2014 Jan Qualification Exam and Solution (Theory Part)

1. The joint density function of X and Y is given by

$$f(x,y) = \frac{e^{-yx^2/2}}{\sqrt{2\pi/y}} \cdot ye^{-y}, \quad -\infty < x < \infty, y > 0.$$

   (a) Find the conditional density $f_{X|Y}(x|y)$ of X given $Y = y$.

   (b) Compute $\mathrm{E}(X|Y)$.

   (c) Compute $\mathrm{Var}(X|Y)$.

   (d) Compute $\mathrm{Var}(X)$.

2. Suppose that X and Y are independent $exp(\lambda)$ random variables with density

$$f(x) = \frac{1}{\lambda} e^{-x/\lambda}, \quad x > 0,$$

   (a) Show that the sum $X + Y$ and the ratio X/Y are independent.

   (b) Let $Z = \frac{X}{X+Y}$, show that for $0 < z < 1$, $F_Z(z) = P(Z \leq z) = z$, i.e. the randome variable Z is uniformly distributed over (0,1).

3. Let $X_1, X_2, \cdots, X_n$ be independent Gaussian random variables, having mean μ and variance σ^2. Define by recurrence the sequence

$$Y_0 = x \in \mathbb{R}, \quad Y_{n+1} = \lambda Y_n + X_{n+1},$$

   for some $\lambda \in (-1, 1)$. Prove that the sequence $\{Y_n\}_{n \in \mathbb{N}}$ converges in distribution and determine the limiting distribution.

4. Let $X_1, \cdots, X_n$ be a random sample from a distribution with the pdf given by

$$f(x|\theta) = \theta^{-c} c x^{c-1} e^{-(x/\theta)^c} I(x > 0),$$

   where $c > 0$ is a known constant.

   (a) Find the maximum likelihood estimator (MLE) of θ.

   (b) Find the uniformly minimum variance unbiased estimator (UMVUE) of θ.

(c) Find the uniformly most powerful (UMP) test of size α for testing

$$H_0 : \theta \leq \theta_0, \quad \text{vs} \quad H_1 : \theta > \theta_0,$$

where θ_0 is a positive constant.

5. Let $X_1, \cdots, X_n$ be a random sample of i.i.d. observations drawn from the following probability density function (pdf)

$$f(x|\theta) = \theta^{-1} x^{(1-\theta)/\theta} I(0 \leq x \leq 1), \quad \theta > 0.$$

(a) Show that $T(\mathbf{X}) = -2\sum_{i=1}^{n} \log(X_i)$ is a minimal sufficient statistic for θ.

(b) Find the distribution of $Y = -2 \log X_1$.

(c) Find a two-sided 95% confidence interval for θ based on T.

(d) Argue or prove that the expected length of the confidence interval obtained in part (c) converges to zero as $n \to \infty$.

6. For regression data $\{(x_i, Y_i\}_{i=1}^{n}$ assume the model

$$Y_i \sim \text{Poisson}(x_i \beta), \quad i = 1, \cdots, n, \quad \text{independent},$$

where $x_1, \cdots, x_n$ are strictly positive and known constants. Let $\mathbf{Y} = (Y_1, \cdots, Y_n)$, $\bar{Y} = \frac{1}{n}\sum_{i=1}^{n} Y_i$, and $\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$.

(a) Show that the MLE of β is $\hat{\beta} = \bar{Y}/\bar{x}$.

(b) Compute the mean and variance of $\hat{\beta}$.

(c) Now assume that β has a gamma prior distribution $\beta \sim \Gamma(wb_0, 1/w)$, where b_0 is our prior best guess and $w > 0$ is a weight attached to this guess. To be specific, β has the prior density

$$\pi(\beta|w, b_0) = \frac{w^{wb_0}}{\Gamma(wb_0)} \beta^{wb_0 - 1} \exp(-w\beta).$$

Find the posterior density of β given $\mathbf{Y}$.

(d) Show that the posterior mean of β is the weighted average of the prior mean and the MLE. What does the posterior mean converge to when the weight $w \to 0$?

