Bayesian model discrimination for partially-observed epidemic models

James N. Walker[a,b,*], Andrew J. Black[a,b], Joshua V. Ross[a,b]

[a] *School of Mathematical Sciences, University of Adelaide, Adelaide, SA 5005, Australia*
[b] *ACEMS, School of Mathematical Sciences, University of Adelaide, Adelaide, SA 5005, Australia*

Abstract

An efficient method for Bayesian model selection is presented for a broad class of continuous-time Markov chain models and is subsequently applied to two important problems in epidemiology. The first problem is to identify the shape of the infectious period distribution; the second problem is to determine whether individuals display symptoms before, at the same time, or after they become infectious. In both cases we show that the correct model can be identified, in the majority of cases, from symptom onset data generated from multiple outbreaks in small populations. The method works by evaluating the likelihood using a particle filter that incorporates a novel importance sampling algorithm designed for partially-observed continuous-time Markov chains. This is combined with another importance sampling method to unbiasedly estimate the model evidence to a given precision. Our method is general and can be applied to a wide range of model selection problems in biological and epidemiological systems with intractable likelihood functions.

Keywords: `Model selection, Epidemic modelling, Importance sampling, Markov chain, Particle filter`

1. Introduction

Biological and epidemiological processes can have highly complex behaviour and in many cases we do not observe all of their important features. As our observations

*Corresponding author
Email address: `james.n.walker@adelaide.edu.au` (James N. Walker)

may depend on features we do not see, it can be difficult to infer properties of this partially-observed process. To add complexity to the problem, there may be several theories about the behaviour of the process and it may not be obvious as to which of these models is most appropriate given some set of observations. For example in epidemiology, it it widely recognised that the shape of the infectious period distribution is critical to understanding historical disease incidence, and also for accurate evaluation of control measures for public health use [1, 2]. Similarly, it has been identified that the relative timing of symptoms and infectiousness is a key determinant of ability to control an outbreak [3]. Investigating these issues is most easily done by encoding certain features in the structure of the model itself rather than just parametrising it, hence to discriminate different possibilities we need to perform some type of model selection.

Bayesian model selection compares models using the probability that each model generated the data given some observations and prior distribution over models and parameters [4]. We consider the model to be inferred as a parameter (typically with a uniform prior) and hence multiplying the model evidences by the prior distribution and normalising gives the probability of each model given the data [5]. While the interpretability and the consistency with the Bayesian paradigm is desirable, Bayesian model selection has some issues: a large amount of data may be required before models can be distinguished; and the evidence is typically difficult to calculate as the likelihood is intractable. This is true for epidemic models, which for small populations are most naturally represented as partially-observed continuous-time Markov chains (CTMCs) [6, 7, 8].

In this paper we describe an efficient way of performing Bayesian model selection for partially-observed CTMCs and apply our method to the two problems in mathematical epidemiology discussed above: infering the shape of the infectious period distribution and identifying the onset of symptoms relative to infectiousness. This method works by calculating unbiased estimates of the likelihood via sequential importance resampling, which in turn is used in another importance sampling algorithm to estimate posterior model probabilities. A novel feature of this method is that the likelihood is estimated via a scheme which is ideally suited for partially-observed CTMCs, where

one component of the state is observed exactly [8]. This works by sampling realisations of the partially-observed process in a way that realisations always match with observations. Our combined method is both computationally efficient and embarrassingly parallelisable and hence well suited to implementation on modern computing hardware. Further, the ability to estimate error bounds and use stopping criterion ensures the accuracy of evidence estimates. Similar CTMC models see wide use in areas of biology such as phylogenetics [9], ecology [10] and cell biology [11], and hence our method may be applied to a broad range of biological model selection problems. Matlab code is provided for all algorithms developed in this paper as part of the EpiStruct project [12].

We first discuss the methods for importance sampling over the parameter space, and the sequential importance resampling algorithm used to estimate the likelihood in Section 2. We then describe the two case studies, their implementation and results in Sections 3 and 4. Lastly we conclude in Section 5 and discuss how our method relates to those already in the literature.

2. Methods

Importance sampling is a method for estimating properties of one distribution by sampling from another distribution; the bias is corrected to obtain unbiased estimates related to the distribution of interest [13]. This technique is useful when the target distribution is difficult to sample from, such as the distribution of latent variables from a partially-observed Markov chain. This method has been effectively applied for parameter inference, for example in particle marginal MCMC schemes [14, 15] and has also been applied to model selection [4, 16, 17, 18]. In this paper we use importance sampling in the space of latent variables to estimate likelihoods via sequential importance resampling and within the parameter space to calculate the model evidence. The novelty of this approach is that the likelihood is estimated via the most suitable kind of importance sampling scheme for these kinds of models, as described in [8].

Here we introduce importance sampling for model selection, detail how likelihood estimates are used in model selection and describe how sequential importance resam-

pling is used for likelihood estimation.

2.1. Importance sampling for evidence estimation

Importance sampling for model selection has been discussed in [4, 16, 17], and in this paper we adopt a similar approach. However, our implementation uses a particle filter to obtain unbiased estimates of the likelihood, which allows us to obtain unbiased estimates of the evidence.

Suppose we have data set y of observations from a model with parameter set θ in parameter space Θ. We may sample n parameter sets $\theta_1, \ldots, \theta_n$ from some arbitrary, importance sampling, density, $q(\cdot|y)$, with support Ψ where $\Theta \subseteq \Psi$. We then compute random variables

$$Z_i = \frac{p(y|\theta_i)p(\theta_i)}{q(\theta_i|y)}, \tag{1}$$

for $i = 1, \ldots n$, where $p(\theta)$ is the prior distribution and $p(y|\theta)$ is the likelihood function. These random variables are chosen so that their expectation is given by

$$\begin{aligned} E[Z_i] &= \int_\Psi \frac{p(y|\theta)p(\theta)}{q(\theta|y)} q(\theta|y) \ \ d\theta \\ &= \int_\Theta p(y|\theta)p(\theta) \ \ d\theta \\ &= p(y). \end{aligned}$$

Hence, the mean of the Z_i's provides an unbiased estimate of the evidence, $\hat{p}(y)$. If the likelihood is intractable we can use unbiased estimates, $\hat{p}(y|\theta)$, instead of $p(y|\theta)$ in Equation (1), by the law of total expectation. Note that using importance sampling to estimate $p(y)$ involves obtaining independent identically distributed samples of Z_i and taking their mean. Therefore the central limit theorem can be applied once a large number of samples are obtained, so we can estimate error bounds on $\hat{p}(y)$ which shrink like $1/\sqrt{n}$. Hence, one can sample until estimates of the evidence are within a specified tolerance. Further, the independence of the sampling procedure allows these computations to be run in parallel and so modern computation hardware can be easily utilised.

Of course, having to estimate the likelihood inflates the variance of $\hat{p}(y)$, so a larger number of samples are required for $\hat{p}(y)$ to converge to a given accuracy.

The variance of this estimate is lowest if the sampling density, $q(\cdot|y)$, is similar to the posterior distribution. Our method requires no parameter inference, but may be made more efficient, especially in higher dimensions, by choosing a sampling distribution similar to the posterior distribution if parameter inference has been performed; for example, the sampling distribution could be a Gaussian with the mean and covariance matrix estimated from MCMC samples [17]. In this paper we do not use MCMC samples to inform the choice of $q(\cdot|y)$, but this could be done in practice to improve the convergence of estimates.

2.2. Sequential importance resampling algorithm for likelihood estimation

Consider a time series, $y = (y_1, \ldots, y_T)$, from a partially-observed CTMC. We can express the likelihood at θ as

$$p(y|\theta) = \prod_{t=1}^{T} p(y_t|y_{1:t-1}, \theta),$$

where we adopt the convention that $p(y_1|y_{1:0}, \theta)$ is taken to mean $p(y_1|\theta)$. Hence, the likelihood can be calculated via a product of the likelihood increments, $p(y_t|y_{1:t-1})$, for $t = 1, \ldots, T$. We apply a sequential importance resampling algorithm, which uses importance sampling to estimate the likelihood increments [19, 20]. This algorithm uses resampling at the end of each time increment to lower the variance of likelihood estimates, and remove realisations that were unlikely to have occurred. Throughout this section all probabilities are calculated with respect to some parameter set, θ, but this notation is suppressed to make statements concise.

