CS3000: Algorithms & Data
Jonathan Ullman

Lecture 22:

- Online Learning

Dec 4^{th}, 2018

Picking Good Experts

- Suppose you have N "experts" making predictions
  - Weather forecasters
  - Financial advisors
  - Recommender systems
  - …

- Most of them are bad, but one might be good!
- Who's predictions should we trust?

Model

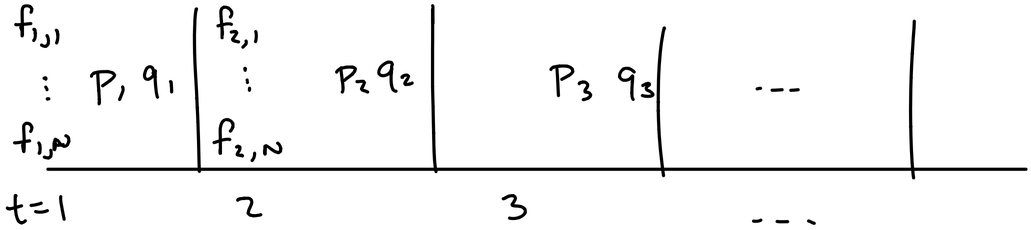

- There are T time periods, in each:
  - experts make predictions $f_{t,1}, \ldots, f_{t,N} \in \{0,1\}$
  - you have to make a decision $p_t \in \{0,1\}$
  - some outcome $q_t \in \{0,1\}$ is revealed (may be adversarial)

binary outcome

- **Goal:** minimize mistakes ($M \approx M^*$)
  - You have a mistake on day t if $p_t \neq q_t$
  - Want # of mistakes $\ll T$
  - M^* is the # of mistakes made by the best expert in hindsight. M is # of mistakes we make

Level I: The Halving Algorithm (HA)

- Assumption: some expert makes 0 mistakes

C_t is the set of experts that have made no mistakes up to day t

Let $C_1 \leftarrow \{1, \ldots, N\}$
For $t = 1, \ldots, T$:
    Let p_t be the majority vote of experts in C_t
    Let C_{t+1} be the experts with no mistakes so far

Level I: The Halving Algorithm

$M^* = 0 \implies M \le \log_2 N$

- **Thm:** If some expert makes 0 mistakes then HA makes $\le \log_2 N$ mistakes.

- Key Idea: Use $|C_t|$ as a "measure of progress"

- $N = |C_1| \ge |C_2| \ge |C_3| \ge \ldots \ge |C_T| \ge 1$

some expert never makes a mistake

① $|C_T| \le N \cdot 2^{-M}$

(because if we make a mistake on day t then $|C_{t+1}| \le |C_t|/2$)

② $|C_T| \ge 1$

$$1 \leq |C_T| \leq N \cdot 2^{-M}$$

$$\Rightarrow \quad 1 \leq N \cdot 2^{-M}$$

$$\Rightarrow \quad 0 \leq \log_2 N - M$$

$$\Rightarrow \quad M \leq \log_2 N \qquad \square$$

Level II: Repeated Halving (RHA)

Let $C_1 \leftarrow \{1, \ldots, n\}$
For $t = 1, \ldots, T$:
  Let p_t be the majority vote of experts in C_t
  Let C_{t+1} be the experts with no mistakes so far
  If $C_{t+1} = \emptyset$, let $C_{t+1} = \{1, \ldots, n\}$

- Suppose some expert makes M^* mistakes

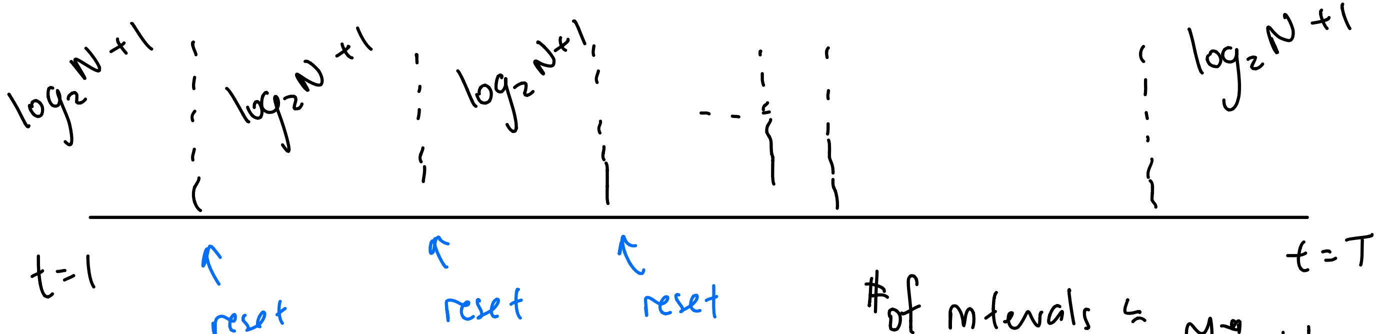

# of intervals $\leq M^* + 1$

Level II: Repeated Halving

- **Thm:** If some expert makes $\leq M^*$ mistakes then RHA makes $\leq (M^* + 1)(\log_2 N + 1)$ mistakes.

