Today

- Deterministic MWU
- Randomized MWU
- General loss/gain + Actions

Recall

$1-\varepsilon \approx \exp(-\varepsilon)$, will use

Claim: $1-\varepsilon \geq \exp(-\varepsilon-\varepsilon^2) \quad \forall\ \varepsilon \in (0, .6)$

Today: how to make decisions when future unknown (or even adversarial)
But: discovered/applications in: ML, solving LPs, Computational geometry, fast algorithms, ...

Deterministic Setup

On day $i = 1, 2, \dots, T$
Each expert $j \in [n]$ recommends $r_{ij} \in \{0,1\}$
We decide $d_i \in \{0,1\}$ via "strategy"
Adversary reveals correct answer $t_i \in \{0,1\}$

Goal: Minimize total # mistakes $M := \sum_i^T |d_i - t_i|$

Can do as bad as $M = T$ if adversary just chooses opposite of us
More achievable: do well if experts give good advice

Let $M_j := \sum_i^T |r_{ij} - t_i|$ and $M^* := \min_j M_j$

Hope Today: $M \approx M^*$

Claim: If $M^* = 0$, then $\exists$ strategy s.t. $M \leq \log n$

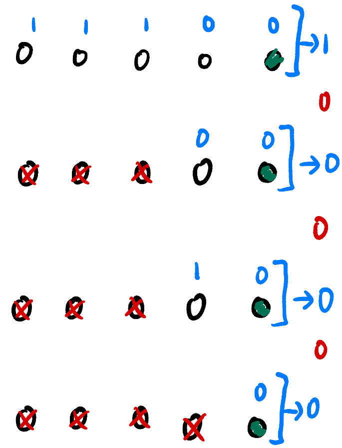

Initially all experts "active"

On day i

$d_i \leftarrow$ Majority vote of active experts

Deactivate any incorrect experts

Each time we make a mistake, # active agents halved
But # active experts ≥ 1 and initially n
So $M \leq \log n$ (A halving argument!)

But what if $\nexists$ perfect agent

Claim: $\exists$ strategy s.t. $M \leq \log n + M^* \cdot \log n$

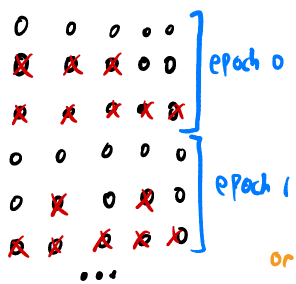

Initially all experts "active"

On day i

$d_i \leftarrow$ Majority vote of active experts

Deactivate any incorrect experts

If all experts inactive, reactivate all experts

or to T if not defined

Let kth epoch be days [kth time activate experts, (k+1)th time reactivate all experts), $K := $ # epochs

In each epoch we make $\leq \log n$ mistakes so $M \leq K \cdot \log n$
But every expert makes a mistake in each epoch so $M^* \geq K$

$\rightarrow M \leq \log n \cdot K$

Previous strategy very harsh; if expert always right except for start of epoch is ignored; want gentler penalty

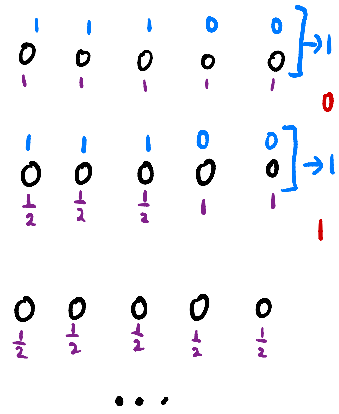

Claim: $\exists$ strategy $\forall \varepsilon \in (0,.6)$ s.t. $M \le (2+\varepsilon)\cdot M^* + O(\log n/\varepsilon)$

Initially expert j gets weight $w_j = 1$

On day i

$d_i \leftarrow$ weighted Majority vote of active experts (d_i is $b \in \{0,1\}$ maxing $\sum_{j:r_{ij}=b} w_j$)

$w_j \leftarrow (1-\varepsilon\cdot|r_{ij}-t_j|)\cdot w_j$ $\leftarrow$ like a "soft" deactivation

