STA360/601 Midterm Exam

Instructions

- Write your name, NetID, and signature below.
- If you need extra space for any problem, continue on the back of the page.

Community Standard

To uphold the Duke Community Standard:

- I will not lie, cheat, or steal in my academic endeavors;
- I will conduct myself honorably in all my endeavors; and
- I will act if the Standard is compromised.

I have adhered to the Duke Community Standard in completing this exam.

Name: ______________________________

NetID: ______________________________

Signature: ______________________________

Score

(For TA use only — leave this section blank.)

1. ____________
2. ____________
3. ____________
4. ____________
5. ____________
6. ____________

Overall: ____________________

List of common distributions

$$\text{Geometric}(x|\theta) = \theta(1-\theta)^x \, \mathbb{1}(x \in \{0,1,2,\ldots\}) \text{ for } 0 < \theta < 1$$

$$\text{Bernoulli}(x|\theta) = \theta^x (1-\theta)^{1-x} \, \mathbb{1}(x \in \{0,1\}) \text{ for } 0 < \theta < 1$$

$$\text{Binomial}(x|n,\theta) = \binom{n}{x} \theta^x (1-\theta)^{n-x} \, \mathbb{1}(x \in \{0,1,\ldots,n\}) \text{ for } 0 < \theta < 1$$

$$\text{Poisson}(x|\theta) = \frac{e^{-\theta}\theta^x}{x!} \, \mathbb{1}(x \in \{0,1,2,\ldots\}) \text{ for } \theta > 0$$

$$\text{Exp}(x|\theta) = \theta e^{-\theta x} \, \mathbb{1}(x > 0) \text{ for } \theta > 0$$

$$\text{Uniform}(x|a,b) = \frac{1}{b-a} \, \mathbb{1}(a < x < b) \text{ for } a < b$$

$$\text{Gamma}(x|a,b) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx} \, \mathbb{1}(x > 0) \text{ for } a, b > 0$$

$$\text{Pareto}(x|\alpha,c) = \frac{\alpha c^\alpha}{x^{\alpha+1}} \, \mathbb{1}(x > c) \text{ for } \alpha, c > 0$$

$$\text{Beta}(x|a,b) = \frac{1}{B(a,b)} x^{a-1} (1-x)^{b-1} \, \mathbb{1}(0 < x < 1) \text{ for } a, b > 0$$

$$\mathcal{N}(x|\mu,\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\big(-\tfrac{1}{2\sigma^2}(x-\mu)^2\big) \text{ for } \mu \in \mathbb{R},\ \sigma^2 > 0$$

$$\mathcal{N}(x|\mu,\lambda^{-1}) = \sqrt{\frac{\lambda}{2\pi}} \exp\big(-\tfrac{1}{2}\lambda(x-\mu)^2\big) \text{ for } \mu \in \mathbb{R},\ \lambda > 0$$

Exponential family form

$$p(x|\theta) = \exp\big(\varphi(\theta)^{\mathtt{T}} t(x) - \kappa(\theta)\big) h(x)$$

List of special functions

$$\text{Beta function: } B(a,b) = \int_0^1 t^{a-1}(1-t)^{b-1} dt \text{ for } a, b > 0$$

$$\text{Gamma function: } \Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt \text{ for } x > 0$$

$$\text{Log function for } a > 1\text{: } \log a = \int_1^a (1/t) dt$$

1. (15 points)

(a) (5 points) What is the formula for Bayes' theorem?

(b) (5 points) You receive x_i telephone calls on day i, for $i = 1, \ldots, n$. You wish to model this as $X_1, \ldots, X_n$ i.i.d. from some distribution. Which of the following distributions would make sense to use? (Circle one.)

i. Beta

ii. Poisson

iii. Bernoulli

iv. Exponential

(c) (5 points) Suppose $X, X_1, \ldots, X_N$ are i.i.d. and assume $\mathbb{E}|X| < \infty$ and $\mathbb{V}(X) < \infty$. What is the standard deviation of

$$\frac{1}{N}\sum_{i=1}^{N} X_i?$$

Hint: It is the same as the RMSE of the Monte Carlo approximation. (Circle one.)

i. $\mathbb{V}(X)/N$

ii. $\mathbb{V}(X)/\sqrt{N}$

iii. $\sigma(X)/N$

iv. $\sigma(X)/\sqrt{N}$

2. (17 points) (Marginal likelihood)
   Suppose $X_1, \ldots, X_n \overset{\text{iid}}{\sim} \text{Geometric}(\theta)$ given θ. Consider a $\text{Beta}(a, b)$ prior on θ. What is the marginal likelihood $p(x_{1:n})$?

   (Your answer must be an explicit expression in terms of $a, b, x_1, \ldots, x_n, n$, and any of the special functions on page 2. You must show your work to receive full credit.)

3. (17 points) (Exponential families, Normal distribution)
Show that the collection of $\mathcal{N}(\mu, \sigma^2)$ distributions, with $\mu \in \mathbb{R}$ and $\sigma^2 > 0$, is a two-parameter exponential family, and identify the sufficient statistics function $t(x) = (t_1(x), t_2(x))^\mathtt{T}$ for your parametrization.

4. (17 points) (Conjugate priors)
Suppose $X_1, \ldots, X_n \overset{\text{iid}}{\sim} \text{Uniform}(0, \theta)$ given θ, that is,

$$p(x_i|\theta) = \frac{1}{\theta} \mathbb{1}(0 < x_i < \theta).$$

You would like to find a conjugate prior for θ. Show that the family of Pareto(α, c) distributions, with $\alpha > 0$ and $c > 0$, is a conjugate prior family.

5. (17 points) (Sampling methods)
Suppose $c > 1$ and

$$p(x) \propto \frac{1}{x}\mathbb{1}(1 < x < c).$$

(Note that $p(x)$ is proportional to this, not equal to this.) Assume you can generate $U \sim \text{Uniform}(0, 1)$. Give an explicit formula, in terms of c and U, for generating a sample from $p(x)$. You must show your work to receive full credit.

6. (17 points) (Decision theory)
Consider a decision problem in which the state is $\theta \in \mathbb{R}$, the observation is x, you must choose an action $\hat{\theta} \in \mathbb{R}$, and the loss function is

$$\ell(\theta, \hat{\theta}) = a\theta^2 + b\theta\hat{\theta} + c\hat{\theta}^2$$

for some known $a, b, c \in \mathbb{R}$ with $c > 0$. Suppose you have computed the posterior distribution and it is $p(\theta|x) = \mathcal{N}(\theta|M, L^{-1})$ for some M and L. What is the Bayes procedure (minimizing posterior expected loss)?

(Your answer must be an explicit expression in terms of a, b, c, M, and L. You must show your work to receive full credit.)