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.

Solution 1. There are several ways in which you can proceed. We provide three different solutions — challenge: can you come up with even more?

Proof 1. Via the variational representation of TV distance, which is Lemma 4.10 of the lecture notes.

$$\begin{aligned}
\delta(\tilde{P}, \tilde{Q}) &= \max_{f:\mathcal{Y}\to[0,1]} \mathbb{E}_{\tilde{P}}[f(Y)] - \mathbb{E}_{\tilde{Q}}[f(Y)] && (1)\\
&= \max_{f:\mathcal{Y}\to[0,1]} \sum_y \tilde{p}(y)f(y) - \sum_y \tilde{q}(y)f(y) && (2)\\
&= \max_{f:\mathcal{Y}\to[0,1]} \sum_y \left(\sum_x W(y|x)p(x)\right) f(y) - \sum_y \left(\sum_x W(y|x)q(x)\right) f(y) && (3)\\
&= \max_{f:\mathcal{Y}\to[0,1]} \sum_x \sum_y W(y|x)p(x)f(y) - W(y|x)q(x)f(y) && (4)\\
&= \max_{f:\mathcal{Y}\to[0,1]} \sum_x \underbrace{\left(\sum_y W(y|x)f(y)\right)}_{=\tilde{f}(x)} p(x) - \left(\sum_y W(y|x)f(y)\right) q(x) && (5)\\
&= \max_{f:\mathcal{Y}\to[0,1]} \mathbb{E}_P[\tilde{f}(X)] - \mathbb{E}_Q[\tilde{f}(X)]. && (6)
\end{aligned}$$

Finally, we observe that if $0 \leq f(y) \leq 1$ for all y, then $0 \leq \tilde{f}(x) \leq 1$ for all x, because

$$\tilde{f}(x) = \sum_y W(y|x)f(y) \leq \sum_y W(y|x) = 1. \tag{7}$$

This is the crucial observation here. How do we exploit this observation? We will now expand the maximization to include *all* possible functions $\tilde{f}(x)$ satisfying $0 \leq \tilde{f}(x) \leq 1$ (not only the ones that are of the form $\tilde{f}(x) = \sum_y W(y|x)f(y)$ for some $f(y)$). Clearly, expanding the maximization can only *increase* the value of the maximum, that is,

$$\begin{aligned}
\delta(\tilde{P}, \tilde{Q}) &= \max_{f:\mathcal{Y}\to[0,1]} \mathbb{E}_P[\tilde{f}(X)] - \mathbb{E}_Q[\tilde{f}(X)] && (8)\\
&\leq \max_{\tilde{f}:\mathcal{X}\to[0,1]} \mathbb{E}_P[\tilde{f}(X)] - \mathbb{E}_Q[\tilde{f}(X)] && (9)\\
&= \delta(P, Q), && (10)
\end{aligned}$$

where in the last step, we have again used the variational representation of TV distance, which is Lemma 4.10 of the lecture notes.

Proof 2. Via the triangle inequality (that is, a "pedestrian" proof).

$$
\begin{aligned}
\delta(\tilde{P}, \tilde{Q}) &= \frac{1}{2} \sum_{y} |\tilde{p}(y) - \tilde{q}(y)| && (11)\\
&= \frac{1}{2} \sum_{y} \left| \sum_{x} W(y|x)p(x) - \sum_{x} W(y|x)q(x) \right| && (12)\\
&= \frac{1}{2} \sum_{y} \left| \sum_{x} W(y|x)(p(x) - q(x)) \right| && (13)\\
&\leq \frac{1}{2} \sum_{y} \sum_{x} |W(y|x)(p(x) - q(x))| && (14)\\
&= \frac{1}{2} \sum_{y} \sum_{x} |W(y|x)| \; |p(x) - q(x)| && (15)\\
&= \frac{1}{2} \sum_{y} \sum_{x} W(y|x)|p(x) - q(x)| && (16)\\
&= \frac{1}{2} \sum_{x} \underbrace{\left(\sum_{y} W(y|x) \right)}_{=1} |p(x) - q(x)| && (17)\\
&= \frac{1}{2} \sum_{x} |p(x) - q(x)| && (18)\\
&= \delta(P, Q), && (19)
\end{aligned}
$$

where the inequality step follows from the triangle inequality $|a + b| \leq |a| + |b|$.

