Department of Statistical Science Duke University

FIRST YEAR EXAM - SPRING 2018

Monday May 7th 2017, 9:00 AM – 1:00 PM

NOTES: PLEASE READ CAREFULLY BEFORE BEGINNING EXAM!

1. Do not write solutions on the exam; please write your solutions on the paper provided.
2. Please use **black pen/ink** (no pencils) to complete your final solutions.
3. Put the problem number and your assigned code on the top of **each page**.
4. Write only on **one side** of the page (solutions on the reverse side of the page will be ignored).
5. Start each problem on a new page.
6. It is to your advantage to show your work and explain your answers.
   Do not erase anything– just draw a line through work you do not want graded.
7. You have 4 hours to finish the written exam.
8. **You may choose to complete five out of the six questions. In case you attempt all six questions, only five will be graded and please clearly indicate which five you choose to be graded.**
9. All five graded questions will carry *equal* weight.
10. This is a closed book exam. No notes are permitted.

1. Two continuous random quantities $x, y \in (0,1)$ have a joint distribution defined by the conditional $x \mid y \sim \mathrm{U}(0,y)$, for any $y \in (0,1)$, and the marginal $y \sim \mathrm{Beta}(2,1)$.

   (a) Show that conditional for $y \mid x$ is $y \mid x \sim \mathrm{U}(x,1)$, for any $x \in (0,1)$.

   (b) Show that the marginal for x is $x \sim \mathrm{Beta}(1,2)$.

   A Gibbs sampler is run using these two complete conditionals, initialised at $x_0 = 0.5$. Specifically, this generates pairs x_t, y_t $(t = 1, 2, \ldots)$, by iterating between samples of

   - y_t from $y_t \mid x_{t-1} \sim \mathrm{U}(x_{t-1}, 1)$, and then
   - x_t from $x_t \mid y_t \sim \mathrm{U}(0, y_t)$.

   Consider now only the first-order Markov process x_t so generated.

   (c) Show that $\mathrm{E}(x_{t+1} \mid x_t) = a + bx_t$ and identify the constants a and b.

   (d) How do you know that this x_t process is ergodic?

   (e) What is the univariate stationary distribution of this marginal x_t process?

   (f) How do you know that the x_t process is reversible?

2. Let $\{X_k\} \sim \text{Exp}(1)$ be independent random variables from the unit exponential distribution. Set

$$S_n = \sum_{k=1}^{n} \mathbb{1}_{\{X_k > k\}} \quad \text{and} \quad T_n = \sum_{k=1}^{n} \mathbb{1}_{\{X_k > 2\}}.$$

(a) State $\mathsf{E}[S_n]$ and $\mathsf{E}[T_n]$.

(b) Does $\{S_n\}$ converge to a random variable as $n \to \infty$? Explain your answer.

(c) Consider the event $T_n = 0$, does this event occur infinitely often as $n \to \infty$? Explain your answer.

3. Suppose $X_1, X_2, \ldots, X_n$ are i.i.d. Normal$(0, \sigma^2)$ and we are interested in estimating σ^2. Consider a loss function of the form $L(\sigma^2, d) = (\sigma^2 - d)^2/\sigma^{2k}$ for some finite constant $k > 1$.

   (a) Adopting a Gamma(α, β) prior on $1/\sigma^2$, what is the Bayes estimator for σ^2?

   (b) Find the uniformly minimum risk unbiased (UMRU) estimator for σ^2.

   (c) Show that when $k = 2$, the estimator $\delta^*(\mathbf{X}) = \frac{1}{n+2} \sum_i X_i^2$ is minimax.

   Hint: The inverse-Gamma(α, β) distribution has pdf

$$p(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{-\alpha-1} e^{-\beta/x}$$

   and its mean is $\beta/(\alpha - 1)$ for $\alpha > 1$. Also, for $Z \sim$ Normal$(0, 1)$, $\mathsf{E}(Z^4) = 3$.

