QUEEN MARY, UNIVERSITY OF LONDON

MTH6102: Bayesian Statistical Methods

Exercise sheet 10

2023-2024

This exercise sheet 10 is assessed and counts for 4% of the module total. The deadline for submission is **Monday the 11th December at 11am**.

Submit the R code used as an R script file (with extension .R). But you need to write the answers in a separate file. This can be a Word document, pdf or a clearly legible image of hand-written work. So you need to submit two files.

1. **50 marks.** Let the observed data be $y = (6, 4, 9, 2, 0, 3)$, a random sample from the Poisson distribution with mean λ, where $\lambda > 0$ is unknown. Suppose that we assume a Gamma(1, 1) prior distribution for λ. The posterior density, $p(\lambda \mid y)$, for λ is Gamma$(1+S, 1+n)$, where $S = \sum_{i=1}^{6} y_i$ and $n = 6$. Suppose that you want to construct a symmetric Metropolis-Hastings on the log-scale to generate a sample from this posterior distribution by using a normal proposal distribution with standard deviation $b = 0.2$.

   (a) Write down the steps in this symmetric Metropolis-Hastings (on the log-scale) to simulate realisations from the posterior density $p(\lambda \mid y)$.

   (b) Implement the algorithm in R and plot the observations as a function of the iterations. Use $M = 5000$ for the number of iterations.

   (c) To assess the accuracy compare the empirical distribution of the sample with the exact posterior density, Gamma$(1 + S, 1 + n)$.

   (d) Rerun the algorithm in R using a smaller $b = 0.01$ and a larger $b = 20$. What are the effects on the behaviour of the algorithm of making b smaller? What are the effects of making it larger?

   (e) Add code to count how many times the proposed value for λ was accepted. Rerun the algorithm using values of $b = 0.01$, $b = 0.2$ and $b = 20$, and each time calculate the proportion of steps that were accepted. Then plot this acceptance probability against b. Examine how the acceptance probability for this algorithm depends b.

2. **50 marks.** Let $y_1, \ldots, y_n$ be a sample from a Poisson distribution with mean λ, where λ is given a Gamma(α, β) prior distribution.

   (a) It is observed that $y_1 = y_2 = \cdots = y_n = 0$, and we take $\alpha = 1, \beta = 1$.

      i. What is the posterior distribution for λ?

      ii. What is the posterior mean?

      iii. What is the posterior median and an equal tail 95% credible interval for λ (without using R)?

(b) Show that if a new data-point x is generated from the same Poisson distribution, the posterior predictive probability that $x = 0$ is

$$p(x = 0 \mid y) = \frac{n+1}{n+2}.$$

(c) Now suppose that we have general $y_1, \ldots, y_n$, α and β; and that again x is a new data-point from the same Poisson distribution.

i. Find the mean and variance of x.

ii. Derive the full posterior predictive distribution for x.

(d) Use R to check as many of these results as you can.