

Adaptive Preferential Sampling in Phylodynamics With an Application to SARS-CoV-2

Lorenzo Cappello & Julia A. Palacios

To cite this article: Lorenzo Cappello & Julia A. Palacios (2021): Adaptive Preferential Sampling in Phylodynamics With an Application to SARS-CoV-2, Journal of Computational and Graphical Statistics, DOI: [10.1080/10618600.2021.1987256](https://doi.org/10.1080/10618600.2021.1987256)

To link to this article: <https://doi.org/10.1080/10618600.2021.1987256>



© 2021 The Author(s). Published with license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 29 Nov 2021.



[Submit your article to this journal](#)



Article views: 338



[View related articles](#)



[View Crossmark data](#)

Adaptive Preferential Sampling in Phylodynamics With an Application to SARS-CoV-2

Lorenzo Cappello^a and Julia A. Palacios^{a,b}

^aDepartment of Statistics, Stanford University, Stanford, CA; ^bDepartment of Biomedical Data Science, Stanford Medicine, Stanford, CA

ABSTRACT

Longitudinal molecular data of rapidly evolving viruses and pathogens provide information about disease spread and complement traditional surveillance approaches based on case count data. The coalescent is used to model the genealogy that represents the sample ancestral relationships. The basic assumption is that coalescent events occur at a rate inversely proportional to the effective population size $N_e(t)$, a time-varying measure of genetic diversity. When the sampling process (collection of samples over time) depends on $N_e(t)$, the coalescent and the sampling processes can be jointly modeled to improve estimation of $N_e(t)$. Failing to do so can lead to bias due to model misspecification. However, the way that the sampling process depends on the effective population size may vary over time. We introduce an approach where the sampling process is modeled as an inhomogeneous Poisson process with rate equal to the product of $N_e(t)$ and a time-varying coefficient, making minimal assumptions on their functional shapes via Markov random field priors. We provide efficient algorithms for inference, show the model performance vis-a-vis alternative methods in a simulation study, and apply our model to SARS-CoV-2 sequences from Los Angeles and Santa Clara counties. The methodology is implemented and available in the R package `adaptref`. Supplementary files for this article are available online.

ARTICLE HISTORY

Received September 2020
Revised July 2021

KEYWORDS

Coalescent process; Markov random fields; Poisson processes; Population size

1. Introduction

Molecular sequence data, within the framework of phylodynamics (Grenfell et al. 2004), is increasingly being used to track disease spread caused by rapidly evolving viruses and pathogens such as Influenza viruses (Rambaut et al. 2008), Zika (Faria et al. 2016), and SARS-CoV-2 (Hadfield et al. 2018). The coalescent process (Kingman 1982a, 1982b), a probability model of gene genealogies, depends on a parameter called effective population size $N_e(t)$, which is a time-varying measure of genetic diversity. When disease dynamics can be modeled by simple epidemiological models such as Susceptible-Infected-Recovered, the coalescent effective population size can be expressed in terms of transmission rates and prevalence (Volz et al. 2009; Frost and Volz 2010). Accurate and efficient inference for $N_e(t)$ is thus relevant to estimate epidemiological parameters of great interest in public health. Although this work is motivated by applications in molecular epidemiology of infectious diseases, estimation of $N_e(t)$ is an active area of research with applications ranging across many other scientific domains such as conservation biology and population genetics (e.g., Shapiro et al. 2004; Huff et al. 2010; Lorenzen et al. 2011).

A common feature in these applications is that genetic data are collected sequentially (heterochronous samples). In viral studies, samples are collected and sequenced when infected individuals attend clinics, hospitals, or testing centers. In ancient DNA studies, specimens are dated according to the time they lived, estimated through radiocarbon dating or

other techniques. The coalescent typically models the gene genealogy conditionally on sampling dates, that is, the sampling dates are treated as censoring information (Felsenstein and Rodrigo 1999). However, in some situations, it is reasonable to assume that samples are collected at a higher frequency when the population is large and at a lower frequency when the population is small: for example, at the onset of an epidemic, as the viral population grows and more people get infected, more resources may be allocated to monitor the viral spread, possibly leading to more molecular sequence collection. The number of SARS-CoV-2 sequences uploaded daily in GISAID offers some evidence of this claim (Shu and McCauley 2017) (see the histogram in the supplementary material).

Karcher et al. (2016) study the scenario in which the sampling process depends on the population size, and show that an estimator of $N_e(t)$ that does not account for this dependence is biased. This issue was first discussed in the spatial statistics literature by Diggle, Menezes and Su (2010), who term preferential sampling a situation in which the process that determines the data locations and the process under study are dependent. In this article, we will introduce a new model that accounts for preferential sampling in a coalescent framework, while making minimal assumptions on $N_e(t)$, the sampling process, and their dependence.

Three estimators that incorporate preferential sampling into the coalescent framework have been proposed. Volz and Frost (2014) propose an estimator in the case that $N_e(t)$ grows

CONTACT Julia A. Palacios  juliapr@stanford.edu  Department of Statistics, Stanford University, 390 Serra Mall, Sequoia Hall, Stanford, CA.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JCGS.