Solutions:

1. (a) By Inspection,

$$f_Y(y) = ye^{-y}, y \geq 0, \quad , f_{X|Y}(x|y) = \frac{e^{-yx^2/2}}{\sqrt{2\pi/y}}.$$

   Therefore $X|Y \sim N(0, 1/y)$.

   (b) $\mathrm{E}[X|Y] = 0$.

   (c) $\mathrm{Var}[X|Y] = 1/Y$.

   (d)

$$\begin{aligned}
\mathrm{Var}(X) &= \mathrm{E}[\mathrm{Var}(X|Y)] + \mathrm{Var}[\mathrm{E}(X|Y)] \\
&= \mathrm{E}(1/Y) + 0 = \int_0^\infty \frac{1}{y} ye^{-y} dy \\
&= -e^{-y}|_0^\infty = 1.
\end{aligned}$$

2. (a) Let $U = X + Y$, $V = X/Y$. Solve for X, Y, $x = \frac{uv}{v+1}, y = \frac{u}{v+1}$. Jacobian Matrix

$$J = \begin{pmatrix} \frac{v}{v+1} & \frac{1}{v+1} \\ \frac{u}{(v+1)^2} & -\frac{u}{(v+1)^2} \end{pmatrix},$$

   hence $|J| = \frac{-u(v+1)}{(v+1)^3} = \frac{-u}{(v+1)^2}$.

$$\begin{aligned}
f_{U,V}(u, v) &= \frac{1}{\lambda^2} e^{-u/\lambda} \frac{u}{(v+1)^2} \\
&= \frac{1}{\lambda^2} e^{-u/\lambda} u \cdot \frac{1}{(v+1)^2}.
\end{aligned}$$

   Therefore, the joint density function of U and V can be written as the product of a function of U and a function of V. Thus U and V are independent.

(b) For $0 < z < 1$,

$$\begin{aligned} P(Z \le z) &= P\left(\frac{X}{X+Y} \le z\right) = P(X \le zX + zY) \\ &= P(X(1-z)/z \le Y) \\ &= \int\int_{x(1-z)/z \le y} \frac{1}{\lambda^2} e^{-(x+y)/\lambda} dx dy \\ &= \int_0^\infty \left(\int_{x(1-z)/z}^\infty \frac{1}{\lambda} e^{-y/\lambda} dy\right) \frac{1}{\lambda} e^{-x/\lambda} dx \\ &= \int_0^\infty e^{-x(1-z)/\lambda z} \frac{1}{\lambda} e^{-x/\lambda} dx \\ &= \int_0^\infty \frac{1}{\lambda} e^{-x/\lambda z} dx = -z e^{-x/\lambda z}|_0^\infty = z. \end{aligned}$$

Therefore Z follows a uniform over $(0, 1)$ distribution.

3. From MGF we know that Y_{n+1} must be normally distributed. To show the convergence in distribution, firstly we note

$$Y_0 = x, Y_n = \lambda^n x + \sum_{i=1}^n \lambda^{n-i} X_i,$$

with mean $\mu \sum_{i=1}^n \lambda^{n-i} + \lambda^n x = \lambda^n x + \mu \frac{1-\lambda^n}{1-\lambda} \to \frac{\mu}{1-\lambda}$,
and variance $\sigma^2 \sum_{i=1}^n \lambda^{2(n-i)} = \sigma^2 \frac{1-\lambda^{2n}}{1-\lambda^2} \to \frac{\sigma^2}{1-\lambda^2}$. So we know that the parameters in the MGFs converge therefore the MGFs of Y_n converges to the MGF of a normal distribution with mean $\frac{\mu}{1-\lambda}$ and variance $\frac{\sigma^2}{1-\lambda^2}$, i.e. $Y_n \xrightarrow{D} N(\frac{\mu}{1-\lambda}, \frac{\sigma^2}{1-\lambda^2})$.

