

Estimating Epidemic Incidence and Prevalence from Genomic Data

Timothy G. Vaughan,^{*,†,1,2,3} Gabriel E. Leventhal,^{†,4,5} David A. Rasmussen,^{2,6,8} Alexei J. Drummond,^{1,7} David Welch,^{1,7} and Tanja Stadler^{2,3}

¹Centre for Computational Evolution, University of Auckland, Auckland, New Zealand

²Department of Biosystems Science and Engineering, ETH Zürich, Basel, Switzerland

³Swiss Institute of Bioinformatics (SIB), Lausanne, Switzerland

⁴Institute of Integrative Biology, ETH Zürich, Zurich, Switzerland

⁵Department of Civil and Environmental Engineering, Massachusetts Institute of Technology (MIT), Cambridge, MA

⁶Department of Entomology and Plant Pathology, North Carolina State University, Raleigh, NC

⁷School of Computer Science, University of Auckland, Auckland, New Zealand

⁸Bioinformatics Research Center, North Carolina State University, Raleigh, NC

[†]These authors contributed equally to this work.

*Corresponding author: E-mail: tgvaughan@gmail.com.

Associate editor: Daniel Falush

Abstract

Modern phylodynamic methods interpret an inferred phylogenetic tree as a partial transmission chain providing information about the dynamic process of transmission and removal (where removal may be due to recovery, death, or behavior change). Birth–death and coalescent processes have been introduced to model the stochastic dynamics of epidemic spread under common epidemiological models such as the SIS and SIR models and are successfully used to infer phylogenetic trees together with transmission (birth) and removal (death) rates. These methods either integrate analytically over past incidence and prevalence to infer rate parameters, and thus cannot explicitly infer past incidence or prevalence, or allow such inference only in the coalescent limit of large population size. Here, we introduce a particle filtering framework to explicitly infer prevalence and incidence trajectories along with phylogenies and epidemiological model parameters from genomic sequences and case count data in a manner consistent with the underlying birth–death model. After demonstrating the accuracy of this method on simulated data, we use it to assess the prevalence through time of the early 2014 Ebola outbreak in Sierra Leone.

Key words: phylodynamics, particle filter, epidemiology, Bayesian phylogenetics.

Introduction

A primary goal of infectious disease epidemiology is to understand epidemic dynamics which are most fully described by the prevalence and incidence of cases through time. Yet most epidemics are only partially observed so their dynamics need to be inferred using statistical methods on incomplete data that can come from a wide variety of sources and over a wide range of scales. A key tool for summarizing and understanding epidemic dynamics are compartmental models—such as the SIR model (Kermack and McKendrick 1927)—which partition the hosts at any time into compartments (e.g., susceptible, infectious, or removed) and describe how the counts in the compartments change. By estimating the parameters of a compartmental model, we can calculate fundamental quantities like the basic reproductive number, R_0 , or simulate prevalence and incidence curves to approximate the true epidemic. However, the reliability of these estimated quantities heavily depends on the adequacy of the model used.

In recent years, several statistical methods have been developed for epidemiological inference from genomic data. These methods lie at the intersection of statistical phylogenetics and epidemiology and exploit the rapid evolution of many pathogens that occurs on the same time-scale as their epidemiological spread. In these cases, pathogens are said to be *measurably evolving* (Drummond et al. 2003) and the use of phylogenetics in this context is termed *phylodynamics* (Grenfell et al. 2004).

Early phylodynamic methods used ad hoc methods to infer epidemiological parameters, incidence, and prevalence. The “skyline plot” (Pybus et al. 2000), based on the mathematical relationship between the effective population size and the time between coalescent events in phylogenetic trees (Kingman 1982), was first used to produce nonparametric estimates of HIV prevalence (Pybus et al. 2000). Later, in the context of Hepatitis C virus, skyline plots were fitted to a parametric epidemiological model to estimate the basic reproduction rate, R_0 (Pybus et al. 2001). A subsequent

© The Author(s) 2019. Published by Oxford University Press on behalf of the Society for Molecular Biology and Evolution.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Open Access

approach combined the estimation of the viral phylogeny and the effective viral population size through time into a joint Bayesian method known as the Bayesian skyline plot (Drummond et al. 2005), but this still lacked an explicit model of the epidemiological process. Another variant of the skyline plot based on the birth–death process (Stadler 2010) allowed for piecewise-constant variation in the birth and death rates (Stadler et al. 2013) from which R_0 could be derived. An important limitation of all of these approaches is that they either do not directly integrate epidemiological modeling into the phylogenetic inference method or use piecewise-constant approximations to changing incidence and prevalence through time.

There have recently been three approaches to incorporate compartmental models into phylodynamic inference. First, Volz et al. (Volz et al. 2009; Volz 2012) showed how to derive prior probability distributions for viral gene trees in the coalescent limit from arbitrary birth–death processes. This method gives a theoretical basis for joint Bayesian inference of epidemic model parameters, prevalence curves, and phylogenetic trees. Inference of model parameters and prevalence curves has been performed using this theory (Rasmussen et al. 2011, 2014; Volz and Siveroni 2018). The coalescent basis of this method requires epidemic curves to either be deterministic or be stochastic as long as the epidemic events are statistically independent from the events that make up the sampled epidemic transmission tree (Rasmussen et al. 2011). Either assumption is justified in the case of large population size (prevalence). But when prevalence is low, the coalescent method is known to lead to biased estimates of the phylogenetic tree and the epidemiological parameters (Boskova et al. 2014; Stadler et al. 2015). Furthermore, large sample fractions may lead to violation of statistical independence assumption, as in this case the majority of epidemic events are present on the sampled phylogeny.

Second, Kühnert et al. (2014) used a parametric compartmental model—specifically, a stochastic SIR model—to produce the piecewise-constant rates of the birth–death skyline plot. Like the coalescent methods of Volz et al. (Volz et al. 2009; Volz 2012), this enables joint inference of epidemiological parameters, epidemic curves, and phylogeny which can be performed using the software package, BDSIR. The stochastic formulation of the epidemiological process does not rest on the assumption of large population sizes but, like the coalescent methods, the tree events and the epidemic events are assumed to be statistically independent.

Third, Leventhal et al. (2014) presented the first inference approach to employ an approximation-free computation of the phylogenetic tree probability under a stochastic epidemiological model. The method involves a tailored numerical algorithm to integrate the master equations of a stochastic epidemiological process that is conditioned on the phylogenetic tree. Although this approach can be extended to full joint inference of epidemic model parameters and the phylogeny, the available implementation assumes a known phylogeny and integrates using differential equations over all possible prevalence curves to infer epidemic model parameters.

In this article, we introduce a new method that uses the Particle Marginal Metropolis-Hastings (PMMH) algorithm (Andrieu et al. 2010) to jointly infer prevalence and incidence curves, phylogenetic trees, and epidemiological parameters under stochastic epidemiological models. Our approach addresses several of the shortcomings of previous methods: 1) It accounts for the dependence of epidemic and tree events, 2) it incorporates stochastic models of epidemic dynamics, 3) it includes “sampled ancestors,” and 4) it provides a natural route to the inclusion of additional (nongenetic) incidence data in full joint phylodynamic analyses. The sampled ancestors (Gavryushkina et al. 2014) mentioned in (3) are samples which appear in the phylogenetic tree as direct ancestors to other samples, meaning a patient transmitted after the time of sampling and one or more patients in the downstream transmission chain were also sampled.

Although particle filtering approaches have been previously applied to phylodynamic inference (Rasmussen et al. 2011, 2014; Li et al. 2017; Smith et al. 2017), our application is distinct. In the case of Rasmussen et al. (2011), this approach has only been used in the diffusion limit where the discrete nature of the compartment occupancies is ignored. This assumption was relaxed in Rasmussen et al. (2014), however the tree density was still computed using a coalescent approximation and inference was conditioned on a known genealogy. Similarly, Li et al. (2017) employed particle filtering to estimate the effect of nongeometric distributions of secondary infection counts on the estimation of reproductive number under a coalescent assumption. In contrast, our particle filter is used to compute the exact probability of a transmission tree under the full stochastic discrete compartmental model and used within a joint inference framework. This distinction is especially important near the start of epidemics where prevalence is low and diffusion or coalescent limits do not hold (Stadler et al. 2015). In the case of Smith et al. (2017), particle filtering is applied to individual-based epidemic models. Such models offer greater flexibility than the compartment-based models we use here at the expense of greater computational complexity and a correspondingly lower limit on the number of samples that can be realistically analyzed.

Note that in this article we use *prevalence* to refer to the absolute number of infectious individuals, as this connects concretely to the discrete population models we employ. The proportion (rather than absolute number) of infected individuals can also be easily derived using the methods we describe, as we will demonstrate.

New Approaches

In this section, we derive a flexible and exact inference method for unstructured stochastic compartmental models.

Stochastic Compartmental Epidemic Models

Compartmental models are the centerpiece of epidemiological modeling. They partition individuals in a population into compartments according to their infection status and describe how they transition between the compartments. For example, in an SIS model, individuals are either susceptible (S)

or infectious (I). Susceptible individuals move to the infectious compartment upon infection, and infectious individuals move back to the susceptible compartment upon recovery. The SIR is similar to the SIS model, except that infectious individuals do not move back to the susceptible compartment but are removed (R) meaning that these individuals cannot move back to the infectious compartment. Removal may be due to, for example, recovery with immunity, or death. Let $S[t]$, $I[t]$, and $R[t]$ (or the relevant set for a given model) represent the number of individuals in the respective compartments at time t and define $\mathcal{A}[t] = (S[t], I[t], R[t])$ to be the state of the epidemic at time t .

