MCMC 2: Lecture 3
SIR models - more topics

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

Contents

- 1. What can be estimated?
- 2. Reparameterisation
- 3. Marginalisation

Contents

- 1. What can be estimated?
- 2. Reparameterisation
- 3. Marginalisation

1. What can be estimated?

As with any attempt to fit a model to data, it is always important to think about how informative the data are about the model parameters which one is trying to estimate.

In some settings it is obvious what can or cannot be estimated; in other settings it can be much less obvious.

1. What can be estimated?

Example: Latent periods

Consider the Markov SIR model with latent periods, say of fixed unknown length = c.

Thus when an individual is infected, they must wait c days until they become infectious.

Latent individuals are called “exposed”.

1. What can be estimated?

Example: Latent periods

Thus we have an SEIR model where
Exposed period = c time units.

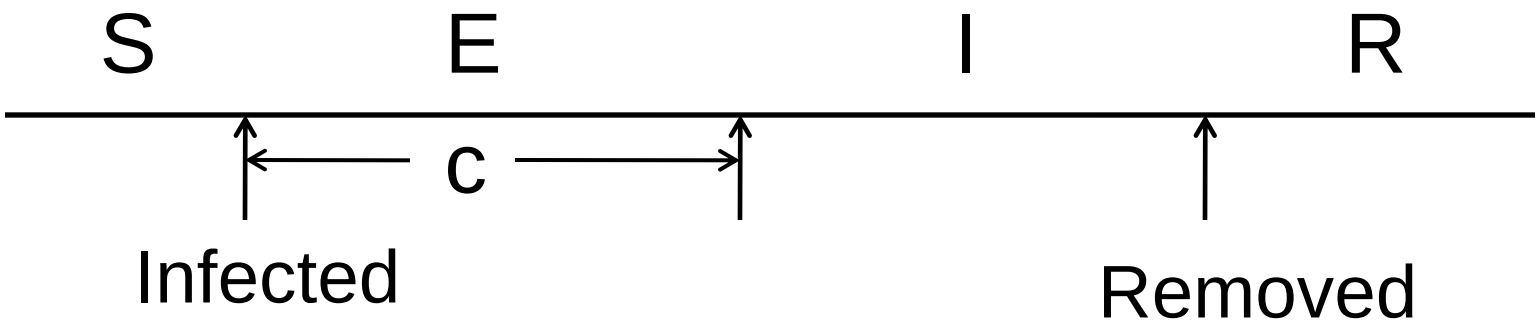

1. What can be estimated?

Example: Latent periods

Introduce "exposure" times (= infection times)

$e_1, e_2, \ldots, e_n$, where $e_k = i_k - c$

and define $\mathbf{e} = (e_2, \ldots, e_n)$.

As before, $\mathbf{r} = (r_1, r_2, \ldots, r_n)$ is observed.

Define $\mathbf{i} = (i_1, i_2, \ldots, i_n)$.

1. What can be estimated?

Example: Latent periods

π (**i**, **r**, **e** | β , γ, c, e_1) =

$$\left(\prod_{j=2}^{n}(\beta/N)S_{ij}I_{ij}\right)\left(\prod_{j=1}^{n}\gamma I_{rj}\right)\exp\left(-\int_{i_1}^{r_n}\{(\beta/N)S_tI_t+\gamma I_t\}dt\right)$$

$$\times 1_{\{i_k-e_k=c\ \ ,\,k=1,2,\ldots,n\}}$$

1. What can be estimated?

Example: Latent periods

It would be straightforward to adapt the standard MCMC algorithm to include c as an extra parameter - e.g. using M-H updates for c.

However such an algorithm would be uninformative about c given removal data alone.

1. What can be estimated?

Example: Latent periods

Roughly speaking, for any value of c, the infection rate β would be estimated accordingly - large c means large β and small c means small β.

So although the MCMC algorithm is correct, the output would need to be carefully interpreted. Here we would see high posterior correlation between c and β.

1. What can be estimated?

Example: Latent periods

In practice, a better strategy would be to fix c to certain (biologically reasonable) values and then perform estimation for β and γ.

1. What can be estimated?

Example: Gamma infectious periods

A common generalisation of the Markov SIR model is to have Gamma-distributed infectious periods.

Thus each infective remains so for a period of time T_I, where

$$T_I \sim \Gamma(c,d), \text{ say } (c = \text{shape}, d = \text{rate}).$$

Note $E(T_I) = c / d$.

1. What can be estimated?

Example: Gamma infectious periods

As seen in lecture 2, the likelihood is

π (i, r | β ,c , d, a, i_a) =

$$\left(\prod_{j\neq a}(\beta/N)\mathrm{I}_{ij}\right)\exp\left(-\int_{i_a}^{r_b}\{(\beta/N)\mathrm{S}_t\mathrm{I}_t\}\mathrm{d}t\right)\prod_{j=1}^{n} f(r_j - i_j \mid c,d)$$

where $f(x \mid c,d) = x^{c-1}d^c \exp(-dx) / \Gamma(c)$

is the p.d.f. of T_I .