© 2021 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

exponentially and samples are collected as an inhomogeneous Poisson process with rate linearly dependent on the effective population size. Karcher et al. (2016) assumed that $N_e(t)$ is a continuous function, and the samples are collected as an inhomogeneous Poisson process with rate $\lambda(t) = \exp(\beta_0)N_e(t)^{\beta_1}$, for $\beta_0, \beta_1 \geq 0$, that is, the dependence between the sampling process and the effective sample size is described by a parametric model. Recent work by Parag, du Plessis, and Pybus (2020) weakened this assumption substantially, allowing for the dependence between the sampling process and effective population size to vary over time. The key assumption in Parag, du Plessis, and Pybus (2020) is that the sampling rate depends linearly on $N_e(t)$ within a given time interval, but the linear coefficient changes across time intervals. The estimator of Parag, du Plessis, and Pybus (2020), termed Epoch skyline plot (ESP), is an extension of the classic skyline plot estimator for $N_e(t)$ (Pybus, Rambaut, and Harvey 2000), in which the sampling rate and the effective population size are both piecewise-constant, and the location and number of change points (boundary points of the time intervals) are either specified or inferred. As it is typical with skyline plots, the estimates are highly variable, rough, and highly dependent on the specification of change points locations and the number of piecewise-constant pieces used. All three works show that under correct model specification, accounting for preferential sampling leads to a more accurate estimation of $N_e(t)$ (in terms of absolute deviations to the true trajectory), and narrower credible regions.

Other noncoalescent approaches for phylodynamics, such as birth–death processes (Stadler 2010), explicitly incorporate the sampling process by modeling the sampling dates as a partially observed death process where only a fraction of the population is observed. Stadler et al. (2013) extended previous work to allow sampling rates to vary through time. Volz and Frost (2014) show that in both, coalescent and birth-death processes alike, statistical power largely depends on the correct specification of the sampling process rate, rather than on the genealogical model. Hence, the need for a flexible modeling approach of the sampling process, adaptive to any possible scenarios encountered in applications.

There are a plethora of nonparametric estimators of $N_e(t)$ following the skyline plot (Pybus, Rambaut, and Harvey 2000). Among others, the generalized skyline plot (Strimmer and Pybus 2001) and the Bayesian skyline plot (Drummond et al. 2005) reduce the high variance and roughness that characterize the skyline plot estimators. These methodologies require either fixing or estimating change-points in $N_e(t)$. A set of models that do not employ change-points but arbitrary discrete grids is based on Markov random fields (MRF): the Gaussian MRF (GMRF) (Minin, Bloomquist and Suchard 2008; Palacios and Minin 2012) allows for the recovery of smooth continuous trajectories; the Horseshoe MRF (HSMRF) (Faulkner et al. 2020) is an alternative to GMRF which is locally adaptive, that is, it can successfully recover sharp changes in a trajectory and it is adaptive to a varying level of smoothness.

In this article, we borrow from this literature and introduce an adaptive preferential sampling framework for phylodynamics, where the adaptivity follows from the fact that the dependence of the sampling process on $N_e(t)$ changes over time. The effective population size $N_e(t)$ is modeled as a latent parameter

included in both, the coalescent and the sampling processes. The latter is assumed to be an inhomogeneous Poisson process with rate $\lambda(t) = \beta(t)N_e(t)$, where $\beta(t)$ is a continuous function controlling the dependence on $N_e(t)$, analogous to that introduced by Parag, du Plessis, and Pybus (2020). We a priori model $N_e(t)$ and $\beta(t)$ as two Markov random fields (MRF), with the flexibility of using either a GMRF or a HSMRF. The prior choice follows from the properties of the fields. The advantage of the proposed adaptive preferential sampling over the ESP estimator is that there is no need to specify (or estimate) the number and location of the change points of β and N_e . Also, the resulting estimates are smooth and the high variability that characterizes skyline estimates disappears.

We develop the methodology assuming that a genealogy is available to the researcher and develop algorithms for inference under this framework. We test our model on simulated data and compare it to alternative methods, including both, estimators that account for preferential sampling and others that do not. We implement our method in the R package `adapref`, available at <https://github.com/lorenzocapp/adapref>, provide two algorithms for posterior approximation: a Hamiltonian MCMC and a Laplace approximation. We apply our method to SARS-CoV-2 sequences from California and study whether there is evidence of preferential sampling.

The rest of the article proceeds as follows. In Section 2, we provide background on the coalescent process, the MRF priors on $N_e(t)$, and previous work on preferential sampling. In Section 3 we introduce the adaptive preferential sampling framework and explain how to approximate the posterior distribution of model parameters. Section 4 includes a simulation study, in which we test our proposal through simulated data and compare it to alternatives. In Section 5, we apply our method to two datasets of SARS-CoV-2 sequences from Santa Clara and Los Angeles counties in California. Section 6 concludes.