In this article, we consider unstructured compartmental models: models in which there is only one class of infected individual, that is, those individuals in the single infectious compartment. This rules out 1) models that include an exposed compartment, often called E, where an individual can be infected but not yet infectious (such as SEIR and SEIS) and 2) structuring of the infectious compartment via space, age, or other factors. The reason for this restriction is that lineages of the transmission tree we discuss below would, under a structured model, require labeling to indicate the compartment each part of the lineage occupies thereby greatly increasing the difficulty of the inference problem.

The overall epidemiological model is defined by the set of compartments, the set of epidemic event types, \mathcal{Q} , and their corresponding rates, $\{\alpha_q : q \in \mathcal{Q}\}$. The transitions of individuals between compartments via the epidemic events can be described by a continuous-time Markov process on the state vector \mathcal{A} with master equation:

$$\frac{d}{dt}f(\mathcal{A}, t) = \sum_{q \in \mathcal{Q}} \{\alpha_q(\mathcal{A} - \mathbf{v}_q)f(\mathcal{A} - \mathbf{v}_q, t) - \alpha_q(\mathcal{A})f(\mathcal{A}, t)\}. \quad (1)$$

Here, $f(\mathcal{A}, t) \equiv P(\mathcal{A}[t] = \mathcal{A} | \mathcal{A}[0])$ is the probability that the system state $\mathcal{A}[t]$ at time t has the particular value \mathcal{A} , $\alpha_q(\mathcal{A})$ is the overall rate at which the epidemiological event of type q occurs when the epidemic is in state \mathcal{A} , and \mathbf{v}_q is the effect of event type q on the state: $\mathcal{A} \rightarrow \mathcal{A} + \mathbf{v}_q$.

This formulation encompasses a broad range of models. For instance, a linear birth–death model consists of just one compartment: $\mathcal{A}[t] = I[t]$, the number of infectious individuals at time t . Possible events are infections and removals, so $\mathcal{Q}_{BD} = \{\text{Infection, Removal}\}$. The infection event produces a single new infection as described by $\mathbf{v}_{\text{Inf}} = +1$, and the overall infection rate is $\alpha_{\text{Inf}}(\mathcal{A}[t]) = \beta I[t]$. Here, β is a constant describing the rate at which infectious individuals produce subsequent infected individuals. Similarly, the removal event removes an individual from the infectious compartment, $\mathbf{v}_{\text{Rem}} = -1$, at overall removal rate $\alpha_{\text{Rem}}(\mathcal{A}[t]) = \gamma I[t]$. The SIS model, $\mathcal{A}[t] = (S[t], I[t])$, has the same event type set as the linear birth–death process, $\mathcal{Q}_{SIS} = \mathcal{Q}_{BD} = \{\text{Infection, Removal}\}$, but different rate functions and event effects. An infection has effect vector $\mathbf{v}_{\text{Inf}} = (-1, 1)$ and occurs at rate $\alpha_{\text{Inf}}(\mathcal{A}[t]) = \beta S[t]I[t]$, whereas a removal event has an effect vector $\mathbf{v}_{\text{Rem}} = (1, -1)$ and occurs at rate $\alpha_{\text{Rem}}(\mathcal{A}[t]) = \gamma I[t]$. The SIR model,

$\mathcal{A}[t] = (S[t], I[t], R[t])$, is similar to the SIS model, only with effect vectors $\mathbf{v}_{\text{Inf}} = (-1, 1, 0)$ and $\mathbf{v}_{\text{Rem}} = (0, -1, 1)$. For brevity, we combine the set of constants into a single variable $\boldsymbol{\eta}$, $\boldsymbol{\eta}_{BD} = \boldsymbol{\eta}_{SIS} = \boldsymbol{\eta}_{SIR} = (\beta, \gamma)$.

A specific realization of an epidemic forward in time—an epidemic trajectory—up to some predetermined maximum time T can be generated as follows: The epidemic starts at time $t_0 = 0$ in state $\mathcal{A}[0]$. Typically, $I[0] = 1$ for the infectious compartment, but other choices are possible. This initial state is modified by a series of events with types e_1, e_2, \dots, e_s at times t_1, t_2, \dots, t_s , where s is a random variable indicating the number of events which occurred before T . The number of the individuals in each compartment after the i th event has occurred at time t_i is denoted by $\mathcal{A}_i = \mathcal{A}[t_i]$. The population trajectory of the epidemic is then just given by $((t_0, \mathcal{A}_0), (t_1, \mathcal{A}_1), \dots, (t_s, \mathcal{A}_s))$. Figure 1a shows an example of the infectious compartment occupancy over time. We can then equivalently expand \mathcal{A}_i as a sum of effect vectors:

$$\mathcal{A}_i = \mathcal{A}_0 + \mathbf{v}_{e_1} + \dots + \mathbf{v}_{e_i} = \mathcal{A}_0 + \left(\sum_{k=1}^i \mathbf{v}_{e_k} \right).$$

An epidemiological trajectory \mathcal{E} is thus well defined by the initial state, \mathcal{A}_0 , the vector of transition events $\mathbf{e} = (e_1, e_2, \dots, e_s)$, and the corresponding event times, $\mathbf{t} = (t_1, t_2, \dots, t_s)$,

$$\mathcal{E} = \{\mathcal{A}_0, \mathbf{E} = (\mathbf{e}, \mathbf{t})\}. \quad (2)$$

As for any time-homogeneous discrete state continuous-time Markov process, the probability density of a particular trajectory is a product of exponentially distributed waiting times between the s events with factors representing the probability density of each event. That is,

$$P(\mathbf{E} | \boldsymbol{\eta}, \mathcal{A}_0, T) = \prod_{i=1}^s \exp\{-a_{i-1}(t_i - t_{i-1})\} \alpha_{e_i}(\mathcal{A}_{i-1}) \times \exp\{-a_s(T - t_s)\}, \quad (3)$$

where $a_i = \sum_{q \in \mathcal{Q}} \alpha_q(\mathcal{A}_i)$ is the sum of the rates of all possible transitions in the interval (t_i, t_{i+1}) . For example, under the SIS model, new infections happen at a rate $\beta S_i I_i$ and infected individuals are removed at a rate γI_i . By defining $\mathcal{I}_{\text{Inf}} \subset \mathcal{I} = \{1, \dots, s\}$ to be the indices of infection events, and $\mathcal{I}_{\text{Rem}} \subset \mathcal{I}$ to be the indices of the removal events, we can write the probability density for an SIS trajectory as

$$P_{\text{SIS}}(\mathbf{E} | \boldsymbol{\eta}, \mathcal{A}_0, T) = \prod_{i \in \mathcal{I}} \exp\{-(\beta S_{i-1} I_{i-1} + \gamma I_{i-1})(t_i - t_{i-1})\} \prod_{i \in \mathcal{I}_{\text{Inf}}} \beta S_{i-1} I_{i-1} \prod_{i \in \mathcal{I}_{\text{Rem}}} \gamma I_{i-1} \times \exp\{-(\beta S_s I_s + \gamma I_s)(T - t_s)\}. \quad (4)$$

Modeling the Sampling Process

Sampling of individuals can be described by expanding \mathcal{Q} to include two additional event types, sampling with and

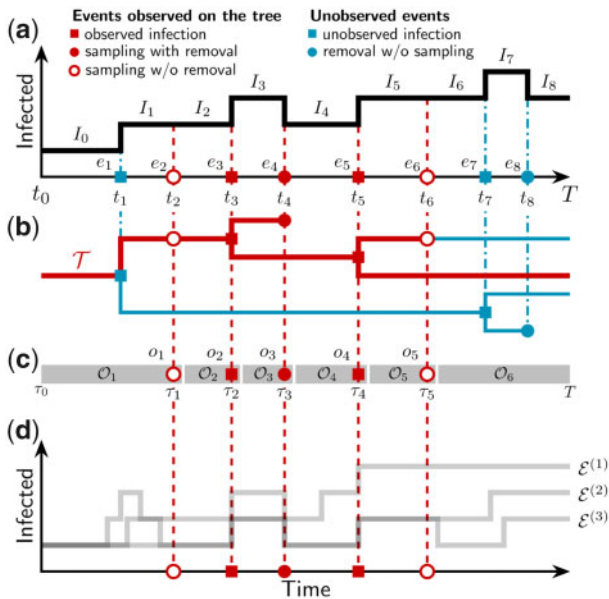


FIG. 1. The true epidemiological trajectory can be inferred from the reconstructed phylogeny. (a) The trajectory \mathcal{E} of an epidemic outbreak consists of a sequence of events (infection, sampling, and recovery) e_i at times t_i that result in a corresponding sequence of compartment occupancies such as the infectious compartment occupancies I_i . (b) The full transmission tree contains information on when infections happened and between which lineages (filled squares) and when individuals were removed (filled circles). The sampled transmission tree \mathcal{T} represents a subset of the full tree (red). The rest of the transmission tree is unobserved (blue). (c) The time ordered observations \mathcal{O}_j consist of the events o_j seen on the tree (infection, sampling with removal, and sampling without removal) at times τ_j , combined with the number of lineages on the sampled tree in the intervals immediately before each of these events. (d) There is an ensemble of trajectories $\mathcal{E}^{(1)}, \mathcal{E}^{(2)}, \dots$ that are compatible with the sampled transmission tree. Note that the sampled transmission tree contains only a subset of the events represented by the full tree and true trajectory \mathcal{E} , and each of these “observed” events must be present in every compatible trajectory.