1. What can be estimated?

Example: Gamma infectious periods

The two parameters c,d can be updated in an MCMC algorithm.

It is not immediately obvious if it is possible to estimate both parameters separately from removal data.

One might expect $E(T_I) = c / d$ to be estimated with reasonable precision.

1. What can be estimated?

Example: Data for Markov SIR model

Another important aspect of estimation is the detail of the data.

For example, suppose we have observations (= removal times) in a population of N=100 susceptibles, of whom n become infected.

Clearly if n=0, no inference can be drawn.

But what if n=1? n=10? n=100?

Contents

- 1. What can be estimated?
- 2. Reparameterisation
- 3. Marginalisation

2. Reparameterisation

Reminder: MCMC theory and practice

A common problem with MCMC algorithm mixing is correlation, i.e. when two (or more) parameters in the target density are highly correlated.

This problem usually means it is hard to update correlated parameters separately.

2. Reparameterisation

Example: Gamma infectious periods

As described above, consider the SIR model with $\Gamma(c,d)$ infectious periods.

Since the mean infectious period is c/d, c and d will be positively correlated.

It therefore makes sense to consider a reparameterisation to (m,v), where

m = mean, v = variance of T_I .

2. Reparameterisation

Example: Gamma infectious periods

Note that this makes the corresponding part of the likelihood more complex:

$$\prod_{j=1}^{n} f(r_j - i_j \mid c, d) = \prod_{j=1}^{n} (r_j - i_j)^{c-1} d^c \exp(-d(r_j - i_j)) / \Gamma(c)$$

$$\prod_{j=1}^{n} f(r_j - i_j \mid m, v)$$

$$= \prod_{j=1}^{n} (r_j - i_j)^{((m^2/v)-1)} (m/v)^{m^2/v} \exp(-(m/v)(r_j - i_j)) / \Gamma(m^2/v)$$

2. Reparameterisation

Example: Gamma infectious periods

In practice it is often the case that there is a trade-off between “elegance” and “efficiency” of MCMC algorithms.

2. Reparameterisation

Example: Non-centering in Markov SIR model

For the Markov SIR model with n (= number infected) large, it can be hard to move quickly around the space of all possible infection times.

A key difficulty is that the infectious period rate parameter, γ, affects the likelihood of proposed new infection times.

2. Reparameterisation

Example: Non-centering in Markov SIR model

e.g. if γ is large, so the mean infectious period is small, proposed updates that give large infectious periods will often be rejected.

One way around this is to try and update infections times and γ simultaneously (i.e. “block updating”).

An alternative is to reparameterise to break the correlation: this is “non-centering”.

2. Reparameterisation

Example: Non-centering in Markov SIR model

Aside: the term “non-centering” arises via the following graphical model representation:

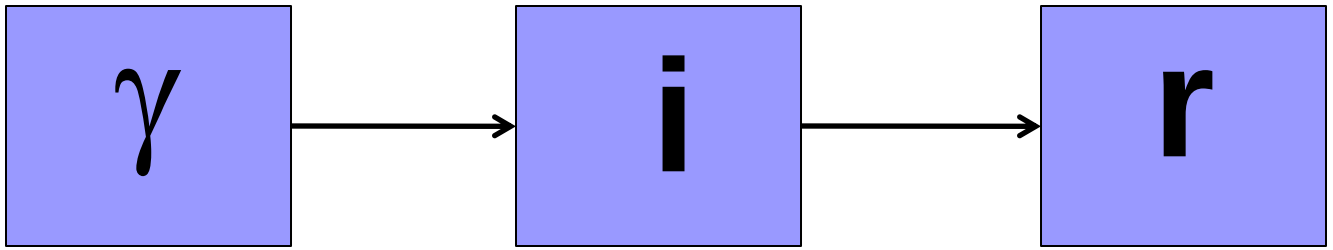

Here, (γ, **i**) is a centered parameterisation.

2. Reparameterisation

Example: Non-centering in Markov SIR model

If the data **r** are relatively uninformative about **i**, then the dependence between γ and **i** can give poor MCMC mixing.