Let x_t denote the state of the partially-observed process at time t. The sequential importance resampling algorithm begins with a set of n particles, each particle is associated with an initial state and a weight, $\{x_0, w\}^{1:n}$. These initial states are distributed according to $p(x_0)$ and the weights begin as 1. Suppose at each iteration of the algorithm we have states $x_{t-1}^{1:n}$ distributed according to $p(x_{t-1}|y_{1:t-1})$, we evolve these over a day according to the importance sampling density, $q(\cdot|y_t, x_{t-1})$, to obtain realisations $\tilde{x}_{(t-1,t]}^{1:n}$ with weights given by

$$w^i = \frac{p\left(\tilde{x}^i_{(t-1,t]}, y_t | x^i_{t-1}\right)}{q\left(\tilde{x}^i_{(t-1,t]} | y_t, x^i_{t-1}\right)}, \qquad \text{for } i = 1, \ldots, n.$$

These are computed such that the expected value of the weights are

$$\begin{aligned} E[w] &= \int\int \frac{p\left(\tilde{x}_{(t-1,t]}, y_t | x_{t-1}\right)}{q\left(\tilde{x}_{(t-1,t]} | y_t, x_{t-1}\right)} q\left(\tilde{x}_{(t-1,t]} | y_t, x_{t-1}\right) \\ &\quad \times p(x_{t-1} | y_{1:t-1}) \; d\tilde{x}_{(t-1,t]} \; dx_{t-1} \\ &= p(y_t | y_{1:t-1}). \end{aligned}$$

Hence, the mean of these weights provides an unbiased estimate of the likelihood increment, $p(y_t|y_{1:t-1})$. Further, resampling states from $\tilde{x}^{1:n}_t$ according to normalised weights, $w_i / \sum_j w_j$, gives a sample, $x^{1:n}_t$, from $p(x_t|y_{1:t})$. Hence, from the initial conditions we can iterate forwards and recursively estimate the likelihood increments.

Thus far we have described the general approach of using importance sampling and sequential importance resampling to estimate posterior model probabilities, however, an important aspect of the sequential importance sampling process is the evolution of particles over each day. The difficulty in evolving particles in the cases considered here is that precise observations are made from models with many latent variables, which can make data-augmentation or rejection-sampling approaches slow. In sections 3.2 and 4.3 we describe the method from [8], which builds upon [21], for generating realisations $\tilde{x}_{(t-1,t]}$ from initial state x_{t-1} for each of the case studies. The main benefit of this approach is that we use an efficient sampling distribution $q(\cdot|y_t, x_{t-1})$ which is tractable and generates $\tilde{x}_{(t-1,t]}$ which match observations almost surely, and, are similar to the partially-observed process. This method makes particles match observations by first generating observation times, then generating events between observations and occasionally forcing events to occur or blocking events to ensure feasibility of the process. For example, if we are yet to observe an infection but a recovery would lead to the epidemic ending, the recovery rate is set to 0. Importance sampling in this way allows for estimation of likelihoods associated with rare events – for example it can be easily used to estimate the tails of the likelihood, or just used to sample when observations are

unlikely – whereas rejection sampling methods tend to perform poorly as simulations are unlikely to match the data.

The overall approach to model selection here is to: (i) sample from the parameter space using some importance sampling distribution; (ii) for each of those samples estimate the likelihood by sequential importance sampling; (iii) plug this likelihood estimate into Equation (1) to obtain importance weights; and, (iv) use the mean of these weights as an unbiased estimate of the model evidence. We can keep sampling and obtaining estimates of the model evidence until they satisfy some stopping criteria; for example, once credible intervals are of a specified width. Once the model evidence is computed for all candidate models we can multiply these by the prior distribution over the models and normalise to obtain the posterior model probabilities.

3. Case study I: Inferring the infectious period distribution

Our first case study uses the importance sampling Bayesian model selection method to infer the appropriate infectious period distribution for an SI(k)R model [22, 23, 24], where symptoms are assumed to coincide with a transition into the infectious class. This study is motivated by influenza outbreaks in which only symptom onset (which is correlated with infection) is observed, and recoveries are not. Here we consider the infectious period to be either exponential, Erlang-2 or Erlang-5 distributed; these represent high, medium and low variance infectious periods respectively. This study aims to answer whether case data at a daily resolution is sufficient for discriminating between these models, how well parameters need to be known in order to discriminate between models, and how much data is required to effectively discriminate between models

The data considered in this paper consists of daily symptom onset counts from completed outbreaks in small populations, which we refer to as households. We let all households be of size 4 for simplicity, though this can easily be extended to allow for a distribution of household sizes. Each outbreak is modelled as a compartmental CTMC, however we only observe a small portion of the epidemic process, so this is a partially-observed CTMC. We assume that all events of symptom onset are observed

until the epidemic fades out in the household and set the time of the first observation within each household as 0.

We describe the epidemic model for households in Section 3.1, the importance sampling method for estimating likelihood increments in Section 3.2, the parameters used for the simulation study in Section 3.3, and show results in Section 3.4.

3.1. SI(k)R model

For a population of size N, the SI(k)R model is a compartmental model that allows each individual to be in one of $k+2$ compartments: they are either susceptible; infectious in phase j, for $j = 1, \ldots, k$; or, recovered. Here the different phases of infectiousness have no physical interpretation; they are introduced in order to allow the overall infectious period to be Erlang-k distributed [22, 23, 24]. Let S, I_j and R denote the number of susceptible, infectious phase j and recovered individuals respectively. As these numbers must always be non-negative and sum to N, we have the state space

$$\mathscr{S} = \left\{ (S, I_1, \ldots, I_k, R) \in \mathbb{N}^{k+2} : S + \sum_{j=1}^{k} I_j + R = N \right\},$$

where we take $\mathbb{N}$ to contain 0. There are three kinds of transitions for this model: infection; phase change; and, recovery. Infectious individuals of all phases make effective contact with other individuals in the population at rate β, and if this contact is with a susceptible individual then that individual becomes infectious, corresponding to a transition into the infectious phase 1 compartment. An infectious phase j individual for $j = 1, \ldots, k-1$, moves into the next phase at rate $k\gamma$; this rate is chosen such that the infectious period is Erlang-k distributed with mean $1/\gamma$. Similarly, an infectious phase k individual recovers at rate $k\gamma$ and is no longer able to spread the disease. The transitions and rates associated with changes in the number of individuals in each compartment are given in Table 1.

For this study we assume that symptom onset corresponds to the infection transition, and we assume that we observe the number of these transitions at a daily resolution. The model is initialised at the time of the first observations, hence the initial state is $(S, I_1, I_2, \ldots, R) = (N-1, 1, 0, \ldots, 0)$.

Table 1: Transitions and rates for the SI(k)R model. Only compartments that change are shown, all other compartments remain the same.

Transition Type	State Change	Rate
Infection	$(S, I_1) \to (S-1, I_1+1)$	$\frac{\beta S \sum_{j=1}^{k} I_j}{N-1}$
Phase change type j	$(I_j, I_{j+1}) \to (I_j - 1, I_{j+1}+1)$	$k\gamma I_j$
Recovery	$(I_k, R) \to (I_k - 1, R+1)$	$k\gamma I_k$

3.2. Importance sampling for SI(k)R outbreaks

Suppose we have observations from M outbreaks, $y^{1:M}$. As these outbreaks occur independently we have a likelihood function which is of the form

$$p\left(y^{1:M}|\theta\right) = \prod_{i=1}^{M} p\left(y^{i}|\theta\right).$$

Hence, the likelihood can be estimated via multiplying estimates of likelihoods for each outbreak, $p(y^i|\theta)$. The sequential importance resampling scheme allows us to calculate each of these iteratively, so we are left to describe how to generate realisations over the day from an initial state, x_{t-1}, in a way that ensures consistency with the observations, and, how to evaluate the importance sampling weights. We do this as per the method of [8].