without removal. Although the particular form of the effect vectors depends on the dimension of the compartmental model, their effect remains the same: $\mathbf{v}_{\text{SampR}}$ removes an individual from the infectious class, whereas $\mathbf{v}_{\text{SampNR}}$ leaves $\mathcal{A}[t]$ unchanged. We explicitly model sampling by augmenting the stochastic process with sampling events and times, and their corresponding rates: $\alpha_{\text{SampR}}(\mathcal{A}[t]) = r\psi I(t)$ and $\alpha_{\text{SampNR}}(\mathcal{A}[t]) = (1-r)\psi I(t)$, where ψ is the per-individual sampling rate parameter and r is the probability of removal following sampling. Additionally, any remaining infected individuals at time T , that is, the end of the process, are sampled with probability ρ . For convenience, we group all parameters related to sampling together in the vector $\boldsymbol{\sigma} = (\psi, r, \rho)$. We then define $P(\mathbf{E}, m | \boldsymbol{\eta}, \boldsymbol{\sigma}, \mathcal{A}_0, T)$ to represent the probability density of this combined process producing a trajectory \mathbf{E} terminated by m contemporaneous samples at time T . For example, in the case of the SIS model, this probability density is

$$P_{\text{SIS}}(\mathbf{E}, m | \boldsymbol{\eta}, \boldsymbol{\sigma}, \mathcal{A}_0, T) =$$

$$\begin{aligned} & \prod_{i \in \mathcal{I}} \exp \{ -(\beta S_{i-1} I_{i-1} + (\gamma + \psi) I_{i-1})(t_i - t_{i-1}) \} \\ & \prod_{i \in \mathcal{I}_{\text{Inf}}} \beta S_{i-1} I_{i-1} \prod_{i \in \mathcal{I}_{\text{Rem}}} \gamma I_{i-1} \prod_{i \in \mathcal{I}_{\text{SampR}}} r \psi I_{i-1} \prod_{i \in \mathcal{I}_{\text{SampNR}}} (1-r) \psi I_{i-1} \\ & \times \exp \{ -(\beta S_T I_T + (\gamma + \psi) I_T)(T - t_s) \} \\ & \times \binom{I_s}{m} \rho^m (1-\rho)^{(I_s-m)}. \end{aligned} \quad (5)$$

From Epidemiological Trajectories to Transmission Trees

By tracking the identity of who infected whom along an epidemiological trajectory, we obtain the transmission tree of the epidemic (full tree in [fig. 1b](#)). All events e_i in the trajectory ([fig. 1a](#)) correspond to nodes in the full tree. The number of extant lineages in the full tree immediately following the event time t_i is l_i .

The *sampled phylogeny*, \mathcal{T} , is the subset of the full tree where only the subtree ancestral to sampling events is retained (red subtree in [fig. 1b](#)). We use k_i to represent the number of lineages present in the sampled phylogeny immediately following time t_i , so $k_i \leq l_i$. The number of lineages remaining in the sampled tree at time T is $k_s = m$.

Because of its relation to the full tree, each node in the sampled phylogeny must correspond to a compatible event in the trajectory for the probability of the sampled phylogeny given the trajectory $P(\mathcal{T} | \mathcal{E}, m)$ to be nonzero. Furthermore, this distribution is independent of the particular epidemiological model. In particular, conditional on the trajectory, the sampled phylogeny can be considered a result of a discrete-time Markov chain proceeding from the most recent sample to the start of the epidemic process. This can be illustrated by defining \mathcal{T}_i to be the portion of the sampled phylogeny \mathcal{T} on the interval $(t_i, t_{i+1}]$, that is, including the tree node (if any) which corresponds to the event e_{i+1} . We assume that lineages in the tree \mathcal{T}_i are labeled, such that the correspondence between lineages in \mathcal{T}_i and \mathcal{T}_{i+1} is unambiguous.

For example, an infection event, $e_i = \text{Infection}$, in the trajectory only produces a branching event in the sampled tree when both the infector and the infected correspond to lineages in the sampled phylogeny, so

$$P(\mathcal{T}_{i-1} | \mathcal{T}_i, l_i, e_i = \text{Infection}) = \begin{cases} \frac{1}{\binom{l_i}{2}} & \text{for } k_i = k_{i+1} - 1, \\ 1 - \frac{\binom{k_i}{2}}{\binom{l_i}{2}} & \text{for } k_i = k_{i+1}, \\ 0 & \text{otherwise,} \end{cases}$$

where l_i is the total number of infected individuals (including the newly infected individual) and thus $\binom{l_i}{2}$ is the total

number of pairs of lineages after the infection event, each of which could have been the pair of lineages involved in the event.

Unsampled removal events do not themselves correspond to any nodes in sampled phylogenies, so if $e_i = \text{Removal}$ we have

$$P(\mathcal{T}_{i-1}|\mathcal{T}_i, l_i, e_i = \text{Removal}) = \begin{cases} 1 & \text{for } k_{i-1} = k_i, \\ 0 & \text{otherwise.} \end{cases}$$

On the other hand, any sampling with removal event corresponds to a leaf node at the time of the event in the sampled phylogeny with probability one:

$$P(\mathcal{T}_{i-1}|\mathcal{T}_i, l_i, e_i = \text{SampR}) = \begin{cases} 1 & \text{for } k_{i-1} = k_i + 1, \\ 0 & \text{otherwise.} \end{cases}$$

In the case of samples that do not coincide with removal of the sampled lineage, there is ambiguity regarding whether the event is represented by an external leaf node or an internal sampled ancestor node in the sampled phylogeny, as this depends on whether or not any descendants of the sample are subsequently sampled:

$$P(\mathcal{T}_{i-1}|\mathcal{T}_i, l_i, e_i = \text{SampNR}) = \begin{cases} \frac{1}{l_i} & \text{for } k_{i-1} = k_i, \\ 1 - \frac{k_i}{l_i} & \text{for } k_{i-1} = k_i + 1, \\ 0 & \text{otherwise.} \end{cases}$$

Combining the probabilities above allows us to calculate the full probability of the sampled phylogeny given a complete compatible trajectory as

$$\begin{aligned} P(\mathcal{T}|\mathcal{E}, m) &= P(\mathcal{T}_s|m) \prod_{i=1}^s P(\mathcal{T}_{i-1}|\mathcal{T}_i, e_i, l_i) \\ &= \delta_{k_s, m} \prod_{i \in \mathcal{I}_{\text{Inf}}} \left(\delta_{k_{i-1}, k_i} \left(1 - \frac{k_i(k_i - 1)}{l_i(l_i - 1)} \right) \right. \\ &\quad \left. + \delta_{k_{i-1}, k_i - 1} \frac{2}{l_i(l_i - 1)} \right) \\ &\quad \times \prod_{i \in \mathcal{I}_{\text{SampNR}}} \left(\delta_{k_{i-1}, k_i} \frac{1}{l_i} + \delta_{k_{i-1}, k_i + 1} \left(1 - \frac{k_i}{l_i} \right) \right), \end{aligned} \tag{6}$$

where δ is the Kronecker delta, and $P(\mathcal{T}_s|m) = 1$ provided $k_s = m$.

Accounting for Unsequenced Samples

We now consider the possibility that samples generated by the birth–death–sampling process may be absent from the sampled phylogeny. These samples, which we refer to here as *unsequenced samples*, arise naturally in epidemiological settings where a large number of pathogen samples may be collected at known times but only a subset are subsequently sequenced. Similarly, doctors’ records can provide evidence

that individuals were carrying a pathogen at a particular time, but without sequencing there is no information about where exactly the pathogen lineages ancestral to these samples attach to a sample phylogeny.

It is possible to directly include unsequenced samples in the phylogeny but their relationship to the rest of the phylogeny would not be informed by data and they would contribute nothing to the inference of relationships between the sequenced samples while increasing the complexity of the overall inference problem.