- $M \approx M^* \cdot \log_2 N$
- Meaningless unless some expert is correct a $\left(1 - \frac{1}{\log_2 N}\right)$ fraction of the time

Level III: Weighted Majority (WM)

Give each expert a weight $w_{t,i} \leftarrow 1$ $\quad W_t = \sum_{i=1}^{N} w_{t,i}$

For $t = 1, \ldots, T$:

Let p_t be the weighted majority vote of experts

For $i = 1, \ldots, N$:

If (expert i made a mistake): $w_{t+1,i} \leftarrow \frac{w_{t,i}}{2}$ can be $(1-\varepsilon) w_{t,i}$

Else: $w_{t+1,i} \leftarrow w_{t,i}$

$$\sum_{i:\, f_{t,i}=1} w_{t,i} \quad \text{vs.} \quad \sum_{i:\, f_{t,i}=0} w_{t,i}$$

One of these is at least $\frac{W_t}{2}$

If we make a mistake on day t, $W_{t+1} \leq \left(\frac{3}{4}\right) W_t$

Level III: Weighted Majority

want this to be ≈ 1

- **Thm:** If some expert makes $\le M^*$ mistakes then WM makes $\le 2.4(M^* + \log_2 N)$ mistakes.

- Proof: "Measure of progress" W_t

- $W_T \le N \cdot \left(\frac{3}{4}\right)^M$

- $W_T \ge w_{T, i^*} \ge \left(\frac{1}{2}\right)^{M^*}$

(i* is the best expert)

- $\left(\frac{1}{2}\right)^{M^*} \le N \cdot \left(\frac{3}{4}\right)^M$

$$M^* \cdot \log_2\left(\frac{1}{2}\right) \leq \log_2 N + M \cdot \log_2\left(\frac{3}{4}\right)$$

$\log\left(\frac{1}{x}\right)$
$= -\log(x)$

$$-M^* \cdot \cancel{\log_2(2)} \leq \log_2 N - M \cdot \log_2\left(\frac{4}{3}\right)$$

$$M \leq \frac{M^* + \log_2 N}{\log_2(4/3)}$$

Level III: Weighted Majority

- **Thm:** If some expert makes $\leq M^*$ mistakes then WM makes $\leq 2.4(M + \log_2 N)$ mistakes.

Some expert has wt $\geq 2^{-M^*}$

Every mistake decreases total wt by a factor of $\left(\frac{1}{4}\right)$

$$2^{-M^*} \leq W_T \leq N \cdot \left(\frac{3}{4}\right)^M$$

$$-M^* \leq \log_2 N + M \log_2 \left(\frac{3}{4}\right)$$

$$M \leq \left(M^* + \log_2 N\right) \frac{-1}{\log_2 (3/4)}$$

$$= \frac{1}{\log_2 (4/3)} \approx 2.4$$

Level III: Weighted Majority

- **Thm:** If some expert makes $\leq M^*$ mistakes then WM makes $\leq 2.4(M + \log_2 N)$ mistakes.

Some expert has wt $\geq (1-\varepsilon)^{M^*}$

Every mistake decreases total wt by a factor of $\left(1-\frac{\varepsilon}{2}\right)$

$$(1-\varepsilon)^{M^*} \leq N \cdot \left(1-\frac{\varepsilon}{2}\right)^M$$

$$M^* \cdot \ln(1-\varepsilon) \leq \ln N + M \cdot \ln\left(1-\frac{\varepsilon}{2}\right)$$

$$-M^* \cdot \ln\left(\frac{1}{1-\varepsilon}\right) \leq \ln N - M \cdot \ln\left(\frac{1}{1-\varepsilon/2}\right)$$

$$M \leq M^* \cdot \left(\frac{\ln(1/1-\varepsilon)}{\ln(1/1-\varepsilon/2)}\right) + \frac{\ln N}{\ln(1/1-\varepsilon)}$$

Level III: Weighted Majority

- **Thm:** If some expert makes $\leq M$ mistakes then WM makes $\leq 2.4(M + \log_2 N)$ mistakes.

- **Thm**: Any **deterministic** strategy can be forced to make at least $2M^*$ mistakes

① we have to put the whole dollar on one expert, would like to split the dollar on many expert

② have to make a single prediction, would like to randomize

Level IV: Randomized Weighted Majority

Give each expert a weight $w_{t,i} \leftarrow 1$, $W_t \leftarrow \sum_i w_{t,i}$

For $t = 1, \ldots, T$:

  Choose i with probability $w_{t,i}/W_t$

  For $i = 1, \ldots, N$:

    If (expert i made a mistake): $w_{t+1,i} \leftarrow (1-\varepsilon) \cdot w_{t,i}$

