Chapter 11

Making Bayesian analysis objective?

11.1 Jeffreys prior for a normal likelihood

Suppose that we are modelling the result of a medical test, which to a suitable approximation can be regarded as being continuous and unbounded. We suppose that a normal probability model is a reasonable sampling model to use here, $X_i \sim \mathcal{N}(\mu, \sigma)$, where μ is unknown but σ is known (perhaps based on the results of many previous tests).

Problem 11.1.1. Write down the likelihood for a single observation.

$$L(\mu|X_i, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(X_i-\mu)^2}{2\sigma^2}} \tag{11.1}$$

Problem 11.1.2. Find the information matrix (here a scalar).

First calculate the log-likelihood for a single observation,

$$l(\mu|X_i, \sigma) = const - \frac{(X_i - \mu)^2}{2\sigma^2}. \tag{11.2}$$

Differentiating this with respect to μ we obtain,

$$\frac{\partial l}{\partial \mu} = const - \frac{\mu}{\sigma^2}, \tag{11.3}$$

which when we differentiate again yields,

$$\frac{\partial^2 l}{\partial \mu^2} = -\frac{1}{\sigma^2}. \tag{11.4}$$

Therefore the information matrix is given by,

$$I(\mu) = \frac{1}{\sigma^2}. \tag{11.5}$$

Problem 11.1.3. Hence calculate the information matrix for a sample of N observations.

Since the only difference is the summation we obtain,

$$I(\mu) = \frac{N}{\sigma^2}. \tag{11.6}$$

Problem 11.1.4. State Jeffreys prior for μ.

The Jeffreys prior is simply the square root of the information matrix,

$$p(\mu) \propto \sqrt{\frac{N}{\sigma^2}} \tag{11.7}$$

Problem 11.1.5. Is Jeffreys prior proper here?

No! It is a uniform prior between $-\infty < \mu < +\infty$.

11.2 The illusion of uninformative priors revisited

Suppose that θ represents the probability that one randomly-chosen individual has a particular disease.

Problem 11.2.1. Suppose that we start by assigning a uniform prior on θ. Use sampling to estimate the prior distribution that in a sample of two one person has the disease and another doesn't. Hence comment on the assumption that a uniform prior is uninformative.

The probability of one of each type in a sample of two is $\phi = \theta(1-\theta)$. So if we first sample θ then transform it as according to this probability we find that there is a maximum probability of 0.25 with this decreasing to 0 as $\phi \rightarrow 0$. This is not uninformative. So a uninformative prior in one frame of reference is not uninformative in another.

Problem 11.2.2. Assume instead that we ascribe a uniform prior to the probability that 2/2 individuals have the disease. What is the implicit prior distribution for the probability that one individual has the disease?

Use sampling again here, but instead transform according to $\theta^{0.5}$. This gives us a linearly increasing line (the exact form is $p(\phi) = 2\phi$).

Bibliography