Instead, we assume that the set of all sampling event indices $\mathcal{I}_{\text{SampNR}} \cup \mathcal{I}_{\text{SampR}}$ is arbitrarily partitioned into subsets \mathcal{I}_{Seq} and $\mathcal{I}_{\text{Unseq}}$ containing indices of sequenced and unsequenced sampling events, respectively. (By allowing this partitioning to be arbitrary, we are choosing not to explicitly model the decision to sequence a given sample, but to instead condition on this decision.) Since this classification then has no effect on the probability density of the stochastic trajectory, we simply exclude the unsequenced sample indices from the final product in the tree probability given by equation (6). This gives the following joint probability for the time tree \mathcal{T} and the unsequenced sample times \mathcal{S} :

$$\begin{aligned} P(\mathcal{T}, \mathcal{S}|\mathcal{E}, m, \mathcal{I}_{\text{Seq}}) &= \delta_{k_s, m} \prod_{i \in \mathcal{I}_{\text{Inf}}} \left(\delta_{k_{i-1}, k_i} \left(1 - \frac{k_i(k_i - 1)}{l_i(l_i - 1)} \right) \right. \\ &\quad \left. + \delta_{k_{i-1}, k_i - 1} \frac{2}{l_i(l_i - 1)} \right) \\ &\quad \times \prod_{i \in \mathcal{I}_{\text{SampNR}} \cap \mathcal{I}_{\text{Seq}}} \left(\delta_{k_{i-1}, k_i} \frac{1}{l_i} + \delta_{k_{i-1}, k_i + 1} \left(1 - \frac{k_i}{l_i} \right) \right). \end{aligned} \tag{7}$$

Again, we emphasize that this expression assumes each event in \mathcal{T} and \mathcal{S} has a corresponding event in the trajectory \mathcal{E} and that otherwise the joint probability is zero.

Bayesian Inference

One of our goals is to perform asymptotically exact Bayesian inference of both the prevalence trajectory and the epidemiological parameters using a set of pathogen sample times, a subset for which genetic sequence data are available, collected throughout an epidemic. To this end, for a given pathogen sequence alignment (with a sampling time associated with each sequence) D and set of times of unsequenced samples \mathcal{S} , we use Bayes’ rule to express the joint posterior distribution for the model parameters and the epidemic trajectories in terms of the conditional distributions composing the full model:

$$\begin{aligned} P(\mathcal{E}, \mathcal{T}, \boldsymbol{\mu}, \boldsymbol{\eta}, \boldsymbol{\sigma}, T|D, \mathcal{S}) &= \frac{1}{P(D, \mathcal{S})} P(D|\mathcal{T}, \boldsymbol{\mu}) \\ &\quad \times P(\mathcal{T}, \mathcal{S}|\mathcal{E}, m, \mathcal{I}_{\text{Seq}}) \\ &\quad \times P(\boldsymbol{E}, m|\mathcal{A}_0, \boldsymbol{\eta}, \boldsymbol{\sigma}, T) P(\mathcal{A}_0, \boldsymbol{\mu}, \boldsymbol{\eta}, \boldsymbol{\sigma}, T). \end{aligned} \tag{8}$$

Here, $P(D, \mathcal{S})$ can be treated as a normalization constant and $P(D|\mathcal{T}, \boldsymbol{\mu})$ is the probability of D evolving down the

sampled transmission tree \mathcal{T} under a substitution model parameterized by μ , also known as the *phylogenetic likelihood*. $P(\mathcal{A}_0, \mu, \eta, \sigma, \mathcal{T})$ represents the joint prior probability distribution for the model parameters.

Several approaches to characterizing this posterior for particular models already exist in the literature, all of which involve using Markov chain Monte Carlo (MCMC) to sample (or maximum likelihood to optimize) a marginalized and/or approximate form of equation (8). For instance, [Stadler et al. \(2012\)](#) analytically marginalize over the trajectory subspace in the case of the linear birth–death model and use MCMC to sample from $(\mathcal{T}, \mu, \eta, \sigma, \mathcal{T})$. Similarly, [Leventhal et al. \(2014\)](#) express the marginalization of equation (8) over trajectories for the nonlinear stochastic SIS model as the solution to a master equation which is then integrated numerically with parameter inferences being drawn by applying MCMC or maximum likelihood.

[Kühnert et al. \(2014\)](#) provide an approximation to the posterior for discretized trajectories under the stochastic SIR model and use MCMC to sample $(\mathcal{E}, \mathcal{T}, \mu, \eta, \sigma, \mathcal{T})$. [Volz et al. \(2009\)](#) and [Volz \(2012\)](#) present an approximation to this posterior under the assumption that the relative amplitude of the stochastic noise in \mathcal{E} is negligible and that $P(\mathcal{E}, m | \eta, \sigma, \mathcal{A}_0, \mathcal{T})$ therefore collapses to a point mass centered on the approximate deterministic solution of the model.

In contrast to these methods, we use the PMMH algorithm ([Andrieu et al. 2010](#)). This has previously been applied in a phylodynamic context by [Rasmussen et al. \(2011, 2014\)](#) using a coalescent approximation to the distribution of sampled phylogenies, but not to sample directly from the exact phylodynamic posterior as we do in the algorithm described below.

Particle Filtering Algorithm

We employ the PMMH algorithm described by [Andrieu et al. \(2010\)](#). In the form presented here, it involves using a bootstrap particle filter to simulate trajectories \mathcal{E} conditional on both a sampled transmission tree \mathcal{T} and the times of unsequenced samples \mathcal{S} .

We call the union of the sampled phylogeny \mathcal{T} and the temporally distributed unsequenced samples \mathcal{S} the *observed process*, \mathcal{O} , and use o_j to represent the j th observation (either a node of the sampled phylogeny or an unsequenced sample) when ordered according to the observation times τ_j , as illustrated in [figure 1c](#). The final (N th) observation represents the contemporaneous sampling of m lineages in the sampled phylogeny, although it is possible for m to be zero.

We divide the time into intervals between observations. The first of these intervals begins at time $\tau_0 = t_0 = 0$, whereas the last ends at time T . We denote the portion of the observed process within interval j using \mathcal{O}_j , which is understood to include both the number of tree lineages extant within the interval and the observation o_j at end of the interval. Similarly, we divide the full trajectory \mathcal{E} into corresponding partial trajectories \mathcal{E}_j which contain only the trajectory events within each observation interval and define

\mathcal{E}_j' to be the partial trajectory excluding the event e_j corresponding to the observation o_j .

The algorithm involves simulating an ensemble of M trajectories or “particles” in each of the N intervals between τ_0 and $\tau_N = T$. The initial condition for each particle is sampled from the ensemble of finishing states of particles simulated in the previous interval, weighted according to the probability of the observation event that divides the intervals.

The algorithm is as follows:

- (1) Set the interval index $j \leftarrow 1$ and define $x_0^{(a)} = \mathcal{A}_0$ to be the starting state of particle a .
- (2) For each $a \in [1 \dots M]$, use Gillespie’s stochastic simulation algorithm ([Gillespie 1976, 1977](#)) or its asymptotically exact equivalent ([Gillespie 2001](#)) to sample a partial trajectory $\mathcal{E}_j^{(a)}$ from the distribution:

$$P(\mathcal{E}_j' | \eta, \sigma, \tau_j - \tau_{j-1}, x_{j-1}^{(a)}), \tag{9}$$

which is a solution to the master equation given in [equation \(1\)](#) conditioned on the initial state $x^{(a)}$ and the interval time $\tau_j - \tau_{j-1}$.

- (3) Each sampled partial trajectory $\mathcal{E}_j^{(a)}$, which is defined as the union of $\mathcal{E}_j'^{(a)}$ and the event corresponding to the observation o_j , is assigned a weight:

$$\omega_j^{(a)} = P(\mathcal{O}_j | \mathcal{E}_j^{(a)}, m, \mathcal{I}_{\text{seq}}) \alpha_{o_j}(y_j^{(a)}). \tag{10}$$

The probability on the right-hand side is given by [equation \(7\)](#) but restricted to include only the epidemic events within the interval and the observation event o_j . The factor $\alpha_{o_j}(y_j^{(a)})$ is the transition rate for the epidemic event corresponding to o_j given the final state of $\mathcal{E}_j^{(a)}$, denoted here $y_j^{(a)}$. (This factor ensures that the particle trajectories are constrained to be consistent with the observation event o_j , as inconsistent trajectories will be assigned a weight of zero.)

- (4) The mean of weights $\Omega_j = (\sum_{a=1}^M \omega_j^{(a)})/M$ is recorded, and a new set of M trajectory states $x_j^{(1)} \dots x_j^{(M)}$ is sampled with replacement from the weighted distribution of the final states of the partial trajectories \mathcal{E}_j .
- (5) If $j < N$, set $j \leftarrow j + 1$ and go to step 2.
- (6) Compute the product $\hat{P}(\mathcal{T}, \mathcal{S} | \mathcal{A}_0, \eta, \sigma, \mathcal{T}) \equiv \prod_{j=1}^M \Omega_j$ which is, as highlighted below, an estimate of the marginal density $P(\mathcal{T}, \mathcal{S} | \mathcal{A}_0, \eta, \sigma, \mathcal{T})$, with the marginalization being over the epidemiological trajectories. Also, sample a single final partial trajectory $\hat{\mathcal{E}}_i$ from the final distribution of weighted partial trajectories and follow the sequence of events back through the observation intervals until $t = 0$, yielding a single sampled full trajectory $\hat{\mathcal{E}}$.

