STAT 489-01: Bayesian Methods of Data Analysis

Problem Set 2

Assigned 2017 January 31
Due 2017 February 7

Show your work on all problems! Be sure to give credit to any collaborators, or outside sources used in solving the problems. Note that if using an outside source to do a calculation, you should use it as a reference for the method, and actually carry out the calculation yourself; it's not sufficient to quote the results of a calculation contained in an outside source.

1 Inference with the Cauchy Distribution

The Cauchy distribution[1]

$$p(y_i|\theta) = \frac{1}{\pi}\frac{1}{1+(y_i-\theta)^2} \qquad -\infty < y_i < \infty \tag{1.1}$$

is often maligned as "pathological" because its mean, variance and other moments cannot be defined, since the integral $\int_{-\infty}^{\infty}(y_i)^n\,|p(y_i|\theta)|$ diverges for $n > 0$. It poses no particular difficulty for Bayesian inference with minimal assumptions, however. Consider the data $\{y_i\} = \{-2.809, 1.034, 5.441, 0.908, 1.223\}$ which are assumed to be a sample from a Cauchy distribution with location parameter θ. Perform the following inferences numerically using the grid approximation, peferably in R. (Note that you will not need to code up the probability distribution; you can use `dcauchy()` with the appropriate location parameter.)

- **(a)** Construct and plot the posterior pdf $p(\theta|\mathbf{y})$ associated with these data, assuming an improper uniform prior $p(\theta) = \text{const.}$
- **(b)** Find the median of the posterior pdf, and the symmetric 90% plausible interval.
- **(c)** Find the maximum a posteriori estimate of θ and the 90% HDR.
- **(d)** Estimate the curvature and plot the Gaussian approximation to the posterior on the same axes as the posterior itself. There are two ways to do this: either you can go back to the form of the posterior and work out the second derivative $\frac{\partial^2 \ln p(\theta|\mathbf{y})}{\partial\theta^2}$ analytically and plug your numerical MAP estimate into it, or you can install the `rethinking` R package which has can numerically estimate the Hessian (in this case the second derivative) as part of its `map()` function.

2 Gelman Chapter 2, Excercise 5

3 Gelman Chapter 2, Excercise 8

[1] which is a Student t-distribution with one degree of freedom

4 Counting Experiments

Consider an experiment which counts the number of events y in a time interval of duration T from a Poisson process with unknown rate $\theta > 0$, whose (discrete) sampling distribution is

$$p(y|\theta, T, I) = \frac{(\theta T)^y}{y!} e^{-\theta T} \qquad y = 0, 1, 2, \ldots \tag{4.1}$$

(a) Suppose the prior is given by a Gamma distribution[2] with parameters $\alpha > 0$ and $\beta > 0$:

$$p(\theta|I_{\alpha,\beta}) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} e^{-\beta\theta} \qquad 0 < \theta < \infty \tag{4.2}$$

Work out the posterior distribution $p(\theta|y, T, I_{\alpha,\beta})$, and show that it is also a Gamma distribution, and note the parameters. (This demonstrates that the Gamma is a conjugate prior family for this problem.)

(b) Show that the improper prior

$$p(\theta|I_c) = \text{constant} \qquad 0 < \theta < \infty \tag{4.3}$$

is a limiting form of the Gamma distribution and give the relevant limits of the parameters.

(c) Using the results of parts (a) and (b), specify what conditions, if any, are needed on y and T so that the posterior $p(\theta|y, T, I_c)$ is normalizable.

(d) Show that the improper prior

$$p(\theta|I_\ell) \propto \frac{1}{\theta} \qquad 0 < \theta < \infty \tag{4.4}$$

is a limiting form of the Gamma distribution and give the relevant limits of the parameters.

(e) Using the results of parts (a) and (b), specify what conditions, if any, are needed on y and T so that the posterior $p(\theta|y, T, I_\ell)$ is normalizable.

(f) Consider the transformation $\lambda = \ln\theta$. Work out the transformed prior distributions $p(\lambda|I_{\alpha,\beta})$, $p(\lambda|I_c)$, and $p(\lambda|I_\ell)$.

(g) Consider an alternative experiment where the data collected are $\mathbf{t} = t_1, \ldots, t_y$, a sample of size y from an exponential distribution of unknown rate parameter $\theta > 0$:

$$p(t_i|\theta, y, I) = \theta e^{-t_i\theta} \qquad 0 < t_i < \infty \tag{4.5}$$

Show that, for an arbitrary prior distribution $p(\theta|I)$, the posterior distribution will be $p(\theta|\mathbf{t}, y, I) = p(\theta|y, T, I)$, where $p(\theta|y, T, I)$ is the posterior for the Poisson counting experiment above, and $T = \sum_{i=1}^{y} t_i$.

[2] Note that the definition of the parameter β is different from what you may be used to, and in particular from the one used in books like Devore and Hogg. This is the convention that Gelman uses, e.g., in Table A.1.