

Bayesian inference

Assessment workshop

Arni Magnusson

Hafro, 16–24 Jan 2012

Outline

- 1 Overview - background, inference, priors, MCMC
- 2 Application - practical Bayesian, MCMC, exercise

Outline

- 1 Overview - background, inference, priors, MCMC
- 2 Application - practical Bayesian, MCMC, exercise

History

Bayes (1763) and Laplace (1774)

Bayesian inference with uniform priors

Fiercely debated in the 1930s–1960s

Dismissed by Fisher and Neyman

Widely used since the 1990s

Gelman et al. (1995), MCMC

<http://cran.r-project.org/web/views/Bayesian.html>

Bayesian theory

$$P(A|B) = \frac{P(B|A) \times P(A)}{P(B)}$$

$$P(\theta|Y) = \frac{P(Y|\theta) \times P(\theta)}{P(Y)}$$

$$P(\theta|Y) \propto L(\theta|Y) \times P(\theta)$$

$$f = -\log L + \text{Penalty}$$

Statistical inference

Significance tests, p values

Given that the null hypothesis is true . . . probability of getting a sample . . .
Sometimes all we need

Uncertainty, confidence intervals

Given that the estimates are the true parameter values . . .
. . . and the experiment was repeated many times . . .

Subjective probability

Probability redefined

Likelihood inference, uniform prior

Support, penalized likelihood

$$f = -\log L + \text{constant}$$

Prior distribution

- Uniform
- Conjugate
- Objective, empirical, meta analysis
- Assist model convergence without fixing a parameter

Markov chain Monte Carlo

Metropolis algorithm

Start from the best fit, and then:

- 1 Sample proposal θ^* from jumping distribution at time t , $J_t(\theta^* | \theta^{t-1})$
- 2 Calculate the ratio of the densities, $r = \frac{p(\theta^* | y)}{p(\theta^{t-1} | y)}$
- 3 Set $\theta^t = \begin{cases} \theta^* & \text{with probability } \min(r, 1) \\ \theta^{t-1} & \text{otherwise} \end{cases}$

Repeat *many* times

Markov chain Monte Carlo

Metropolis-Hastings algorithm

Extends the Metropolis algorithm:

- Jumping distribution $J_t(\theta^* | \theta^{t-1})$ can be asymmetric
- To accommodate the asymmetry, the ratio becomes

$$r = \frac{p(\theta^* | y) / J_t(\theta^* | \theta^{t-1})}{p(\theta^{t-1} | y) / J_t(\theta^{t-1} | \theta^*)}$$

Outline

- 1 Overview - background, inference, priors, MCMC
- 2 Application - practical Bayesian, MCMC, exercise

In practice

Frequentist and Bayesian methods address different questions

How likely is $\theta = 160$ or $\theta = 220$?

What is the probability of $B < B_{\text{lim}}$?

MCMC is practical

Evaluate uncertainty

Find global minimum

Diagnose model behavior

Exercise

Use a prior for the Beverton-Holt parameter a (1/initial slope)

$$R = \frac{S}{a + bS}$$

Imagine that we are having model convergence problems, but know from other data-rich stocks that for this species, a tends to be around 0.12

Use normal prior distribution with $\mu=0.12$. Try (1) $\sigma=0.02$, (2) $\sigma=0.06$, and (3) shrinking the dataset down to the first $n=30$ observations.

Hint: $f = -\log L + \text{Penalty}$