

- sub-Gaussian R.V.s

A mean-zero R.V. Z is σ^2 -sub-Gaussian if $\mathbb{E}[\exp(\lambda Z)] \leq \exp(\sigma^2 \lambda^2 / 2) \quad \forall \lambda \in \mathbb{R}$

- Chernoff bound w/ sub-Gaussian R.V.s:

If X_1, \dots, X_n are IID R.V. w/ $\mathbb{E}[X_i] = \mu$, and $X_i - \mu$ is σ^2 -sub-Gaussian, then

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n X_i - \mu > \varepsilon\right) \leq \exp\left(-\frac{n\varepsilon^2}{2\sigma^2}\right) = \delta$$


- Hoeffding's lemma: (If $X \in [a, b]$, then take $a=0, b=1 \Rightarrow X - \mathbb{E}[X]$ is $\frac{1}{4}$ -sub-Gaussian)

If $X \in [a, b]$ then $X - \mathbb{E}[X]$ is $\frac{(b-a)^2}{4}$ -sub-Gaussian

- Union Bound (or, sub-additivity of probability)

Let A_1, \dots, A_n be events in a well-defined probability space,

then $\mathbb{P}\left(\bigcup_{i=1}^n \{A_i\}\right) \leq \sum_{i=1}^n \mathbb{P}(A_i)$



Given n treatments.

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

Patient t arrives w/ illness

Agent prescribes treatment $I_t \in [n] = \{1, \dots, n\}$

Observe $X_{I_t, t} = \mathbb{1}\{\text{patient } t \text{ recovers given } I_t\}$

Let $X_{i,t} = \mathbb{1}\{\text{patient } t \text{ recovers given } i\}$.

Assume for each $i \in [n]$, $\{X_{i,t}\}_{t=1}^T$ are IID w/ $\mathbb{E}[X_{i,1}] = \theta_i$

perhaps $X_{i,t} = \mathbb{1}\{e_{i,t} > 1\}$

Questions:

- How long does it take to identify the best treatment $\arg \max_{i \in [n]} \theta_i$ with probability at least $1 - \delta$?

- If $R_T = \max_{i \in [n]} \mathbb{E}\left[\sum_{t=1}^T X_{i,t} - X_{I_t, t}\right]$, how can we minimize R_T ?

- Expected Regret

$$R_T = \max_{i \in [n]} \mathbb{E} \left[\sum_{t=1}^T X_{i,t} - X_{I_t,t} \right] = \sum_{i=1}^n \Delta_i \mathbb{E}[T_i]$$

where $\Delta_i = \max_{j \in [n]} \theta_j - \theta_i$, $T_i = \sum_{t=1}^T \mathbb{1}\{I_t = i\}$

Warm-up: $n=2$. Consider a strategy that plays action $i \in [n]$ round-robin until $T_i = 3$, then plays $\operatorname{argmax}_{i \in [n]} \hat{\theta}_{i,3}$ for the remainder $T - n3$ times.

↖ empirical mean of first 3 plays.

$i \in [n]$
 Define $\mathcal{E}_i = \left\{ |\hat{\theta}_{i,3} - \theta_i| \leq \sqrt{\frac{2\sigma^2 \log(2n/\delta)}{3}} \right\}$, $\sigma^2 = \frac{1}{4}$

Lemma $\mathbb{P} \left(\bigcap_{i=1}^n \mathcal{E}_i \right) = 1 - \mathbb{P} \left(\bigcup_{i=1}^n \mathcal{E}_i^c \right) \geq 1 - \delta$.

