Lecture #17

<u>Last Time:</u> Interior Point Methods

(Newton's method, self concordance)

<u>Today:</u> Multiplicative Weights Method

Rediscovered many times, central to learning and algorithm design

<u>Setup:</u> n experts whose advice you can take

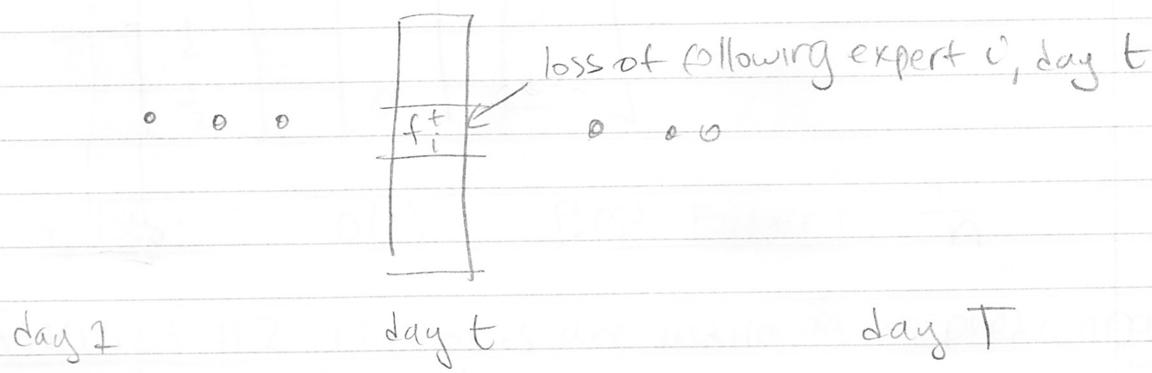

<u>Important:</u> These losses are bdded ($\|f^t\|_\infty \le 1 \ \forall t$), but are otherwise <u>arbitrary</u>

<u>Goal:</u> perform almost as well as best expert in hindsight — i.e. minimize

$$\text{Regret} := \sum_{t=1}^{T} \langle p^t, f^t \rangle - \min_i \sum_{t=1}^{T} f_i^t$$

indicator of expert chosen on day t or <u>distribution</u> on experts sampled from

Applications: Picking stocks, playing strategic games, building algorithms/predictors

Any ideas for how to pick experts?

Attempt #1: Choose best expert so far

But what if the losses are:

$$\begin{bmatrix} 0 \\ \frac{1}{2} \\ \vdots \\ \frac{1}{2} \\ \frac{1}{2} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix} \cdots \qquad \text{"round robin"}$$

Our loss: $T - o(1)$ Best Expert: $\frac{T}{n}$

Attempt #2: (If losses are indicators of predicting correctly) Weighted Majority

Can show our loss $\lessapprox$ 2x Best Expert

Attempt #3: Multiplicative Weights

Initialize $w_i^1 = 1 \ \forall i$ (equal weights)

For $t = 1$ to T

Follow expert i with probability $\frac{w_i^t}{\sum_j w_j^t} := p_i^t$

Update each $w_i^{t+1} = w_i^t(1 - \varepsilon f_i^t)$

Theorem 1: If $0 < \varepsilon \leq \frac{1}{2}$ then MWU satisfies

$$\underbrace{\sum_{t=1}^{T} \langle p^t, f^t \rangle}_{\text{expected loss in day } t} - \min_i \sum_{t=1}^{T} f_i^t \leq \frac{\ln n}{\varepsilon} + \varepsilon T$$

If we know T can set $\varepsilon = \sqrt{\frac{\ln n}{T}}$ to get

regret at most $2\sqrt{T \ln n}$ (avg regret $\xrightarrow[T \to \infty]{} 0$)

Intuition: The weights focus quickly on best experts

Let $\Phi^t := \sum_{i=1}^{n} w_i^t$, will be our potential fctn

Proof Plan: (a) when we incur loss, Φ decreases
(b) If there is a good expert, Φ cannot be too small

Proof: We have $\Phi^{t+1} = \sum_i w_i^{t+1} = \sum_i w_i^t (1 - \varepsilon f_i^t)$

Now $w_i^t = p_i^t \Phi^t$, hence

$$\Phi^{t+1} = \sum_i p_i^t \Phi^t (1 - \varepsilon f_i^t) = \Phi^t (1 - \varepsilon \langle p^t, f^t \rangle)$$

$$\underset{(\text{using } 1 - x \leq e^{-x})}{\leq} \Phi^t e^{-\varepsilon \langle p^t, f^t \rangle}$$

Thus we have $\Phi^{T+1} \leq \underbrace{\Phi^1}_{n} e^{-\varepsilon \sum_{t=1}^{T} \langle p^t, f^t \rangle}$

