Serial No.: 1718A

| | | |
|---|---|---|
| Unique Paper Code | : | 2372014801 |
| Name of the Paper | : | Bayesian Inference |
| Name of the Course | : | B.Sc. (Hons.) Statistics under NEP-UGCF |
| Semester | : | VIII |
| Duration | : | 3 hours |
| Maximum Marks | : | 90 Marks |

**Instructions for Candidates**

- Write your Roll No. on the top immediately on receipt of this question paper.
- Attempt six questions in all.
- Use of a non-programmable scientific calculator is allowed.

1. (a) State and prove Bayes' theorem for random variables. Hence, explain the notions of likelihood function, prior distribution, posterior distribution and marginal distribution.

(b) Let $\underset{\sim}{X} = (X_1, X_2, \ldots, X_n)$ be a random sample from a Poisson distribution with parameter $\theta$, where $\theta$ is unknown. Assume that the prior distribution of $\theta$ is $Gamma(\alpha, \beta)$. Obtain the posterior distribution of $\theta$ given $\underset{\sim}{x}$, and identify the resulting posterior distribution.

(7, 8)

2. (a) Let X be a random variable taking values 0 and 1 with probability mass function

$$f(x|\theta) = \begin{cases} 1-\theta, & x = 0 \\ \theta, & x = 1. \end{cases}$$

Assume that the prior distribution for $\theta$ is given by

$$g(\theta) = \begin{cases} \frac{1}{0.2}, & 0.4 \le \theta \le 0.6 \\ 0, & otherwise. \end{cases}$$

Obtain the marginal probability mass function of x and the posterior distribution of $\theta$ given an observed value of x.

(b) Suppose $X_1, X_2, \ldots, X_n$ is a random sample from a normal distribution with known mean $\theta$ and unknown precision $r$. Construct a conjugate prior for $r$ and identify the prior distribution. Also, obtain the posterior distribution of $r$ given the sample.

(7, 8)

3. (a) Let a random variable $X \sim Bernoulli(\theta)$, where $\theta$ is unknown. Obtain Jeffreys' non-informative prior for $\theta$ and identify the resulting prior distribution. Also, obtain the posterior distribution of $\theta$ given $x$.

(b) Define generalized maximum likelihood estimate. Suppose

$$f(x|\theta) = \begin{cases} e^{-(x-\theta)}, & x > \theta \\ 0, & otherwise. \end{cases}$$

Assume that the prior distribution of $\theta$ is $g(\theta) = \frac{1}{\pi(1+\theta^2)}$. Obtain the generalised maximum likelihood estimate of $\theta$.

(8, 7)

4. (a) Explain how the Bayes estimate of the unknown parameter $\theta$ is obtained under a specified loss function $L(\theta, a)$ using the extensive form of Bayesian analysis. Hence, derive the Bayes estimator of $\theta$ under weighted squared error loss function $L(\theta, a) = \omega(\theta)(\theta - a)^2, \omega(\theta) > 0$.

(b) Consider the simple linear regression model $Y_i = \beta X_i + U_i, i = 1,2, \dots, n$, where $Y_i$ is the dependent variable, $X_i$ is a fixed (non-stochastic) independent variable, and $U_i$ represents the random error term. The regression parameter $\beta$ is unknown. Assume that the errors $U_i$ are independently and identically distributed as $N(0, \sigma^2)$, with $\sigma^2$ known. Further, suppose the prior distribution of $\beta$ is $g(\beta) = 1, \beta \in (-\infty, \infty)$. Derive the posterior distribution of $\beta$. Obtain the Bayes estimator of $\beta$ under squared error loss function, and compare it with the ordinary least squares estimator of $\beta$.

(6, 9)

5. (a) Define the $100(1-\alpha)\%$ highest posterior density (HPD) credible interval for a parameter $\theta$. Let $X_1, X_2, \dots, X_n$ be a random sample from $N(\theta, 1)$, where the mean $\theta$ is unknown. Assume that the prior distribution of $\theta$ is $N(0,1)$. Derive the 95% HPD credible interval of $\theta$. Further, compare this interval with the corresponding 95% classical confidence interval and comment on the differences.

(b) Let $X \sim Bernoulli(\theta)$, where $\theta$ is unknown, and suppose the prior distribution for $\theta$ is $Beta(\alpha, \beta)$. Given the loss function

$$L(\theta, \hat{\theta}) = \theta^{a-1}(1-\theta)^{b-a-1}(\theta - \hat{\theta})^2, a > 1, b > a + 1$$

find the Bayes estimate of $\theta$ using the duality between the prior distribution and the loss function.

(7, 8)

6. (a) Let $x = (x_1, x_2)$ be two independent observations from a normal distribution $N(\theta, 1)$. Suppose we wish to obtain the Bayes estimate of $\theta$ as a linear estimator of the form $\delta(x) = c_1 x_1 + c_2 x_2$. Assume that the loss function is $L(\theta, \delta(x)) = (c_1 x_1 + c_2 x_2 - \theta)^2$, and the prior distribution is uniform on (0,1). Obtain the risk function and the Bayes risk for the given estimator. Also, obtain the Bayes estimate of $\theta$ using the normal form of Bayesian analysis.

(b) Suppose $X_1, X_2, \dots, X_n$ is a random sample from $N(\theta, 1)$, where $\theta$ is unknown. Consider the hypotheses: $H_0: \theta = 0$ Vs $H_1: \theta = 1$. Assume that, a-priori, both hypotheses $H_0$ and $H_1$ are equally likely. Find the posterior odds in favour of $H_0$ after observing the sample, and also compute the Bayes factor in favour of $H_0$.

(8, 7)

7. (a) Consider testing two composite hypotheses $H_0: \theta \in \Theta_0$ against $H_1: \theta \in \Theta_1$. Show that the Bayes factor is the ratio of weighted (or marginal) likelihoods over the parameter spaces $\Theta_0$ and $\Theta_1$, where the likelihoods are averaged with respect to their respective prior distributions.

(b) Consider testing the hypothesis $H_0: \theta = 2$ against $H_1: \theta \neq 2$, where the observed data follow a Poisson distribution with parameter $\theta$. Assume that the prior distribution of $\theta$ under $H_1$ is Gamma with parameters $\alpha$ and $\beta$, i.e., $g_1(\theta) \sim Gamma(\alpha, \beta)$. Given that an observation (x = 4) is recorded, derive the expression for the Bayes factor in favour of $H_0$. Evaluate the Bayes factor when $\alpha = \beta = 1$. Hence, find the posterior probability of $H_0$ assuming equal prior probabilities for $H_0$ and $H_1$. Comment on the behaviour of the Bayes factor when the prior becomes non-informative (i.e., as $\beta \to 0$).

(6, 9)