2. Background

2.1. Coalescent Model

Coalescent models are continuous-time Markov chains used to model the set of ancestral relationships of a sample of n individuals from a large population. This set of ancestral relationships is called gene genealogy. In the context of molecular epidemiology, a genealogy is a subset of the transmission history among the samples (Figure 1). Starting from the original work of Kingman (1982a), several extensions to the standard coalescent have been developed to incorporate more realistic population and sampling features, such as variable population size (Slatkin and Hudson 1991), longitudinal sampling (also called heterochronous sampling) (Felsenstein and Rodrigo 1999), and population structure (Hudson 1990). Wakeley (2009) provided a good introduction to the subject. Coalescent processes can be characterized by two underlying processes: a jump chain defining the ancestral relationships represented by a binary tree topology and a pure death process that defines the timing of the coalescent events, that is, the times when pairs of lineages meet their common ancestors. This sequence of holding times defines the branch lengths of the corresponding tree topology.

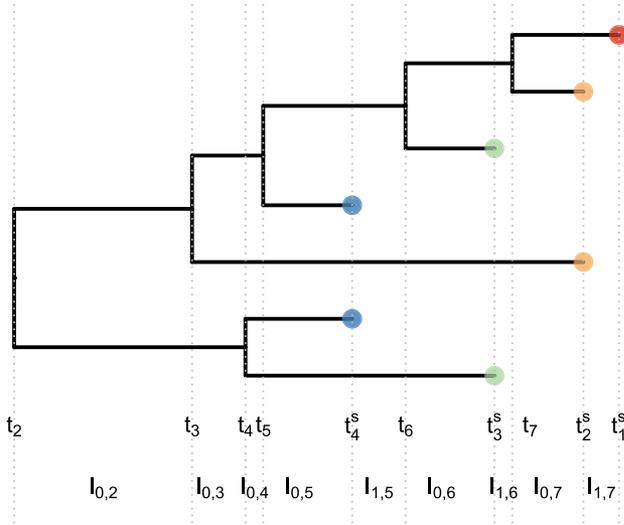


Figure 1. Example of a heterochronous genealogy. A genealogy of 7 individuals sampled at 4 different times (color of tips) with multiplicities ($n_1 = 1, n_2 = 2, n_3 = 2, n_4 = 2$). Sampling times are denoted by $(t_k^s)_{1,4}$, coalescent times are denoted by $(t_k)_{2,7}$ and I_{ij} denoted the interval lengths delimited by coalescent times and/or sampling times, that is, every time there is a change in the number of lineages.

Let the vector $\mathbf{n} = (n_1, \dots, n_m)$ denote the sample sizes at times $\mathbf{t}^s = (t_1^s, \dots, t_m^s)$, with m number of sampling points and n the total sample size. The process goes backward in time (from present toward the past): with $t_1^s = 0$ denoting the present time, and $t_j^s > t_{j-1}^s$ for $j = 2, \dots, m$. Let $\mathbf{t} = (t_{n+1}, \dots, t_2)$ be the vector of coalescent times with $t_{n+1} = 0 < t_n < \dots < t_2$. Note that the subscript in t_k is not the current number of extant lineages (often a convention in the coalescent literature) but the number of lineages that have yet to coalesce. Starting from $t = 0$, the vectors \mathbf{t} and \mathbf{t}^s partition time into intervals (Figure 1). An interval ending with a coalescent event, say t_k , is denoted by $I_{0,k}$; the intervals that end with a sampling time within the interval (t_{k+1}, t_k) are denoted as $I_{i,k}$, where $i \geq 1$ indexes all the sampling events in (t_{k+1}, t_k) . Formally, for every $k \in \{2, \dots, n\}$, we define

$$I_{0,k} = [\max\{t_{k+1}, t_j^s\}, t_k), \quad \text{where the maximum is taken over all } t_j^s < t_k,$$

and for every $i \geq 1$ we set

$$I_{i,k} = [\max\{t_{k+1}, t_{j-i}^s\}, t_{j-i+1}^s) \text{ with the max taken over all } t_{j-i+1}^s > t_{k+1} \text{ and } t_j^s < t_k.$$

With $n_{i,k}$ denoting the number of extant lineages during the time interval $I_{i,k}$, Figure 1 plots an example of a heterochronous genealogy with $\mathbf{n} = (1, 2, 2, 2)$, at times $\mathbf{t}^s = (t_1^s, \dots, t_4^s)$ with $t_1^s = 0$. In the interval (t_6, t_5) there are two intervals: $I_{1,5} = [t_6, t_4^s)$, $I_{0,5} = [t_4^s, t_5)$.

