LECTURE 19

Last Lecture

Inequalities for Tail Probabilities.

X : r.v

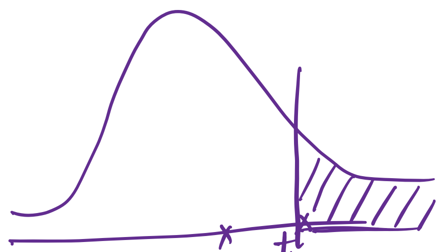

μ : mean of X

$\mathbb{P}\{X \geq t\}$

$t \geq \mu$

$\downarrow$

Right tail probability

Chernoff Bound:

$$\mathbb{P}\{X \geq t\} \leq \min_{\lambda \geq 0} \frac{\mathbb{E}\, e^{\lambda X}}{e^{\lambda t}}$$

} Chernoff

Left Tail:

$$\mathbb{P}\{X \leq t\} \qquad t \leq \mu$$

$$= \mathbb{P}\{-X \geq -t\} \qquad -t \geq -\mu$$

$$\leq \min_{\lambda \geq 0} \frac{\mathbb{E}\, e^{\lambda(-X)}}{e^{\lambda(-t)}}$$

---

① Binomial $\boxed{X \sim \text{Bin}(n,p)}$

mean $= np$

$$\mathbb{P}\{X \geq t\} \leq \boxed{\exp\left[-n\, D\left(\tfrac{t}{n}, p\right)\right]} \qquad \leq \delta$$

$t \geq np \qquad D(\rho, p) = \rho \log \frac{\rho}{p} + (1-\rho)\log\frac{1-\rho}{1-p}$

② Hoeffding Bound:

$$\mathbb{P}[\underbrace{X}_{\mathrm{Bin}(n,p)} \geq \underbrace{t}_{t}] \leq \exp\left[-2n\left(\frac{t}{n}-p\right)^2\right]$$

$$= \underbrace{\exp\left[-\frac{2(t-np)^2}{n}\right]}_{\delta}$$

Solve

$$\exp\left[-2\frac{(t-np)^2}{n}\right] = \delta$$

$$\frac{(t-np)^2}{n} = \frac{1}{2}\log\frac{1}{\delta}$$

$t \geq np$

$$t - np = \sqrt{\frac{n}{2}\log\frac{1}{\delta}}$$

$$t = np + \sqrt{\frac{n}{2}\log\frac{1}{\delta}}$$

$\therefore$ $X \leq np + \sqrt{\frac{n}{2}\log\frac{1}{\delta}}$ with probability at least $1-\delta$

③ General Hoeffding:

Binomial: $X = X_1 + \cdots + X_n$

$X = X_1 + \cdots + X_n$

for some $X_1 \ldots X_n$

such that ⓐ each X_i is bounded between a & b.

ⓑ $X_1 \ldots X_n$ are independent

$\mathbb{P}[X \geq t]$

$t \geq \mathbb{E}X = \mathbb{E}(X_1) + \cdots + \mathbb{E}(X_n)$

notation: $\mu_i = \mathbb{E}X_i$

$\mathbb{E}X = \sum \mu_i$

$$t \geq \sum \mu_i$$

$$\leq \exp\left[- \frac{2 (t - \sum \mu_i)^2}{n (b-a)^2} \right]$$

$$0.05$$

Note: $X \sim \text{Bin}(n, p)$ $\quad t \leq np$

$$\mathbb{P}(X \leq t) ,$$

$$= \mathbb{P}\left[\sum_{i=1}^{n} (-X_i) \geq -t \right]$$

$$\leq \exp\left[- \frac{2 (t - np)^2}{n} \right]$$

$$= \delta$$

$$\frac{2 (t - np)^2}{n} = \log \frac{1}{\delta}$$

$$np - t = \sqrt{\frac{n}{2} \log \frac{1}{\delta}}$$

$$t = np - \sqrt{\frac{n}{2} \log \frac{1}{\delta}}$$

Deduced:

$$\mathbb{P}\left[X \leq np - \sqrt{\frac{n}{2} \log \frac{1}{\delta}} \right] \leq \delta$$

or $\mathbb{P}\left[X \geq np - \sqrt{\frac{n}{2} \log \frac{1}{\delta}} \right] \geq 1 - \delta$

$$\mathbb{P}\left[p \leq \frac{X}{n} + \sqrt{\frac{\log 1/\delta}{2n}} \right] \geq 1 - \delta$$

$$\boxed{p} \leq \underbrace{\frac{X}{n} + \sqrt{\frac{\log 1/\delta}{2n}}}_{UCB} \quad w.p \geq 1-\delta$$

UCB = Upper Confidence Bound

MULTI-ARMED BANDIT

I have K coins:
probability of success: $p_1, p_2, \ldots, p_K$

Each round: You pick a coin that I will toss.

