FINAL EXAMINATION
Statistics 461

Friday 20th April, 2007
Time: 8:30am - 11:00am

Notes:

• This exam has 4 problems. Each exercise is worth 25 points.

• The amount each part of each question is worth is shown in []
on the left-hand side.

• Clear and complete solutions are essential; little partial credit will be given.

• This is a closed-book exam.

• Calculators are not allowed.

• No devices that can store text or send/receive messages are allowed.

Exercise 1. Let $X_1, X_2, \ldots, X_n$ be n independent observations from an exponential distribution with density

$$f(x|\theta) = \theta^{-1} \exp\left(-\theta^{-1} x\right) 1_{[0,\infty)}(x)$$

where $\theta \geq 0$. We introduced the following estimators of θ :

$$\begin{aligned} T_{1,n} &= \frac{\sum_{i=1}^{n} X_i}{n}, \\ T_{2,n} &= \frac{\sum_{i=1}^{n} X_i}{n+1}, \\ T_{3,n} &= \min(X_1, ..., X_n). \end{aligned}$$

(a) [**3 points**] Which of these estimators are functions of sufficient statistics?

(b) [**7 points**] Which of these estimators are unbiased?

(c) [**7 points**] Show that $T_{1,n}$ and $T_{2,n}$ are consistent estimates of θ; that is $T_{1,n} \to \theta$ and $T_{2,n} \to \theta$ in probability as $n \to \infty$. [Hint: Remember that $T_n \to \theta$ in probability means that $P_\theta(|T_n - \theta| > c) \to 0$ as $n \to \infty$ for any $c > 0$.]

(d) [**8 points**] Show that $T_{3,n}$ is not a consistent estimate of θ.

Exercise 2. Let $X_1, X_2, \ldots, X_n$ be n independent observations from a zero mean normal of variance σ^2 denoted $\mathcal{N}\left(0, \sigma^2\right)$. One can find an M-estimator of σ by minimizing the following objective function

$$\frac{1}{n}\sum_{i=1}^{n} \rho\left(X_i/\sigma\right) + \log\sigma \tag{1}$$

for some function ρ. Suppose ρ is differentiable except perhaps at a finite number of points.

(a) [**3 points**] By differentiating (1) with respect to σ, establish the equation satisfied by the M-estimate.

(b) [**4 points**] Show that if

$$\rho\left(u\right) = |u|$$

then the M-estimator is given by

$$\widehat{\sigma} = \frac{1}{n}\sum_{i=1}^{n} |X_i| \,.$$

(b) [**6 points**] Is this M-estimator consistent?

(c) [**7 points**] Establish the asymptotic distribution of this estimator.

(d) [**5 points**] How does it compare with the MLE of σ in terms of consistency, efficiency and robustness?

Exercise 3. Let $X_1, X_2, \ldots, X_n$ be n independent observations from a zero mean normal of variance τ denoted $\mathcal{N}(0, \tau)$.

(a) [**3 points**] Find the MLE for τ.

(b) [**3 points**] Find the Fisher's information of τ.

(c) [**4 points**] Detail the likelihood ratio test for $H_0 : \tau = 1$ versus $H_1 : \tau \neq 1$.

(d) [**5 points**] Find the Wald test of (asymptotic) size α. Express the result in terms of a z_β where $\Pr(Z > z_\beta) = \beta$ for $Z \sim \mathcal{N}(0, 1)$.

(e) [**7 points**] Compute the power function of the Wald test. Express the result in terms of the cdf of the standard normal denoted $\Phi(\cdot)$.

(f) [**3 points**] In a Bayesian framework, how would you test $H_0 : \tau = 1$ versus $H_1 : \tau \neq 1$.

Exercise 4. Let $X_1, X_2, \ldots, X_n$ be n independent observations from a mixture of two Poisson distributions with probability mass function

$$f(X = k|\theta) = \pi \exp(-\lambda_1) \frac{\lambda_1^k}{k!} + (1-\pi) \exp(-\lambda_2) \frac{\lambda_2^k}{k!}$$

where $\theta = (\pi, \lambda_1, \lambda_2)$ with $\pi \in (0,1)$, $\lambda_1 > 0$ and $\lambda_2 > 0$.

(a) [**3 points**] Write down the log-likelihood of the observations $X_1, X_2, \ldots, X_n$. Is it possible to obtain an expression for the MLE estimate in closed-form? Suggest a numerical method which could be used to find the MLE without having to introduce any missing data.

(b) [**7 points**] Describe precisely how you would implement the Expectation-Maximization to find the MLE estimate.

(c) [**7 points**] Assume now that the parameter θ follows the following prior distribution

$$p(\theta) = p(\pi) p(\lambda_1) p(\lambda_2)$$

with

$$\begin{aligned} p(\pi) &= \mathcal{B}eta(\pi; \alpha, \beta), \\ p(\lambda_1) &= \mathcal{G}amma(\lambda_1; a, b), \\ p(\lambda_2) &= \mathcal{G}amma(\lambda_2; a, b), \end{aligned}$$

where α, β, a, b are fixed hyper-parameters.

Describe precisely how you could implement the Gibbs sampler to sample from the posterior distribution of θ given $X_{1:n} = (X_1, X_2, \ldots, X_n)$.

(d) [**5 points**] Show that, whatever being the observations, we have $\mathbb{E}[\lambda_1 | X_{1:n}] = \mathbb{E}[\lambda_2 | X_{1:n}]$.

(e) [**3 points**] Suggest an alternative to the Gibbs sampler to sample from the posterior distribution of θ given $X_{1:n}$. Discuss briefly the advantages and drawbacks of this alternative scheme.

Hints.

- The Beta distribution is given for $a, b > 0$ by

$$\mathcal{B}eta(x; a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1} (1-x)^{b-1} 1_{[0,1]}(x).$$

We have $\mathbb{E}[X] = \frac{a}{a+b}$.

- The Gamma distribution is given for $a, b > 0$ by

$$\mathcal{G}amma(x; a, b) = \frac{b^a}{\Gamma(a)} x^{a-1} \exp(-bx) 1_{(0,\infty)}(x).$$

We have $\mathbb{E}[X] = \frac{a}{b}$.