The vector of coalescent times \mathbf{t} is a random vector whose density with respect to the Lebesgue measure on \mathbb{R}_+^{n-1} depends on two quantities: the coalescent factor $C_{i,k} := \binom{n_{i,k}}{2}$, and the effective population size $N_e(t)$. The coalescent density can be factorized as the product of the conditional densities of t_{k-1} given t_k , that is,

$$p(\mathbf{t}^s, \mathbf{n}, N_e(t)) = \prod_{k=n+1}^3 p(t_{k-1}|t_k, \mathbf{t}^s, \mathbf{n}, N_e(t)) = \prod_{k=n+1}^3 \frac{C_{0,k-1}}{N_e(t_{k-1})} \exp \left\{ - \int_{I_{0,k-1}} \frac{C_{0,k-1}}{N_e(t)} dt + \sum_{i=1}^m \int_{I_{i,k-1}} \frac{C_{i,k-1}}{N_e(t)} dt \right\}, \quad (1)$$

where the conditional density $p(t_{k-1}|t_k, \mathbf{t}^s, \mathbf{n}, N_e(t))$ is the extension of Kingman n -coalescent holding time density to account for variable effective population size and heterochronous samples (Felsenstein and Rodrigo 1999). The integral over $I_{i,k-1}$ accounts for the probability of no coalescence during $I_{i,k-1}$. It is zero if there are less than i sampling times between t_k and t_{k-1} , and t_k appears in the definition of the intervals (either $I_{0,k-1}$ or $I_{1,k-1}$). In addition, conditionally on \mathbf{t}^s, \mathbf{n} , and t_k , the coalescent factors can be computed exactly and $N_e(t)$ is the only unknown parameter. Sampling times are assumed fixed.

Coalescent times can be alternatively viewed as the realization of an inhomogeneous point process with rate $C(t)N_e(t)^{-1}$, with the coalescent factor $C(t)$ being defined for all $t \geq 0$ by the notation above. This alternative view allows us to frame the problem of inferring $N_e(t)$ as that of inferring the intensity function of an inhomogeneous point process. Palacios and Minin (2013) is an example of how this representation is useful in inference and simulations.

2.2. Some Priors for the Effective Population Size

Markov random field-based priors on the log effective population trajectory allows the recovery of smooth trajectories. They are computationally tractable thanks to the sparsity assumption in the covariance matrix of the field (Rue and Held 2005). All MRF-based priors for phylodynamic inference share the assumption that the trajectory $N_e(t)$ is an unknown continuous function. The integral in Equation (1) is numerically approximated by the Riemann sum at a regular grid of $M + 1$ points $(k_i)_{1:M+1}$, and one assumes that the trajectory $N_e(t)$ is well approximated by $\sum_{i=1}^{M+1} \exp \theta_i 1(t \in (k_i, k_{i+1}))$, with $\boldsymbol{\theta} = (\theta_i)_{1:M}$. We stress that neither the grid cell boundaries $(k_i)_{1:M+1}$ nor M depend on \mathbf{t} and $N_e(t)$, with the choice of M commonly based on n (Faulkner et al. 2020). A description of the discretized coalescent log-likelihood $\mathcal{L}(\boldsymbol{\theta}|\mathbf{t})$ is given in detail in Palacios and Minin (2012) and Faulkner et al. (2020).

The Horseshoe Markov random field prior (HSMRF) for $\boldsymbol{\theta}$ (Faulkner et al. 2020) assumes that the p th-order forward differences of $\boldsymbol{\theta}$ are independent and Horseshoe distributed (Carvalho, Polson, and Scott 2010), that is,

$$\Delta^p \theta_i | \tau_i \sim N(0, \tau_i^2) \quad \tau_i | \gamma \sim C^+(0, \gamma) \\ \gamma | \zeta \sim C^+(0, \zeta) \quad \text{for } p+1 \leq i \leq M-1, \quad (2)$$

where $C^+(0, a)$ is the standard half-Cauchy distribution with positive support with scale parameter a , τ_i are the local shrinkage parameters and γ is the global smoothing parameter. To completely specify the prior, one sets $\theta_1 \sim N(\mu, \sigma_1^2)$, and for $p \geq 2$, the first p values of the field have running order q difference priors as follows:

$$\Delta^q \theta_q | a_q \tau_q \sim N(0, a_q \tau_q^2) \quad \tau_q | \gamma \sim C^+(0, \gamma) \\ \gamma | \zeta \sim C^+(0, \zeta) \quad \text{for } 1 \leq q \leq p-1,$$

with $a_q = 2^{-(p-q)/2}$. As it is common in the trend filtering literature (Kim et al. 2009), only orders 1 and 2 are typically employed in applications.

A related prior consists in assuming that the p th-order forward difference of θ , more precisely the vector $(\theta_1, \Delta^1\theta_1, \dots, \Delta^p\theta_p, \Delta^p\theta_{M-1})$ is distributed as a GMRF

$$\begin{aligned} \Delta^p\theta_i|\gamma &\sim N(0, \gamma^2) \quad \gamma|\zeta \sim C^+(0, \zeta) \\ \text{for } p \leq i \leq M-1, \end{aligned} \quad (3)$$

$\theta_1 \sim N(\mu, \sigma_1^2)$, and for $1 \leq q \leq p-1$ we set $\Delta^q\theta_q|a_q\gamma \sim N(0, a_q\gamma^2)$. That is, the vector $(\theta_1, \Delta^1\theta_1, \dots, \Delta^p\theta_p, \Delta^p\theta_{M-1})$ is normally distributed with a covariance matrix that depends on the order p considered. A common alternative to the half-Cauchy distribution on γ is a Gamma prior, for example, as in Palacios and Minin (2012). We will employ both formulations in our implementations.