    Else: $w_{t+1,i} \leftarrow w_{t,i}$, $W_{t+1} \leftarrow \sum_i w_{t+1,i}$

On day t, I lose

$$F_t = \sum_{i:\, f_{t,i} \neq q_t} w_{t,i}/W_t$$

In total, I lose

$$M = \sum_{t=1}^{T} F_t$$

Level IV: Randomized Weighted Majority

- **Thm:** If some expert makes $\leq M^*$ mistakes then RWM makes $\leq (1+\varepsilon)\cdot M^* + \frac{\log_2 N}{\varepsilon}$ mistakes

Set $\varepsilon = \sqrt{\frac{\log_2 N}{T}}$

Then
$$M \leq M^* + M^* \cdot \sqrt{\frac{\log_2 N}{T}} + \sqrt{T \log_2 N}$$
$$\leq M^* + \sqrt{T \log_2 N} + \sqrt{T \log_2 N}$$
$$= M^* + 2\sqrt{T \log_2 N}$$

Level IV: Randomized Weighted Majority

- **Thm:** If some expert makes $\leq M$ mistakes then RWM makes $\leq (1+\varepsilon)\cdot M + \frac{\log_2 N}{\varepsilon}$ mistakes

Proof: measure of progress is W_t

goal is to bound $\sum_{t=1}^{T} F_t$

- $W_T \geq (1-\varepsilon)^{M^*}$

- If I make a mistake on day t

then $W_{t+1} = (1-\varepsilon F_t)\cdot W_t$

- $W_T \leq N \cdot \prod_{t=1}^{T} (1-\varepsilon F_t)$

$$(1-\varepsilon)^{M^*} \le N \cdot \prod_{t=1}^{T} (1-\varepsilon F_t)$$

$$M^* \cdot \ln(1-\varepsilon) \le \ln(N) + \sum_{t=1}^{T} \ln(1-\varepsilon F_t)$$

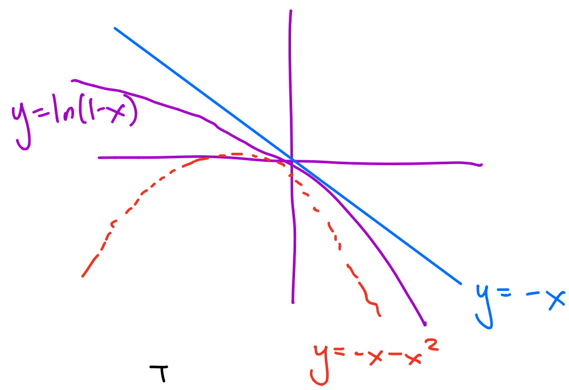

$$M^* \cdot \ln(1-\varepsilon) \le \ln(N) - \varepsilon \sum_{t=1}^{T} F_t$$

$$\sum_{t=1}^{T} F_t \le \frac{-M^* \cdot \ln(1-\varepsilon) + \ln(N)}{\varepsilon}$$

$$= M^* \left(\frac{-\ln(1-\varepsilon)}{\varepsilon} \right) + \frac{\ln(N)}{\varepsilon}$$

$$\le M^* \left(\frac{\varepsilon + \varepsilon^2}{\varepsilon} \right) + \frac{\ln(N)}{\varepsilon}$$

$$= M^* (1+\varepsilon) + \frac{\ln(N)}{\varepsilon}$$

Level IV: Randomized Weighted Majority

- **Thm:** If some expert makes $\leq M$ mistakes then RWM makes $\leq (1+\varepsilon) \cdot M + \frac{\log_2 N}{\varepsilon}$ mistakes

Let F_t be the fraction of mistakes we made on day t

Want to bound $\sum_t F_t$ $\qquad F_t \propto \sum_i w_t m_{t,i}$

$$(1-\varepsilon)^{M^*} \leq W_T \leq N \cdot \prod_{t=1}^{T} (1-\varepsilon F_t)$$

$$M^* \cdot \ln(1-\varepsilon) \leq \ln N + \sum_{t=1}^{T} \ln(1-\varepsilon F_t)$$

$$\leq \ln N - \sum_{t=1}^{T} \varepsilon F_t$$

$$\sum_{t=1}^{T} F_t \leq \frac{\ln N}{\varepsilon} + \frac{M^*}{\varepsilon \ln(1/1-\varepsilon)}$$

Why I love this algorithm

- **Endless applications:**
  - continuous optimization / linear programming
    - including maximum flow!
  - machine learning
    - training machine learning models
    - combining weak models into strong models
    - online learning: updating models with more data
  - probability theory
  - game theory
    - how to play zero-sum games
  - theory of computation

Why I care so much about this

- We often teach algorithms as a set of *ad hoc* tricks
  - These algorithms are easier to deploy
  - These algorithms are used often
  - These algorithms came first historically
  - These algorithms require less mathematical background
- Algorithms research today is much more systematic
  - More powerful and unified techniques
  - But requires more mathematical sophistication
    - Randomization / Probability / Statistics
    - Continuous Mathematics / Linear Algebra
  - **But beautiful and worth studying!**