

STA531 Midterm Exam 2

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

1. _____

2. _____

3. _____

4. _____

Overall: _____

List of common distributions

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

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

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

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

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

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

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

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

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

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

$$\mathcal{N}(x|\mu, C) = \frac{1}{(2\pi)^{d/2}|C|^{1/2}} \exp\left(-\frac{1}{2}(x - \mu)^T C^{-1}(x - \mu)\right) \text{ for } \mu \in \mathbb{R}^d, C \in \mathbb{R}^{d \times d} \text{ symmetric positive definite.}$$

Exponential family form

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


List of special functions

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

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

1. (25 points) Graphical models

Suppose $(X_1, \dots, X_8) \sim p(x_1, \dots, x_8)$, and p respects the following DAG.



(a) Write down the factorization of $p(x_1, \dots, x_8)$ implied by this graph.


(b) For each question below, circle either Yes, No, or Indeterminate. Indeterminate means that the answer cannot be determined from the information given. (You do not need to justify your answer.)

- | | | | |
|--|-----|----|---------------|
| i. Is $X_1 \perp X_3$? | Yes | No | Indeterminate |
| ii. Is $X_1 \perp X_3 \mid X_7$? | Yes | No | Indeterminate |
| iii. Is $X_1 \perp X_3 \mid X_6$? | Yes | No | Indeterminate |
| iv. Is $X_1 \perp X_3 \mid X_4, X_5$? | Yes | No | Indeterminate |


(c) Draw the moral graph (i.e., the moralization) of the DAG above.

2. (25 points) Markov chains and graphical models

Let's refer to the following directed graph as G_D :



Let's refer to the following undirected graph as G_U :



Suppose $p(x_1, \dots, x_n)$ is a probability distribution.

(a) Show that if p respects G_D , then p respects G_U .

(b) Show that if p respects G_U , then p respects G_D .

3. (25 points) Modeling the data collection process

Doctor Drastic has each of his patients fill out a questionnaire on their first visit. One of the questions is whether or not the patient has a certain sexually transmitted disease (STD), but some of the patients might be embarrassed and might not want to answer this question. Consider the following model:

$$\theta \sim \text{Beta}(a, b)$$

$$Y_1, \dots, Y_n | \theta \text{ i.i.d. } \sim \text{Bernoulli}(\theta)$$

(where $Y_j = 1$ if patient j has the STD, and $Y_j = 0$ otherwise)

$$I_j | y, \theta \sim \text{Bernoulli}(\phi_{y_j}) \text{ for each } j = 1, \dots, n$$

(where $I_j = 1$ if patient j answers the question, and $I_j = 0$ otherwise).

For simplicity, suppose the parameters ϕ_0 and ϕ_1 are known, and suppose that if the patient answers, they always answer truthfully.

The doctor has had n patients, and so far, all of them have answered the STD question (i.e., $I_j = 1$ for all $j = 1, \dots, n$). What is the posterior on θ ? Express your answer in terms of a well-known distribution.

4. (25 points) Viterbi algorithm

In class, in order to derive the Viterbi algorithm, we first devised an algorithm for computing $\max_{z_{1:n}} p(x_{1:n}, z_{1:n})$ in which we visited the hidden variables in the “forward” direction: z_1, z_2, \dots, z_n .

Devise an algorithm for computing $\max_{z_{1:n}} p(x_{1:n}, z_{1:n})$ in which the variables are visited in the “backward” direction: z_n, z_{n-1}, \dots, z_1 . (Your algorithm should take nm^2 time, like the forward version. You do not need to include steps for computing the argmax.)