max = OPT$.

Let $K = \{x \geq 0,\ c^T x = OPT\}$

~~Say we know~~

~~Say only one constraint: $\max c^T x$ $\alpha x \leq \beta$, $x \geq 0$.~~

~~Let's say $K = \{x : c^T x = OPT, x \geq 0\}$. So want $x \in K, \alpha x \leq \beta$~~

~~choose e_{j^*} where $j^* = \arg\max_j c_j/\alpha_j$ and we'll set $x =$~~

- Sps only one constraint: $\alpha x \leq \beta$. so want $x \in K$ st $\alpha x \leq \beta$. (if feasible).

  Basically want $j^* = \arg\max_j c_j/\alpha_j$ and set $\begin{cases} \bar{x} = \bar{e}_{j^*} \cdot \min\{OPT/c_{j^*}, \beta/\alpha_{j^*}\} \\ \text{if } \bar{x} \notin K \text{ return infeasible.} \end{cases}$

- So one constraint is easy. What about multiple constraints?

  Let's combine them together using weights. Say $p^{(t)}$ is how we combine them, $p^{(t)} \in \Delta_m$.

  Set $\alpha^t = p^{(t)T} A$, $\beta^t = (p^t)^T b$ — if $\exists$ a feasible sol^n to $x \in K$, $Ax \leq b$ $\Rightarrow$ x is also a sol^n to $x \in K$, $\alpha x \leq \beta$.

  So we get sol^{n}s $x^1, x^2 \ldots x^T$.

  define $\hat{x} = \frac{1}{T}\sum_{i=1}^{T} x^i$.

  BTW, what are the loss vectors? (gain) $\ell_i^{(t)} = b_i - (Ax^t)_i$ // increase weights on unsatisfied constraints!

  or $g_i^{(t)} = (Ax^{(t)})_i - b_i$

  $= \langle a_i, x^{(t)} \rangle - b_i$

Clearly if $x^{(t)} \in K \ \forall t \Rightarrow \hat{x} \in K$.

So what is $\langle a_i, \hat{x} \rangle - b_i$?

by Hedge, $\frac{1}{T}\sum_t \langle p^t, g^t \rangle \geq \frac{1}{T}\sum_t \langle e_i, g^t \rangle - \varepsilon$.

LHS: $\langle p^t, g^t \rangle = \langle p^t, Ax^t - b \rangle = \langle \alpha^t x^t - \beta^t \rangle \leq 0$

$\Rightarrow$ ~~4~~ RHS $= \frac{1}{T}\sum_t \langle g^t, e_i\rangle - \varepsilon \leq 0$

but $g^t = Ax^t - b$

$\Rightarrow \frac{1}{T}\sum g^t = A\hat{x} - b. \Rightarrow (A\hat{x}-b)_i \leq \varepsilon. \Rightarrow a_i x - b_i \leq \varepsilon.$

i.e. sol^n is ε-feasible! As claimed.

BTW: all answers were $\in [-\rho, \rho]$ where $\rho = \max_{x\in K} |(a_i x - b_i)|$

width of the "oracle".

Nb: If we could ensure that all answers we give have small width then all is OK. But in general width could become large.

Also: since we're getting an apx already, could actually solve the oracle itself approximately. if get $\alpha x \leq \beta + \delta$ (and have small width) then the δ just gives $A\hat{x} \leq b + \epsilon + \delta$ and get faster runtime

---

Example of maxflow: Say all edge capacities $=1$
OPT $= F$

Now: could send all F flow on the shortest path. May get width F.

$\Rightarrow$ #rounds $= O\left(\frac{F^2 \log m}{\epsilon^2}\right)$.

Can get it down to single F.

How? the width is asymmetric. When we are satisfied the gain is at ~~most~~ least -1 else it goes up to F. So we can use that in a more careful analysis of Hedge to get $O\left(\frac{F\log m}{\epsilon^2}\right)$ iters, each is a shortest path computation!

$$\max \sum f_P$$

$$\text{st } \sum_{P\ni e} f_P \leq 1 \ \forall e.$$

$$f_P \geq 0.$$

$$\Downarrow$$

$$\max \sum f_P$$

$$\underbrace{\text{st } \sum_e p_e \cdot \sum_{P\ni e} f_P \leq 1}$$

$$\sum_P f_P \underbrace{\sum_{e\in P} p_e}_{\text{len}(P)} \leq 1.$$