

MCMC2: Lab Session 4: Households

Theo Kypraios and Phil O'Neill

Background

In Lecture 4 we were concerned with models for disease transmission which incorporate households. Although we looked at two distinct cases with regards to what sort of data are available, in this lab session we will focus on the case where there is only *final outcome* data available. In other words, the available data consist only of the final number of cases in each household.

The main objective of this lab session is to draw Bayesian inference for the model parameters by implementing a Markov Chain Monte Carlo algorithm. Below, there is a brief description of the model (see the lecture notes for more details).

The model, the data and the likelihood

We assume that a population of N individuals is partitioned into households. Each individual in the population has, independently, a constant “risk” per unit time of becoming infected from the community. In addition, within a household, the disease spreads according to the mechanism of an SIR model.

Suppose that the data consist of the set of numbers $\mathbf{n} = \{n(j, k)\}$ where $n(j, k)$ = number of households in which j out of k initial susceptibles become infected. Therefore, the likelihood takes the form

$$\pi(\mathbf{n}|\alpha, p) = \prod_j \prod_k q(j, k)^{n(j, k)}$$

where $q(j, k)$ is the probability that in a household of size k the number of individuals who ever become infected is j .

In the lecture we have derived an expression for the probability $q(j, k)$ and showed that it depends on two parameters:

- p which is the probability that an individual avoids infection from outside the household and
- α which is the person-to-person infection rate between individuals within the same household.

Preliminaries

Download the file `households.R` from the module’s website

The file contains two functions:

- `compute_conditional_prob` which computes the probability that j susceptibles become infected in an SIR model with constant infectious period of length c , a initial infectives and m initial susceptibles.

```
# this is a function to compute the final size P(T=k|Y=y) = probability  
# that k-y susceptibles become infected with $y$ initial infectives  
# and n-y susceptibles  
  
# m initial susceptibles  
# a initial infectives
```

```

# alpha infection rate
# the length of the fixed infectious period
compute.conditional.prob <- function(m, j, a, alpha, c){
  if (j==0) {
    res <- (exp(-alpha*m*c))^a
  }
  else {
    part.one <- exp(-alpha*(m-j)*c)^(j+a)
    part.two <- choose(m, j)

    sum <- 0;
    for (k in 0:(j-1)) {
      sum <- sum + choose(m-k, j-k)*compute.conditional.prob(m,k,a,alpha,c)/(exp(-alpha*(m-j)*c)^(k+a))
    }
    res <- part.one*(part.two - sum)
  }
  return(res)
}

```

- `compute.marginal.probab` which computes the probability that in a household of size n the number of individuals who ever become infected is k .

```

# We wish to compute  $P(T = k)$ , for  $k = 0, \dots, n$ 
# "n" is the size of the household
# "alpha" is the person to person infection rate
# "c" is the length of the (fixed) infectious period
compute.marginal.prob <- function(k, n, c, alpha, p) {
  prob <- 0;
  for (y in 0:n) {
    if (k >= y) {
      cond.prob <- compute.conditional.prob(n - y, k - y, y, alpha, c)
    }
    else {
      cond.prob <- 0
    }
    prob <- prob + cond.prob*dbinom(y, n, 1 - p)
  }
  prob
}

```

In addition, Table 1 presents some Asian influenza epidemic household data taken from Longini & Koopman (1982). There is a community which consists of 42 households each of them of size 3.

```

cases <- c(0,1,2,3)
households <- c(29,9,2,2)
mat <- matrix(NA, nrow=4, 2)
mat[,1] <- cases
mat[,2] <- households
mat <- as.data.frame(mat)
colnames(mat) <- c("No.of.Cases", "No.of.Households")
mat

```

```

##   No.of.Cases No.of.Households
## 1           0                29
## 2           1                 9

```

```
## 3          2          2
## 4          3          2
```

Exercises

Exercise 1

Go through each of the functions in the file `households.R` and make sure that you understand what is going on.

Exercise 2

Write a function in R which computes the log-likelihood of the data given the parameters p and α assuming that the infectious period is constant and of length one unit.

Exercise 3

Suppose that we are interested in drawing Bayesian inference for the model parameters and therefore, we could employ an MCMC algorithm to draw samples from the posterior distribution:

$$\pi(p, \alpha | \mathbf{n}) \propto \pi(\mathbf{n} | \alpha, p) \times \pi(p) \times \pi(\alpha).$$

Assume that *a priori*

$$p \sim \text{Beta}(\mu_p, \nu_p)$$

and

$$\alpha \sim \text{Gamma}(\mu_\alpha, \nu_\alpha).$$

What values should we choose for the hyper-parameters μ_p, ν_p, μ_α and ν_α such that we end up with fairly uninformative priors for the model parameters p and α ?

Exercise 4

Write a function in R which draws samples from $\pi(p, \alpha | \mathbf{n})$ using MCMC and by adopting uninformative priors for the parameters of interest p and α .

Hint: In the lecture we discussed two different ways to update the parameter p . Either using an independence sampler or a random walk Metropolis. For the purposes of this exercise, update p by an independence sampler.

Exercise 5

Investigate how correlated the posterior samples of α and p are.

Exercise 6

Draw samples from the posterior distribution of the probability that a susceptible individual avoids infection from one infected household member.

Exercise 7

In the current version of your MCMC algorithm the parameter p is updated via an independence sampler. Modify your code such that p is now updated using a random walk Metropolis. Does the mixing of your algorithm improve?