STK 4021 - Applied Bayesian Analysis and Numerical Methods
Thordis L. Thorarinsdottir
Fall 2014

Problem Set 12

Problem 1 (binomial and multinomial models). Suppose data $(y_1, \dots, y_J)$ follow a multinomial distribution with parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_J)$. Also suppose that $\boldsymbol{\theta}$ has a Dirichlet prior distribution. Let $\alpha = \frac{\theta_1}{\theta_1+\theta_2}$.

(a) Write the marginal posterior distribution for α.

(b) Show that the distribution in (a) is identical to the posterior distribution for α obtained by treating y_1 as an observation from the binomial distribution with probability α and sample size $y_1 + y_2$, ignoring the data $y_3, \dots, y_J$.

Problem 2 (Poisson models).

(a) Suppose $y|\theta \sim \text{Po}(\theta)$. Find Jeffreys' prior density for θ, and then find α and β for which the $\Gamma(\alpha, \beta)$ density is a close match to Jeffreys' density.

(b) Suppose $y|\theta \sim \text{Po}(\theta)$ and $\theta \sim \Gamma(\alpha, \beta)$. Then the marginal (prior predictive) distribution of y is negative binomial with parameters α and β. Use the formulas

$$\mathbb{E}(\theta) = \mathbb{E}(\mathbb{E}(\theta|y))$$
$$\text{var}(\theta) = \mathbb{E}(\text{var}(\theta|y)) + \text{var}(\mathbb{E}(\theta|y))$$

to derive the mean and the variance of this marginal distribution.

Problem 3 (Poisson and binomial distributions). A student sits on a street corner for an hour and records the number of bicycles b and the number of other vehicles v that go by. Two models are considered:

- The outcomes b and v have independent Poisson distributions, with unknown means θ_b and θ_v.
- The outcome b has a binomial distribution, with unknown probability p and sample size $b + v$.

Show that the two models have the same likelihood if we define $p = \frac{\theta_b}{\theta_a+\theta_b}$.

Problem 4 (discrete mixture models).

(a) If $p_m(\theta)$, for $m = 1, \dots, M$ are conjugate prior densities for the sampling model $p(y|\theta)$, show that the class of finite mixture prior densities given by

$$p(\theta) = \sum_{m=1}^{M} \lambda_m p_m(\theta)$$

is also a conjugate class, where the λ_m's are nonnegative weights that sum to 1.

(b) Use the mixture form to create a bimodal prior density for a normal mean, that is thought to be near 1, with a standard deviation of 0.5, but has a small probability of being near -1, with the same standard deviation. If the variance of each observation $y_1, \ldots, y_{10}$ is known to be 1, and their observed mean is $\bar{y} = -0.25$, derive your posterior distribution for the mean.

Solutions will be discussed in class on December 5.