A fully nonparametric prior on $\log N_e(t)$ has been studied by Palacios and Minin (2013), who proposed a Gaussian process prior on the log effective population size. The advantage of this approach is that no grid needs to be specified a priori. In applications, we believe that the GMRF, the discretized version of this prior, achieves a comparable empirical performance.

2.3. Preferential Sampling

Preferential sampling arises when the process that determines the locations of the data (i.e., sampling process) and the process under study are stochastically dependent. The notion was introduced by Diggle, Menezes and Su (2010) who show that not accounting for this effect leads to biased inference as a result of the model misspecification. On the other hand, a correctly specified sampling model can lead to more accurate estimates.

In phylodynamics, preferential sampling arises when the sampling process depends on $N_e(t)$. Volz and Frost (2014) provide the first evidence that coalescent-based inference under a misspecified sampling process can be biased. They propose a new estimator tailored to a coalescent process with exponentially growing effective population size and a sampling process with rate linearly dependent on $N_e(t)$. They show that the estimator obtained by correctly modeling the sampling process is more accurate than the standard coalescent estimator.

Karcher et al. (2016) assumed that $N_e(t)$ is a continuous function and the sampling process is a Poisson process with rate $\lambda(t) = \exp(\beta_0)N_e(t)^{\beta_1}$, for $\beta_0, \beta_1 \geq 0$, that is, the rate $\lambda(t)$ is proportional to the effective population size. This model is parsimonious, capturing a variety of scenarios with two parameters: with $\beta_1 = 1$, the rate is a constant times $N_e(t)$, on the opposite side of the spectrum, with $\beta_1 = 0$, one models uniform sampling. Another advantage is that little assumptions are made on $N_e(t)$. However, the parametric assumptions on $\lambda(t)$ make the sampling dates strongly informative about $N_e(t)$. This situation can be problematic in the case of sampling dates errors. Moreover, under no preferential sampling or under a different rate $\lambda(t)$ (model misspecified), it constitutes a relevant model misspecification, and leads to estimation biases; see, for example, Figure 3 in Section 4. Karcher et al. (2020) addressed some of the limitations of the parametric model by including time-varying covariates into the Poisson process rate: $\lambda(t) =$

$\exp(\beta_0)N_e(t)^{\beta_1} + \beta' \mathbf{X}(t)$, where \mathbf{X} is a vector of covariates and β' the corresponding linear coefficients. Here a covariate can be for example a dummy variable indicating a change in sampling protocols, or when a new sampling center joined the study. The term $\beta' \mathbf{X}$ adds more flexibility to the parametric dependence enforced by $\exp(\beta_0)N_e(t)^{\beta_1}$. Clearly, this extension requires the availability of covariates informative on the sampling design.

Parag, du Plessis, and Pybus (2020) introduced the epoch sampling skyline plot (ESP) estimator that allows for a more flexible dependence of the sampling process on the effective population size. More specifically, the ESP method assumes that $N_e(t)$ is a piecewise-constant function with r segments described by the vector (N_1, \dots, N_r) of r parameters, and time is further partitioned in d epochs, such that in epoch i and segment r the sampling process is a Poisson process with rate $\beta_i N_j$, where $(\beta_1, \dots, \beta_d)$ is a vector of d parameters. The vector $(\beta_1, \dots, \beta_d)$ modulates the dependence of the sampling process on the effective population size, assuming that the dependence changes across d epochs. This is a notable advantage over the parametric model of Karcher et al. (2016): one can model a variety of realistic time-varying sampling protocols, or simply deal with sampling discontinuities typical of outbreaks. We conjecture that the higher flexibility in the ESP preferential model reduces the risk of model misspecification and bias when the preferential sampling assumption does not hold. Furthermore, the model remains identifiable: the authors show that the expected Fisher information matrix is nonsingular, a sufficient condition for parameter identifiability in exponential families (Rothenberg 1971).

In the ESP, the endpoints of the r segments coincide with a subset of the coalescent times \mathbf{t} . Similarly, the boundary points of the d epochs are determined by a subset of the sampling times. The number of segments r and epochs d , as well as their lengths, need to be determined or inferred. A downside of this approach is that these choices affect the ESP estimates heavily. The authors implement a frequentist and a Bayesian version with independence assumptions in (N_1, \dots, N_p) and $(\beta_1, \dots, \beta_d)$, leading to estimates with high variance, a characteristic feature of skyline plot-type estimators.