Proof 3. Via the original definition of TV distance, Definition 4.2 in the lecture notes.

$$
\begin{aligned}
\delta(\tilde{P}, \tilde{Q}) &= \max_{\mathcal{S} \subseteq \mathcal{Y}} \tilde{P}(S) - \tilde{Q}(S) && (20)\\
&= \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{y \in \mathcal{S}} \tilde{p}(y) - \sum_{y \in \mathcal{S}} \tilde{q}(y) && (21)\\
&= \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{y \in \mathcal{S}} (\tilde{p}(y) - \tilde{q}(y)) && (22)\\
&= \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{y \in \mathcal{S}} \left(\sum_{x \in \mathcal{X}} W(y|x)(p(x) - q(x)) \right) && (23)\\
&= \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{S}} W(y|x)(p(x) - q(x)) && (24)
\end{aligned}
$$

Now comes the tricky step: Like in the lecture notes, let $\mathcal{A} = \{x \in \mathcal{X} : p(x) \geq q(x)\}$. Consider the last sum over $x \in \mathcal{X}$ and observe that $W(y|x) \geq 0$. Now, we note that all

$x \in \mathcal{A}$ contribute non-negatively. But all $x \notin \mathcal{A}$ contribute strictly negatively to the sum. So, dropping all $x \notin \mathcal{A}$ can only make the sum larger. That is,

$$\begin{aligned}
\delta(\tilde{P}, \tilde{Q}) &= \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{S}} W(y|x)(p(x) - q(x)) && (25) \\
&\leq \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{x \in \mathcal{A}} \sum_{y \in \mathcal{S}} W(y|x)(p(x) - q(x)) && (26) \\
&= \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{x \in \mathcal{A}} \left(\underbrace{\sum_{y \in \mathcal{S}} W(y|x)}_{\leq 1} \right) (p(x) - q(x)) && (27) \\
&= \max_{\mathcal{S} \subseteq \mathcal{Y}} \sum_{x \in \mathcal{A}} (p(x) - q(x)) && (28) \\
&= \sum_{x \in \mathcal{A}} (p(x) - q(x)). && (29)
\end{aligned}$$

To complete the proof, we recall from class that this last sum, for the very particular choice of the set $\mathcal{A}$ that we have made, is precisely the TV distance between P and Q.

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 $t = 1, 2, \dots$:

(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 $m > 0$

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

(Hint: decompose according to whether $T_a(t) < m$ or $T_a(t) \geq m$, 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 $T_a(t) \geq m$,

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

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 $G_a := \left\{\left|\hat{\mu}_{a,t} - \mu_a\right| \leq \frac{\varepsilon}{4}\right\}$.

(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

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

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

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

Solution 2. (a) Model and conjugacy. For arm $a \in \{1,2\}$ put prior $\mu_a \sim \mathrm{Beta}(1,1)$ with density

$$p(\mu_a) = 1, \qquad \mu_a \in (0,1).$$

If by time t arm a has been pulled $N_a(t)$ times with $S_a(t)$ successes and $F_a(t)$ failures, the likelihood is

$$p(\text{data} \mid \mu_a) = \mu_a^{S_a(t)}(1-\mu_a)^{F_a(t)}.$$

Bayes' rule gives

$$p(\mu_a \mid \text{data}) = \frac{p(\text{data} \mid \mu_a)\, p(\mu_a)}{\int_0^1 p(\text{data} \mid \mu)\, p(\mu)\, d\mu}.$$

Substituting the prior and likelihood,

$$p(\mu_a \mid \text{data}) = \frac{\mu_a^{S_a(t)}(1-\mu_a)^{F_a(t)}}{\displaystyle\int_0^1 \mu^{S_a(t)}(1-\mu)^{F_a(t)}\, d\mu}.$$

The denominator is the normalizing constant (partition function):

$$Z = \int_0^1 \mu^{S_a(t)}(1-\mu)^{F_a(t)}\, d\mu = B(1+S_a(t),\, 1+F_a(t)),$$

where $B(\cdot,\cdot)$ is the Beta function.