4. Let $\{X_1, ..., X_n\} \overset{iid}{\sim} \mathrm{N}(\theta, 1)$ where $\theta \in \mathbb{R}$. Suppose we think θ might be zero, and so are considering using the estimator δ defined by

$$\delta(\bar{X}) = \begin{cases} \bar{X} & \text{if } |\bar{X}| > c_n \\ 0 & \text{if } |\bar{X}| < c_n \end{cases}$$

where $c_n = 1/n^{1/4}$. More compactly, $\delta(\bar{X}) = \bar{X}W_n$, where $W_n = 1(|\bar{X}| > c_n)$. Define $\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$.

(a) State the probability $\Pr(W_n = 1)$ as $n \to \infty$, for both $\theta \neq 0$ and $\theta = 0$.

(b) Obtain the limiting distribution of $\sqrt{n}(\delta - \theta)$ for both $\theta \neq 0$ and $\theta = 0$.

(c) Show that $\Pr(\delta = 0) \to 1$ if $\theta = 0$.

5. The inverse Gaussian distribution can be used to model neurons firing, specifically the inter-spike interval (ISI) is a measure of the amount of time between firing. We consider a model where we measure ISI as random variable $\{X_i\} \overset{iid}{\sim} \text{IG}(\lambda, \mu)$.

The density function of the two parameter inverse Gaussian is

$$f(x \mid \mu, \lambda) = \left(\frac{\lambda}{2\pi x^3}\right)^{1/2} \exp\left(-\frac{\lambda(x-\mu)^2}{2\mu^2 x}\right). \quad x > 0, \lambda > 0, \mu > 0.$$

(a) Is the inverse Gaussian an example of a two parameter exponential family distribution ? Show why or why not.

(b) What are the sufficient statistics of the inverse Gaussian distribution and what are the natural parameters ?

(c) Let $\mathbf{x} = (x_1, \ldots, x_{12})$ be a random sample of size 12 from $\text{IG}(\mu, \lambda)$.
Assume two independent Gamma distributions for the priors on μ and λ

$$\begin{aligned} \pi_1(\mu) &= \frac{b^a}{\Gamma(a)} \mu^{a-1} \exp(-b\mu), \quad a = 1, b = 2 > 0 \\ \pi_2(\lambda) &= \frac{c^d}{\Gamma(c)} \lambda^{c-1} \exp(-d\lambda), \quad c = 3, d = 4 > 0 \end{aligned}$$

State the conditional posterior distribution $\lambda \mid \mu, \mathbf{x}$.
I now tell you $\sum_{i=1}^{12} x_i = 9.17$, $\sum_{i=1}^{12} (x_i)^{-1} = 20.7$, and $\mu = 10$. State the distribution of

$$\lambda \mid \mu, \mathbf{x},$$

simplify as much as possible.

6. Bivariate measurements $X = (X_m, X_w)$ on student math and writing scores are assumedly independent normal $X \sim \mathrm{N}(\mu, \Sigma)$ with Σ known.

   (a) The prior for $\mu = (\mu_m, \mu_w)$ is also bivariate normal, $\mu \sim \mathrm{N}(\mu_{\text{prior}}, \Sigma_{\text{prior}})$, defined implicitly by

   $$\mu_{\mathrm{m}} \mid \mu_{\mathrm{w}} \sim \mathrm{N}(2\mu_{\mathrm{w}}, 1), \quad \mu_{\mathrm{w}} \sim \mathrm{N}(0, 1).$$

   In numbers, what are the prior mean vector and covariance matrix of μ, and what is the prior correlation $C(\mu_{\mathrm{m}}, \mu_{\mathrm{w}})$?

   (b) Observations $X_i = x_i$, $(i = 1, \ldots, n)$, on $n = 100$ students have sample mean $\bar{x} = (0, 0)$.

      i. State the posterior distribution of μ in terms of Σ, Σ_{prior}, μ_{prior} and the sample mean $\bar{x}$. Give a simplified expression for the posterior mean.

      ii. What is the mean of the posterior predictive distribution for scores of a future student?

      iii. What is the posterior predictive probability of the event $\{X_{\mathrm{m}} < X_{\mathrm{w}}\}$ for a future student?

      iv. In the observed sample of n students the observed frequency of the event $\{X_{\mathrm{m}} < X_{\mathrm{w}}\}$ is quite different from the posterior predictive value. What does this tell you about the model and prior?