ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE
School of Computer and Communication Sciences

Foundations of Data Science — Assignment date: Wednesday, November 12th, 2025, 13:15
Fall 2025 — Due date: Wednesday, November 12th, 2025, 14:45

Midterm Exam – AAC 231

This exam is open book. No electronic devices of any kind are allowed. There are three problems. Good luck!

Only answers given on this handout count.

Name: ______________________________

SCIPER: ______________________________

Problem 1	/ 8
Problem 2	/ 12
Problem 3	/ 10
Total	/30

Recall: The Beta distribution with parameters α and β is the probability density function given by

$$p_{\alpha,\beta}(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}$$

where $B(\alpha,\beta)$ is the beta function defined as

$$B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}.$$

The mean of a random variable X with probability density function $p_{\alpha,\beta}(x)$ is $\mathbb{E}[X] = \frac{\alpha}{\alpha+\beta}$.

Problem 1 (DPI for TV distance – 8 pts)**.** Show that the Total Variation Distance satisfies the Data Processing Inequality. That is, using the definitions given in Lemma 4.3 of the Lecture Notes: a probability kernel $W(y|x)$, an original distribution P, and the distribution $\tilde{P}$ whose probability mass function is $\tilde{p}(y) = \sum_x W(y|x)p(x)$, show that

$$\delta(\tilde{P}, \tilde{Q}) \leq \delta(P, Q),$$

where $\delta(P, Q)$ denotes the Total Variation distance between P and Q.

Remark: If you refer to class materials, be precise (Theorem or equation numbers, Homework problem identifiers and so on.) Your overall argument must be complete.

(space for problem 2)

(space for problem 2)

Problem 2 (Thompson Sampling for Two-Armed Bernoulli Bandits – 12 pts)**.** Consider the following two-arm setting:

- Let μ_1 and μ_2 be two independent samples from the uniform distribution on the interval $[0, 1]$. As in class, denote $\Delta = |\mu_1 - \mu_2|$. But note that we do not know which arm is better.
- Now, pulling arm $a \in \{1, 2\}$ at round t yields a reward
$$X_t \sim \text{Bernoulli}(\mu_a).$$
- Let $S_a(t-1)$ and $F_a(t-1)$, respectively, denote the number of observed successes (reward=1) and failures (reward=0) of arm a up to round $t-1$ (for $a = 1$ and $a = 2$).

We use the following algorithm (*Thompson Sampling*):

> At round $t = 1, 2, \ldots$:
>
> (a) Sample independently
> $$\tilde{\mu}_1^{(t)} \sim \text{Beta}(1 + S_1(t-1),\, 1 + F_1(t-1))$$
> $$\tilde{\mu}_2^{(t)} \sim \text{Beta}(1 + S_2(t-1),\, 1 + F_2(t-1))$$
>
> **Note:** See the exam's cover page for a reminder about the Beta distribution.
>
> (b) Pull the arm with the larger sample:
> $$A_t = \arg\max_{a \in \{1,2\}} \tilde{\mu}_a^{(t)}.$$
>
> (c) Observe reward $X_t \in \{0, 1\}$ from arm A_t.

As in class, let $T_{\text{bad}}(n)$ denote the number of times the bad arm (the arm with the smaller value μ_a) is pulled up to horizon n. As in class, the regret can then be written as $R_n = \Delta\, \mathbb{E}[T_{\text{bad}}(n)]$.

Tasks. Answer the following:

(a) [2 points] **(Formulation)** Recall that the means μ_a were drawn from a uniform prior distribution. Suppose by time t, arm a has been pulled $T_a(t)$ times with $S_a(t)$ successes and $F_a(t)$ failures. Give an expression for the posterior distribution $p(\mu_a | S_a(t), F_a(t))$. What type of a distribution is it?

(b) [2 points] **(Partition of the time axis)** Show that for $m > 0$

$$\mathbb{E}[T_a(n)] \le m + \mathbb{E}\big[\# \text{ of pulls of arm } a \text{ after it has already been pulled } m \text{ times}\big].$$

(Hint: decompose according to whether $T_a(t) < m$ or $T_a(t) \ge m$, where $T_a(t)$ is the number of pulls of arm a before round t.)