↑ The "multiplicative weights" strategy

Analysis intuition: before, tracking # active agents; now weight ~ how active an agent is so track it

Let $W_i := \sum_j w_j$ at start of day i

$W_{T+1} \le n \cdot \exp(-\frac{\varepsilon}{2}\cdot M)$

$W_{i+1} \le \frac{1}{2}W_i + \frac{1}{2}(1-\varepsilon)\cdot W_i = (1-\frac{\varepsilon}{2})\cdot W_i$

If we make a mistake on day i then $\sum_{j:r_{ij}\neq t_i} w_j \ge \sum_{j:r_{ij}=t_i} w_j$ but $W_i = \sum_{j:r_{ij}\neq t_i} w_j + \sum_{j:r_{ij}=t_i} w_j \le 2\cdot \sum_{j:r_{ij}\neq t_i} w_j \to \frac{W_i}{2} \le \sum_{j:r_{ij}\neq t_i} w_j$

But $W_{i+1} = \sum_{j:r_{ij}=t_i} w_j + (1-\varepsilon)\sum_{j:r_{ij}\neq t_i} w_j = W_i - \varepsilon\cdot\sum_{j:r_{ij}\neq t_i} w_j \le W_i - \varepsilon\cdot\frac{W_i}{2} = (1-\frac{\varepsilon}{2})\cdot W_i$

So $W_{T+1} \le (1-\frac{\varepsilon}{2})^M \cdot W_1 = (1-\frac{\varepsilon}{2})^M\cdot n \underset{1-x\le e^{-x}}{\le} \exp(-\frac{\varepsilon}{2}\cdot M)\cdot n$

$W_{T+1} \ge \exp(-\varepsilon M^* - \varepsilon^2 M^*)$

Best expert has weight $(1-\varepsilon)^{M^*}$ at end so $W_{T+1} \ge (1-\varepsilon)^{M^*}$

But $1-\varepsilon \ge \exp(-\varepsilon-\varepsilon^2)$ for $\varepsilon\in(0,.6)$

So $W_{T+1} \ge \exp(-\varepsilon\cdot M^* - \varepsilon^2\cdot M^*)$

So $\exp(-\varepsilon M^* - \varepsilon^2 M^*) \le n\cdot \exp(-\frac{\varepsilon}{2}\cdot M) \iff -\varepsilon M^* - \varepsilon^2 M^* \le \ln n - \frac{\varepsilon}{2}\cdot M$

$\iff M \le 2\frac{\ln n}{\varepsilon} + 2M^* + \varepsilon M^* = (2+\varepsilon)M^* + O(\frac{\log n}{\varepsilon})$

Above factor of 2 basically tight:

Claim: $\exists$ adversary+experts s.t. any deterministic strategy has $M \ge 2\cdot M^*$

Consider 2 experts, $r_{i1} = 0$ and $r_{i2} = 1$ $\forall i$

Adversary outputs $t_i := \bar{d}_i$ $\forall i$

$M = T$

By averaging $\exists j \in \{1,2\}$ s.t. $M_j \le \frac{T}{2}$ so $M^* \le \frac{T}{2}$ so $M \ge 2\cdot M^*$

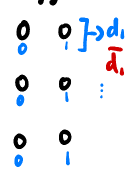

But hope was $M \approx M^*$; will improve by using randomization

Randomized Setup

On day $i = 1, 2, \ldots, T$

Each expert $j \in [n]$ recommends $r_{ij} \in \{0,1\}$

We fix a distribution P_i over experts

Adversary reveals correct answer $t_i \in \{0,1\}$

We decide $d_i = r_{ij}$ w/ Probability P_{ij} (adversary can see P_i, but not d_i)

Goal: Minimize total expected # mistakes $\mathbb{E}[M] = \sum_i \sum_j P_{ij} \cdot M_{ij}$ where $M_{ij} := |r_{ij} - t_i|$

