MCMC 2: Lecture 2
SIR models (1)

Phil O'Neill Theo Kypraios
School of Mathematical Sciences
University of Nottingham

Contents

- 1. The modelling process
- 2. Gamma distribution example
- 3. General (Markov) SIR epidemic model

Contents

- 1. The modelling process
- 2. Gamma distribution example
- 3. General (Markov) SIR epidemic model

1. The modelling process

The general procedure we use is as follows.

- We have some observed data y
- We formulate a model with parameters θ
- In a Bayesian framework, we are interested in the posterior density $\pi(\theta \mid y)$

1. The modelling process

- By Bayes' Theorem, we have

  $\pi(\theta \mid y) \propto \pi(y \mid \theta)\,\pi(\theta)$

  posterior $\propto$ likelihood $\times$ prior
- $\pi(\theta \mid y)$ is usually not known explicitly (i.e. we typically only know it up to proportionality), so we then use an MCMC algorithm to get samples from $\pi(\theta \mid y)$

1. The modelling process

- Typically, θ is multi-dimensional. If, say, $\theta = (\theta_1, \ldots, \theta_n)$, then we usually need a way of updating each θ_k.

  If the full conditional density

  $$\pi(\theta_k \mid \theta_1, \ldots, \theta_{k-1}, \theta_{k+1}, \ldots, \theta_n, y)$$

  is known explicitly* then we can use it to perform the update for θ_k ("Gibbs step")

  * i.e. not just up to proportionality

1. The modelling process

- Conversely, if the full conditional density is not known explicitly then we can update θ_k using a Metropolis-Hastings step.

1. The modelling process

What if the likelihood $\pi(y \mid \theta)$ is unknown?

(Meaning – hard/complicated to compute)

Two possible solutions are

- Data augmentation: introduce extra quantities x such that $\pi(x, y \mid \theta)$ is tractable
- Give up on MCMC and do something else… (e.g. Approximate Bayesian Computation = ABC)

Contents

- 1. The modelling process
- 2. Gamma distribution example
- 3. General (Markov) SIR epidemic model

2. Gamma distribution example

Example

Suppose we have data on incubation periods

$$y = (y_1, \ldots, y_n)$$

and we wish to fit a Gamma distribution to these data.

Day	2	3	4	5	6	7
Freq	2	6	11	3	0	1

$$y_1 = 2,\ y_2 = 2,\ y_3 = 3,\ \ldots,\ y_{23} = 7$$

(Campylobacter data from Evans et al. 1996)

2. Gamma distribution example

The gamma distribution has probability density function

$$f(x \mid \alpha, \beta) = \beta^{\alpha} x^{\alpha-1} \exp(-\beta x) / \Gamma(\alpha)$$

(where $x > 0$, $\alpha > 0$, $\beta > 0$).

Assuming the data are independent draws from this distribution, the likelihood is

$$\pi(y \mid \alpha, \beta)$$

$$= f(y_1 \mid \alpha, \beta)\, f(y_2 \mid \alpha, \beta) \ldots f(y_n \mid \alpha, \beta)$$

2. Gamma distribution example

The likelihood simplifies to

$\pi(y \mid \alpha, \beta)$

$= \beta^{n\alpha} \prod_k y_k^{\alpha-1} \exp(-\beta \sum y_k) / \{\Gamma(\alpha)\}^n$

We assign independent priors as

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

$\beta \sim \text{Gamma}(\lambda_\beta, \nu_\beta)$

So

$\pi(\alpha) = f(\alpha \mid \lambda_\alpha, \nu_\alpha)$ and $\pi(\beta) = f(\beta \mid \lambda_\beta, \nu_\beta)$

2. Gamma distribution example

The posterior density of interest is

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

So to define an MCMC algorithm to sample from this target density we need a way to update α and β.

One option is to update them separately.

2. Gamma distribution example

<u>For separate updates:</u>

Find the full conditional densities

$$\pi(\alpha \mid \beta, y) \quad \text{and} \quad \pi(\beta \mid \alpha, y)$$

In this case (see lab exercise) we find that

(i) $\pi(\alpha \mid \beta, y)$ is only known up to proportionality;

(ii) $\pi(\beta \mid \alpha, y)$ is the density of a Gamma distribution

2. Gamma distribution example

Therefore we

(i) update α using a Metropolis-Hastings step, i.e. we propose a new value α* from a proposal density $q(\alpha^* \mid \alpha)$ and accept α* with probability min(p,1) where

$$p = \frac{\pi(y \mid \alpha^*, \beta)\, \pi(\alpha^*)\, q(\alpha \mid \alpha^*)}{\pi(y \mid \alpha, \beta)\, \pi(\alpha)\, q(\alpha^* \mid \alpha)}$$

2. Gamma distribution example

The choice of $q(\alpha^* \mid \alpha)$ is fairly arbitrary; possible options include

- Propose $\alpha^* \sim N(\alpha, \sigma^2)$

  (Gaussian random walk)

- Propose $\alpha^* \sim \text{Gamma}(a,b)$

  (independence sampler – needs a good choice of a and b to work well)

2. Gamma distribution example

(ii) update β using a Gamma distribution

(see lab exercise for details!)