3. Adaptive Preferential Sampling

In the adaptive preferential sampling framework, the sampling times are determined by the jumps of an inhomogenous Poisson process with rate $\lambda(t) = \beta(t)N_e(t)$, with $N_e(t)$ effective population size, and $\beta(t)$, a function modulating the linear dependence between $\lambda(t)$ and $N_e(t)$. Let s_i denote the i th sampling time generated by this process, for $i = 1, \dots, n$ and s_1 fixed to 0. The notation $(s_i)_{1:n}$ is added to highlight the difference between a realization from the Poisson simple point process $(s_i)_{1:n}$ of one event at a time, and the observed vector \mathbf{t}^s introduced in Section 2 in which multiple events can be observed at a given time point. As we will show later in this section, this discrepancy is resolved by discretizing the observed time window into regular intervals. We assume that both $\beta(t)$ and $N_e(t)$ are unknown continuous functions. To numerically approximate the integrals in (1), we resort to the approximation sketched in Section 2.2 and detailed in Palacios and Minin (2012). We employ the regular

grid $(k_i)_{1:M+1}$ and assume that $N_e(t)$ is governed by parameters $\theta = (\theta_i)_{1:M'}$. Similarly, to model the time-varying rate $\lambda(t)$, we assume that $\beta(t)$ is governed by parameters $\alpha = (\alpha_i)_{1:M'}$, where $M' = \min_i\{k_{i+1} : k_{i+1} > s_n\}$ and $\beta(t) \approx \exp\{\alpha_i\}$ for $t \in (k_i, k_{i+1}]$.

By the independence of nonoverlapping intervals of the Poisson process, we can write the log-likelihood contribution of the sampling process as

$$\mathcal{L}(\alpha, \theta | \mathbf{s}) = \sum_{i=1}^{M'} \left[\left[\{s_i : s_i \in (k_i, k_{i+1}]\} \mid (\alpha_i + \theta_i + \log \Delta_i) - \exp\{\alpha_i \theta_i\} \Delta_i \right], \quad (4)$$

where $\Delta_i = k_{i+1} - k_i$ and the first interval $[k_1, k_2]$ is closed to include s_1 . Equation (4) is the log-likelihood of M' independent Poisson distributed random variables with rates $(\exp\{\alpha_i \theta_i\} \Delta_i)_{1:M'}$ and number of events $\left(\left[\{s_i : s_i \in (k_i, k_{i+1}]\} \right] \right)_{1:M'}$.

A few preliminary remarks. M' is not related to the number of epochs d in ESP: it is solely determined by s_n and the grid $(k_i)_{1:M+1}$, which in turn does not depend on \mathbf{t} . We have by definition $M' < M$ to ensure that $\beta(t)$ is not modeled after the last sampling time (there is no information in the sample to estimate it). In addition, through the term $\left[\{s_i : s_i \in (k_i, k_{i+1}]\} \right]$, the discretized log-likelihood (4) allows to naturally account for multiple sampling times collected at once, reconciling this model with the description of the heterochronous coalescent given in Section 2.1. Lastly, the identifiability of the model follows by the result in Parag, du Plessis, and Pybus (2020) since our likelihood and the ESP likelihood only differ by the locations of the breakpoints. Intuitively, while the two parameters would not be identifiable from the sampling process alone, the parameter $N_e(t)$ appears in both the coalescent and the sampling likelihoods, making the joint model parameters $N_e(t)$ and $\beta(t)$ identifiable.

Here, we model both α and θ through Markov random field priors, either HSMRFs or GMRFs. This allows us to make minimal assumptions on $N_e(t)$ and $\beta(t)$: the choice of the grid practically depends solely on the sample size and no major assumptions are made on the underlying sampling process. The choice of prior for $N_e(t)$ follows from the well-studied characteristics of the two priors discussed in Section 2.2. Under the HSMRF prior on $\beta(t)$, one can model situations in which there are sharp changes in the sampling design (both first and second orders). Under the GMRF prior on $\beta(t)$, one favors smooth sampling designs, a situation which is also desirable when one does not have exact knowledge of the underlying sampling protocol. Note that the choice of field and order of the priors can be disjoint: for example, one can place a HSMRF of order 1 prior on $N_e(t)$ and a GMRF of order 2 prior on $\beta(t)$. We study in details the properties of the prior distributions in the supplementary material.

To formalize, Bayesian phylodynamic inference under adaptive preferential sampling can be written in the most general form as

$$\begin{aligned} \mathbf{t} | \theta, \mathbf{n}, \mathbf{s} &\sim \text{Coalescent model} & \mathbf{s} | \theta, \alpha, n &\sim \text{Poisson process} \\ \theta | \tau, \gamma &\sim \text{HSMRF-}p_1 \text{ or } \theta | \gamma &\sim \text{GMRF-}p_1 \\ \alpha | \psi, \xi &\sim \text{HSMRF-}p_2 \text{ or } \alpha | \xi &\sim \text{GMRF-}p_2, \end{aligned} \quad (5)$$

where ξ is the global smoothing parameter of the MRF on α , ψ is the vector of local shrinkage parameter of the HSMRF prior on α , p_1 , and p_2 are the orders of the respective MRFs. We will refer to any combination of priors above as the adaptive preferential model.

Note that the adaptive preferential model differs notably from the framework of the ESP estimator by the fact that the parameter vectors θ and α are each dependent, the grid at which they are defined does not depend on \mathbf{t} , and these priors favor smooth estimates.