Suppose for each outbreak we have a dataset $y = (y_1, \ldots, y_T)$ where y_t gives the cumulative number of infection events over $(t-1, t]$ for $t = 1, \ldots, T$; note that this does not include the initial infectious individual in the population. To estimate $p(y_t|y_{1:t-1})$, we begin by uniformly generating y_t observation times over $(t-1, t]$. These times are ordered, so the joint density is that of y_t order statistics of Uniform$(t-1, t)$ random variables, so we initialise the importance weights by $w = 1/y_t!$. Then, beginning from time $\tau = t-1$ in state,

$$\tilde{x}_\tau = (S(\tau), E_1(\tau), \ldots, R(\tau)),$$

determined by the final state of the previous iteration, we generate events between observation times. However, we only allow phase change or recovery transitions to occur, as the observations (and hence infections) have already been generated. Let a be a vector of rates,

$$a = \left(\frac{\beta S(\tau) \sum_{j=1}^{k} I_j(\tau)}{N-1}, k\gamma I_1(\tau), \ldots, k\gamma I_k(\tau) \right),$$

as per Table 1, associated with state $\tilde{x}_\tau$. We consider a process with modified rates b, given by

$$\begin{aligned} b_1 &= 0 \\ b_{2:k} &= a_{2:k} \\ b_{k+1} &= 1_{\{\sum_j I_j > 1\}} a_{k+1}; \end{aligned}$$

where $1_{\{A\}}$ is an indicator function which takes the value 1 if logical statement A holds or 0 otherwise. The modified rates are constructed so that no further observations (infections) can occur, as $b_1 = 0$, and no recoveries can occur if it would lead to epidemic fade-out. Let τ' denote either the next observation (or if there are no further observations over the day let $\tau' = t$). We sample an Exponential($\sum_j b_j$) candidate time increment, $\Delta\tau$. Then one of three kinds of updates occur:

(i) if $\tau + \Delta\tau < \tau'$ we generate an event at time $\tau + \Delta\tau$ and let the event be of type i with probability $b_i / \sum_j b_j$, the importance weight is updated to

$$w \leftarrow w \times \frac{a_i e^{-\sum_j a_j \Delta\tau}}{b_i e^{-\sum_j b_j \Delta\tau}},$$

we update the time $\tau \leftarrow \tau + \Delta\tau$, and we update the state to $\tilde{x}_\tau$ according to transition type i;

(ii) if $\tau + \Delta\tau > \tau'$ and $\tau' \neq t$ no event occurs in (τ, τ') and the next event is an observation, so we update the weights to

$$w \leftarrow w \times \frac{a_1 e^{-\sum_j a_j(\tau'-\tau)}}{e^{-\sum_j b_j(\tau'-\tau)}},$$

we update the time, $\tau \leftarrow \tau'$, and update the state $\tilde{x}_\tau$ according to an infection transition; or,

(iii) if $\tau + \Delta\tau > \tau'$ and $\tau' = t$ then no event occurs before the end of the day and the weights are updated to

$$w \leftarrow w \times \frac{e^{-\sum_j a_j(\tau'-\tau)}}{e^{-\sum_j b_j(\tau'-\tau)}}.$$

We repeatedly recalculate the rates and make updates until an update of type (iii) occurs, at which time the process ends for this time increment.

Once all particles have gone through this process the weights are averaged to form an estimate of $p(y_t|y_{1:t-1})$ and the states are resampled, using systematic resampling [25], according to the normalised weights. These samples provide initial states for calculating the next likelihood increment. Once the last observation in the time series has been generated, we generate from a modified process where $b = (0, a_{2:k+1})$ and update weights according to (i) until the epidemic dies out; this allows us to calculate the portion of the likelihood associated with the assumption that the epidemic died out after our last observation.

3.3. Implementation

We simulate 20 data sets of multiple completed outbreaks in households of size 4 with parameters $(\beta, \gamma) = (0.933, 2/3)$ under each of the three models. We use 500 particles per likelihood calculation, as this is found to be sufficient for low variance likelihood estimates, and sample from the parameter space until 95% credible intervals of the model evidence are non-overlapping. The number of particles chosen for the likelihood calculation only effects the speed of convergence of estimates, not their accuracy. We calculate posterior model probabilities using data from 50, 100 and 150

complete household outbreaks. For each of these data sets we consider two cases: where the mean infectious period and the reproduction number are known to a high level of accuracy *a priori*; and, where the prior distribution on model parameters is relatively uninformative. We refer to these as the assumption of tight priors and loose priors, respectively. For the tight priors we suppose that $1/\gamma$ has a gamma distributed prior with mean $1/\gamma$ and variance 0.01, and that β/γ has a Uniform$(\beta/\gamma - 0.03, \beta/\gamma + 0.03)$ prior. For the loose priors we suppose that $1/\gamma$ has a gamma distributed prior with mean 2 and variance 0.75, and we assume that β/γ has a Uniform$(1,2)$ prior; which is a typical range of β/γ for influenza.

3.4. Results

Our results are given in terms of box plots of the difference in posterior model probabilities of the true model and other candidate models in Figures 1 and 2, and in terms of the proportion of times the correct model was identified in Figure 3. If the box plots are near one it means the posterior model probability for the correct model is nearly one; values that are negative represent times at which the other candidate models have higher posterior model probabilities.

For tight priors, most often the correct model was associated with the highest posterior model probabilities for each of the models considered (Figure 3). Note that the lowest box in the panels on Figures 1 and 2 always correspond to an adjacent model. This fits with the intuition that adjacent models are most often misidentified as each other, that is, the models with infectious period distributions which are most similar are more difficult to discriminate. As the bars in Figure 3 are lowest for the SI(2)R model we conclude that the SI(2)R model is most often misidentified; this agrees with intuition as this model's infectious period has a shape parameter between the other two. Results were similar under loose priors, (Figure 2), however the box plots tended to have a larger range, showing that there was less certainty in the correct model. There was also very little change in the proportion of correct times the model was identified under loose and tight priors (Figure 3). Specifically, the proportion of times the correct model was identified under tight priors for data on 50, 100 and 150 outbreaks for the SIR model was 0.75, 0.85 and 0.8 respectively; for the SI(2)R model was 0.55, 0.8

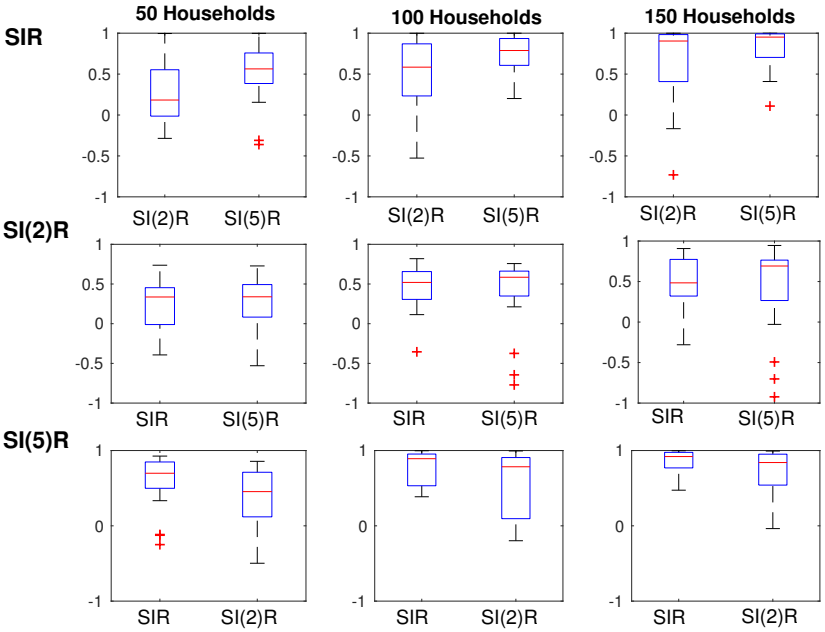

Figure 1: Box plots of the difference in posterior model probability of the true model and the other candidate models for the SI(k)R models with tight priors based on 20 simulated data sets. Rows from top to bottom show results from data sets generated from the SIR, SI(2)R and SI(5)R models. Columns from left to right represent data sets containing 50, 100 and 150 outbreaks.

and 0.7 respectively; and, for the SI(5)R model was 0.8, 0.75 and 0.95 respectively. With loose priors the proportion of times the correct model was identified for data on 50, 100 and 150 outbreaks for the SIR model was 0.75, 0.7 and 0.8 respectively; for the SI(2)R model was 0.55, 0.75 and 0.55 respectively; and, for the SI(5)R model was 0.65, 0.7 and 0.9 respectively. Even if the parameters are not well known, with multiple completed outbreaks (from small populations), infectious case data can be used to distinguish between these models, but the SI(2)R model is the most difficult to identify.

