Lecture 11: Intro to Online Learning

Batch Statistical Learning $(x_1, y_1)\ (x_2, y_2) \cdots (x_n, y_n) \sim D$

Realizable Setting or $x \sim D$, $y = h^*(x)$ $h^*: X \to \{0,1\}$

$h \in \mathcal{H}$ Hypothesis class

Goal find h with low error $err_D(h) = \mathbb{E}_{x \sim D} \mathbb{1}(h^*(x) \neq h(x)) \to 0$

Online Learning you see the data point one at a time $(x_1, y_1) \cdots (x_t, y_t)$

Alg:
h_1: [illegible] x_1
$h_2 \in (x_1, y_1), x_2$
$h_3 \in (x_1, y_1) (x_2, y_2), x_3$
$\vdots$
$h_{n+1} \in (x_1, y_1), \ldots, (x_n, y_n), x_{n+1}$

How to evaluate this algorithm?

Mistake Bound. M:

Alg A learns class H if A makes $\leq M$ mistakes on any sequence of examples consistent w/ some $h \in \mathcal{H}$

Example. $\mathcal{X} = \{0,1\}^d$

$\mathcal{H}$: the class of all disjunctions e.g. $\{X[5]$ or $\bar{X}[8]$ or $X[16]\}$

Alg 1: ERM

① $V_1 = \mathcal{H}$
$t = 1, 2, \ldots$
receive x_t, pick any $h \in V_t$
Predict $\hat{y}_t = h(x_t)$
Receive $y_t = h^*(x_t)$
update $V_{t+1} = \{h \in V_t : h(x_t) = y_t\}$

$\exists$ sequence of x_t, h^* and sequence of consistent h s.t. $M_{(ERM)} = |\mathcal{H}| - 1$

Bad!

Alg 2. Halving

① $V_1 = \mathcal{H}$
$t = 1, 2, \ldots$
Receive x_t
Vote using all h in V_t
$\hat{y}_t = \arg\max_{r \in \{0,1\}} |\{h \in V_t ; h(x_t) = r\}|$

Receive y_t,
~~if mistake do~~ Update $V_{t+1} = \{h \in V_t : h(x_t) = y_t\}$

Observation: Each mistake we make

Proof: At least $\frac{|V_t|}{2}$ hyp-theses are wrong

$$1 \le |V_{T+1}| \le |\mathcal{H}| 2^{-M} \Rightarrow M_{(\text{Halving})} \le \log_2(|\mathcal{H}|)$$

Is this optimal? ~~...~~ Consider a sequence with you make mistake
M bits of information to describe the unique hypothesis in $\mathcal{H}$

What if Not realizable? Regret $= \sum_{t=1}^{n} \mathbb{1}(h_t(x_t) \neq y_t) - \min_{h \in \mathcal{H}} \sum_{t=1}^{n} \mathbb{1}(h(x_t) \neq y_t)$

Example: Stock prediction with Experts

	Exp 1 (Chang)	Exp 2 (Yugang)	Exp 3 (Omid)	Exp 4 (Octopus Paul) throwing	Outcome
Day 1	Down	Up	Up	Down	Down
Day 2	Up	Up	Down	Down	Down
Day 3	Up	Down	Up	Up	Up

WM: Weighted Majority:

Start weight	1	$\frac{1}{2}$	$\frac{1}{2}$	1	(Change Ring) Day 1	
	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{2}$	1	choose Up	~~Down~~
					c	Day 2
	$\frac{1}{2}$	$\frac{1}{8}$	$\frac{1}{2}$	1		Day 3

Claim

Proof: M = # of mistakes so far

m = # of mistakes best expert made so far

W = total weight (start at n, n = # of experts)

Every mistake ⇒ W drops at least 25% (why?)

weight of best expert is $(1/2)^m$

So $\left(\frac{1}{2}\right)^m \le n\left(\frac{3}{4}\right)^M$, $\left(\frac{4}{3}\right)^M \le n\,2^m$

take log $M \le \underbrace{2.4}_{\text{constant ratio}}(m + \log n)$

(still, pretty good!)