3.1. Inference

Posterior distributions. Under the assumption that \mathbf{t} and \mathbf{s} are known, and that we place HSMRF priors on both α and θ , the posterior distribution of model parameters could be readily computed

$$\begin{aligned} \pi(\alpha, \theta, \psi, \tau, \gamma, \xi | \mathbf{t}, \mathbf{s}) &\propto \mathcal{L}(\alpha, \theta | \mathbf{s}) \mathcal{L}(\theta | \mathbf{t}) \pi(\theta | \tau) \pi(\tau | \gamma) \\ &\quad \pi(\gamma | \zeta_1) \pi(\alpha | \psi) \pi(\psi | \xi) \pi(\xi | \zeta_2), \end{aligned}$$

where $\mathcal{L}(\theta | \mathbf{t})$ is the discretized coalescent log-likelihood. Under GMRF priors on α and θ , the posterior would be

$$\begin{aligned} \pi(\alpha, \theta, \gamma, \xi | \mathbf{t}, \mathbf{s}) &\propto \mathcal{L}(\alpha, \theta | \mathbf{s}) \mathcal{L}(\theta | \mathbf{t}) \pi(\theta | \gamma) \\ &\quad \pi(\gamma | \zeta_1) \pi(\alpha | \xi) \pi(\xi | \zeta_2). \end{aligned}$$

For our analysis, we fixed the pair (g, \mathbf{t}) , which can be estimated by other methods such as the Maximum clade credibility tree of the posterior distribution of the genealogy. In order to approximate the posterior distribution we use two methods: Hamiltonian MCMC and Integrated Nested Laplace approximation (INLA; Rue, Martino, and Chopin 2009). The choice of the two algorithms is based on the recent results in the phylodynamics literature: Lan et al. (2015) and Faulkner (2020) study the use of HMC in the context of phylodynamics with GP, GMRF, and HSMRF priors on $N_e(t)$; Palacios and Minin (2012) and Karcher et al. (2016) study INLA with a GMRF-1 prior on $N_e(t)$. Both INLA and HMC are shown to accurately approximate the posterior distributions and to efficiently handle large sample sizes (in the order of thousands of samples).

For Hamiltonian MCMC, we rely on Stan (Carpenter et al. 2017). The hyperparameters ζ_1 and ζ_2 are as described in Faulkner et al. (2020, app. B) (their method is suitably adapted to ζ_2 , the global smoothing parameter of the MRF on α).

INLA approximation. Posterior inference from latent Gaussian models can be achieved by approximating posterior marginal distributions via Laplace approximations. INLA allows us to replace MCMC entirely and approximate the posterior marginals of model parameters when our model is based on GMRF priors. What follows is largely based on Palacios and Minin (2012), who discuss INLA for GMRFs in phylodynamics. We extend it here to include the adaptive preferential sampling priors.

INLA approximates posterior marginals $\pi(\gamma, \xi | \mathbf{t}, \mathbf{s})$, $\pi(\theta_i | \mathbf{t}, \mathbf{s})$ for $1 \leq i \leq M$, and $\pi(\alpha_j | \mathbf{t}, \mathbf{s})$ for $1 \leq j \leq M'$. The posterior marginal distribution of hyperparameters is

$$\widehat{\pi}(\gamma, \xi | \mathbf{t}, \mathbf{s}) \propto \frac{\pi(\gamma, \xi, \theta, \alpha, \mathbf{t}, \mathbf{s})}{\widehat{\pi}_G(\theta, \alpha | \gamma, \xi, \mathbf{t}, \mathbf{s})} \Big|_{\alpha = \alpha^*(\xi, \gamma), \theta = \theta^*(\xi, \gamma)},$$

where $\widehat{\pi}_G(\boldsymbol{\theta}, \boldsymbol{\alpha} | \gamma, \xi, \mathbf{t}, \mathbf{s})$ is the Gaussian approximation of $\pi(\boldsymbol{\theta}, \boldsymbol{\alpha} | \gamma, \xi, \mathbf{t}, \mathbf{s})$ obtained from a Taylor expansion around its modes $\boldsymbol{\theta}^*(\xi, \gamma)$ and $\boldsymbol{\alpha}^*(\xi, \gamma)$ (modes can be computed through any optimization algorithm, e.g., Newton–Raphson).

The approximation of the marginal distributions of the MRFs $\pi(\theta_i | \mathbf{t}, \mathbf{s})$ and $\pi(\alpha_j | \mathbf{t}, \mathbf{s})$ are

$$\widehat{\pi}(\theta_i | \gamma, \xi, \mathbf{t}, \mathbf{s}) \propto \frac{\pi(\gamma, \xi, \boldsymbol{\theta}, \boldsymbol{\alpha}, \mathbf{t}, \mathbf{s})}{\widehat{\pi}_{GG}(\boldsymbol{\theta}_{-i}, \boldsymbol{\alpha} | \gamma, \xi, \mathbf{t}, \mathbf{s})} \Big|_{\boldsymbol{\alpha}=\boldsymbol{\alpha}^*, \boldsymbol{\theta}_{-i}=\boldsymbol{\theta}_{-i}^*} \quad \text{and}$$