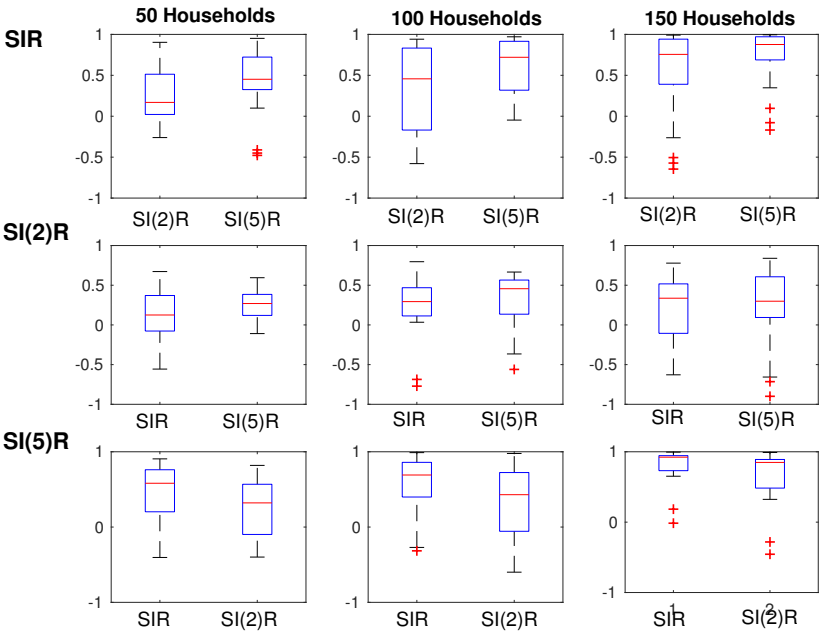

Figure 2: Box plots of the difference between the posterior model probability of the true model and the other candidate models for the SI(k)R model for 20 simulations based on loose priors. Posterior model probabilities are calculated after 50, 100 and 150 households have completed outbreaks.

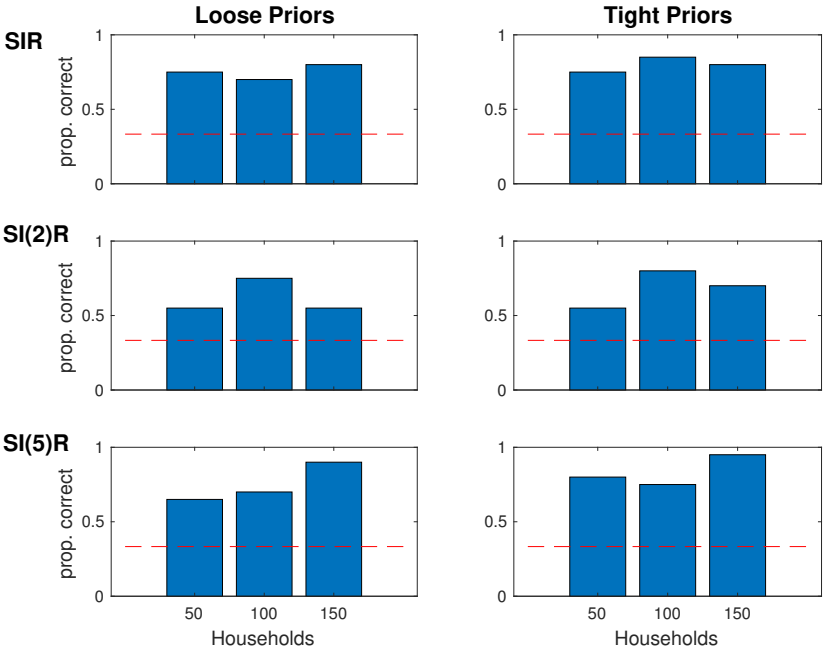

Figure 3: The proportion of times the correct model was identified from 20 data sets generated from each of the SIR, SI(2)R and SI(5)R models for loose (left panels) and tight (right panels) priors for data sets based on 50, 100 and 150 households (bars from left to right). The red dotted line is at 1/3, which represents the proportion of times the correct model would be chosen if models were indistinguishable.

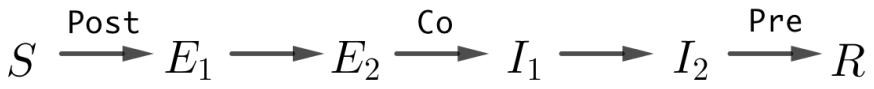

Figure 4: A diagram showing how we relate post-symptomatic (Post), coincidental-symptomatic (Co) and pre-symptomatic (Pre) infections to transitions in the SE(2)I(2)R model.

4. Case study II: Inferring symptom onset relative to infectiousness

This study considers inferring the time of infectiousness relative to symptom onset, where infectiousness could either be pre-symptomatic (Pre), coincidental-symptomatic (Co) or post-symptomatic (Post). For example, influenza may have symptoms that largely coincide with infectiousness [26], SARS is largely infectious post-symptom onset [27] and HIV is infectious pre-symptom onset [3]. We consider diseases that have a lag between exposure of individuals and infectiousness, specifically an SE(2)I(2)R model, which is detailed in Section 4.1. This kind of model has been used in previous work on inference using early outbreak data; it has realistic features, such as non-exponential exposed and infectious periods, while being simple enough for inference [28, 29, 30, 31, 7]. We model the observations of symptom onset as either a transition into an exposed, infectious or recovered state for the Post, Co and Pre models respectively (Figure 4).

4.1. SE(2)I(2)R model

For a population of size N, the SE(2)I(2)R model is a compartmental model that allows each individual to be in one of six compartments: they are either susceptible; exposed phase 1 or 2; infectious phase 1 or 2; or, recovered. The key differences between this model and the SI(k)R model is that the infectious period is assumed to be Erlang-2 distributed and there is an Erlang-2 distributed lag between being exposed to the disease and being able to spread it. Let S, E_i, I_j and R denote the number of susceptible, exposed phase i, infectious phase j and recovered individuals respectively. Again, as these numbers must always be non-negative and sum to N, we have state space

$$\mathscr{S} = \left\{ (S, E_1, E_2, I_1, I_2, R) \in \mathbb{N}^6 : S + E_1 + E_2 + I_1 + I_2 + R = N \right\}.$$

There are five kinds of transitions: exposure; exposed phase change; infectious; infectious phase change; and, recovery. These transitions are similar to those described for the SI(k)R model and are shown in Table 2.

Table 2: Transitions and rates for the SE(2)I(2)R model. Only compartments that change are shown, all other remain the same.

Transition Type	State Change	Rate
(1) Exposure	$(S, E_1) \to (S-1, E_1+1)$	$\frac{\beta S(I_1+I_2)}{N-1}$
(2) Exposed phase change	$(E_1, E_2) \to (E_1-1, E_2+1)$	$2\sigma E_1$
(3) Infectious	$(E_2, I_1) \to (E_2-1, I_1+1)$	$2\sigma E_2$
(4) Infectious phase change	$(I_1, I_2) \to (I_1-1, I_2+1)$	$2\gamma I_1$
(5) Recovery	$(I_2, R) \to (I_2-1, R+1)$	$2\gamma I_2$

We assume that symptom onset corresponds to a transition into either the exposed phase 1, infectious phase 1 or recovered class for the Post, Co and Pre models respectively, as shown in Figure 4. Hence, the process has an initial state that is either

$$(S, E_1, E_2, I_1, I_2, R) = (N-1, 1, 0, 0, 0, 0),$$

for exposure observations,

$$(S, E_1, E_2, I_1, I_2, R) = (N-1, 0, 0, 1, 0, 0),$$

for infectious observations, or a stochastic initial state for the recovery observations; this last case is discussed in the following section.

