

Office hours

Daniel : Tue after lecture until 4:35pm (426 Mudd)

Noah : Mon 3-5pm (467 Mudd)

Nikhil : TBD

See website / Edstem for announcements, e.g., OH changes

PSØ: Problem 1(b): Prove $\Pr(X=0) \cdot E(X^2) \leq \text{var}(X)$

ONLINE ALLOCATION

① In expert advice scenario w/ K -mistake assume

WM has M.B.

$$A \cdot K + B \cdot \log N$$

Can we change things so we can get $A \approx 1$?

② O.A. is a "central" or "core" problem in ML.

③ Alg. for OA gives a way to prove other theorems

A. Issue in E.A.S.

WM: on day t : $Z_{t,0}$ = sum of WM weights on experts predicting 0
 $Z_{t,1}$ = ... predicting 1

	weight on "0"	weight on "1"	WM predicts	y_t True outcome
Day 1	1 X	1 ✓	0	1
Day 2	$\frac{1}{2}$ ✓	1 X	1	0
Day 3	$\frac{1}{2}$	$\frac{1}{2}$	0	1

Suppose WM randomizes its predictions by choosing

$$a_t = \begin{cases} 1 & \text{w.p. } \frac{Z_{t,1}}{Z_t} \\ 0 & \text{o.w.} \end{cases}$$

$$\Pr[a_t \neq y_t] = 1 - \frac{Z_{t,y_t}}{Z_t}$$

This is OK if adversary doesn't see the coins being tossed...

i.e. adversary commits to y_t before coins are tossed

OBVIOUS ADVERSARY ASSUMPTION

This is enough to let Randomized version of WM to get

$$E[\text{Mistakes}] \leq AK + B \log N$$

with $A \approx 1$

③ Online allocation problem in $\mathbb{R}^N$

For $t=1, 2, \dots, T$ $\rightarrow$ horizon

• Learner chooses an allocation p_t
an N -vector whose components are ≥ 0
and sum to 1 $\left. \vphantom{\begin{matrix} \text{an } N\text{-vector} \\ \text{whose components are } \geq 0 \\ \text{and sum to 1} \end{matrix}} \right\} \Delta^{N-1}$
(probability simplex)

• learner observes loss vector l_t
an N -vector whose components are ≥ 0 $\left. \vphantom{\begin{matrix} \text{an } N\text{-vector} \\ \text{whose components are } \geq 0 \end{matrix}} \right\} \mathbb{R}_+^N$

• learner incurs loss $\langle p_t, l_t \rangle$

$$L_T = \sum_{t=1}^T \langle p_t, l_t \rangle \quad \text{cumulative loss of learner after } T \text{ rounds}$$

For any allocation vector $q \in \Delta^{N-1}$

$$L_{T,q} = \sum_{t=1}^T \langle q, l_t \rangle \quad \text{cumulative loss of fixed alloc. vector } q.$$
$$= \langle q, \sum_{t=1}^T l_t \rangle$$



$$R_{T,q} = L_T - L_{T,q} \quad \text{regret of learner to alloc. vector } q.$$

$$R_T = \max_{q \in \Delta^{N-1}} R_{T,q} \quad \text{regret to the best allocation}$$
$$= \max_{q \in \Delta^{N-1}} L_T - L_{T,q}$$

$$= L_T - \min_{q \in \Delta^{N-1}} L_{T,q}$$

$$\min_{q \in \Delta^{N-1}} L_{T,q} = \min_{q \in \Delta^{N-1}} \left\langle q, \sum_{t=1}^T l_t \right\rangle = \min_{i \in [N]} \sum_{t=1}^T l_{t,i}$$

$$l_t = (l_{t,1}, \dots, l_{t,N})$$



[C] Hodge algorithm

Hyperparameter $\eta > 0$ (step size, learning rate)

Initialize: $W_1 = (1, 1, \dots, 1)$

For $t = 1, 2, \dots, T$

Choose $p_{t,i} = \frac{W_{t,i}}{Z_t} \quad \forall i \in [N]$
 where $Z_t = \sum_{i=1}^N W_{t,i}$

Get $l_t = (l_{t,1}, \dots, l_{t,N}) \in \mathbb{R}_+^N$

Update

$$W_{t+1,i} = W_{t,i} \exp(-\eta l_{t,i}) \quad \forall i \in [N]$$

Theorem: In O.A. for $\mathbb{R}^N$, for any T and loss vectors $l_1, \dots, l_T \in [0, 1]^N$

$$\begin{aligned} \textcircled{1} \quad R_T &\leq \frac{\ln N}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i=1}^N p_{t,i} l_{t,i}^2 \\ &\leq \frac{\ln N}{\eta} + \frac{\eta}{2} \cdot T \end{aligned}$$

$\underbrace{\langle p_t, l_t^2 \rangle}_{\leq \langle p_t, l_t \rangle} \leq \langle p_t, l_t \rangle$

If η is chosen to balance two terms $\frac{\ln N}{\eta}$ and $\frac{\eta T}{2}$,

$$\text{then } R_T \leq 2 \cdot \sqrt{\ln N \cdot \frac{T}{2}} = \sqrt{2T \ln N}$$

$$\frac{R_T}{T} \leq \sqrt{\frac{2 \ln N}{T}} \rightarrow 0 \text{ as } T \rightarrow \infty$$

"No-regret": R_T is sublinear in T

$$\textcircled{2} \quad (1 - \frac{\eta}{2}) L_T \leq \min_{i \in [N]} L_{T,i} + \frac{\ln N}{\eta}$$

If $\eta \in (0, 2]$,

$$\frac{1}{1-x} \approx 1+x \quad L_T \leq \frac{1}{1-\eta/2} \left(\min_{i \in [N]} L_{T,i} + \frac{\ln N}{\eta} \right)$$

$$\frac{1}{1-x} = 1+x+x^2+\dots \approx (1+\eta) \min_{i \in [N]} L_{T,i} + \frac{1+\eta}{\eta} \ln N$$

Choosing η in some way to balance $\eta \cdot \min_{i \in [N]} L_{T,i}$

and $\frac{1+\eta}{\eta} \ln N$ gives

$$R_T \leq \sqrt{2L^* \ln N} + L^* N + \ln N$$

← "small loss regret bound" or "L-star bound"

WM:

$$\text{Mistakes} \leq AK + B \log N \quad A \approx 2$$

$$\text{Mistakes} - K \leq (A-1)K + B \log N$$

D Proof of Hedge Theorem

Using concepts from information theory

Entropy $H(q) = \sum_{i=1}^N q_i \ln\left(\frac{1}{q_i}\right)$

measure of 'surprise' you feel when i^{th} outcome occurs : $\ln\left(\frac{1}{q_i}\right)$

Entropy is expected surprise you feel when outcome is $\sim q$.

Cross entropy $CE(q, p) = \sum_{i=1}^N q_i \ln\left(\frac{1}{p_i}\right)$

Relative entropy $RE(q, p) = CE(q, p) - H(q)$

Fact $RE(q, p) \geq 0$, equality iff $q = p$

"Gibbs' inequality"