2. Gamma distribution example

Block updating

An alternative to separate updates for α and β is to update them simultaneously in a "block".

In this case we could do this using M-H, e.g.

propose (α^*, β^*) from $q(\alpha^*, \beta^* | \alpha, \beta)$ and then accept/reject accordingly.

One reason to do this is if α and β are strongly correlated, i.e. it is hard to move one without the other.

Contents

- 1. The modelling process
- 2. Gamma distribution example
- 3. General (Markov) SIR epidemic model

3. General epidemic model

Population of N individuals

At time t there are:

S(t) susceptibles

I(t) infectives

R(t) recovered / immune individuals

Thus $S(t)+I(t)+R(t) = N$ for all t.

Initially $(S(0),I(0),R(0)) = (N-1,1,0)$.

3. General epidemic model

- Each infectious individual remains so for a length of time $T_I \sim \text{Exp}(\gamma)$.
- During this time, infectious contacts occur with each susceptible according to a Poisson process of rate β / N.
- Thus overall infection rate is $\beta S(t) I(t) / N$.
- Two model parameters, β and γ.

3. General epidemic model

- We call this model the “Markov SIR model”.
- This is because the process

$\{(S(t), I(t)): t \geq 0\}$

is a bivariate Markov chain. This follows from the fact that the infectious periods are exponentially distributed.

- If infectious periods are not exponential we have a “non-Markov” SIR model.

3. General epidemic model

- Suppose we observe n removals at times

$$r_1 \leq r_2 \leq r_3 \leq \ldots \leq r_n.$$

- The inference problem is to estimate the model parameters β and γ.
- In a Bayesian setting, this means we need to find (or sample from) the posterior density

$$\pi(\beta, \gamma \mid r_1, r_2, \ldots, r_n)$$

3. General epidemic model

- However, the likelihood

$$\pi(r_1, r_2, \ldots, r_n \mid \beta, \gamma)$$

is very hard to compute.

- A solution (as discussed in MCMC I) is to introduce infection times as extra variables to give a tractable augmented likelihood.

3. General epidemic model

Augmented Likelihood

- Let b be the label of the last removal time, i.e. $r_b \geq r_k$ for all k = 1, ..., n.
- Given removal data, b is observed and fixed for any given labelling.
- Define a as the label of the first infection time, i.e. $i_a < i_k$ for all $k \neq a$.
- Given removal data, a is unknown.

3. General epidemic model

- Define $\mathbf{r} = (r_1, r_2, \ldots, r_n)$
- Define $\mathbf{i} = (i_1, i_2, \ldots, i_{a-1}, i_{a+1}, \ldots, i_n)$
- Let

$$f(x \mid \gamma) = \gamma \exp(-\gamma x) \qquad (x > 0)$$

denote the probability density function of the infectious period distribution (Exp(γ)).

3. General epidemic model

- The augmented likelihood is

$$\pi(\mathbf{i}, \mathbf{r} \mid \beta, \gamma, i_a, a) =$$

$$\prod_{j \neq a} \beta N^{-1} I(i_j-) \times \exp\left(-\beta N^{-1} \int S(t) I(t)\, dt\right) \times \prod_{1 \leq j \leq n} f(r_j - i_j \mid \gamma)$$

$$= \prod_{j \neq a} \beta N^{-1} I(i_j-) \times \exp\left(-\beta N^{-1} \int S(t) I(t)\, dt\right) \times \gamma^n \exp\{-\gamma \sum (r_j - i_j)\}$$

where I(t-) means I(t) <u>just before</u> time t

3. General epidemic model

- The target posterior density is

  $\pi(\beta, \gamma, \mathbf{i}, i_a, a \mid \mathbf{r})$

  $\propto \pi(\mathbf{i}, \mathbf{r} \mid \beta, \gamma, i_a, a)\, \pi(\beta, \gamma, i_a, a)$

- Set independent priors as

  $\beta \sim \text{Gamma}(m_\beta, \lambda_\beta)$

  $\gamma \sim \text{Gamma}(m_\gamma, \lambda_\gamma)$

  $a \sim$ uniform on $\{1, \ldots, n\}$

  $i_a \sim$ uniform on $(-\infty, r_1)$

3. General epidemic model

- So an MCMC algorithm that targets the posterior density needs ways of updating the parameters β, γ, $\mathbf{i}$, i_a and a.
- In each case, we can first find the full conditional distribution to see if we get a standard distribution. If so, we can update using that distribution. If not, we need a Metropolis-Hastings step.