It can be shown ([Del Moral 2004](#)) that the value of $\hat{P}(\mathcal{T}, \mathcal{S} | \mathcal{A}_0, \eta, \sigma, \mathcal{T})$ is an unbiased and consistent estimate of the marginal probability density for the sampled phylogeny and unsequenced samples $P(\mathcal{T}, \mathcal{S} | \mathcal{A}_0, \eta, \sigma, \mathcal{T})$. (This probability

density is sometimes called the *phylogenetic likelihood*, and below we simply write “likelihood,” although the implicit classification of \mathcal{T} as “data” should not be understood to mean that phylogenies are physically observed.) As shown by Andrieu et al. (2010), this implies that by using this estimate in place of the terms $P(\mathcal{T}, \mathcal{S} | \mathcal{E}, m, \mathcal{I}_{\text{seq}})P(\mathcal{E}, m | \mathcal{A}_0, \boldsymbol{\eta}, \boldsymbol{\sigma}, T)$ in the posterior given by equation (8), and using the resulting expression as the target distribution for a MCMC algorithm, we obtain an algorithm for sampling from the joint posterior marginalized over the epidemic trajectories. Furthermore, by recording the sampled trajectories $\hat{\mathcal{E}}$ generated by the particle filter alongside the parameter values and sampled phylogenies visited by the MCMC procedure, the algorithm generates samples from the full (unmarginalized) joint posterior.

The use of particle filtering to condition the epidemic trajectories on the tree is potentially confusing, due to the (backward-time) correlations between the observations that make up the sampled phylogeny. Despite these correlations, the PMMH algorithm remains applicable since the joint probability of the observations and hidden state, $P(\mathcal{T}, \mathcal{S}, \mathcal{E} | \mathcal{A}_0, \boldsymbol{\eta}, \boldsymbol{\sigma}, T)$, can be expressed in precisely the same form as the weighted sequence of conditional probabilities generated by a standard hidden Markov model. This is shown in the [supplementary text, Supplementary Material](#) online, along with a simple demonstration that the resulting algorithm does indeed produce samples from the required marginal density of the observations given the phylodynamic model parameters.

Results

Implementation and Validation

We have implemented the algorithm described above as a BEAST 2 (Bouckaert et al. 2014) package. This allows the algorithm to be used in conjunction with standard phylogenetic models such as those describing the nucleotide substitution process as well as existing algorithms for performing the MCMC sampling of the phylogenetic tree space. The package is released under the GNU General Public License and instructions for installing and using it can be found, along with source code, at <http://tgvaughan.github.io/EpilInf>.

All of the BEAST 2 input files necessary to reproduce the results described in this section, together with instructions on how to use them, may be downloaded from

<http://github.com/tgvaughan/ParticleFilterResults>.

Direct Likelihood Comparison

We validated our algorithm and its implementation by comparing the likelihoods generated by the particle filter with those computed analytically under the linear birth–death model (Stadler 2010) and numerically under the nonlinear stochastic SIS model (Leventhal et al. 2014). These comparisons were performed for a variety of parameter combinations and in all cases yielded perfect agreement (fig. 2).

Comparison of Tree-Based and Incidence-Based Sampling

The joint tree and sample time prior defined in equation (7) has the property that marginalizing over the time tree yields a quantity which is independent of which samples are

sequenced and which samples are not. In other words, if the sequence data from the sampled individuals provide no information about the phylogenetic tree then the only information we have are the sample times: our estimates of the epidemiological model parameters should therefore not depend on which samples were sequenced. This suggests the following test for the consistency of the joint posterior:

- (1) Fix a set of sampling times.
- (2) Assign a fraction f of these times to be associated with tree leaves (i.e., play the role of “sequenced” sample times),
- (3) Sample from the joint posterior defined in equation (8) without sequence data (i.e., setting $P(D | \mathcal{T}, \boldsymbol{\mu})$ to a constant).

Provided the unsequenced sampling times are being handled consistently by the sampler, the posteriors for model parameters should be identical regardless of f .

We performed this test using a set of 83 sample times simulated using a birth–death–sampling process and using these times, via the procedure above, to produce the posterior for the birth rate parameter β as a function of f . The lack of variation in this posterior as with respect to f , shown in figure 3, is strong evidence that our treatment of unsequenced samples is indeed consistent with our treatment of sequenced samples.

Inference from Simulated Data

In order to assess the capability of the sampler to recover prevalence trajectories, we simulated trajectories under each of the three models supported by our implementation: linear birth–death ($\beta = 1.2$, $\gamma = 0.1$, $\psi/(\psi + \gamma) = 0.5$, $T = 7.0$), stochastic SIS ($\beta = 0.02$, $\gamma = 1.0$, $\psi/(\psi + \gamma) = 0.1$, $T = 5$, $S_0 = 199$), and stochastic SIR ($\beta = 0.2$, $\gamma = 1.0$, $\psi/(\psi + \gamma) = 0.1$, $T = 5$, $S_0 = 199$). In all cases, we fixed the removal probability $r = 1$, the present-day sampling probability $\rho = 0$ and set $I_0 = 1$. Sampled transmission trees were then simulated from each of these trajectories, which were in turn used to simulate 2-kb genetic sequence alignments under a simple Jukes–Cantor model with a substitution rate of 5×10^{-3} per site per unit time. For each of these three alignments, we then used our algorithm to sample from the joint posterior for the transmission tree, epidemic trajectory, and the model parameters β , γ , T , and (in the case of SIS and SIR) S_0 . (The remaining parameters ψ , r , and ψ were fixed to the truth.) For the continuous parameters, we employed improper priors $P(\beta) = 1/\beta$, $P(\gamma) = 1/\gamma$ and $P(T) = 1/T$. For the discrete S_0 parameter, we used $P(S_0) = \text{Unif}(0, 300)$.

Figure 4 illustrates the agreement between the posterior prevalence distributions obtained from each of these analyses (red lines) and the true prevalence curves (black lines). Also shown is the distribution of prevalence curves generated directly from the posterior samples of the model parameters (blue lines). Prior to our PMMH algorithm, the blue lines were the best estimates obtained for prevalence under compartmental models (unless coalescent approximations were

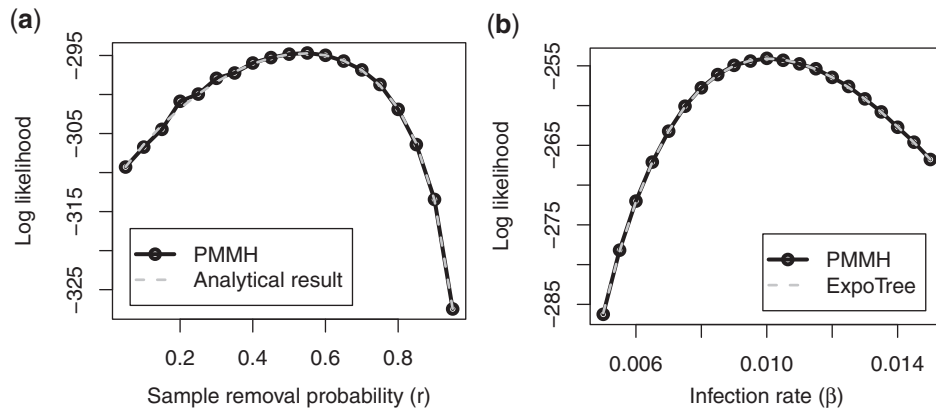


FIG. 2. Comparison between values of the phylodynamic likelihoods computed using the PMMH algorithm with those calculated using other approaches: (a) likelihood of r under the linear birth–death model from PMMH compared with the analytical result (Stadler 2010) and (b) likelihood of β under the stochastic SIS model from PMMH compared with a numerical result from ExpoTree (Leventhal et al. 2014).

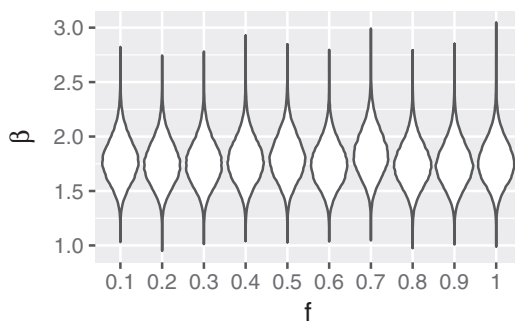


FIG. 3. Marginal posteriors for the infection rate as a function of the fraction f of samples regarded as “sequenced” when no data besides the sampling times are available. The invariance of this distribution with respect to f shows that the treatment of unsequenced samples is consistent with the treatment of sequenced samples.

appropriate in the particular application). As these blue trajectories are not explicitly conditioned on the corresponding sampled transmission trees however, there is a significantly greater variance in their distribution.

Quantitative Validation of Trajectory Inference

Although agreement between simulated and subsequently inferred trajectories is encouraging, we use a well-calibrated approach (Dawid 1982) for a more robust quantitative validation of the inference algorithm. The steps of this approach are as follows.

- (1) Under each model (linear birth–death, SIS and SIR) and a chosen set of parameters (table 1), we simulate 200 trajectories and sampled trees.
- (2) A random DNA sequence is simulated down each sampled tree, resulting in a unique simulated sequence alignment.
- (3) For each simulated sequence alignment, infer the corresponding trajectory conditional on the true model parameters using our inference algorithm.
- (4) We compute the proportion of analyses for which the true prevalence at a particular time falls within the 100 α % highest posterior density (HPD) interval of

the sampled posterior distribution for the prevalence at this time. This is repeated for a range of times and α values.

Figure 5 shows, for each model, the perfectly linear relationship between α and the proportion of analyses for which the 100 α % HPD includes the truth. This relationship provides strong evidence that our implementation of the algorithm correctly samples from the true distribution of epidemiological trajectories.

