QUEEN MARY, UNIVERSITY OF LONDON

MTH6102: Bayesian Statistical Methods

Solutions of exercise sheet 5

2023-2024

1. If the data follow a normal distribution with unknown mean μ and known standard deviation, and μ is assigned a normal prior distribution, show that the posterior mean for μ can be written as a weighted sum of the maximum likelihood estimate and the prior mean.

   Solution

   In the notation used in the lectures, the posterior distribution for μ is

   $$\mu \sim N(\mu_1, \sigma_1^2)$$

   where

   $$\mu_1 = \left(\frac{\mu_0}{\sigma_0^2} + \frac{n\bar{y}}{\sigma^2}\right) \Big/ \left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}\right)$$

   So the posterior mean μ_1 is

   $$\begin{aligned}\mu_1 &= \frac{\mu_0}{\sigma_0^2}\Big/\left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}\right) + \frac{n\bar{y}}{\sigma^2}\Big/\left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}\right) \\ &= (1-w)\mu_0 + w\bar{y},\end{aligned}$$

   where

   $$w = \frac{n}{\sigma^2}\Big/\left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}\right) = \frac{1}{1 + \dfrac{\sigma^2}{n\sigma_0^2}}.$$

2. The column in the exercise sheet 5 dataset, labelled `x`, contains the observed data be $x_1, \ldots, x_n$. Suppose that each data-point x_i is normally distributed, with unknown mean θ and standard deviation assumed to be equal to 1.

   As a prior distribution for θ, we assign a normal distribution with mean 5 and standard deviation 10.

   (a) What is the posterior distribution for θ?

   (b) What is the posterior mean for θ?

   (c) What is the posterior median for θ?

   (d) What is a 95% equal tail credible interval for θ?

   Solution:

(a) This is the problem with unknown mean, where the left-hand column below is the symbol used in the lectures and the right-hand column is the corresponding symbol or numerical value in this question:

μ	θ
σ	1
μ_0	5
σ_0	10

Let the sample mean be $\bar{x}$. The posterior distribution for θ is normal,

$$\theta \mid x \sim N(\mu_1, \sigma_1^2)$$

where

$$\mu_1 = \left(\frac{\mu_0}{\sigma_0^2} + \frac{n\bar{x}}{\sigma^2}\right) \Big/ \left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}\right)$$

$$\sigma_1^2 = 1 \Big/ \left(\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}\right)$$

(b) The posterior mean is μ_1.

(c) The posterior median is also μ_1, since the normal distribution is symmetric about its mean.

(d) For a 95% equal tail credible interval, we can either use the `qnorm` function in R (the quantile function of the normal distribution), after defining the posterior parameters `mu1` and `sigma1`

```
qnorm(c(0.025, 0.975), mean=mu1, sd=sigma1)
```

Or we can remember that the 2.5% and 97.5% percentiles of a normal distribution $N(\mu_1, \sigma_1^2)$ are given by

$$\mu_1 - 1.96\,\sigma_1,\ \mu_1 + 1.96\,\sigma_1,$$

and so these points form the credible interval.

3. Let $x_1, \ldots, x_n$ be iid Poisson(λ), and let λ have a gamma(α, β) distribution.

(a) Show that gamma(α, β) is conjugate to the Poisson likelihood.

(b) Calculate the posterior mean and variance.

(c) Show how to find a 95% equal tail credible interval for λ?

(d) Show how to find a 95% HPD credible interval for λ.

Solution

(a) Let $x = (x_1, \ldots, x_n)$. The posterior $p(\lambda \mid x)$ is

$$\begin{aligned} p(\lambda \mid x) &\propto \text{likelihood} \times \text{prior} \\ &= \frac{\beta^\alpha \lambda^{\alpha-1} \exp(-\beta\lambda)}{\Gamma(\alpha)} \cdot \exp(-n\lambda)\lambda^{\sum y_i} \\ &\propto \lambda^{\alpha+S-1} \exp(-(\beta+n)\lambda). \end{aligned}$$

Hence, the posterior distribution is Gamma(α_1, β_1) with $\alpha_1 = \alpha + S$ and $\beta_1 = \beta + n$, where n is the sample size and $S = \sum y_i$ is the sum of the data.