Proof $\hat{\theta}_{i,3} = \frac{1}{3} \sum_{s=1}^3 X_{i,s}$, $X_{i,s} \in \{0,1\}$ are iid $\mathbb{E}[X_{i,1}] = \theta_i$
 $\Rightarrow (X_{i,s} - \theta_i)$ is σ^2 -sub-Gaussian with $\sigma^2 = 1/4$.

$$\begin{aligned} \mathbb{P}(\mathcal{E}_i^c) &= \mathbb{P} \left(|\hat{\theta}_{i,3} - \theta_i| > \sqrt{\frac{2\sigma^2 \log(2n/\delta)}{3}} \right) \\ &\leq 2 \exp \left(-\frac{3}{2\sigma^2} \left(\sqrt{\frac{2\sigma^2 \log(2n/\delta)}{3}} \right)^2 \right) \\ &= 2 \exp \left(-\log(2n/\delta) \right) \\ &= \frac{\delta}{n} \end{aligned}$$

$$P\left(\bigcup_{i=1}^n \{\mathcal{E}_i^c\}\right) \leq \sum_{i=1}^n P(\mathcal{E}_i^c) \leq \sum_{i=1}^n \frac{\delta}{n} = \delta.$$

Because $P\left(\bigcap_{i=1}^n \{\mathcal{E}_i\}\right) \geq 1 - \delta$, assume $\bigcap_{i=1}^n \mathcal{E}_i$ occurs.

Assume WLOG $\theta_1 \geq \theta_2 \geq \dots \geq \theta_n$

If $\mathfrak{I} = \lceil 8\sigma^2 \Delta_2^2 \log(2n/\delta) \rceil$ then

$$\hat{\theta}_{1,\mathfrak{I}} > \theta_1 - \sqrt{\frac{2\sigma^2 \log(2n/\delta)}{\mathfrak{I}}} \quad (\mathcal{E}_1)$$

$$> \theta_1 - \frac{\Delta_2}{2}$$

$$= \theta_2 + \frac{\Delta_2}{2}$$

$$\geq \hat{\theta}_{2,\mathfrak{I}} - \sqrt{\frac{2\sigma^2 \log(2n/\delta)}{\mathfrak{I}}} + \frac{\Delta_2}{2} \quad (\mathcal{E}_2)$$

$$\geq \hat{\theta}_{2,\mathfrak{I}}.$$

\Rightarrow With probability at least $1 - \delta$, strategy identifies $\operatorname{argmax}_{i \in [n]} \theta_i$.

$$\theta_1 \geq \theta_2 \geq \dots \geq \theta_n \quad (n=2 \quad \theta_1 > \theta_2)$$

$$R_T = \sum_{i=1}^n \Delta_i \mathbb{E}[T_i] = \Delta_2 \mathbb{E}[T_2] \quad \text{since } n=2 \text{ and } \Delta_1=0.$$

$$= \Delta_2 \mathbb{E}[T_2 (\mathbb{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}} + \mathbb{1}_{\{\mathcal{E}_1^c \cup \mathcal{E}_2^c\}})]$$

$$= \Delta_2 \mathbb{E}[T_2 \mathbb{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}}] + \Delta_2 \mathbb{E}[T_2 \mathbb{1}_{\{\mathcal{E}_1^c \cup \mathcal{E}_2^c\}}]$$

$$\leq \Delta_2 \mathfrak{I} \mathbb{E}[\mathbb{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}}] + \Delta_2 (T - \mathfrak{I}) \mathbb{E}[\mathbb{1}_{\{\mathcal{E}_1^c \cup \mathcal{E}_2^c\}}]$$

$$\leq 1$$

$$\begin{aligned} &= \mathbb{P}(\cup_{i=1}^n \{E_i^c\}) \\ &\leq \delta \end{aligned}$$

$$\leq \Delta_2 \mathfrak{S} + \Delta_2 T \delta$$

$$\leq 1 + 8\Delta_2^{-1} \sigma^2 \log(2n/\delta) + \Delta_2 T \delta$$

Idea: choose $\delta = 1/T$ ($\Delta_2 \leq 1$)

$$\hookrightarrow \leq 1 + 8\Delta_2^{-1} \sigma^2 \log(2nT) + 1$$


We wanted $R_T = o(T)$ - we have it! $\frac{\log(T)}{T} \rightarrow 0$

Note: If treatments were nearly identical so that $\Delta_2 = 10^{-10}$ then our regret bound says $R_T \leq 10^{10}$.