4.2. Importance sampling for initial state generation

For the SE(2)I(2)R model, if we observe the daily number of recovery transitions, then a population with a single exposed individual initially may be in one of several states at the time of the first observation, as multiple exposures may have occurred before the first recovery. If we see no secondary transmission in the household we can calculate the likelihood exactly as

$$p(y|\beta,\gamma,\sigma) = \left(\frac{2\gamma}{2\gamma+\beta}\right)^2.$$

If the epidemic does not die out after the initial infection we can efficiently generate weighted initial states of the process via importance sampling.

If our data set for the household has daily observation vector $y_{1:T} = (y_1,\ldots,y_T)$ we know that each type of transition can occur at most $\psi = 1+\sum_{t=1}^{T} y_t$ times and that recovery cannot occur if it would lead to epidemic fade out; note that the plus 1 is because the initial condition is not included in $y_{1:T}$. So if the SE(2)I(2)R process has rate vector

$$a = \left(\frac{\beta S(\tau)\,(I_1(\tau)+I_2(\tau))}{N-1}, 2\sigma E_1(\tau), 2\sigma E_2(\tau), 2\gamma I_1(\tau), 2\gamma I_2(\tau)\right),$$

as per Table 2, we consider a modified process with rates that never lead to inconsistencies in our data. The modified rates are

$$b = a \circ \left(1_{\{z_1<\psi\}}, 1_{\{z_2<\psi\}}, 1_{\{z_3<\psi\}}, 1_{\{z_4<\psi\}}, 1_{\{z_4-z_1>1\}}\right);$$

where '$\circ$' denotes an elementwise product, z_i denotes the cumulative number of individuals that entered the ith compartment, and $1_{\{\cdot\}}$ denotes the indicator function. The modified process is used to generate transitions until an observation occurs; the bias from simulating from this modified process is corrected for by using importance sampling weights.

The process begins with particles in state

$$(S, E_1, E_2, I_1, I_2, R) = (N-1, 0, 0, 1, 0, 0),$$

with importance weight $w = 1$. Each particle makes transitions according to probabilities $b/\sum b$, and with each transition the particles weight is updated to

$$w \leftarrow w \times \frac{a_i \sum_j b_j}{b_i \sum_j a_j},$$

where i is the type of the transition that occurred (type is numbered as per Table 2). When the first recovery (transition type 5) occurs the process ends. Once we have a set of particles distributed according to the initial state they are resampled according to the normalised importance weights to give initial particles for the sequential importance resampling algorithm.

4.3. Importance sampling for SE(2)I(2)R outbreaks

The importance sampling scheme for the SE(2)I(2)R model is similar to the SI(k)R model except that the modified rates are slightly different, and, at times, we need to force certain events to occur to ensure feasibility of samples [8]. Again, we uniformly generate observation times over $(t-1, t]$ and from time $t-1$ generate rates between observations according to a modified process with rate vector, b. We set modified rates $b_j = a_j$ for all j that correspond to rates which are not set to 0 (those set to 0 will be specified in this section). For the Post, Co and Pre models we set the modified rates $b_1 = 0$, $b_3 = 0$ and $b_5 = 0$ respectively. We also set b_5 to 0 if a recovery would lead to epidemic fadeout without the correct number of observations occurring. For these models we may also need to force events to occur to ensure that the particle generated, $\tilde{x}_{(t-1,t]}$, is a feasible realisation from an SE(2)I(2)R process. For example, for the Co model, if there are no exposed phase 2 individuals in the population and an observation is yet to occur, an exposed phase change would need to occur. If there were also no exposed phase 1 individuals then an exposure time would need to be generated before the exposed phase change; we refer to these kinds of events, and observation events, as *forced events*. More generally, if the particle is a realisation generated up to time τ, $\hat{x}_{(t-1,\tau]}$, where the next forced event occurs at time τ' and would lead to an infeasible

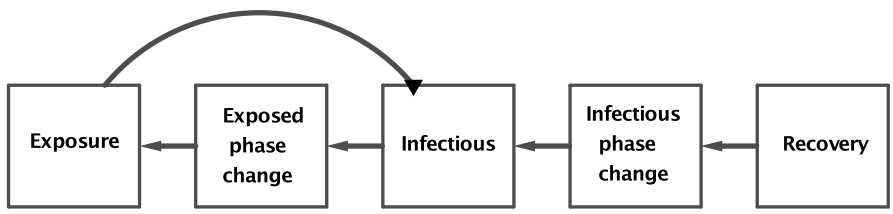

Figure 5: A flow chart showing the order in which transitions are proposed as forced events in order to ensure feasibility of the particle. Starting from the type of the next forced event, propose a new forced events according to arrows in the flow chart.

realisation, we generate a new forced event which would allow the next forced event to be feasible. The next forced event is chosen by proposing transitions further back in the chain, or forcing an infectious event (transition type 3), until an event with a positive rate is proposed. The order in which to propose events is shown via a flowchart in Figure 5. For other CTMC's this can be done by proposing forced events according to a logic tree which is model specific. If we need to generate a forced event of type i, we generate an inter-arrival time, s, according to a $\text{TruncExp}(a_i, \tau' - \tau)$ distribution, set $b_i = 0$, set the next forced event as type i at time $\tau' = \tau + s$, where $\text{TruncExp}(a,t)$ refers to the truncated exponential distribution with rate a, truncated to $[0,t]$. The truncated exponential distribution is chosen as it has appropriate support and generates events with a similar distribution to the true process. The weights are updated according to

$$w \leftarrow w \times \frac{a_i e^{-a_i s}}{1 - e^{-a_i(\tau' - \tau)}}.$$

If no more forced events are necessary we continue to propose candidate events from time τ according to the new modified rates; as in Section 3.1, we let these occur, move to the next forced event, or move to the end of the time increment as per (i), (ii) and (iii) respectively. Once the cumulative number of events of type i equals the number of observations in total we set $b_i = 0$; as the epidemics are completed within households the final state must be $(N - \psi, 0, 0, 0, 0, \psi)$, so no event can occur more than ψ times. For more details and a more general description of the importance sampling scheme see [8].

4.4. Implementation

Again we simulate 20 data sets of multiple completed outbreaks in households of size 4 with parameters $(\beta, \sigma, \gamma) = (0.933, 0.5, 2/3)$ under each of the three models and calculate posterior model probabilities with data from 50, 100 and 150 completed household outbreaks. In this case we used 2000 particles per likelihood calculation; a larger number of particles were chosen because the sampling distribution is less like the true process, due to the possibility of needing to force events. We consider inference based upon tight and loose priors on all model parameters. For the tight priors we suppose that $1/\gamma$ has a gamma distributed prior with mean $1/\gamma$ and variance 0.01, similarly we assume that $1/\sigma$ has a gamma distributed prior with mean $1/\sigma$ and variance 0.01 and we assume that β/γ has a Uniform$(\beta/\gamma - 0.03, \beta/\gamma + 0.03)$ prior. For the loose priors we suppose that $1/\gamma$ and $1/\sigma$ have gamma distributed priors with mean 2 and variance 0.75 and we assume that $\beta/\gamma - 1$ has a gamma distributed prior with mean 1 and variance 0.75.