Inference of Ebola Prevalence in Sierra Leone

In order to demonstrate the applicability of our method, we analyzed 101 full Ebola virus (EBOV) genomes collected from the Kailahun district in eastern Sierra Leone during the 2014 west-African epidemic (Gire et al. 2014; Bell et al. 2015; Carroll et al. 2015; Park et al. 2015), as curated and aligned by Dudas et al. (2017). These sequences were analyzed jointly with the temporal distribution of unsequenced Kailahun cases (World Health Organization 2016). To assess the degree to which the inclusion of unsequenced data affected the inferred trajectory distributions, we conducted a separate analysis based solely on sequence data collected during the first 4 weeks. Later, sequences were excluded from the latter analysis to avoid introducing bias due to the sequencing fraction being skewed toward earlier weeks (fig. 6f).

We assumed a standard neutral model of sequence evolution allowing for distinct transition/transversion rates and nonequilibrium base frequencies (Hasegawa et al. 1985), together with Gamma-distributed rate heterogeneity among sites (Yang 1994). We further assumed a strict clock rate whose value was jointly estimated using an informative prior derived from a recent meta-analysis (Holmes et al. 2016).

We assumed a stochastic SIR epidemiological model in which each sample (whether sequenced or unsequenced) is assumed to be generated by a linear sampling process with fixed rate ψ between the times of the most recent and earliest samples. Importantly, although the temporal distribution of sample collection times is determined by this model, the choice of which samples to sequence is not. We feel that

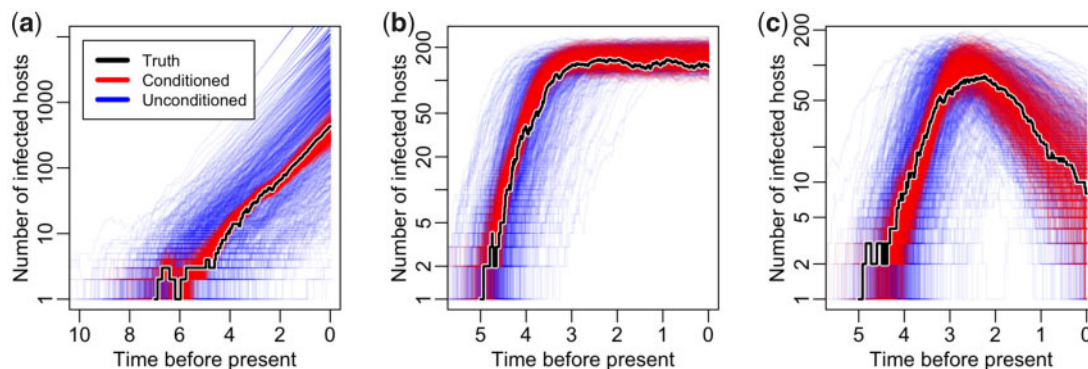


FIG. 4. Inference of prevalence dynamics from sequence data simulated under (a) linear birth–death, (b) stochastic SIS, and (c) stochastic SIR model. Samples from the posterior of the prevalence trajectory are shown in red, whereas the black line represents the truth. The blue lines are prevalence trajectories simulated from the posterior samples of the compartmental model parameters.

Table 1. Fixed Parameter Values Used for Well-Calibrated Trajectory Inference Validation.

Model	β	γ	S_0	ψ	r	ρ	T
Linear birth–death	0.5	0.1	—	0.25	0.0	0.0	10.0
SIS	0.02	1.0	199	0.1	0.0	0.0	5.0
SIR	0.02	1.0	199	0.1	0.0	0.0	5.0

this is a sensible decision, given the nonlinear relationship between the sequenced and unsequenced cases.

The total removal rate γ was fixed at 25 removals per infectious individual per year, corresponding to an expected infectious period of ~ 15 days. Similarly, the removal probability at sampling r was fixed to 0, meaning that sampling was not assumed to affect infectious potential. All other epidemiological parameters were estimated from the data. The complete list of prior distributions used for these analyses is presented in the second column of [table 2](#).

For the full analysis and the sequence-only analysis, a total of 30 independent MCMC chains were run for 2×10^7 steps each and compared with assess convergence. The initial 10% of each chain was removed to account for burn-in and the remaining samples combined into two long chains (one for each analysis type) from which the final results were derived.

The 95% HPD intervals for each of the estimated compartmental model parameters are presented in the right-most columns of [table 2](#). Interestingly, despite the broad uniform prior, the initial size of the susceptible population is inferred to be very low: on the order of one or two thousand individuals. This is likely due to the effects of population structure, with the fitted value representing the effective magnitude of the susceptible population rather than a demographic count. Additionally, we find that the overall rate of sampling is comparable to the removal rate γ , suggesting a relatively high sampling fraction $\psi/(\psi + \gamma)$ of 39–60% (95%HPD interval) during the period that sampling was taking place, that is, between the first and the last sample recorded for this region.

The posterior distributions for the absolute number of infectious hosts, $I(t)$, and effective reproduction number, $R_e(t) = \beta S(t)/\gamma$, trajectories are shown as the distributions of red curves in [figures 6a and b](#), respectively. The blue curves shown alongside are trajectories simulated under the model

using the sampled epidemiological parameter values and not explicitly conditioned on the observed sample data nor inferred transmission trees, hence their broader variance.

[Figure 6c](#) shows the posterior for the prevalence in terms of the number of infectious hosts per 10^5 initially susceptible hosts in the population. Since the SIR model is a constant population size model, this is also just the proportion of the population at any time which is inferred to be infected. Furthermore, since the initial number of susceptible hosts S_0 is jointly estimated, the shape of the estimated curve differs subtly from the absolute infected host count trajectories shown in [figure 6a](#) due to correlations between this shape and the susceptible host count.

[Figure 6d](#) shows the posterior for the rate of incidence. Specifically, it shows the inferred rate of new infections per susceptible host per week, with time measured in weeks.

The comparison between analysis of the full data set and the sequence-only analysis ([fig. 6e](#)) clearly displays the advantage of including the additional unsequenced case count data. In particular, it is clear that the unsequenced samples ([fig. 6f](#)) provide a wealth of information regarding the peak prevalence of the epidemic, a value that is almost completely unresolved in the sequence-only analysis.

Discussion

The primary strengths of the inference method and associated software presented here are their versatility and exactness. The method jointly samples from the exact posterior of transmission trees, epidemic trajectories, and model parameters under compartmental models without needing to make assumptions about the size of the epidemic or the size of the host population. (In contrast, coalescent methods are usually only applicable when population sizes are large.) The current implementation treats SI, SIS, and SIR epidemic models but, with only minor modifications, it can be used under any unstructured stochastic compartmental model whose dynamics can be described by [equation \(1\)](#).

There is also versatility in the type of data the method accepts. Many phylodynamic methods have relied solely on sequence data to inform their models which, while increasingly available, is more costly and scarce than simple case

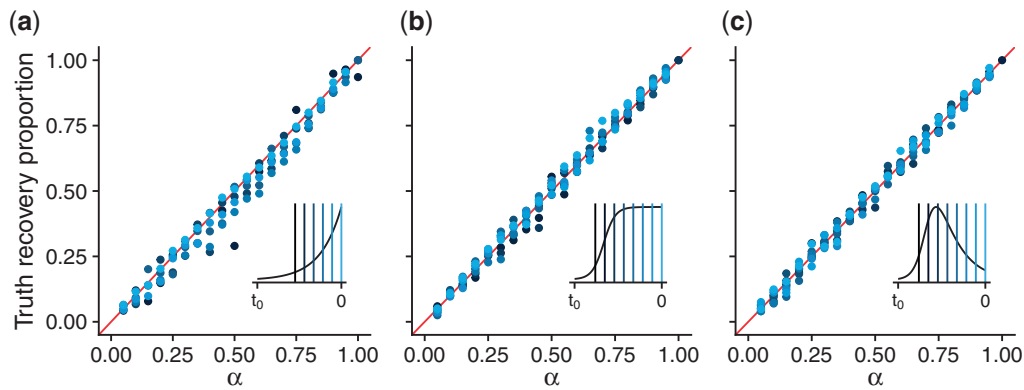


FIG. 5. Proportion of simulated data analyses which included the true prevalence in their 100% HPD intervals, for alignments simulated under each of the (a) linear birth–death, (b) SIS, and (c) SIR models. Colors represent the distinct times at which the coverage fractions were computed, and the insets indicate where these times fall in relation to the approximate deterministic prevalence curves. The linear relationship between the relative inclusion frequencies and α indicates that the PMMH algorithm is correctly sampling from the posterior prevalence distribution under each of these models.

reports. Our method can use cases reports and sequences together. The benefits of including case reports (unsequenced samples) to improving prevalence estimation are clearly shown in the Ebola analysis where the time of the epidemic peak is much more tightly estimated than when the sequences are analyzed alone. We also expect that including the case reports could inform the dating of the tree in data sets where the case reports are numerous and only a small number of sequences are available.

