QUEEN MARY, UNIVERSITY OF LONDON

MTH6102: Bayesian Statistical Methods

Exercise sheet 11

2023-2024

1. If the data is normally distributed

$$y_1, \dots, y_n \sim N(\mu, \sigma^2)$$

where σ is known and μ has prior distribution $\mu \sim N(\mu_0, \sigma_0^2)$, we saw that the posterior distribution $p(\mu \mid y)$ is

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

where

$$\mu_1 = \frac{\mu_0/\sigma_0^2 + n\bar{y}/\sigma^2}{1/\sigma_0^2 + n/\sigma^2},\ \sigma_1^2 = \frac{1}{1/\sigma_0^2 + n/\sigma^2}.$$

Suppose we take this model as model M_2, and also consider model M_1 in which μ is known to be equal to μ_0:

$$M_1 : y_i \sim N(\mu_0, \sigma^2),\ i = 1, \dots, n$$
$$M_2 : y_i \sim N(\mu, \sigma^2),\ i = 1, \dots, n\ ;\ \mu \sim N(\mu_0, \sigma_0^2)$$

In model M_1, μ_0 is fixed; for both models, σ is known.

It can be shown that

$$p(y \mid M_2) = \frac{\sigma_1\, p(y \mid M_1)}{\sigma_0\, exp\left(-\dfrac{(\mu_1 - \mu_0)^2}{2\sigma_1^2}\right)}.$$

(a) What is the Bayes factor B_{12} for comparing models 1 and 2?

(b) Show that for sufficiently large values of σ_0

$$B_{12} \approx \frac{\sqrt{n}\sigma_0}{\sigma}\, exp\left(-\frac{n(\bar{y} - \mu_0)^2}{2\sigma^2}\right).$$

For the data, let ABC be the last three digits of your 9 digit ID number. Take $y_1, \dots, y_n = (20 + A, 20 + B, 20 + C, 28, 30)$. Take $\mu_0 = 25$ and $\sigma = 3$.

(c) Calculate the posterior mean and standard deviation of μ under model M_2. Do this for $\sigma_0 = 10$, and then repeat with $\sigma_0 = 100$.

(d) For the same values of σ_0 as in (1c), calculate the Bayes factor B_{12}, using both the exact formula and the approximation.

(e) Assuming that the models have equal prior probabilities, calculate the posterior probabilities that each model is the correct model, $p(M_1 \mid y)$ and $p(M_2 \mid y)$. Do this for each of the two values of σ_0. Use the exact Bayes factor, not the approximation.

(f) Comment briefly on the results of parts (1c) and (1e).

2. Suppose that the observed data is $y_1, \ldots, y_n$, which we assume is a sample from a Poisson distribution. Two models are under consideration: model M_1 is that the distribution that generated the data is Poisson with mean 1; model M_2 is that the data were generated by a Poisson distribution with mean λ, where λ has a Gamma(α, β) prior distribution.

   (a) Suppose that $y_1 = \cdots = y_n = 0$ and that $\alpha = 1$, $\beta = 1$. Show that in this case the Bayes factor B_{12} for comparing the two models is given by

   $$B_{12} = (n+1)e^{-n}.$$

   If $n = 10$, what is the Bayes factor? If we also take as prior model probabilities $p(M_1) = 1/2$ and $p(M_1) = 1/2$, what is the posterior probability of model M_1?

   (b) Show that for general data and $\alpha = 1$, $\beta = 1$, the Bayes factor is given by

   $$B_{12} = \frac{(n+1)^{S+1}e^{-n}}{S!}, \text{ where } S = \sum_{i=1}^{n} y_i.$$

   [This could be done by direct integration, which would involve manipulating gamma functions. Alternatively one could rearrange the Bayes theorem formula as was done for the normal example result that is quoted in question (1).]