Thus,

$$p(\mu_a \mid \text{data}) = \frac{1}{B(1+S_a(t),\, 1+F_a(t))}\mu_a^{S_a(t)}(1-\mu_a)^{F_a(t)}.$$

which is the density of $\mathrm{Beta}(1+S_a(t), 1+F_a(t))$.

(b) Time partition. Count separately the pulls of arm 2 made before it has been pulled m times, and those after. Before reaching m pulls, arm 2 can be pulled at most m times. Hence

$$N_2(n) \le m + \sum_{t: N_2(t-1) \ge m} \mathbf{1}\{A_t = 2\}.$$

Taking expectations gives the desired inequality.

(c) Concentration. After $N_2(t) = m$ pulls of arm 2, its posterior is

$$\mu_2 \mid \mathcal{F}_t \sim \mathrm{Beta}(1+S_2(t), 1+F_2(t)),$$

with $S_2(t) + F_2(t) = m$.

Our goal is to show that there exist universal constants $c_1 = 3e$ and $c_2 = \frac{1}{8}$ such that for every $\varepsilon \in (0,1)$ and every $m \ge 1$,

$$\mathbb{P}\big(\tilde{\mu}_2 > \mu_2 + \varepsilon\big) \;\le\; c_1 \exp\big(-c_2 m \varepsilon^2\big), \tag{30}$$

and the same holds for arm 1.

Proof. Fix $\varepsilon \in (0,1)$ and $m \geq 1$. Write

$$\hat{\mu}_m := \frac{S_2(t)}{m} \quad \text{and} \quad \mu_m := \frac{1 + S_2(t)}{m+2}.$$

Note that μ_m is exactly the posterior mean of $\mathrm{Beta}(1 + S_2(t), 1 + F_2(t))$.

We now split according to a good empirical event. Define the "good" event

$$G_m := \left\{ \left| \hat{\mu}_m - \mu_2 \right| \leq \frac{\varepsilon}{4} \right\}.$$

By Hoeffding's inequality for Bernoulli variables,

$$\mathbb{P}(G_m^c) = \mathbb{P}\left(\left| \hat{\mu}_m - \mu_2 \right| > \frac{\varepsilon}{4} \right) \leq 2 \exp\left(-2m \left(\frac{\varepsilon}{4} \right)^2 \right) = 2 \exp\left(-\frac{m\varepsilon^2}{8} \right). \tag{31}$$

We want to bound $\mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon)$. Split:

$$\begin{aligned}
\mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon) &= \mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon,\ G_m) + \mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon,\ G_m^c) \\
&\leq \mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon,\ G_m) + \mathbb{P}(G_m^c).
\end{aligned} \tag{32}$$

Next, we control the posterior mean on the good event. On G_m we have

$$|\hat{\mu}_m - \mu_2| \leq \frac{\varepsilon}{4}.$$

Since

$$\mu_m = \frac{1 + S_2(t)}{m+2} = \frac{m}{m+2} \hat{\mu}_m + \frac{1}{m+2},$$

we can bound, on G_m,

$$\begin{aligned}
|\mu_m - \mu_2| &= \left| \frac{m}{m+2} \hat{\mu}_m + \frac{1}{m+2} - \mu_2 \right| \\
&\leq \frac{m}{m+2} \left| \hat{\mu}_m - \mu_2 \right| + \frac{1}{m+2} \\
&\leq \frac{\varepsilon}{4} + \frac{1}{m+2}.
\end{aligned}$$

Now distinguish two sub-cases on G_m.

Case A: $m \geq \frac{4}{\varepsilon}$. Then $\frac{1}{m+2} \leq \frac{\varepsilon}{4}$, hence on G_m,

$$|\mu_m - \mu_2| \leq \frac{\varepsilon}{2}. \tag{33}$$

In particular,

$$\mu_2 + \varepsilon \ \geq \ \mu_m + \frac{\varepsilon}{2}.$$

Next, we use the fact that Beta is sub-Gaussian around its mean. Conditioned on $S_2(t)$ (equivalently on $\mathcal{F}_{t-1}$ in the bandit notation), we have

$$\tilde{\mu}_2 \mid S_2(t) \sim \mathrm{Beta}\big(1 + S_2(t),\, 1 + F_2(t)\big),$$

with total mass

$$(1 + S_2(t)) + (1 + F_2(t)) = m + 2.$$

A standard Chernoff/Hoeffding-type bound for the Beta distribution (it can be proved, e.g., by viewing Beta as a normalized Gamma/Binomial, or by log-concavity) states that there exists an absolute constant $C > 0$ such that for all $u \in (0, 1)$,

$$\mathbb{P}\big(\tilde{\mu}_2 - \mu_m \geq u \,\big|\, S_2(t)\big) \;\leq\; \exp\big(-C(m+2)u^2\big). \tag{34}$$