4.5. Results

Our results are given in terms of box plots of the difference in posterior model probabilities of the true model and other candidate models in Figures 6 and 7, and in terms of the proportion of times the correct model was identified in Figure 8. Each model is easily identifiable whether or not the priors are informative, as shown in Figure 8, and by the time 150 households have had completed outbreaks the correct model is identified in almost all cases. Specifically, the proportion of times the correct model was identified under tight priors for data on 50, 100 and 150 outbreaks for the Post model was 0.9, 1 and 1 respectively; for the Co model was 0.8, 0.85 and 0.95 respectively; and, for the Pre model was 0.75, 0.8 and 1 respectively. With loose priors the proportion of times the correct model was identified for data on 50, 100 and 150 outbreaks for the Post model was 0.95, 1 and 1 respectively; for the Co model was 0.85, 0.8 and 0.9 respectively; and, for the Pre model was 0.7, 0.8 and 1 respectively. Figure 6 shows that even when 50 households are infected most posterior model probabilities of the true model are near 1, and by the time 150 households are infected effectively all posterior model probabilities are close to 1. Figure 7 shows that loose priors reduces the

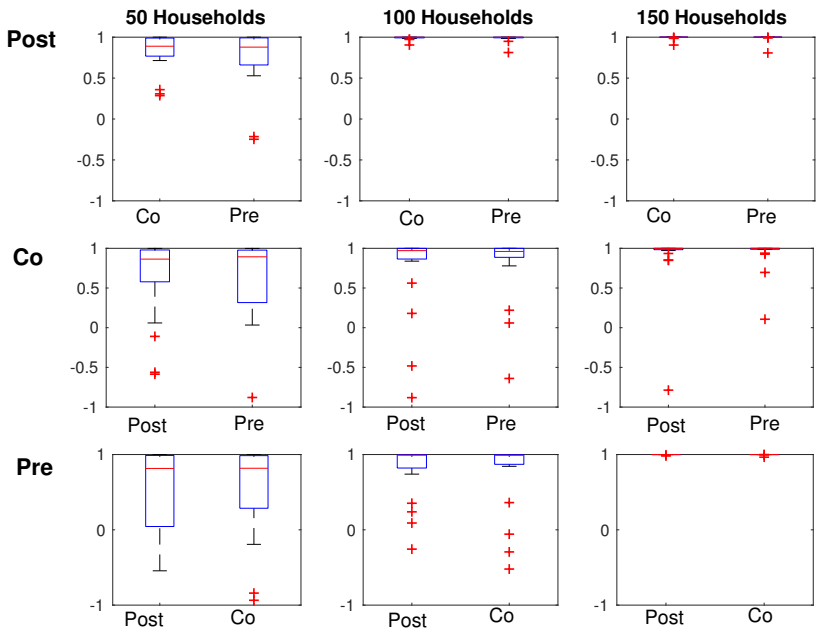

Figure 6: Box plots of the difference in posterior model probability of the true model and the other candidate models for the Post, Co and Pre models with tight priors based on 20 simulated data sets. Rows from top to bottom show results from data sets generated from the Post, Pre and Co models. Columns from left to right represent data sets containing 50, 100 and 150 outbreaks.

certainty of the correct model. However, Figure 8 shows that this has little impact on whether the correct model is chosen. Figure 7 also shows that even with loose priors the posterior model probabilities for the correct model are almost always near 1 for Post and Pre models once 150 households are infected. It also shows that the Co model tended to be chosen with less certainty when priors were loose, with one outlier choosing an incorrect model with posterior model probability near 1. Figure 8 shows that the Pre model is the most difficult model to identify with data on 50 outbreaks. However, every model except the Co model was selected correctly every time with data on 150 outbreaks.

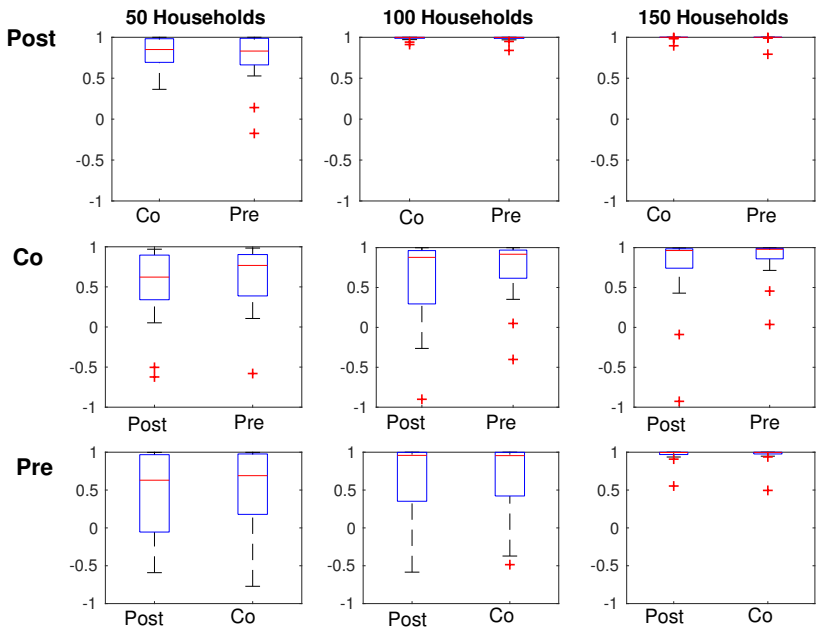

Figure 7: Box plots of the difference in posterior model probability of the true model and the other candidate models for the Post, Co and Pre models with loose priors based on 20 simulated data sets. Rows from top to bottom show results from data sets generated from the Post, Pre and Co models. Columns from left to right represent data sets containing 50, 100 and 150 outbreaks.

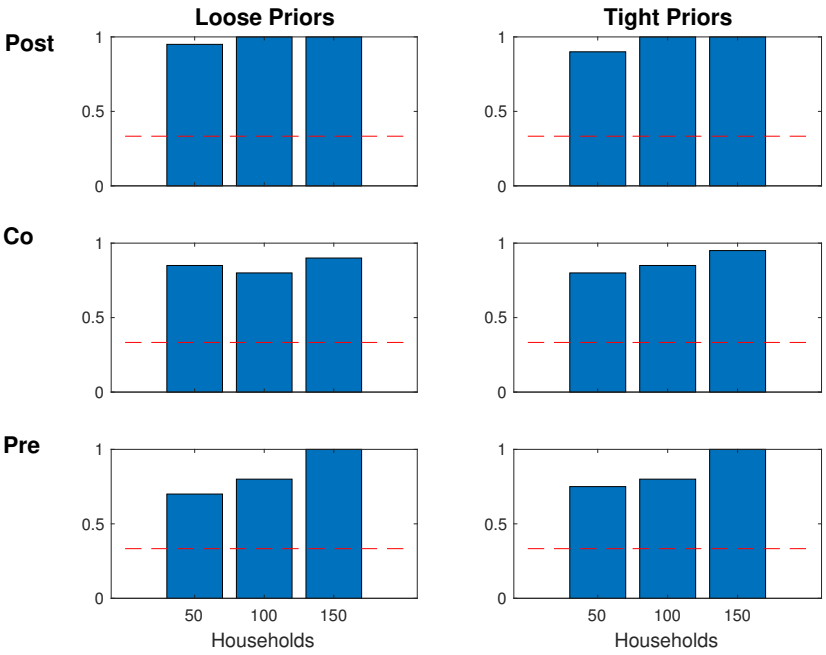

Figure 8: Bar graphs of the proportion of times the correct model was identified from 20 data sets generated from each of the post-symptomatic, coincidental-symptomatic and pre-symptomatic observation models for loose priors (left panels) and tight priors (right panels) for data sets based on 50, 100 and 150 households. The red dotted line is at 1/3, which represents the expected proportion of times the correct model would be chosen if models were indistinguishable.

5. Discussion

This paper has introduced an exact method for Bayesian model selection which used importance sampling for estimating the likelihood function as well as for estimating the evidence. The novelty of this method is the use of an efficient importance sampling scheme ideal for partially-observed state space models used to estimate the likelihood function [8]. Implementation of this scheme has some overhead compared to Doob-Gillespie simulations; however, by construction all simulations agree with the observed data, hence all strictly contribute to the likelihood estimate. This is in sharp contrast to rejection-sampling approaches which do not force samples to fit with observations, such as in approximate Bayesian computation [32] or the Alive particle filter [33, 16]. Hence, it is able to greatly improve computational efficiency, particularly in estimating the tails of the likelihood. This is also an *exact* method, in that it computes unbiased estimates which converge to the target as the number of iterations increases. In contrast, other methods may not be guaranteed to converge; for example, approximate Bayesian computation, which accepts samples that do not perfectly agree with data. Further, methods may have an exact implementation, but in practice approximations are needed for computational feasibility; for example, the Alive particle filter [33, 16] can be implemented as an exact method but may require too many simulations to be practical, so the likelihood function is set to 0 in the tails instead.