3. General epidemic model

- To update β: first find the full conditional density $\pi(\beta \mid \gamma, \mathbf{i}, i_a, a, \mathbf{r})$.
- We do this by finding all the terms involving β in the posterior (=likelihood × prior)
- We find:

$$\beta^{n} \exp(-\beta N^{-1} \int S(t)I(t)\,dt)\, \beta^{m_\beta - 1} \exp(-\beta \lambda_\beta)$$

$$= \beta^{n+m_\beta-1} \exp\{-\beta (N^{-1} \int S(t)I(t)\,dt + \lambda_\beta)\}$$

- Thus $\beta \mid \ldots$

$$\sim \text{Gamma}(m_\beta + n - 1, \lambda_\beta + N^{-1} \int S(t)I(t)\,dt)$$

3. General epidemic model

- So to update β we just need to sample from

  Gamma ($m_\beta + n - 1, \lambda_\beta + N^{-1} \int S(t)I(t)\,dt$)
- Only problem is: how to evaluate the integral? (We will address this shortly…)
- Sampling directly from the full conditional distribution is often called a “Gibbs step”

3. General epidemic model

- Going through the same steps for γ we find the full conditional distribution

  $\gamma \mid \ldots \sim \text{Gamma}\left(m_\gamma + n, \lambda_\gamma + \sum (r_j - i_j)\right)$

  and so we can update γ using this distribution.

3. General epidemic model

- For the infection times, the full conditional distribution turns out to be non-standard.
- We therefore require a Metropolis-Hastings step. One option is to update infection times one-at-a-time, either in order or at random.
- For instance, we might update 10% of the infection times in between each update of β and γ.

3. General epidemic model

- So for infection time i_k , we propose a new time using proposal density $q(i_k^* \mid i_k)$
- Accept with probability min(p,1) where

$$p = \frac{\pi(\mathbf{i}^*, \mathbf{r} \mid \beta, \gamma, i_{a^*}, a^*)\, q(i_k \mid i_k^*)}{\pi(\mathbf{i}, \mathbf{r} \mid \beta, \gamma, i_a, a)\, q(i_k^* \mid i_k)}$$

- Note that if $i_k^* < i_a$ then $a^*=k$, otherwise a is unchanged.
- Note also we need to be able to evaluate

$$\pi(\mathbf{i}, \mathbf{r} \mid \beta, \gamma, i_a, a)$$

3. General epidemic model

- Choices for $q(i_k^* \mid i_k)$ could include

(i) Propose $i_k^* = r_k - \text{Exp}(\gamma)$

(ii) Propose $i_k^* = r_k - \text{Exp}(\mu)$, where μ is fixed throughout

(iii) Propose $i_k^* \sim N(i_k, \sigma^2)$

3. General epidemic model

Evaluating the likelihood

$$\pi(\mathbf{i}, \mathbf{r} \mid \beta, \gamma, i_a, a) = \prod_{j \neq a} \beta N^{-1} I(i_j-) \times \exp(-\beta N^{-1} \int S(t) I(t)\, dt) \times \gamma^n \exp\{-\gamma \sum (r_j - i_j)\}$$

- The parts that are not straightforward to evaluate are the product term and the integral term.

3. General epidemic model

- Product term:

$$\prod_{j \neq a} \beta N^{-1} I(i_j-) = (\beta N^{-1})^{n-1} \prod_{j \neq a} I(i_j-)$$

- No easy way to simplify further
- Need to write code to evaluate number of infectives at each infection event

3. General epidemic model

Integral term

- $\int S(t)\, I(t)\, dt = \Sigma_{1 \le k \le n} \Sigma_{1 \le j \le N} [(r_k \wedge i_j) - (i_k \wedge i_j)]$

Here, "$a \wedge b$" denotes "minimum of a,b".

Also $i_j = \infty$ for $j > n$, i.e. for those individuals never infected.

Formula is easy to code.

3. General epidemic model

Integral term

Explanation:

$$S(t) = \sum_{1 \le j \le N} 1_{\{j \text{ is susceptible at time } t\}}$$

$$= \sum_{1 \le j \le N} 1_{\{i_j > t\}}$$

where $1_A = 1$ if event A occurs

$= 0$ otherwise

3. General epidemic model

Similarly,

$$I(t) = \sum_{1 \le k \le n} 1_{\{i_k < t < r_k\}}$$

$$\int S(t)I(t)\,dt = \int \sum_{1 \le j \le N} \sum_{1 \le k \le n} 1_{\{i_k < t < r_k\}}\, 1_{\{i_j > t\}}\,dt$$

$$= \sum_{1 \le j \le N} \sum_{1 \le k \le n} \int 1_{\{i_k < t < r_k \text{ and } i_j > t\}}\,dt$$

Total time that j susceptible, k infective

$$= (r_k \wedge i_j) - (i_k \wedge i_j)$$

3. General epidemic model

Summary: MCMC algorithm updates

- Update β and γ using their full conditional distributions (Gamma, in this case)
- Update infection times using a Metropolis-Hastings step. We might typically update 10% of the infection times in between each update of β and γ.

3. General epidemic model

Summary: MCMC algorithm

Initialise β, γ, $\mathbf{i}$, i_a, a

Loop:

- Update β
- Update γ
- Update some of the infection times
- Record current values of β, γ

3. General epidemic model

Summary: MCMC algorithm

The output is a sequence

$$(\beta_1,\gamma_1),\ (\beta_2,\gamma_2),\ \ldots,\ (\beta_M,\gamma_M)$$

where M is number of iterations in loop.

Note that we don't usually record **i,** i_a and a since they are not our main focus, and also **i** is typically high-dimensional so costly to store.