We now wrap up Case A. On G_m and for $m \geq 4/\varepsilon$, we can use (33) and (34) with $u = \varepsilon/2$ to get

$$\mathbb{P}\big(\tilde{\mu}_2 > \mu_2 + \varepsilon \,\big|\, S_2(t)\big) \leq \mathbb{P}\big(\tilde{\mu}_2 > \mu_m + \varepsilon/2 \,\big|\, S_2(t)\big) \leq \exp\big(-C(m+2)(\varepsilon/2)^2\big) \leq \exp\big(-C' m\varepsilon^2\big),$$

for some $C' = C/4$. Taking expectation over $S_2(t)$ but keeping G_m,

$$\mathbb{P}\big(\tilde{\mu}_2 > \mu_2 + \varepsilon,\, G_m\big) \leq \exp\big(-C' m\varepsilon^2\big). \tag{35}$$

Plugging (35) and (31) into (32), we get for $m \geq 4/\varepsilon$,

$$\mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon) \leq 2\exp\left(-\frac{m\varepsilon^2}{8}\right) + \exp\big(-C' m\varepsilon^2\big) \leq 3\exp\left(-\frac{m\varepsilon^2}{8}\right),$$

after adjusting constants (take $c_2 = 1/8$).

Case B: $1 \leq m < \frac{4}{\varepsilon}$. In this regime we simply use the trivial bound

$$\mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon) \leq 1.$$

At the same time, since $m < 4/\varepsilon$,

$$\exp\left(-\frac{m\varepsilon^2}{8}\right) \geq \exp\left(-\frac{(4/\varepsilon)\varepsilon^2}{8}\right) = \exp\left(-\frac{\varepsilon}{2}\right) \geq e^{-1/2}.$$

Hence

$$1 \;\leq\; 3e \cdot \exp\left(-\frac{m\varepsilon^2}{8}\right).$$

Therefore, for all $1 \leq m < 4/\varepsilon$ the desired bound

$$\mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon) \leq 3e \cdot \exp\left(-\frac{m\varepsilon^2}{8}\right)$$

also holds.

Finally, we unify the two cases. Combining Case A and Case B, we have shown that for every $m \geq 1$ and every $\varepsilon \in (0,1)$,

$$\mathbb{P}(\tilde{\mu}_2 > \mu_2 + \varepsilon) \leq 3e \cdot \exp\left(-\frac{m\varepsilon^2}{8}\right).$$

This proves (30) with $c_1 = 3e$ and $c_2 = \frac{1}{8}$. $\square$

(d) Bounding suboptimal selection. Choose $\varepsilon = \Delta/2$. Then, using

$$\{\tilde{\mu}_2^{(t)} > \tilde{\mu}_1^{(t)}\} \subseteq \{\tilde{\mu}_2^{(t)} > \mu_2 + \Delta/2\} \cup \{\tilde{\mu}_1^{(t)} < \mu_1 - \Delta/2\},$$

we get

$$\mathbb{P}(A_t = 2 \mid \mathcal{F}_{t-1}) \leq c_1 \exp\left(-c_2 m \frac{\Delta^2}{4}\right) + c_1 \exp\left(-c_2 m \frac{\Delta^2}{4}\right).$$

Since $m \asymp (\log n)/\Delta^2$, the first term is at most $n^{-c_1'}$ for some $c_1' > 0$. Under TS, the optimal arm 1 is also pulled $\Omega(\log n)$ times, so the second term is at most $n^{-c_2'}$. Hence

$$\mathbb{P}(A_t = 2 \mid \mathcal{F}_{t-1}) \leq n^{-c'}$$

for some $c' > 0$.