If it is possible to evaluate conditional densities, a different approach is to directly calculate the evidence after use of a Gibbs sampler [34, 35]. This is problematic in that typically the Gibbs sampler for partially-observed CTMCs will require data augmentation [36, 6], making mixing slow and convergence problematic when the amount of missing data to be imputed becomes large [21, 37, 6]. Further, the conditional densities are usually only analytically tractable if priors are conjugate. In [17] data-augmentation is used to estimate the evidence, but this required a time discretisation of the model so that the forward-filtering backwards-sampling algorithm could be applied to sample from the latent process. If the model is not discretised and it is not possible to sample exactly from the latent process, data-augmentation may still be applied. However, for each point estimate of the likelihood function this requires imputing latent variables

until the process reaches stationarity. Imputation is undesirable as convergence can be slow for models with high dimensional latent processes and likelihood estimates may still have a high variance. Our method circumvents this difficult imputation by instead marginalising over the latent variables in estimating the likelihood. A weakness of our approach is that we are currently restricted to Markovian models where a single event type is observed whereas data-augmented approaches are much more flexible. Current research is looking at how our method can be extended to situations where multiple types, and combinations, of events are observed.

Other methods for Bayesian model selection include reversible-jump MCMC (RJ-MCMC) [5], nested sampling [38, 39] and sequential Monte Carlo-squared (SMC^2) [40, 16, 41]. RJ-MCMC is an algorithm in which the model is considered as one of the parameters to be sampled [5]. One issue with such an approach is that if models are non-nested it can be difficult to make sensible proposals between the different parameter spaces, particularly if data-augmentation techniques are used to address unobserved transitions [17]. Hence, at times, proposals are accepted with low probability and the mixing of the algorithm can be prohibitively slow. Nested sampling is an approach to model selection which does not require MCMC samples, however it requires a tractable likelihood [38, 39], which is uncommon for most dynamical mechanistic models. SMC^2 is an alternative that jointly infers model parameters and the posterior model probabilities by using particle filters in the parameter space and in the state space [40, 16, 41]. As noted by [16], the estimates that allow for model selection may have large Monte Carlo error, but use of importance sampling post-SMC^2 can be an effective way of reducing this error.

The method of importance sampling for estimating the evidence has typically only been considered for cases where the likelihood function is known [4, 18], for inference on continuous-state models [18], where data augmentation is used to estimate the likelihood function [17, 18], or in conjunction with another method which is not suitable for processes with highly variable observations [16]. This sort of implementation is inefficient for models with high dimensional parameter spaces. One way to overcome this is by using particle-marginal MCMC steps to inform a sampling distribution over the parameter space prior to model selection [17].

In this paper we described the general method used but made no attempt to optimise the algorithm. The choice of the number of particles for estimating the likelihood was chosen for reasonable, rather than optimal, performance so there is potential to improve efficiency by informing the number of particles based on some optimality criteria, such as minimising $\mathrm{var}(\hat{p}(y))/t$ for a fixed number of samples from the parameter space, where t is the amount of computation time taken to estimate $\hat{p}(y)$.

A number of other approaches to model selection exist. Competing models are commonly discriminated using information criteria which are based on maximum likelihood estimates, such as AIC [42], AICc [43], BIC [44] or DIC [45]; note none of these account for prior information about model parameters, they depend on asymptotic results or depend on distributions being approximately Gaussian. Further, interpretation of these quantities is non-intuitive other than that they represent a function of maximum likelihood values which are penalised for each parameter in the model, and comparison of Bayesian model selection and DIC found DIC to be unreliable [46].

Our method was applied to two case studies: inferring the shape of the infectious period distribution of an SI(k)R, model and inferring the time of symptom onset relative to infectiousness for an SE(2)I(2)R model. In each of the studies we considered daily symptom onset data from completed outbreaks in multiple households. The first study showed that the data was sufficiently informative to select the correct model much more often than by randomly guessing. The SI(2)R model was the most difficult model to select correctly, intuitively this seems to be because the disease has an infectious period distribution that has a variance between the other two models. It is worth noting that the posterior model probabilities were often near 1 when the correct model was chosen and most often multiple models had reasonable support when the incorrect model was chosen. This shows that if there was insufficient information to identify the correct model, the correct model was often still given some posterior support. We find that this method does well even with datasets from 50 outbreaks in households.

In the second case study we found that the data was highly informative for being able to determine the time of symptom onset relative to infectiousness. In the case of the post-symptomatic infection model we selected the correct model every time once 100 households were infected and for the pre-symptomatic infection model we

selected the correct model every time once 150 households were infected. In all cases we generally saw an increased ability to select the correct model as more data was obtained. With data on 50 outbreaks and loose priors the Pre model had the lowest proportion of correct selections at 0.7. With data on 150 outbreaks and loose priors the Co model had the lowest proportion of correct selections at 0.9. This shows that symptom onset data from multiple outbreaks is highly informative for choosing the time of symptom onset relative to infectiousness. The results of both studies show that these kinds of data sets are sufficient for discriminating between these models in the early stages of an outbreak of a novel disease, which has important implications for informing public health response [1, 3, 2].

Acknowledgements

JNW acknowledges the support provided by the Dr Michael and Heather Dunne PhD Scholarship in Applied Mathematics and an Australian Government Research Training Program Stipend Scholarship. JVR acknowledges the support of an Australian Research Council Future Fellowship (FT130100254). AJB was supported by an Australian Research Council Discovery Early Career Researcher Award (DE160100690). All authors also acknowledge support from both the ARC Centre of Excellence for Mathematical and Statistical Frontiers (CoE ACEMS), and the Australian Government NHMRC Centre for Research Excellence in Policy Relevant Infectious diseases Simulation and Mathematical Modelling (CRE PRISM2). This work was supported with supercomputing resources provided by the Phoenix HPC service at the University of Adelaide.

References

[1] A. L. Lloyd, Realistic distributions of infectious periods in epidemic models: Changing patterns of persistence and dynamics, Theoretical Population Biology 60 (1) (2001) 59 – 71. doi:https://doi.org/10.1006/tpbi.2001.1525.

[2] H. J. Wearing, P. Rohani, M. J. Keeling, Appropriate models for the management

of infectious diseases, PLOS Medicine 2 (7). doi:10.1371/journal.pmed.0020174.

[3] C. Fraser, S. Riley, R. M. Anderson, N. M. Ferguson, Factors that make an infectious disease outbreak controllable., Proceedings of the National Academy of Sciences of the United States of America 101 16 (2004) 6146–51.

[4] A. E. Gelfand, D. K. Dey, Bayesian model choice: Asymptotics and exact calculations, Journal of the Royal Statistical Society. Series B (Methodological) 56 (3) (1994) 501–514.

[5] P. J. Green, Reversible jump markov chain monte carlo computation and bayesian model determination, Biometrika 82 (4) (1995) 711–732. doi:10.1093/biomet/82.4.711.

[6] J. N. Walker, J. V. Ross, A. J. Black, Inference of epidemiological parameters from household stratified data, PLOS ONE 12 (10) (2017) 1–21. doi:10.1371/journal.pone.0185910.

[7] A. J. Black, N. Geard, J. M. McCaw, J. McVernon, J. V. Ross, Characterising pandemic severity and transmissibility from data collected during first few hundred studies, Epidemics 19 (2017) 61 – 73. doi:https://doi.org/10.1016/j.epidem.2017.01.004.

[8] A. J. Black, Importance sampling for partially observed temporal epidemic models, Statistics and Computingdoi:10.1007/s11222-018-9827-1.