But this is nonsense, b/c it shouldn't matter what is prescribed, they act nearly the same.

$$\text{We know: } R_T = \Delta_2 \mathbb{E}[T_2] \leq \Delta_2 T$$

$$R_T \leq \min \left\{ 2 + 8\Delta_2^{-1} \sigma^2 \log(2nT), \Delta_2 T \right\}$$



maximize over Δ_2

$$\Rightarrow R_T \leq Z + \max_{\Delta \geq 0} \min \{ 8\Delta^{-1} \sigma^2 \log(2nT), \Delta T \}$$

$$\Delta = \sqrt{\frac{8\sigma^2 \log(2nT)}{T}}$$

$$= Z + \sqrt{8\sigma^2 T \log(2nT)}$$


Note $R_T/T \rightarrow 0$

↑
"instance-indep" regret bound

Multi-armed Bandit problem.

Imagine Casino

← arm that you pull to play. "1-armed bandit"



Each machine i pays out random amount w/ mean θ_i

Elimination-style algorithm

Input: n arms $\mathcal{X} = \{1, \dots, n\}$, confidence level $\delta \in (0, 1)$.

Let $\hat{\mathcal{X}}_1 \leftarrow \mathcal{X}, l \leftarrow 1$

while $|\mathcal{X}_l| > 1$ **do**

$$\epsilon_l = 2^{-l}$$

Pull each arm in \mathcal{X}_l exactly $\tau_l = \lceil 2\epsilon_l^{-2} \log(\frac{4l^2|\mathcal{X}_l|}{\delta}) \rceil$ times

Compute the empirical mean of these rewards $\hat{\theta}_{i,l}$ for all $i \in \mathcal{X}_l$

$$\mathcal{X}_{l+1} \leftarrow \mathcal{X}_l \setminus \{i \in \mathcal{X}_l : \max_{j \in \mathcal{X}_l} \hat{\theta}_{j,l} - \hat{\theta}_{i,l} > 2\epsilon_l\}$$

$$l \leftarrow l + 1$$

Output: $\hat{\mathcal{X}}_{l+1}$ (or play the last arm forever in the regret setting)

Assume observation $X_{i,t}$ at time t is ^{from arm i} 1-sub-Gaussian

and iid and $\mathbb{E}[X_{i,t}] \in [0, 1]$

Assume wLOG $\theta_1 \geq \theta_2 \geq \dots \geq \theta_n$, $\Delta_i = \theta_1 - \theta_i$

Lemma With probability at least $1 - \delta$,

arm $1 \in \mathcal{X}_l$ and $\max_{i \in \mathcal{X}_l} \Delta_i \leq 8\epsilon_l \quad \forall l \in \mathbb{N}$.

Proof For all $l \in \mathbb{N}$ and $i \in [n]$ define

$$\mathcal{E}_{i,l} = \{|\hat{\theta}_i - \theta_i| \leq \epsilon_l\}$$

Note that $\epsilon_l \geq \sqrt{\frac{2 \log(4l^2|\mathcal{X}_l|/\delta)}{5_l}}$. Thus

$$\mathbb{P}\left(\bigcup_{l=1}^{\infty} \bigcup_{i \in \mathcal{X}_l} \mathcal{E}_{i,l}^c\right) \leq \mathbb{P}\left(\bigcup_{l=1}^{\infty} \bigcup_{i=1}^n \mathcal{E}_{i,l}^c\right)$$

$$\leq \sum_{l=1}^{\infty} \sum_{i=1}^n \mathbb{P}(\mathcal{E}_{i,l}^c)$$

$$\leq \sum_{l=1}^{\infty} \sum_{i=1}^n \frac{\delta}{2l^2|\mathcal{X}_l|} \leq \delta \quad (|\mathcal{X}_l| = n)$$

1
2
3