(e) Final bound. Summing over t,

$$\mathbb{E}\big[\# \text{ of pulls of arm 2 after } m\big] \leq \sum_{t=1}^{n} n^{-c'} = O(1) \quad (\text{for } c' > 1).$$

Therefore

$$\mathbb{E}[N_2(n)] \leq m + O(1) = O\left(\frac{\log n}{\Delta^2}\right),$$

and

$$\mathbb{E}[\text{Regret}(n)] = \Delta\, \mathbb{E}[N_2(n)] = O\left(\frac{\log n}{\Delta}\right).$$

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 $\alpha \in (0,1)$,

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

Solution 3. First, we explicitly derive the density function of the Beta(α, m) distribution. Since $\Gamma(\alpha+1) = \alpha\Gamma(\alpha)$, we have $\Gamma(\alpha+m) = \prod_{i=0}^{m-1}(\alpha+i)\Gamma(\alpha)$. Therefore,

$$B(\alpha, m) = \frac{\Gamma(m)}{\prod_{i=0}^{m-1}(\alpha+i)}.$$

Thus,

$$p_{\alpha,2} = \frac{1}{\Gamma(m)}\prod_{i=0}^{m-1}(\alpha+i)\cdot x^{\alpha-1}(1-x)^m. \tag{36}$$

(a) Using (36), we have

$$\begin{aligned}
E_{\alpha,m}\left[\frac{d}{d\alpha}\log p_{\alpha,m}(X)\right] =& E_{\alpha,m}\left[\frac{d}{d\alpha}\left((\alpha-1)\log X + m\log(1-X) + \sum_{i=0}^{m-1}\log(\alpha+i) - \Gamma(m)\right)\right] \\
=& E_{\alpha,m}\left[\log X + \sum_{i=0}^{m-1}\frac{1}{\alpha+i}\right] \\
=& E_{\alpha,m}[\log X] + \sum_{i=0}^{m-1}\frac{1}{\alpha+i}.
\end{aligned}$$

Notice that $\Gamma(m)$ is a constant w.r.t. α and therefore will disappear. By Lemma 6.3 in the lecture notes, the LHS above is equal to 0 (since it is the expected value of the score function $\ell(\alpha) = E_{\alpha,m}[d \log p_{\alpha,m}(X)/dx]$, and therefore the result is proved.

Alternate solution: We can also proceed directly without the help of the hint, but the approach is a little more tricky and tedious.

From the definition of the expectation,

$$
\begin{aligned}
E_{\alpha,2}[\log X] &= \int_{x\in[0,1]} \log x \cdot p_{\alpha,m}(x) dx \\
&= \int_{x\in[0,1]} \log x \cdot \frac{1}{\Gamma(m)} \prod_{i=0}^{m-1} (\alpha+i) \cdot x^{\alpha-1}(1-x)^m dx \\
&= \int_{x\in[0,1]} \frac{1}{\Gamma(m)} \cdot (1-x)^m \left(\prod_{i=0}^{m-1} (\alpha+i) \cdot (\log x \cdot x^{\alpha-1}) \right) dx \\
&= \int_{x\in[0,1]} \frac{1}{\Gamma(m)} \cdot (1-x)^m \left(\prod_{i=0}^{m-1} (\alpha+i) \cdot \frac{d}{d\alpha} x^{\alpha-1} \right) dx \\
&\overset{(a)}{=} \int_{x\in[0,1]} \frac{1}{\Gamma(m)} \cdot (1-x)^m \left(\frac{d}{d\alpha} \left(\prod_{i=0}^{m-1} (\alpha+i) \cdot x^{\alpha-1} \right) - x^{\alpha-1} \cdot \frac{d}{d\alpha} \prod_{i=0}^{m-1} (\alpha+i) \right) dx \\
&= \int_{x\in[0,1]} \frac{d}{d\alpha} \left(\frac{1}{\Gamma(m)} \cdot \prod_{i=0}^{m-1} (\alpha+i) \cdot x^{\alpha-1}(1-x)^m \right) dx \\
&\qquad - \int_{x\in[0,1]} \frac{1}{\Gamma(m)} \cdot x^{\alpha-1}(1-x)^m \cdot \frac{d}{d\alpha} \prod_{i=0}^{m-1} (\alpha+i) \cdot dx \\
&\overset{(b)}{=} \frac{d}{d\alpha} \int_{x\in[0,1]} \frac{1}{\Gamma(m)} \cdot \prod_{i=0}^{m-1} (\alpha+i) \cdot x^{\alpha-1}(1-x)^m dx \\
&\qquad - \frac{d}{d\alpha} \prod_{i=0}^{m-1} (\alpha+i) \cdot \int_{x\in[0,1]} \frac{1}{\Gamma(m)} \cdot x^{\alpha-1}(1-x)^m dx \\
&\overset{(c)}{=} \frac{d}{d\alpha}(1) - \frac{1}{\prod_{i=0}^{m-1}(\alpha+i)} \cdot \sum_{i=0}^{m-1} \prod_{j\in[0:m-1]\backslash\{i\}} (\alpha+j) \\
&= -\sum_{i=0}^{m-1} \frac{1}{\alpha+i}.
\end{aligned}
$$