This is step (a)

Now for step (b),

$$\Phi^{T+1} \geq w_i^{T+1} = \prod_{t=1}^{T} (1 - \varepsilon f_i^t) \quad \forall i$$

Using $1 - x \geq e^{-x-x^2}$ we have

$$\Phi^{T+1} \geq \prod_{t=1}^{T} e^{-\varepsilon f_i^t - \varepsilon^2 (f_i^t)^2} = e^{-\varepsilon \sum_{t=1}^{T} f_i^t - \varepsilon^2 \sum_{t=1}^{T} (f_i^t)^2}$$

Now putting it all together

$$\ln\left(n\, e^{-\varepsilon \sum_{t=1}^{T} \langle p^t, f^t \rangle}\right) \geq \ln\left(e^{-\varepsilon \sum_{t=1}^{T} f_i^t - \varepsilon^2 \sum_{t=1}^{T} (f_i^t)^2}\right)$$

$$\Rightarrow \ln n + \varepsilon^2 \underbrace{\sum_{t=1}^{T} (f_i^t)^2}_{\leq T} \geq \varepsilon \sum_{t=1}^{T} \langle p^t, f^t \rangle - \varepsilon \sum_{t=1}^{T} f_i^t$$

$$\Rightarrow \frac{\ln n}{\varepsilon} + \varepsilon T \geq \text{regret} \quad \blacksquare$$

---

what assumptions about f^t's did we need?
Stochastic? Adversarial?

In fact the adversary can even know our p^t's and our strategy (as long as he doesn't see our coin flips)

What if the losses satisfy $\|f^t\|_\infty \le \rho$ instead?

width

We can set $w_i^{t+1} = w_i^t\left(1 - \frac{\varepsilon}{\rho} f_i^t\right)$, to get

$$\text{regret} \le \frac{\rho^2 \ln n}{\varepsilon} + \varepsilon T$$

Now let's apply what we've learned to zero sum games.

Let A be an $m \times n$ <u>payoff</u> <u>matrix</u>

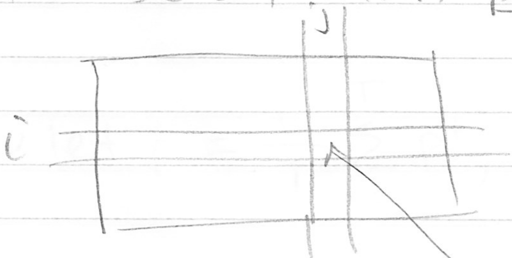

<u>assume</u>: $\max_{i,j} |A_{i,j}| \le 1$

$A_{i,j}$ = "amt Alice wins, Bob loses if they play i and j"

Suppose Alice and Bob play MWU to get randomized strategies:

$$(p^1, q^1), (p^2, q^2) \ldots (p^T, q^T)$$

$$(1) \quad \overbrace{-\sum_{t=1}^{T} \langle p^t, Aq^t \rangle + \max_i \sum_{t=1}^{T} (Aq^t)_i}^{\text{Alice's expected regret}} \le \frac{\ln n}{\varepsilon} + \varepsilon T$$

Similarly for Bob we have:

$$(2)\quad \underbrace{\sum_{t=1}^{T} \langle p^t, Aq^t \rangle - \min_j \sum_{t=1}^{T} ((p^t)^T A)_j}_{\text{Bob's expected regret}} \le \frac{\ln n}{\varepsilon} + \varepsilon T$$

Then (1) + (2) yields

$$\max_i \sum_{t=1}^{T} (Aq^t)_i - \min_j \sum_{t=1}^{T} ((p^t)^T A)_j \le \frac{2 \ln n}{\varepsilon} + 2\varepsilon T$$

$$\left(\text{setting } \varepsilon = \sqrt{\frac{\ln n}{T}}\right) \le 4\sqrt{\frac{\ln n}{T}} := \delta$$

Now let $\bar{p} = \frac{1}{T}\sum_{t=1}^{T} p^t$, $\bar{q} = \frac{1}{T}\sum_{t=1}^{T} p^t$ then

$$\underbrace{\max_i (A\bar{q})_i}_{\substack{\text{Bob plays first,}\\ \text{then Alice}}} - \underbrace{\min_j (\bar{p}^T A)_j}_{\substack{\text{Alice plays first,}\\ \text{then Bob}}} \le \delta$$

Thus we have

$$\min_j (\bar{p}^T A)_j \le \bar{p}^T A \bar{q} \le \max_i (A\bar{q})_i$$

where the l.b. and u.b. are off by at most $+\delta$,

Thus the strategies $(\bar{p}, \bar{q})$ are approximately in equilibrium, and no player can do better by more than δ by changing their strategy