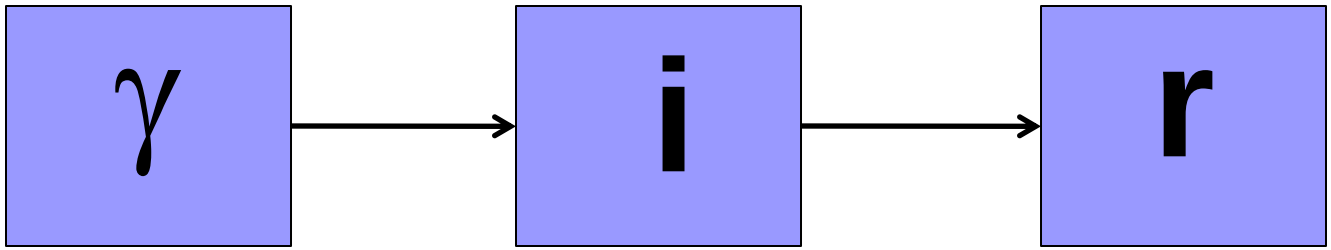

2. Reparameterisation

Example: Non-centering in Markov SIR model

A natural alternative is to find a non-centered parameterisation **u** such that γ and **u** are independent:

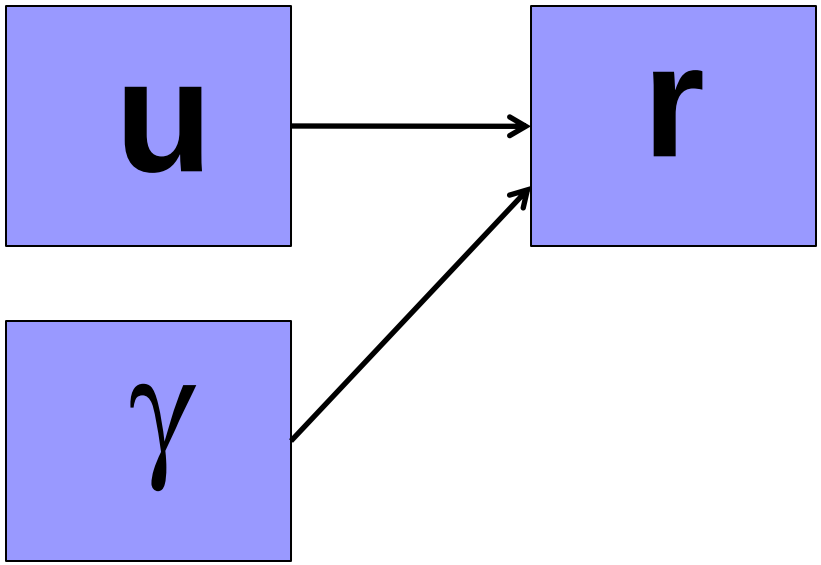

2. Reparameterisation

Example: Non-centering in Markov SIR model

Specifically: observe that if

$$X \sim \text{Exp}(a)$$

then we can write X as

$$X = (1/a)\, Y$$

where

$$Y \sim \text{Exp}(1).$$

2. Reparameterisation

Example: Non-centering in Markov SIR model

Proof:

$$P((1/a)Y > t) = P(Y > at)$$
$$= \exp(-at)$$
$$= P(X > t) \text{ since } X \sim \text{Exp}(a).$$

2. Reparameterisation

Example: Non-centering in Markov SIR model

Now in the model, the infectious period of individual k $\sim$ Exp(γ).

As for the non-Markov case, adopt the labelling where i_k corresponds to r_k .

Thus $r_k - i_k \sim \text{Exp}(\gamma) = (1/\gamma)\text{Exp}(1)$

$= (1/\gamma)\, U_k$, say.

2. Reparameterisation

Example: Non-centering in Markov SIR model

This suggests a reparameterisation in which the infection times $i_1, i_2, \ldots, i_n$ are replaced with $U_1, U_2, \ldots, U_n$.

Note that i_k can be recovered from U_k and γ.

$r_k - i_k = (1/\gamma)\, U_k$, so $i_k = r_k - (1/\gamma)\, U_k$.

2. Reparameterisation

Example: Non-centering in Markov SIR model

Note now that if γ is updated, then all of the infection times are simultaneously updated.

This allows for faster mixing around the space of possible infection times.

2. Reparameterisation

Example: Non-centering in Markov SIR model

Consider now the likelihood in the new parameterisation.

We are making the change of variable

$$u_k = \gamma (r_k - i_k), \qquad k=1, ..., n$$

2. Reparameterisation

Example: Non-centering in Markov SIR model

To evaluate the likelihood we simply write i_k as a function of u_k in the current likelihood, and multiply by the Jacobian of the transformation $J = \left| \frac{\partial(i_1, i_2, \ldots, i_n)}{\partial(u_1, u_2, \ldots, u_n)} \right| = \gamma^{-n}$.