4. (a) The joint likelihood is

$$L(\theta) = -nc \log \theta + n \log(c) + (c-1) \sum_{i=1}^n \log x_i - \frac{\sum_{i=1}^n x_i^c}{\theta^c}.$$

Taking the derivative with respect to θ, we set

$$\frac{dL}{d\theta} = c \frac{\sum_{i=1}^n x_i^c}{\theta^{c+1}} - \frac{nc}{\theta} == 0.$$

This leads to $\hat{\theta} = [\frac{\sum_{i=1}^n x_i^c}{n}]^{1/c}$. Check the second derivative $\frac{d^2 L}{d\theta^2}|_{\hat{\theta}} < 0$. Thus $\hat{\theta}_{MLE} = \left(\frac{T}{n}\right)^{1/c}$, where $T = \sum_{i=1}^n x_i^c$.

(b) This is one-parameter exponential family, with $T = \sum_{i=1}^n X_i^c$ being the complete and sufficient statistic. Define $Y_i = X_i^c$ for $i = 1, \cdots, n$. Using the transformation, we compute the pdf of Y as $f(y|\theta) = \frac{1}{\theta^c} \exp -y/\theta^c$. So $Y = X^c \sim exp(\theta^c)$, so $T \sim Gamma(n, \theta^c)$. Then

$$\begin{aligned} E(\hat{\theta}_{MLE}) &= \int_0^\infty \left(\frac{t}{n}\right)^{\frac{1}{c}} \frac{1}{\Gamma(n)\theta^{cn}} t^{n-1} e^{-\frac{t}{\theta^c}} dt \\ &= \frac{\Gamma(n+\frac{1}{c})}{\Gamma(n)n^{1/c}}\theta. \end{aligned}$$

The UMVUE of θ is $\frac{\Gamma(n)n^{1/c}}{\Gamma(n+\frac{1}{c})}\hat{\theta}_{MLE}$.

(c) For $\theta_2 > \theta_1 > 0$, the density ratio

$$\frac{f(\mathbf{x}|\theta_2)}{f(\mathbf{x}|\theta_1)} = \left(\frac{\theta_1}{\theta_2}\right)^{nc} \exp\left\{\left(\frac{1}{\theta_1^c} - \frac{1}{\theta_2^c}\right)\sum_{i=1}^n x_i^c\right\}$$

is an increasing function of T. So the distribution family of X has an MLR in T. By Karlin-Rubin Theorem, the UMP test of size α is: reject H_0 if $T(\mathbf{X}) > t_0$where t_0 satisfies $P_{\theta_0}(T > t_0) = \alpha$. Note that when $\theta = \theta_0$, $\frac{2T}{\theta_0^c} \sim \chi^2_{2n}$. Then

$$P(T > t_0) = P(\frac{2T}{\theta_0^c} > \frac{2t_0}{\theta_0^c}) = P(\chi^2_{2n} > \frac{2t_0}{\theta_0^c}) = \alpha.$$

So $t_0 = \chi^2_{2n,\alpha}\theta_0^c/2$.

5. (a) This is one-parameter exponential family. It follows that $T(\mathbf{X}) = -2\sum_{i=1}^n \log(X_i)$ is a sufficient statistic for θ. Furthermore,

$$\frac{f(\mathbf{x}|\theta)}{f(\mathbf{y}|\theta)} = \frac{I(0 \le x_{(1)} \le x_{(n)} \le 1)}{I(0 \le y_{(1)}) \le y_{(n)} \le 1} \exp\{\frac{1-\theta}{\theta}\left(\sum_{i=1}^n \log x_i - \sum_{i=1}^n \log y_i\right)\}.$$

The ratio is free of θ if and only if $\sum_{i=1}^n \log X_i = \sum_{i=1}^n \log Y_i$. So T is minimal and sufficient.

(b) Using the variable transformation, the pdf of $Y = -2\log X_1$ is

$$f(y|\theta) = \frac{1}{2\theta}\exp(-\frac{y}{2\theta}).$$

So $Y \sim \exp(2\theta)$.