Notice adversary now cannot just do opposite of us so hope for beating $2M^*$

Claim: $\exists$ randomized strategy $\forall\ \varepsilon \in (0, .6)$ s.t. $\mathbb{E}[M] \leq (1+\varepsilon) \cdot M^* + \frac{\ln n}{\varepsilon}$

Initially expert j gets weight $w_j = 1$

On day i

$P_{ij} \leftarrow w_j / \sum_j w_j$

$w_j \leftarrow (1 - \varepsilon \cdot |r_{ij} - t_j|) \cdot w_j$

Let $W_i := \sum_j w_j$ at start of day i

Let $m_i := |d_i - t_i|$ so $\mathbb{E}[M] = \sum_i \mathbb{E}[m_i]$ (by LoE)

$W_{T+1} \leq n \cdot \exp(-\varepsilon \cdot \mathbb{E}[M])$

$$\mathbb{E}[m_i] = \sum_j P_{ij} |r_{ij} - t_i| = \sum_j \frac{w_j}{W_i} \cdot |r_{ij} - t_i| = \sum_{j : r_{ij} \neq t_i} w_j / W_i \quad \text{so} \quad \sum_{j : r_{ij} \neq t_i} w_j = \mathbb{E}[m_i] \cdot W_i$$

$$W_{i+1} \leftarrow \sum_{j : r_{ij} = t_i} w_j + (1-\varepsilon) \cdot \sum_{j : r_{ij} \neq t_i} w_j = W_i - \varepsilon \cdot \sum_{j : r_{ij} \neq t_i} w_j = W_i (1 - \varepsilon \cdot \mathbb{E}[m_i])$$

$$\text{so} \quad W_{T+1} = W_1 \cdot \prod_{i=1}^{T} (1 - \varepsilon \cdot \mathbb{E}[m_i]) \underset{1-x \leq e^{-x},\ W_1 = n}{\leq} n \cdot \exp\left(-\sum_i \varepsilon \cdot \mathbb{E}[m_i]\right) = n \cdot \exp(-\varepsilon\, \mathbb{E}[M])$$

$W_{T+1} \geq \exp(-\varepsilon M^* - \varepsilon^2 M^*)$ by same argument as deterministic case

$$\text{so} \quad \exp(-\varepsilon M^* - \varepsilon^2 M^*) \leq n \cdot \exp(-\varepsilon \cdot \mathbb{E}[M]) \iff -\varepsilon M^* - \varepsilon^2 M^* \leq \ln n - \varepsilon\, \mathbb{E}[M]$$

$$\iff \mathbb{E}[M] \leq (1+\varepsilon) \cdot M^* + \frac{\ln n}{\varepsilon}$$

2 Big Generalizations

1) Instead of $d_i \in \{0,1\}$, $d_i \in [n]$ where $d_i = j$ ~ do what expert j does

2) Continuous loss + gain

General Setup

On day $i = 1, 2, \dots, T$

We fix a distribution P_i over experts

Adversary reveals gain/loss vector $\ell_i \in [-1,1]^n$

We decide $d_i = j$ w/ Probability P_{ij}

Goal: Minimize expected loss $\mathbb{E}[L] = \sum_i^T \sum_j^n P_{ij} \cdot \ell_{ij}$

Let $L^* := \min_j \sum_i^T \ell_{ij}$

Claim: $\exists$ randomized strategy $\forall\, \varepsilon \in (0, .6)$ s.t. $\mathbb{E}[L] \leq L^* + \varepsilon \cdot T + \frac{\ln n}{\varepsilon}$

Initially expert j gets weight $w_j = 1$

On day i

$P_{ij} \leftarrow w_j / \sum_j^n w_j$

$w_j \leftarrow (1 - \varepsilon \cdot \ell_{ij}) \cdot w_j$

Possibly > 0!

Rest of analysis basically same as before

"Experts" is Very Flexible

"Weak" Supervised Learner → "boost" weak to strong learners

Strategies in Zero-Sum Game → Proof of Minimax

Elements of Probability Distribution → Derandomization

Constraints of an LP → APx. Solve LPs in Poly-time

↳ More on this next time to solve flow!