[9] J. S. Sinsheimer, M. A. Suchard, R. E. Weiss, Bayesian Selection of Continuous-Time Markov Chain Evolutionary Models, Molecular Biology and Evolution 18 (6) (2001) 1001–1013. doi:10.1093/oxfordjournals.molbev.a003872.

[10] E. M. Hanks, M. B. Hooten, M. W. Alldredge, et al., Continuous-time discrete-space models for animal movement, The Annals of Applied Statistics 9 (1) (2015) 145–165.

[11] Y. Fong, P. Guttorp, J. Abkowitz, Bayesian inference and model choice in a hidden stochastic two-compartment model of hematopoietic stem cell fate decisions, Ann. Appl. Stat. 3 (4) (2009) 1695–1709. doi:10.1214/09-AOAS269.

[12] Epistruct, Code supporting the paper Walker, Black and Ross 2019. Available: https://github.com/EpiStruct/WBR-2019., Accessed 2019 April 3 .

[13] D. Kroese, T. Taimre, Z. Botev, Handbook of Monte Carlo Methods, Wiley, 2011.

[14] M. A. Beaumont, Estimation of population growth or decline in genetically monitored populations, Genetics 164 (3) (2003) 1139–1160.

[15] C. Andrieu, G. O. Roberts, The pseudo-marginal approach for efficient Monte Carlo computations, Ann. Statist. 37 (2) (2009) 697–725. doi:10.1214/07-AOS574.

[16] C. C. Drovandi, R. A. McCutchan, Alive SMC2: Bayesian model selection for lowcount time series models with intractable likelihoods, Biometrics 72 (2) (2016) 344–353. doi:10.1111/biom.12449.

[17] P. Touloupou, N. Alzahrani, P. Neal, S. E. F. Spencer, T. J. McKinley, Efficient model comparison techniques for models requiring large scale data augmentation, Bayesian Anal. 13 (2) (2018) 437–459. doi:10.1214/17-BA1057.

[18] N. Alzahrani, P. Neal, S. E. Spencer, T. J. McKinley, P. Touloupou, Model selection for time series of count data, Computational Statistics & Data Analysis 122 (2018) 33 – 44. doi:https://doi.org/10.1016/j.csda.2018.01.002.

[19] N. Gordon, D. Salmond, A. Smith, Novel approach to nonlinear/non-gaussian bayesian state estimation, Radar and Signal Processing, IEE Proceedings F 140 (1993) 107 – 113. doi:10.1049/ip-f-2.1993.0015.

[20] J. S. Liu, R. Chen, T. Logvinenko, A theoretical framework for sequential importance sampling with resampling, in: Sequential Monte Carlo methods in practice, Springer, 2001, pp. 225–246.

[21] T. J. McKinley, J. V. Ross, R. Deardon, A. R. Cook, Simulation-based Bayesian inference for epidemic models, Comput. Stat. Data Anal. 71 (2014) 434–447.

[22] M. J. Keeling, P. Rohani, Modeling Infectious Diseases in Humans and Animals, Princeton University Press, New Jersey, 2007.

[23] A. J. Black, A. J. McKane, A. Nunes, A. Parisi, Stochastic fluctuations in the susceptible-infective-recovered model with distributed infectious periods, Phys. Rev. E 80 (2009) 021922. doi:10.1103/PhysRevE.80.021922.

[24] A. J. Black, J. V. Ross, Computation of epidemic final size distributions, Journal of theoretical biology 367 (2015) 159–165.

[25] R. Douc, O. Cappé, Comparison of resampling schemes for particle filtering, in: ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis, 2005., IEEE, 2005, pp. 64–69.

[26] L. L. Lau, B. J. Cowling, V. J. Fang, K.-H. Chan, E. H. Lau, M. Lipsitch, C. K. Cheng, P. M. Houck, T. M. Uyeki, J. M. Peiris, et al., Viral shedding and clinical illness in naturally acquired influenza virus infections, The Journal of Infectious Diseases 201 (10) (2010) 1509–1516.

[27] R. M. May, A. R. McLean, J. Pattison, R. A. Weiss, R. M. Anderson, C. Fraser, A. C. Ghani, C. A. Donnelly, S. Riley, N. M. Ferguson, G. M. Leung, T. H. Lam, A. J. Hedley, Epidemiology, transmission dynamics and control of sars: the 2002–2003 epidemic, Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences 359 (1447) (2004) 1091–1105. doi:10.1098/rstb.2004.1490.

[28] T. House, M. Baguelin, A. J. Van Hoek, P. J. White, Z. Sadique, K. Eames, J. M. Read, N. Hens, A. Melegaro, W. J. Edmunds, et al., Modelling the impact of local reactive school closures on critical care provision during an influenza pandemic, Proceedings of the Royal Society B: Biological Sciences 278 (1719) (2011) 2753–2760.

[29] A. J. Black, T. House, M. J. Keeling, J. V. Ross, Epidemiological consequences of household-based antiviral prophylaxis for pandemic influenza, Journal of The Royal Society Interface 10 (81) (2013) 20121019. doi:10.1098/rsif.2012.1019.

[30] A. J. Black, J. V. Ross, Estimating a markovian epidemic model using household serial interval data from the early phase of an epidemic, PLOS ONE 8 (8) (2013) 1–8. doi:10.1371/journal.pone.0073420.

[31] A. J. Black, J. V. Ross, Contact tracing and antiviral prophylaxis in the early stages of a pandemic: the probability of a major outbreak, Mathematical Medicine and Biology: A Journal of the IMA 32 (3) (2014) 331–343. doi:10.1093/imammb/dqu014.

[32] D. B. Rubin, Bayesianly justifiable and relevant frequency calculations for the applied statistician, Ann. Statist. 12 (4) (1984) 1151–1172. doi:10.1214/aos/1176346785.

[33] P. D. Moral, A. Jasra, A. Lee, C. Yau, X. Zhang, The alive particle filter and its use in particle markov chain monte carlo, Stochastic Analysis and Applications 33 (6) (2015) 943–974. doi:10.1080/07362994.2015.1060892.

[34] S. Geman, D. Geman, Stochastic relaxation, gibbs distributions, and the bayesian restoration of images, IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-6 (6) (1984) 721–741. doi:10.1109/TPAMI.1984.4767596.

[35] S. Chib, Marginal likelihood from the gibbs output, Journal of the American Statistical Association 90 (432) (1995) 1313–1321.

[36] P. D. O'Neill, G. O. Roberts, Bayesian inference for partially observed stochastic epidemics, J. R. Stat. Soc. A 162 (1999) 121–130.

[37] C. M. Pooley, S. C. Bishop, G. Marion, Using model-based proposals for fast parameter inference on discrete state space, continuous-time Markov processes, J. R. Soc. Interface 12 (2015) 20150225.

[38] J. Skilling, Nested sampling, AIP Conference Proceedings 735 (1) (2004) 395–405. doi:10.1063/1.1835238.

[39] R. Salomone, L. South, C. Drovandi, D. Kroese, Unbiased and consistent nested sampling via sequential monte carlo, arXiv preprint arXiv:1805.03924.

[40] N. Chopin, P. E. Jacob, O. Papaspiliopoulos, SMC2: an efficient algorithm for sequential analysis of state space models, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 75 (3) (2013) 397–426.

[41] A. Golightly, T. Kypraios, Efficient SMC2 schemes for stochastic kinetic models, (submitted).

[42] H. Akaike, A new look at the statistical model identification, IEEE Transactions on Automatic Control 19 (6) (1974) 716–723.

[43] C. M. Hurvich, C.-L. Tsai, Regression and time series model selection in small samples, Biometrika 76 (2) (1989) 297–307. doi:10.1093/biomet/76.2.297.

[44] G. Schwarz, Estimating the dimension of a model, The Annals of Statistics 6 (2) (1978) 461–464.

[45] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, A. Van Der Linde, Bayesian measures of model complexity and fit, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 64 (4) (2002) 583–639. doi:10.1111/1467-9868.00353.

[46] C. M. Pooley, G. Marion, Bayesian model evidence as a practical alternative to deviance information criterion, Royal Society Open Science 5 (3) (2018) 171519. doi:10.1098/rsos.171519.