SEMESTER 2 EXAMINATION 2012/13

MATH6025 Bayesian Methods

Duration: 120 min

Answer ALL questions.

Formula Sheet FS/MATH6025/13 will be available

Only University approved calculators may be used.

A foreign language dictionary (paper version) is permitted provided it contains no notes, additions or annotations.

1. (a) Define the difference between a subjective and an objective Bayesian and give an example of what kind of prior is used in each case. Explain the concepts of a conjugate prior and a non-informative prior.

[7 marks]

Suppose that $\underline{x} = (x_1, \ldots, x_n)$ are independent and identically distributed observations from an Exponential distribution with probability density function

$$f(x|\theta) = \theta e^{-\theta x}, \quad x > 0\,, \theta > 0\,.$$

(b) Show that a Gamma(β, m) prior, the form of which can be found in the Poisson section of the summary sheet, is conjugate for θ.

[5 marks]

(c) Find the Jeffreys prior for θ and obtain the posterior density for $\theta|\underline{x}$ under this prior, up to a normalising constant.

[6 marks]

(d) Find an appropriate normal approximation to the posterior distribution of $\theta|\underline{x}$.

[7 marks]

2. (a) Show that the mean of the posterior distribution is the Bayes estimator $\tilde{\theta}$ corresponding to the loss function

$$L(\theta, \tilde{\theta}) = (\theta - \tilde{\theta})^2$$

[4 marks]

Suppose that we have a single measurement $X = x_1$ from a Normal distribution $N(\mu, \sigma^2)$, so

$$f(x|\mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x-\mu)^2\right)$$

Also assume that we have a Normal prior for $\mu \sim N(\theta, \tau^2)$ and that σ^2 is known.

(b) Working to a constant of proportionality and completing the square, show that the posterior distribution is $f(\mu|x) \sim N(\mu^*, \tau^{*2})$ where

$$\mu^* = \frac{\tau^2 x + \sigma^2\theta}{\sigma^2 + \tau^2}, \qquad \tau^{*2} = \frac{\sigma^2\tau^2}{\sigma^2 + \tau^2}$$

[7 marks]

(c) Consider a loss function

$$L(\theta, \tilde{\theta}) = \frac{(\theta - \tilde{\theta})^2}{\theta^2}$$

Derive the Bayes estimator that minimises the posterior expectation of this loss function in terms of expectations of the form $E(g(\theta)|\underline{x})$. [4 marks]

(d) Suppose we now have multiple measurements $X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n$ from the Normal distribution.

(i) Using the fact that the exponential of a quadratic form in x_{n+1} and θ is proportional to a bivariate normal distribution and the marginal distributions of a bivariate normal are normal, calculate the density of X_{n+1} with parameters $E(x_{n+1}|\underline{x})$ and Var$(x_{n+1}|\underline{x})$, given the previous n observations.

(ii) If $\mu = 50$, $\tau = 2$, $\sigma = 5$ and we have $X = (49, 48, 45, 40, 50)$ Calculate the mean and variance of X_{n+1} to 1dp.

(iii) Using the values in (ii) find the distribution of the posterior predictive sample mean of m further draws, what happens as $m \to \infty$? [10 marks]

3. Suppose that we have an industrial process, in which the numbers of failures per batch (X) follows a Poisson distribution with mean $\theta = 1$. Each batch is independent and the company wishes to assess whether this mean has changed.

A change in mean level would mean that the distribution had changed to Poission(θ), $\quad \theta \neq 1$.

If the mean level has changed, a Gamma(α, β) prior is assumed on θ.

$$X \sim \text{Poisson}(\theta) \quad \text{and} \quad f(\underline{x}) = \prod_{i=1}^{n} \frac{\theta^{x_i}}{x_i!} e^{-\theta} \qquad \theta > 0, x_i \in (0, 1, 2, 3, \ldots)$$

The problem of determining if a change in the mean failure rate has occured can be thought of as choosing between the two models.

$$\begin{aligned} m_1 &: X \sim \text{Po}(\theta) \quad \theta = 1\,, \\ m_2 &: X \sim \text{Po}(\theta) \quad \theta \sim \text{Gamma}(\alpha, \beta)\,. \end{aligned}$$

(a) Define the marginal likelihood for a model m, the posterior odds, the Bayes factor and the prior odds in favour of model m_2.

[4 marks]

(b) Obtain expressions for the marginal likelihood for models m_1 and m_2.

[7 marks]

(c) Write down the posterior odds in favour of m_2, given that $f(m = 1) = p$.

[4 marks]

(d) Suppose that $\alpha = 2, \beta = 2, \underline{X} = (1, 1, 2)$ and that the two models are equally likely a priori. Find the posterior odds. What answer would you give the company about whether the mean has changed?

[10 marks]

4. Consider the following hierarchical model,

$$X|\theta \sim \text{Binomial}(n, \theta)$$

$$\theta \sim \text{Beta}(\alpha, \beta)$$

We know that

$$f(x|\theta) = \binom{n}{x} \theta^x (1-\theta)^{n-x}$$

and the form of the Beta distribution is given as the prior for the Bernoulli distribution on the formula sheet.

(a) Using the above information calculate the joint distribution of $f(X, \theta)$.

[3 marks]

(b) Derive the conditional distributions of $X|\theta$ and $\theta|X$.

[3 marks]

(c) (i) Define a Gibbs sampler giving a general algorithm and discussing why the algorithm would be used.

(ii) Write an example Gibbs sampling routine which could be used to sample from the joint distribution of X and θ.

(iii) Discuss how the sampler would be initialised and how the algorithm would be analysed, include examples of diagnostic plots which should be obtained and show the main features of a Gibbs sampling algorithm, describing the features depicted.

[16 marks]

(d) Describe how the generated samples could be used to estimate $E(\theta|X)$ and Var$(\theta|X)$.

[3 marks]

END OF PAPER