If m is 0 the proof ... no ... fixed

Idea: use $1-\varepsilon$ instead of $\frac{1}{2}$, choose ε, carefully

$(1-\varepsilon)^m \le n\left(1-\frac{1}{2}(1+\varepsilon)\right)^M$ take log move things around

$n\left(\frac{1}{2}+\frac{\varepsilon}{2}\right)^M$

~~choose $\varepsilon = \frac{\sqrt{\log n}}{\sqrt{m}} \Rightarrow M \le m + O(\sqrt{m \log n})$~~ $\le m + O(\sqrt{T \log n})$

$m \log(1-\varepsilon) \le \log n + M\left(\log\left(\frac{1}{2}+\frac{\varepsilon}{2}\right)\right)$

$$M \le \frac{\log(1-\varepsilon)}{\log\left(\frac{1}{2}+\frac{\varepsilon}{2}\right)} m - \frac{\log n}{\log\left(\frac{1}{2}+\frac{\varepsilon}{2}\right)}$$

$$\le 2(1+\varepsilon)m + O(\log n)$$

Can be improved

Randomized Weighted Majority

1. Set weight $w_i^{(1)} = 1$ for all i

2. At time t, output $\begin{cases} 1 & \text{with probability } \dfrac{\sum_{i=1}^{n} w_i^{(t)} \mathbb{1}(\hat{y}_i = 1)}{\sum_{i=1}^{n} w_i^{(t)}} \\ 0 & \text{otherwise} \end{cases}$

3. $w_i^{(t+1)} \leftarrow w_i \cdot (1-\varepsilon)$ if i makes a mistake.

Analysis: At time t, F_t is the fraction of ~~the fraction~~ weights on experts that make mistake

$W = n(1-\varepsilon F_1)(1-\varepsilon F_2) \cdots$

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

when there is a mistake

$\log(W_{final}) = \log n + \sum_{t=1} \log(1-\varepsilon F_t) \leq \log n - \varepsilon \sum_t F_t$

$(\log(1-x) \leq -x)$

$= \ln(n) - \varepsilon M$

$(\sum F_t = E[\# \text{ number of mistakes}] = M)$

$\log(W_{final}) \geq \log((1-\varepsilon)^m)$

$\log(n) - \varepsilon M > m \log(1-\varepsilon)$

$M < \frac{\log(1-\varepsilon)}{\varepsilon} m + \frac{\log n}{\varepsilon} \approx (1+\frac{\varepsilon}{2}) m + \frac{1}{\varepsilon} \log n$

$= OPT + \varepsilon m + \frac{\log n}{\varepsilon}$

Regret $= \sqrt{m \log n} \leq \sqrt{T \log n}$

~~Optimization View~~

~~$\min \sum f_i(\theta)$ only~~

~~$\theta \in$ probability~~

~~More generally~~

~~$\min \sum \langle l_t, \theta \rangle \in$ Bounded~~

~~$\theta \in$ probability $0 \le V \le 1$~~

Hedge alg

① Pick expert i with prob $\propto \frac{W_t(i)}{\sum W_t(i)}$

② incur loss $\ell_t(i)$

③ $W_{t+1} = W_t(i)\, e^{-\varepsilon \ell_t(i)}$ for all $i \in [n]$

same analysis

Optimization View $\min \sum f_i(\theta)$ online f_i is linear $\langle x_i, \theta \rangle$

$\theta \in$ probability $0 \le x_i \le 1$

① Make decision θ
② Receive feedback
③ Incur loss

More generally f_i is convex? f_i is strongly convex?

θ is general convex set?

all low rank matrices? Online Convex Optimization

- adversarial learner (no distribution assumption)
- Do not need training data. learn & make decision on the fly
- Usually come with theoretical guarantees

Homework due on thursday