(c) Based on part (b), we have $T = \sum_{i=1}^n Y_i \sim \Gamma(n, 2\theta)$. Consider the pivotal quantity $\frac{T}{\theta} \sim \chi^2_{2n}$. Then

$$P\left(\chi^2_{2n,0.975} \leq \frac{T}{\theta} \leq \chi^2_{2n,0.025}\right) = 0.95,$$

which implies

$$P\left(\frac{T}{\chi^2_{2n,0.025}} \leq \theta \leq \frac{T}{\chi^2_{2n,0.975}}\right) = 0.95,$$

so a two-sided 95% confidence interval for θ is $\left(\frac{T}{\chi^2_{2n,0.025}}, \frac{T}{\chi^2_{2n,0.975}}\right)$.

(d) The expected length of the confidence interval obtained in part (c) is

$$2n\theta\left(\frac{1}{\chi^2_{2n,0.025}} - \frac{1}{\chi^2_{2n,0.975}}\right).$$

Using the normal approximation, we have

$$\frac{\chi^2_{2n,0.975}}{n} \approx \frac{2n + 1.96\sqrt{4n}}{n} \to 2,$$

and

$$\frac{\chi^2_{2n,0.025}}{n} \approx \frac{2n - 1.96\sqrt{4n}}{n} \to 2,$$

So the length converges to zero as $n \to \infty$.

6. (a) The log-likelihood is

$$L(\beta) = -\beta\sum_{i=1}^n x_i + \sum_{i=1}^n y_i \log(\beta x_i) - \sum_{i=1}^n \log(y_i!).$$

Then setting $\frac{dL}{d\beta} = n\bar{x} + n\bar{y}/\beta$ to zero, we have $\hat{\beta} = \bar{Y}/\bar{x}$. Note that $\frac{d^L}{d\beta^2}|_{\hat{\beta}_{MLE}} < 0$, so $\hat{\beta}_{MLE} = \bar{Y}/\bar{x}$.

(b) Compute

$$E(\hat{\beta}_{MLE}) = \frac{E\bar{Y}}{\bar{x}} = \frac{\beta\bar{x}}{\bar{x}} = \beta,$$

and, by independence,

$$\begin{aligned}\text{Var}(\hat{\beta}_{MLE}) &= \frac{Var(\bar{Y})}{\bar{x}^2} = \frac{\sum_{i=1}^n \text{Var}(Y_i)}{n^2\bar{x}^2} \\ &= \frac{\beta\sum_{i=1}^n x_i}{n^2\bar{x}^2} = \frac{\beta}{n\bar{x}}.\end{aligned}$$

(c) Now assume that β has gamma prior distribution $\beta \sim \Gamma(wb_0, 1/w)$, where b_0 is our prior best guess and $w > 0$ is a weight attached to this guess. To be specific, β has the prior density

$$\pi(\beta|w, b_0) = \frac{w^{wb_0}}{\Gamma(wb_0)} \beta^{wb_0 - 1} \exp(-w\beta).$$

Then the posterior density of β given $\mathbf{Y}$ is

$$\begin{aligned}
\pi(\beta|\mathbf{Y}) &\propto f(\mathbf{y}|\beta)\pi(\beta|w, b_0) \\
&= e^{-\beta \sum_{i=1}^{n} x_i} \prod_{i=1}^{n} (\beta x_i)^{y_i} \frac{w^{wb_0}}{\Gamma(wb_0)} \beta^{wb_0 - 1} \exp(-w\beta) \\
&= e^{-\beta(\sum_{i=1}^{n} x_i + w)} \beta^{wb_0 + \sum_{i=1}^{n} y_i - 1} \\
&\sim \Gamma(wb_0 + n\bar{y}, 1/(w + n\bar{x})).
\end{aligned}$$

(d) The posterior mean of β is

$$\begin{aligned}
E(\beta|\mathbf{Y}) &= \frac{wb_0 + n\bar{Y}}{w + n\bar{x}} \\
&= \frac{wb_0}{w + n\bar{x}} + \frac{n\bar{Y}}{w + n\bar{x}} \\
&= b_0 \frac{w}{w + n\bar{x}} + \frac{\bar{Y}}{\bar{x}} \frac{n\bar{x}}{w + n\bar{x}}.
\end{aligned}$$

Thus the weighted average of the prior mean b_0 and the MLE $\bar{Y}/\bar{x}$. The posterior mean converges to the MLE $\bar{Y}/\bar{x}$ when the weight $w \to 0$.