The method described here is also applicable to the field of macroevolution where past species richness, that is, the number of species through time, is a quantity of much interest. Estimates are typically obtained by using sequences from extant species to estimate past speciation and extinction rates which are then used to simulate unconditioned trajectories (Stadler and Bonhoeffer 2013). As is the case with epidemic trajectories, using our particle filtering tool to fit conditioned trajectories should improve these estimates and make quantification of species richness more precise. Fossil occurrence data have been shown to greatly improve macroevolutionary estimates (Gavryushkina et al. 2017) and are analogous to unsequenced samples, so can be directly incorporated into analyses with our method.

The sampling model we use is relatively simple, with infected samples uniformly taken at a constant rate through the epidemic and the possibility of burst of sampling at the end. This overly simple approach means that data needed to be discarded in the Ebola analysis so as not to bias results. It is feasible to extend the sampling model to more closely reflect how the data is actually collected, for example, by modeling changes in collection effort or having multiple bursts of intense sampling and so avoid potential biases introduced by the current model.

The software implementation of the method within the Beast 2 framework means that the default is to estimate the tree along with other parameters, and the full range of standard phylogenetic models can be used to model sequence evolution along the tree.

The flexibility and exactness of the inference relies on simulation to compute Monte Carlo estimates of the probability density of the transmission tree under the model and so comes at a heavy computational cost. Although a single density estimate can be made very quickly, when it is run as part of a larger MCMC analysis, estimates must be computed many times for each MCMC step and for hundreds of millions of steps. The number of simulations run at each step is a tunable parameter of PMMH and does not, in theory, alter the accuracy of the result. But there is a trade-off in that reducing the number of stochastic simulations that make up a density estimate increases the variance of the estimate with the result the Markov chain can become “stuck” after an extreme estimate is made, and the mixing rate of the chain is drastically reduced to the point that independent draws from the target posterior are not being produced. There is potential to parallelize the density estimate by running simulations in parallel at each step though with overheads the benefit of this may be marginal. Overall, joint analysis under this method is currently limited to hundreds of sequences.

Another obvious shortcoming of the present algorithm is its inability to handle structure in the population. Structure can originate from spatial segmentation of the host population or from the infection having distinct phases, for example, varying degrees of transmissibility or a noninfectious period (such as in the SEIR model). This issue is addressed by Rasmussen et al. (2014), although in an approximate way that assumes events in the epidemic trajectory are independent of the events observed in the phylogeny.

Despite these difficulties, we have presented what is to our knowledge the first algorithm capable of exactly inferring epidemiological trajectories jointly with compartmental model parameters using a combination of pathogen sequencing data and case count records. Our method also enables estimates of species richness through time by combining extant species data and fossil occurrences. A focus for future work will be extending this tool to account for population structure and to allow for the analysis of larger data sets in a mathematically exact framework.

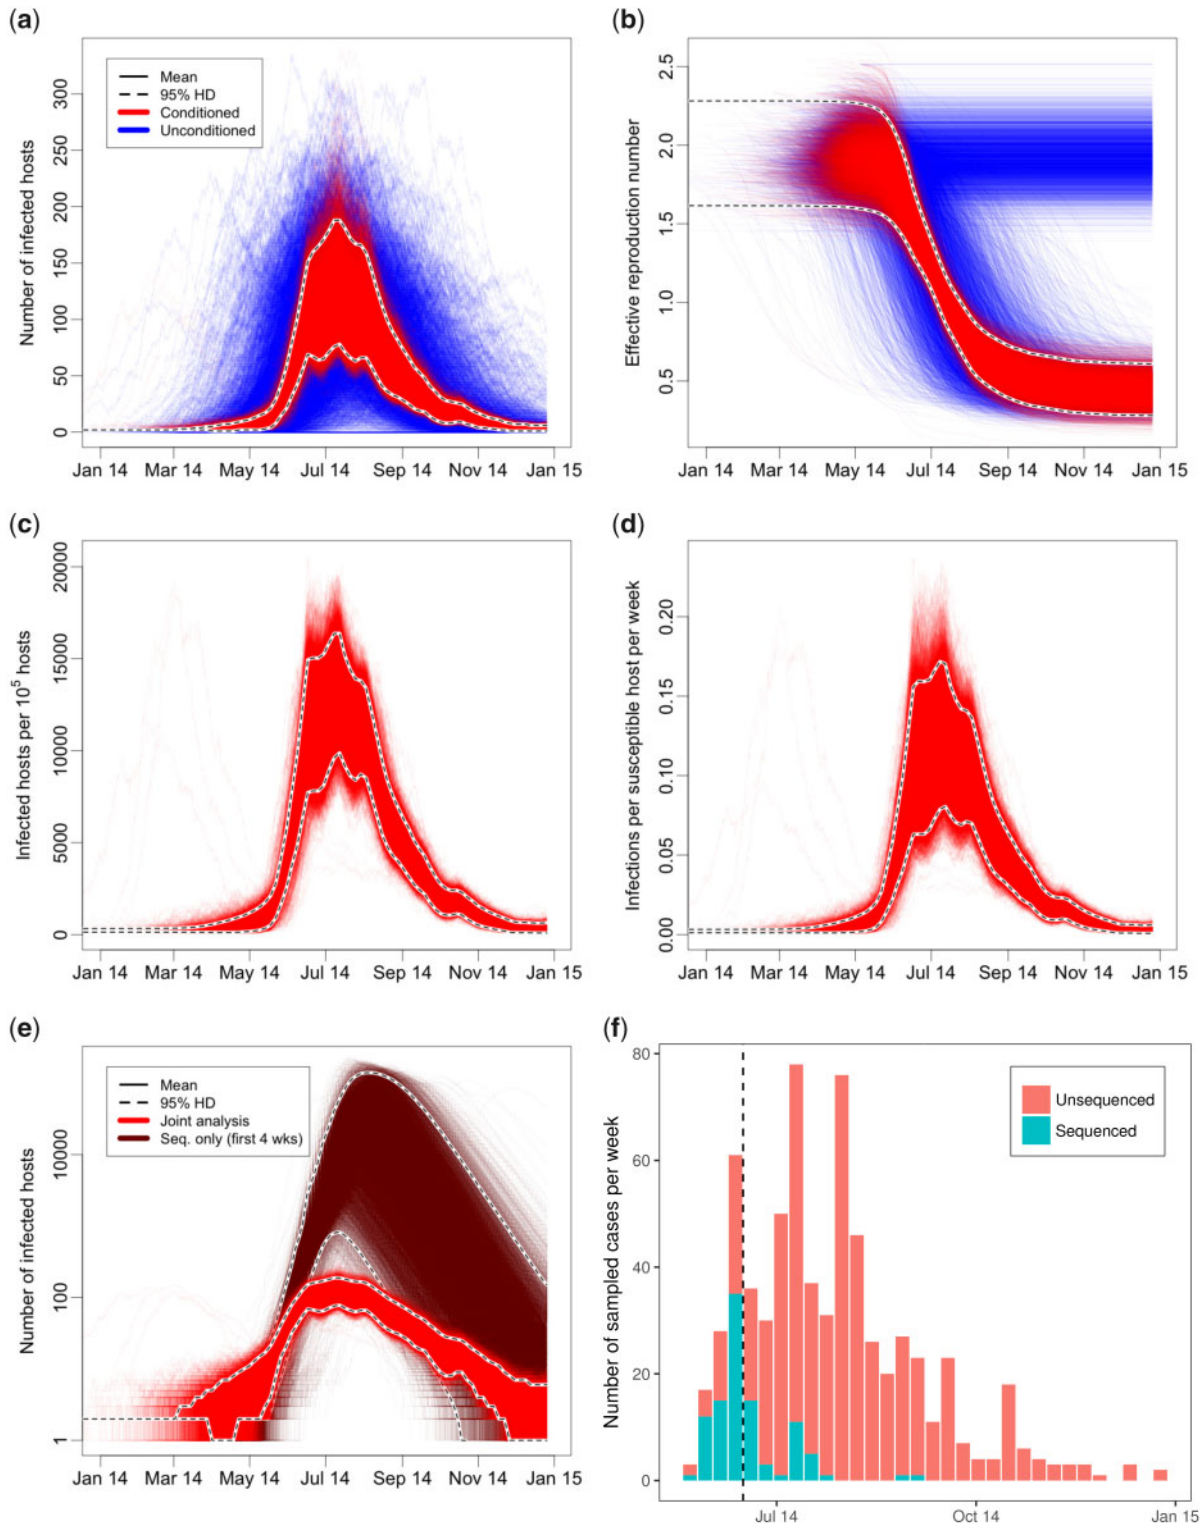


FIG. 6. (a, b) Jointly inferred posterior distributions (red) and unconditioned simulated distributions (blue) for (a) infected host count and (b) effective reproduction number during the Kailahun EVD outbreak. (c) Posterior distribution of infected host count per 10^5 hosts (prevalence). (d) Expected number of new EVD infections per susceptible host per week (incidence). (e) Comparison of inferred number of infected hosts using all data (red curves) and only the first 4 weeks of sequence data (brown curves). (f) Temporal distribution of EBOV cases used in the full analysis, both sequenced (turquoise) and unsequenced (orange). The vertical dashed line in (f) indicates the end of the 4-week period of sequence data used to infer the brown trajectories in (e).

Table 2. Parameter Priors Distributions Used in and 95% HPD Intervals Derived from Our Analysis of EBOV Genomes Sampled from the 2014 EVD Outbreak in Kailahun.