(b) The posterior mean is

$$\frac{\alpha_1}{\beta_1} = \frac{\alpha + S}{\beta + n}.$$

The posterior variance is

$$\frac{\alpha_1}{\beta_1^2} = \frac{\alpha + S}{(\beta + n)^2}.$$

(c) If the posterior parameters are `alpha1` and `beta1`, then the median and 95% credible interval are given by

```
qgamma(c(0.5, 0.025, 0.975), shape=alpha1, rate=beta1)
```

(d) We want to find (θ_L, θ_U) such that

$$\int_{\theta_L}^{\theta_U} p(\lambda \mid x)\, d\lambda = 0.95$$
$$p(\theta_L \mid x) = p(\theta_U \mid x).$$

where $p(\lambda \mid x) \sim \text{Gamma}(\alpha_1, \beta_1)$.

4. In an investigation into the size of the errors produced by a new measurement instrument, n measurements are taken of a standard sample of mass 1000 grams. The measurements (in grams), $y_1, \ldots, y_n$ can be modelled as a random sample from a normal distribution with known mean $\mu = 1000$ and unknown precision τ (reciprocal of the variance). Assume the prior gamma(α, β) distribution on τ, where $\alpha = 5$ and $\beta = 0.05$

(a) Six measurements are taken and the data is

$$1000.11, 999.96, 999.84, 999.89, 999.80, 1000.09.$$

Given these measurements, what is the posterior distribution of τ?

(b) Use R to find the posterior median and a 95% equal tail credible interval for τ.

(c) Can you find the posterior median for $\sigma = 1/\sqrt{\tau}$? Can you find a 95% credible interval for σ? [Hint: these do not need the derivation of the posterior distribution for σ, or any extensive calculations.]

Solution:

(a) The posterior distribution for τ is

$$\tau \sim \text{Gamma}\left(\alpha + \frac{n}{2}, \beta + \frac{\sum_{i=1}^{n}(y_i - \mu)^2}{2}\right)$$

Using the measurements, the posterior distribution for τ is $\text{Gamma}(\alpha_1, \beta_1)$ with

$$\alpha_1 = 5 + 6/2 = 8, \quad \beta_1 = 0.05 + 0.0995/2 = 0.09975$$

(b) Using the `qgamma` command gives 76.9 for the posterior median and $(34.6, 144.6)$ for the 95% credible interval.

(c) Now put $\sigma = 1/\sqrt{\tau}$.
For any $0 < c < 1$ and $t \geq 0$

$$P(\tau \leq t) = c \implies P(1/\sqrt{\tau} \geq 1/\sqrt{t}) = c$$

In particular, if τ_M is the posterior median for τ,

$$0.5 = P(\tau \leq \tau_M) = P(\sigma \geq 1/\sqrt{\tau_M})$$

$$P(\sigma \leq 1/\sqrt{\tau_M}) = 0.5$$

So $1/\sqrt{\tau_M}$ is the posterior median for σ.
If (τ_L, τ_U) is an equal tail 95% credible interval for τ,

$$P(\tau \leq \tau_L) = 0.025 \text{ and } P(\sigma \geq 1/\sqrt{\tau_L}) = 0.025$$

$$P(\tau \leq \tau_U) = 0.975 \text{ and } P(\sigma \geq 1/\sqrt{\tau_U}) = 0.975$$

So

$$P(\sigma \leq 1/\sqrt{\tau_U}) = 0.025,\ P(\sigma \geq 1/\sqrt{\tau_L}) = 0.025$$

Hence $(1/\sqrt{\tau_U}, 1/\sqrt{\tau_L})$ is an equal tail 95% credible interval for σ.