2. Reparameterisation

Example: Non-centering in Markov SIR model

When we do this, then the "infection" part is essentially unchanged:

$$\left(\prod_{j\neq a}(\beta/N)\mathrm{I}_{ij}\right)\exp\left(-\int_{i_a}^{r_b}\{(\beta/N)\mathrm{S}_t\mathrm{I}_t\}\mathrm{d}t\right)\prod_{j=1}^{n} f(r_j - i_j \mid c,d)$$

Unaffected by reparameterisation

2. Reparameterisation

Example: Non-centering in Markov SIR model

Conversely, the remaining product term simplifies:

$$\left(\prod_{j\neq a}(\beta/N)\mathrm{I}_{ij}\right)\exp\left(-\int_{i_a}^{r_b}\{(\beta/N)\mathrm{S}_t\mathrm{I}_t\}\mathrm{d}t\right)\underbrace{\prod_{j=1}^{n} f(r_j - i_j \mid c,d)}$$

2. Reparameterisation

Example: Non-centering in Markov SIR model

Specifically,

$$f(r_k - i_k) = \gamma \exp(-\gamma (r_k - i_k))$$
$$= \gamma \exp(-\gamma (r_k - (r_k - \gamma^{-1} u_k)))$$
$$= \gamma \exp(-u_k),$$

and including the Jacobian means that the whole product becomes $\exp(-\Sigma_{1 \le j \le n} u_j)$.

2. Reparameterisation

Example: Non-centering in Markov SIR model

MCMC algorithm:

- Update β using Gamma full conditional distribution as before
- Update γ using Metropolis-Hastings step (e.g. Gaussian random walk). Note that when γ is updated, so are infection times
- Update u_k's using M-H step (e.g. propose u_k from Exp(1) distribution)

Contents

- 1. What can be estimated?
- 2. Reparameterisation
- 3. Marginalisation

3. Marginalisation

Suppose we have an MCMC algorithm with target density

$$\pi(\theta_1, \theta_2, ..., \theta_n).$$

The idea of marginalisation is to instead run MCMC on a marginal target density, e.g.

$$\pi(\theta_1, \theta_2).$$

This is achieved by integrating out the other parameters.

3. Marginalisation

The obvious reason to do this is that the resulting MCMC chain has less parameters to update.

Note that samples from the integrated-out parameters can be recovered from the output of the reduced chain.

3. Marginalisation

Example: Markov SIR model

Recall that π (**i**, **r** | β , γ, i_1) =

$$\left(\prod_{j=2}^{n}(\beta / N)S_{ij}I_{ij}\right)\left(\prod_{j=1}^{n}\gamma I_{rj}\right)\exp\left(-\int_{i_1}^{r_n}\{(\beta / N)S_tI_t+\gamma I_t\}dt\right)$$

The standard MCMC algorithm updates β , γ, and the infection times **i**.

3. Marginalisation

Example: Markov SIR model

However, it is possible to integrate out β and γ, as follows:

$$\pi(\mathbf{i}, \mathbf{r} \mid i_1) = \iint \pi(\mathbf{i}, \mathbf{r} \mid \beta, \gamma, i_1)\, \pi(\beta, \gamma)\, d\beta\, d\gamma,$$

where $\pi(\beta, \gamma)$ is the joint prior density of (β, γ).

Standard choice is $\pi(\beta, \gamma) = \pi(\beta)\, \pi(\gamma)$.

3. Marginalisation

Example: Markov SIR model

Specifically, suppose $\beta \sim \Gamma(m, \lambda)$ a priori.

$$\int_{(0,\infty)} \pi(\mathbf{i}, \mathbf{r} \mid \beta, \gamma, i_1)\, \pi(\beta)\, d\beta$$

$$\propto \int_{(0,\infty)} \beta^{(n-1)+m-1} \exp(-\beta A)\, d\beta,$$

where $A = \lambda + \int S_t I_t / N \, dt$.

3. Marginalisation

Example: Markov SIR model

However, recalling that

$$\int_{(0,\infty)} y^{n-1} \exp(-y\,b)\, dy = \Gamma(n)\, b^{-n} ,$$

$$\int_{(0,\infty)} \beta^{(n-1)+m-1} \exp(-\beta A)\, d\beta$$

$$= \Gamma(n+m-1)\, A^{-(n+m-1)}$$

$$\propto A^{-(n+m-1)}$$

3. Marginalisation

Example: Markov SIR model

The parameter γ can be integrated out in a similar fashion.

We can thus run an MCMC algorithm on the target density $\pi(\mathbf{i}, \mathbf{r} \mid i_1)$.

Note that only **i** updates are required.

3. Marginalisation

Example: Markov SIR model

To recover samples from the marginal posterior distribution of β, simply note that

$$\pi(\beta \mid \mathbf{i}, \mathbf{r}, \gamma, i_1) \sim \Gamma(m+n-1, A).$$

Each sample for **i** yields a sample for A, from which a sample from β can be obtained.

References

- Neal, P. and Roberts, G.O. (2005) A case study in non-centering for data augmentation: stochastic epidemics. *Statistics and Computing* 15, 315-327.