Parameter	Unit	Prior Distribution	Posterior 95% HPD	
			Lower	Upper
β	Year ⁻¹	Unif(0, 1)	2.9×10^{-2}	8.2×10^{-2}
S_0	—	Unif(0, 5×10^5)	576	1,390
ψ	Year ⁻¹	Unif(1, 365)	16	36
T	Year	Unif(0, 2)	0.64 (May 5)	0.83 (Feb 25)

NOTE.—While T is the time difference between the start of the outbreak and the end of the observation period, for a given time of cessation of observation it implies the absolute time of the start of the outbreak, which we provide in the bracketed (2014) dates.

Supplementary Material

Supplementary data are available at *Molecular Biology and Evolution* online.

Acknowledgments

The authors thank Louis du Plessis for helpful suggestions. They also thank the New Zealand eScience Infrastructure for access to high-performance computing facilities (<http://www.nesi.org.nz>). This work was supported by Marsden grant UOA1324 from the Royal Society of New Zealand. T.S. is supported in part by the European Research Council under the Seventh Framework Program of the European Commission (PhyPD: Grant Agreement No. 335529). G.E.L. was supported by the Swiss National Science Foundation (162251) and the Human Frontiers Science Program (LT000643/2016-L).

References

- Andrieu C, Doucet A, Holenstein R. 2010. Particle Markov chain Monte Carlo methods. *J R Stat Soc B* 72(3):269–342.
- Bell A, Lewandowski K, Myers R, Wooldridge D, Aarons E, Simpson A, Vipond R, Jacobs M, Gharbia S, Zambon M, et al. 2015. Genome sequence analysis of Ebola virus in clinical samples from three British healthcare workers, August 2014 to March 2015. *Eurosurveillance* 20(20):21131.
- Boskova V, Bonhoeffer S, Stadler T. 2014. Inference of epidemiological dynamics based on simulated phylogenies using birth–death and coalescent models. *PLoS Comput Biol*. 10(11):e1003913.
- Bouckaert R, Heled J, Kühnert D, Vaughan T, Wu C-H, Xie D, Suchard MA, Rambaut A, Drummond AJ. 2014. BEAST 2: a software platform for Bayesian evolutionary analysis. *PLoS Comput Biol*. 10(4):e1003537.
- Carroll MW, Matthews DA, Hiscox JA, Elmore MJ, Pollakis G, Rambaut A, Hewson R, García-Dorival I, Bore JA, Koundouno R, et al. 2015. Temporal and spatial analysis of the 2014–2015 Ebola virus outbreak in West Africa. *Nature* 524(7563):97–101.
- Dawid AP. 1982. The well-calibrated Bayesian. *J Am Stat Assoc*. 77(379):605–610.
- Del Moral P. 2004. Feynman-Kac formulae (Hb). New York: Springer-Verlag.
- Drummond AJ, Pybus OG, Rambaut A, Forsberg R, Rodrigo AG. 2003. Measurably evolving populations. *Trends in Ecology & Evolution*. 18(9):481–488.
- Drummond AJ, Rambaut A, Shapiro B, Pybus OG. 2005. Bayesian coalescent inference of past population dynamics from molecular sequences. *Mol Biol Evol*. 22(5):1185–1192.
- Dudas G, Carvalho LM, Bedford T, Tatem AJ, Baele G, Faria NR, Park DJ, Ladner JT, Arias A, Asogun D, et al. 2017. Virus genomes reveal factors that spread and sustained the Ebola epidemic. *Nature* 544(7650):309–315.
- Gavryushkina A, Heath TA, Ksepka DT, Stadler T, Welch D, Drummond AJ. 2017. Bayesian total-evidence dating reveals the recent crown radiation of penguins. *Syst Biol*. 66(1):57–73.
- Gavryushkina A, Welch D, Stadler T, Drummond AJ. 2014. Bayesian inference of sampled ancestor trees for epidemiology and fossil calibration. *PLoS Comput Biol*. 10(12):e1003919.
- Gillespie DT. 1976. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J Comp Phys*. 22(4):403.
- Gillespie DT. 1977. Stochastic simulation of coupled chemical reactions. *J Phys Chem*. 81(25):2340.
- Gillespie DT. 2001. Approximate accelerated stochastic simulation of chemically reacting systems. *J Chem Phys*. 115(4):1716.
- Gire SK, Goba A, Andersen KG, Sealfon RSG, Park DJ, Kanneh L, Jalloh S, Momoh M, Fullah M, Dudas G, et al. 2014. Genomic surveillance elucidates Ebola virus origin and transmission during the 2014 outbreak. *Science* 345(6202):1369–1372.
- Grenfell BT, Pybus OG, Gog JR, Wood JLN, Daly JM, Mumford JA, Holmes EC. 2004. Unifying the epidemiological and evolutionary dynamics of pathogens. *Science* 303(5656):327–332.
- Hasegawa M, Kishino H, Yano T. 1985. Dating of the human–ape splitting by a molecular clock of mitochondrial DNA. *J Mol Evol*. 22(2):160–174.
- Holmes EC, Dudas G, Rambaut A, Andersen KG. 2016. The evolution of Ebola virus: insights from the 2013–2016 epidemic. *Nature* 538(7624):193–200.
- Kermack WO, McKendrick AG. 1927. A contribution to the mathematical theory of epidemics. *Proc R Soc Lond A* 115(772):700.
- Kingman J. 1982. The coalescent. *Stoch Proc Appl*. 13(3):235–248.
- Kühnert D, Stadler T, Vaughan TG, Drummond AJ. 2014. Simultaneous reconstruction of evolutionary history and epidemiological dynamics from viral sequences with the birth–death sir model. *J R Soc Interface* 11(94):20131106.
- Leventhal GE, Günthard HF, Bonhoeffer S, Stadler T. 2014. Using an epidemiological model for phylogenetic inference reveals density dependence in HIV transmission. *Mol Biol Evol*. 31(1):6–17.
- Li LM, Grassly NC, Fraser C. 2017. Quantifying transmission heterogeneity using both pathogen phylogenies and incidence time series. *Mol Biol Evol*. 34(11):2982–2995.
- Park DJ, Dudas G, Wohl S, Goba A, Whitmer SLM, Andersen KG, Sealfon RS, Ladner JT, Kugelman JR, Matranga CB, et al. 2015. Ebola virus epidemiology, transmission, and evolution during seven months in Sierra Leone. *Cell* 161(7):1516–1526.
- Pybus OG, Charleston MA, Gupta S, Rambaut A, Holmes EC, Harvey PH. 2001. The epidemic behavior of the Hepatitis C Virus. *Science* 292(5525):2323–2325.
- Pybus OG, Rambaut A, Harvey PH. 2000. An integrated framework for the inference of viral population history from reconstructed genealogies. *Genetics* 155(3):1429–1437.
- Rasmussen DA, Ratmann O, Koelle K. 2011. Inference for nonlinear epidemiological models using genealogies and time series. *PLoS Comput Biol*. 7(8):e1002136.
- Rasmussen DA, Volz EM, Koelle K. 2014. Phylodynamic inference for structured epidemiological models. *PLoS Comput Biol*. 10(4):e1003570.
- Smith RA, Ionides EL, King AA. 2017. Infectious disease dynamics inferred from genetic data via sequential Monte Carlo. *Mol Biol Evol*. 34(8):2065–2084.
- Stadler T. 2010. Sampling-through-time in birth–death trees. *J Theor Biol*. 267(3):396–404.
- Stadler T, Bonhoeffer S. 2013. Uncovering epidemiological dynamics in heterogeneous host populations using phylogenetic methods. *Philos Trans R Soc Lond B Biol Sci*. 368(1614):20120198.
- Stadler T, Kouyos R, von Wyl V, Yerly S, Böni J, Bürgisser P, Klimkait T, Joos B, Rieder P, Xie D, et al. 2012. Estimating the basic reproductive number from viral sequence data. *Mol Biol Evol*. 29(1):347–357.

- Stadler T, Kühnert D, Bonhoeffer S, Drummond AJ. 2013. Birth–death skyline plot reveals temporal changes of epidemic spread in HIV and Hepatitis C Virus (HCV). *Proc Natl Acad Sci U S A*. 110(1):228–233.
- Stadler T, Vaughan TG, Gavryushkin A, Guindon S, Kühnert D, Leventhal GE, Drummond AJ. 2015. How well can the exponential-growth coalescent approximate constant-rate birth–death population dynamics? *Proc Biol Sci*. 282(1806):20150420.
- Volz E, Siveroni I. 2018. Bayesian phylodynamic inference with complex models. *PLoS Comput Biol*. 14(11):e1006546.
- Volz EM. 2012. Complex population dynamics and the coalescent under neutrality. *Genetics* 190(1):187–201.
- Volz EM, Kosakovsky Pond SL, Ward MJ, Leigh Brown AJ, Frost SDW. 2009. Phylodynamics of infectious disease epidemics. *Genetics* 183(4):1421–1430.
- World Health Organization. 2016. Sierra Leone Ebola case data (Ebola data and statistics website, <http://apps.who.int/gho/data/node.ebola-sitrep>, last accessed September 1, 2017), Geneva, Switzerland.
- Yang Z. 1994. Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: approximate methods. *J Mol Evol*. 39(3):306–314.