Heads → you get \$1
Tails → you get \$0

We play for T=1000 rounds.

Qn: How would you play?

① Explore then $\boxed{\text{Commit}}$ (ETC)

Fix a number $m \geq 1$.

Exploration
ⓐ Toss each coin $\boxed{m}$ times. ($m \times K$ rounds)
ⓑ $\hat{p}_i$ = proportion of observed heads for coin i
ⓒ Pick the coin where $\hat{p}_i$ is the largest

Keep tossing that coin.

$K = 9$

$\underbrace{p_1 = 0.1, \; p_2 = 0.2, \; p_3 = 0.3, \ldots,}$ $\boxed{p_8 = 0.8}$, $\boxed{p_9 = 0.9}$

Best possible Reward = \$900

Qn: What value of m would make it highly unlikely for us to pick 8 instead of 9?

Ans: $\underbrace{\mathbb{P}\left[\text{Bin}(m, 0.8) \geq \text{Bin}(m, 0.9) \right]} \leq \delta$

δ: 0.1 or 0.05

$= \mathbb{P}\left[\sum Y_i \geq \sum X_i \right]$

$Y_i \sim \text{Ber}(0.8)$

$X_i \sim \text{Ber}(0.9)$

$= \mathbb{P}\left[\sum_{i=1}^{m} (Y_i - X_i) \geq 0 \right]$

Right tail

$\boxed{t = 0}$

$\underbrace{-1}_{a} \leq Y_i - X_i \leq \underbrace{1}_{b}$

$\mathbb{E}\left(\sum_{i=1}^{m} (Y_i - X_i) \right) = -m(0.1)$

$\leq \exp\left[-\frac{2[m(0.1)]^2}{m \, 4} \right]$

$= \exp\left[-\frac{m(0.01)}{2} \right]$

$= \exp\left[-\frac{m}{200} \right] = \delta$

$m = 200 \times \log\left(\frac{1}{\delta}\right)$

$\delta = 0.1$ $\quad$ $m = 460$

UCB (Upper Confidence Bound Algorithm)

Version 1:

(a) Basic Exploration. ($m=1$)

(b) For $t = K+1, \dots$

$\underbrace{\hat{p}_1}, \underbrace{\hat{p}_2}, \dots, \underbrace{\hat{p}_K}$: Historical data up to round $t-1$.

Pick the coin with highest $\hat{p}_i$.

eg: Exploration:

1	2	3	. .			K
0	0	$\boxed{1}$	0	$\boxed{1}$	.	$\boxed{1}$
0	0	$\frac{1}{2}$ (0)	0	1	. .	$\boxed{1}$
0 . .	0	$\frac{1}{2}$	0	1	. .	1

$\frac{0.1}{0}$ . . . 0 $\frac{0.7}{1}$ $\frac{0.8}{1}$ $\frac{0.9}{0}$

UCB

① Basic exploration $m = 1$

② For $t = K+1, \dots$

$$\hat{p}_i^{(t)} + \sqrt{\frac{\log 1/\delta}{2\, n_i(t)}}$$

& pick coin with highest value of ↑

δ : some small number.

n_i : number of times the i^{th} coin was tossed in previous round.

Thompson Sampling

Prior: $p_1 \dots p_K \overset{iid}{\sim} \text{Unif}[0,1]$

After t rounds of the game,

$n_i(t)$ = # tosses for i^{th} coin upto round t.

$X_i(t)$ = # heads for i^{th} coin upto round t

Posterior for coin i:

$$\text{Beta}\left(X_i(t) + 1,\ n_i(t) - X_i(t) + 1\right)$$

Sample from these K posterior distributions.

$$\boxed{p_1^{(t)}, \; p_2^{(t)}, \; \ldots, \; p_K^{(t)}}$$

Pick coin for which $p_i^{(t)}$ is the largest.

Start the process at $t = 0$.

(Initially $n_i(0) = 0$, $X_i(0) = 0$)