

MCMC2: Lab Session 2: SIR Models (1)

Theo Kypraios and Phil O'Neill

Background

In Lecture 2 we discussed about the modelling process in general and what steps does one have to follow to perform Bayesian inference. In this lab session we look at how to implement an MCMC algorithm to infer the parameters that govern the incubation period distribution that was discussed in the first part of the lecture.

The Data and the Model

Suppose that we have data on incubation periods $y = (y_1, \dots, y_n)$. We assume that the data are *independent* draws from a Gamma distribution with shape α and rate β , i.e.

$$y_i \sim Ga(\alpha, \beta), \quad (\text{i.i.d.}) \quad i = 1, \dots, n.$$

where the Gamma distribution has probability density function:

$$f(x|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp\{-\beta x\}, \quad \alpha > 0, \beta > 0.$$

The goal is to make (sampling-based) Bayesian inference for the parameters α and β .

Exercises

Exercise 1

Assume that α and β are *a priori* independent and assign Gamma distributions with parameters λ_α and ν_α , and λ_β and ν_β , respectively. Write down (i.e. express in mathematical terms) the posterior density of interest, i.e.

$$\pi(\alpha, \beta|y) \propto \pi(y|\alpha, \beta)\pi(\alpha)\pi(\beta)$$

where $\pi(y|\alpha, \beta)$ denotes the likelihood function.

Exercise 2

Having obtained the joint posterior density in Exercise 1, first derive the densities of the posterior distribution of the parameters α and β up to proportionality, i.e. $\pi(\alpha|\beta, y)$ and $\pi(\beta|\alpha, y)$.

Is any of these densities of a known form (i.e. the density of a standard/well-known distribution, e.g. Gamma, Normal etc)?

Exercise 3

Write a R function to implement an MCMC algorithm (from scratch!) which samples from the joint posterior distribution $\pi(\alpha, \beta|y)$.

First, think about what the input (i.e. arguments) the function should have, e.g. a vector corresponding to data y , hyper-parameter values $\lambda_\alpha, \nu_\alpha$ etc. Then also think about what the output should be; surely the function should output the posterior samples of α and β . Anything else?

To get you started, a pseudo-code of this MCMC algorithm is given below:

- i. Choose initial values for α and β ;
- ii. Update α by sampling from $\pi(\alpha|\beta, y)$ using a Gaussian random-walk;
- iii. Update β by sampling from $\pi(\beta|\alpha, y)$ directly.
- iv. Goto to Step (b)

Look at the trace plots and convince yourself that the chain has reached stationarity and is mixing well.

Exercise 4

Test your algorithm is producing sensible results.

One way to do this is to first simulate a dataset from a Gamma distribution with some specific values for α and β , eg. `y <- rgamma(100, 4, 2)` will simulate a vector of 100 draws from a $\text{Ga}(4, 2)$.

Then use your MCMC algorithm (in Exercise 3) to sample from the joint posterior distribution of α and β assuming vague priors, e.g. $\lambda_\alpha = \lambda_\beta = 1$ and $\nu_\alpha = \nu_\beta = 10^{-3}$.

In principle, your marginal posterior distributions ($\pi(\alpha|y)$ and $\pi(\beta|y)$) should be centered around the values you have chosen to simulate the data from (e.g. 4 and 2 in this case).

Exercise 5

If you have convinced yourself that your MCMC algorithm is doing what is supposed to be doing, then use it to fit the model to the *Campylobacter* data (24 observations in total) from Evans et al. 1996) by

```
y <- c(rep(2,2), rep(3, 6), rep(4, 11), rep(5, 3), rep(7,7))
```

Does the model offer a good fit? How would you go assessing that?

Exercise 6* [Optional]

Now that you have got a working algorithm, can you improve its mixing?

One way to do this is to update both α and β at the same time (in a block). This can be done, for example, and as discussed in the lecture, by proposing values drawn from some (bivariate distribution) with density $q(\alpha, \beta)$ and accept/reject according to the Metropolis-Hastings ratio. There are many choices here, but not all of them will lead to an efficient sampler.

Here are a few choices to try:

- Gaussian random-walk in a block, i.e. propose values for α and β drawn from a bivariate Normal distribution with mean the current values of α and β and some (2×2) variance-covariance matrix, and then accept/reject using the Metropolis-Hastings ratio.
- Gaussian approximation to the posterior distribution $\pi(\alpha, \beta|y)$. (Hint: use the command `optim` to find the values of α and β that maximise the posterior density).