Steps (a) and (c) follow from the product rule of derivatives, step (b) follows after interchange of the derivative and the integral, and we use the fact that $p_{\alpha,m}(x)$ integrates to 1 in step (c).

(b) From class, we know that the score is defined as

$$\ell(\theta) = \frac{d}{d\theta}\log p_\theta(x) = \frac{p'_\theta}{p_\theta} \tag{37}$$

Now, taking the second derivative, we can write by the usual chain rule for derivatives (and more precisely, the quotient rule)

$$\begin{aligned} \frac{d^2}{d\theta^2}\log p_\theta(x) &= \frac{p''_\theta p_\theta - (p'_\theta)^2}{(p_\theta)^2} && (38) \\ &= \frac{p''_\theta}{p_\theta} - \left(\frac{p'_\theta}{p_\theta}\right)^2 && (39) \end{aligned}$$

Taking expected values, and using the linearity of expectation,

$$\mathbb{E}\left[\frac{d^2}{d\theta^2}\log p_\theta(x)\right] = \mathbb{E}\left[\frac{p''_\theta}{p_\theta}\right] - \mathbb{E}\left[\left(\frac{p'_\theta}{p_\theta}\right)^2\right] \tag{40}$$

For the first summand, we may proceed like in class,

$$\begin{aligned} \mathbb{E}\left[\frac{p''_\theta}{p_\theta}\right] &= \int p_\theta(x)\frac{p''_\theta(x)}{p_\theta(x)}dx && (41) \\ &= \int \frac{d^2}{d\theta^2}p_\theta(x)dx && (42) \\ &= \frac{d^2}{d\theta^2}\int p_\theta(x)dx && (43) \\ &= \frac{d^2}{d\theta^2}1 && (44) \\ &= 0. && (45) \end{aligned}$$

For the second summand, we recall that the score can be written as $\ell(\theta) = \frac{p'_\theta}{p_\theta}$. Therefore, we know that

$$\mathbb{E}\left[\frac{d^2}{d\theta^2}\log p_\theta(x)\right] = -\mathbb{E}\left[\left(\frac{p'_\theta}{p_\theta}\right)^2\right] = -\mathbb{E}\left[\ell^2(\theta)\right] = -I_\theta, \tag{46}$$

(As a side note, we point out that the multi-dimensional version of the lemma can be found as Lemma 6.5 in the lecture notes, but we did not cover this part in class.)

(c) We already know that $\frac{d}{d\alpha}\log p_{\alpha,2}(x) = \log x + \sum_{i=0}^{m-1}\frac{1}{\alpha+i}$. Taking the second derivative, we find that it does not, in fact, depend on x, and therefore, its expectation is also simply

$$\begin{aligned} E_{\alpha,m}\left[-\frac{d^2}{d\alpha^2}\log p_{\alpha,2}(X)\right] &= \sum_{i=0}^{m-1}\frac{1}{(\alpha+i)^2} \le \sum_{i=0}^{\infty}\frac{1}{(\alpha+i)^2} \\ &\le \frac{1}{\alpha^2} + \sum_{i=1}^{\infty}\frac{1}{i^2} \le \frac{1}{\alpha^2} + \frac{\pi^2}{6}. \end{aligned}$$