(c) [3 points] **(Concentration of the Beta posterior)** *This part is a little more tedious. You may want to save it for last and first tackle parts (d) and (e).* For a constant $c > 0$ to be chosen later, define

$$m := \left\lceil \frac{c \log n}{\Delta^2} \right\rceil.$$

Argue that once arm a has been pulled m times, the posterior distribution you found in Part *(a)* is concentrated around the true μ_a, in the sense that for any $\varepsilon > 0$, there exist constants $c_1, c_2 > 0$ such that, for all t with $T_a(t) \ge m$,

$$\mathbb{P}\big(\tilde{\mu}_a^{(t)} > \mu_a + \varepsilon \mid H_{t-1}\big) \le c_1 \exp\big(-c_2\, T_a(t)\, \varepsilon^2\big),$$
$$\mathbb{P}\big(\tilde{\mu}_a^{(t)} < \mu_a - \varepsilon \mid H_{t-1}\big) \le c_1 \exp\big(-c_2\, T_a(t)\, \varepsilon^2\big),$$

where as in class, we are using H_{t-1} to denote the entire history up to time $t-1$, that is, $H_{t-1} = (A_1, X_1, A_2, X_2, \ldots, A_{t-1}, X_{t-1})$.

Hint: Calculate $\hat{\mu}_{a,t} := \mathbb{E}[\tilde{\mu}_a^{(t)} \mid H_{t-1}]$. Then, as in class, split by conditioning onto the events G_a and G_a^c, where $G_a := \left\{ \left|\hat{\mu}_{a,t} - \mu_a\right| \le \frac{\varepsilon}{4} \right\}$.

(d) [3 points] **(Bounding the selection of the suboptimal arm)** Use the previous item with $\varepsilon = \Delta/2$. For large enough t (when both arms have been pulled m times), provide an upper bound for the probability that we select the bad arm,

$$\mathbb{P}(A_t = \text{bad} \mid H_{t-1})$$

in terms of some constant $c' > 0, \Delta$ and n.

(e) [2 points] **(Final bound on $\mathbb{E}[T_{\text{bad}}(n)]$)** Pick a large enough c and conclude that

$$\mathbb{E}[T_{\text{bad}}(n)] \le \frac{c \log n}{\Delta^2} + O(1),$$

and hence the regret of Thompson Sampling in this two-armed Bernoulli bandit is bounded by

$$R_n = \Delta\, \mathbb{E}[T_{\text{bad}}(n)] = O\left(\frac{\log n}{\Delta}\right).$$

(space for problem 2)

(space for problem 2)

Problem 3 (Fisher goes Beta – 10 pts)**.** In this problem, we will establish a few properties of the Beta(α, m) distribution when m is an integer. See the exam's cover page for a reminder about the Beta distribution.

Hint: Recall from Homework 4 that for any $x > 0$, we have $x\Gamma(x) = \Gamma(x+1)$, and that $\Gamma(0) = 1$.

(a) [3 points] Let $X \sim \text{Beta}(\alpha, m)$, where m is a positive integer. For $\alpha > 0$, find

$$\mathbb{E}\left[\log X\right].$$

Hint: Try finding $\mathbb{E}_{\alpha,m}\left[\frac{d}{d\alpha}\log p_{\alpha,m}(X)\right]$. *If finding the expectation for a general* m *seems too difficult, try* $m = 1, 2$ *first.*

(b) [3 points] Show that, for any parametric family of density functions $p_\theta(x)$ such that the support is independent of θ, the Fisher information

$$I(\theta) = -\mathbb{E}_\theta\left[\frac{d^2}{d\theta^2}\log p_\theta(X)\right].$$

(c) [4 points] Using part (b), show that for $X \sim \text{Beta}(\alpha, m)$ for a fixed integer $m > 0$, whenever $\alpha \in (0, 1)$,

$$I(\alpha) \leq \frac{1}{\alpha^2} + \frac{\pi^2}{6}.$$

(space for problem 5)

(space for problem 5)