$$\widehat{\pi}(\alpha_i | \gamma, \xi, \mathbf{t}, \mathbf{s}) \propto \frac{\pi(\gamma, \xi, \boldsymbol{\theta}, \boldsymbol{\alpha}, \mathbf{t}, \mathbf{s})}{\widehat{\pi}_{GG}(\boldsymbol{\theta}, \boldsymbol{\alpha}_{-i} | \gamma, \xi, \mathbf{t}, \mathbf{s})} \Big|_{\boldsymbol{\alpha}_{-i}=\boldsymbol{\alpha}_{-i}^*, \boldsymbol{\theta}=\boldsymbol{\theta}^*},$$

where $\widehat{\pi}_{GG}(\boldsymbol{\theta}, \boldsymbol{\alpha}_{-i} | \gamma, \xi, \mathbf{t}, \mathbf{s})$ is the Gaussian approximation of $\pi(\boldsymbol{\theta}, \boldsymbol{\alpha}_{-i} | \gamma, \xi, \mathbf{t}, \mathbf{s})$ obtained from a Taylor expansion at $(\boldsymbol{\alpha}_{-i}, \boldsymbol{\theta}) = E_G[\boldsymbol{\theta}, \boldsymbol{\alpha}_{-i} | \gamma, \xi, \mathbf{t}, \mathbf{s}]$, where E_G denotes the expected value w.r.t. $\widehat{\pi}_G(\boldsymbol{\theta}, \boldsymbol{\alpha} | \gamma, \xi, \mathbf{t}, \mathbf{s})$. Analogously, we can define the same term for $\widehat{\pi}_{GG}(\boldsymbol{\theta}_{-i}, \boldsymbol{\alpha} | \gamma, \xi, \mathbf{t}, \mathbf{s})$. Given $\widehat{\pi}(\theta_i | \gamma, \xi, \mathbf{t}, \mathbf{s})$ and $\widehat{\pi}(\alpha_i | \gamma, \xi, \mathbf{t}, \mathbf{s})$, we use $\widehat{\pi}(\gamma, \xi | \mathbf{t}, \mathbf{s})$ to integrate out γ and ξ and obtained the desired distributions.

4. Simulations

We rely on simulations to evaluate the performance of the adaptive preferential sampling (adaPref) method in estimating the effective population size trajectory. We compare the performance of adaPref to alternative methods with and without preferential sampling. This simulation study is designed to test the benefit of the adaptive framework. We consider two settings: “strong” preferential sampling (sampling process always depends on $N_e(t)$), and “weak” preferential sampling (sampling is preferential only during some time periods).

For parsimony, we address the following questions in the supplementary material: (i) we evaluate the sensitivity of adaPref posterior distributions to the choice of the MRF (GMRF vs HSMRF) priors and their orders (1 vs 2), (ii) we study the approximation error incurred using INLA in place of MCMC, (iii) we study how well our model infer $\beta(t)$.

Simulation setup. In order to exhibit the benefit of the adaPref framework, we consider adaPref with GMRF order 1 priors (GMRF1) on both $N_e(t)$ and $\beta(t)$ and posterior distributions approximated by INLA. We compare adaPref with the preferential sampling method of Karcher et al. (2016) (parPref) and the method without preferential sampling of (Palacios and Minin 2012) (noPref), available in the R package `phyloodyn` (Karcher et al. 2017). Both noPref and parPref rely on a GMRF1 prior on $N_e(t)$ and the INLA approximation. All implementations rely on R-INLA (Rue, Martino, and Chopin 2009).

For each dataset, we test the performance of all models through a set of commonly used summary statistics. For a regular grid of time points $(v_i)_{1:K}$, as a measure of bias we consider the sum of relative errors: $\text{DEV} = \sum_{i=1}^K \frac{|\widehat{N}_e(v_i) - N_e(v_i)|}{N_e(v_i)}$, where $\widehat{N}_e(v_i)$ is the posterior median of N_e at time v_i ; to quantify the uncertainty in the estimate we use the mean relative width: $\text{RWD} = \frac{1}{K} \sum_{i=1}^K \frac{|\widehat{N}_{97.5}(v_i) - \widehat{N}_{2.5}(v_i)|}{N_e(v_i)}$, where $\widehat{N}_{97.5}(v_i)$ and $\widehat{N}_{2.5}(v_i)$ are, respectively, the 97.5% and 2.5% quantiles of the posterior distribution of $N_e(v_i)$; last, the envelope measure $\text{ENV} = \frac{1}{K} \sum_{i=1}^K \mathbf{1}_{\{\widehat{N}_{2.5}(v_i) \leq N_e(v_i) \leq \widehat{N}_{97.5}(v_i)\}}$, which measures the proportion of the curve that is covered by the 95% credible region,

that is, it is a proxy for coverage. We fix $K = 100$, $v_1 = 0$ and $v_K = .8 t_2$.

Data. We simulate genealogies under two population size trajectories: a piece-wise constant and exponential trajectory (CE), and a bottleneck trajectory (B). For the sampling protocols, we simulate sampling trajectories that resemble situations encountered in applications: a sampling protocol proportional to $N_e(t)$ (PP), uniform (U), a “lagged” response to changes in $N_e(t)$ (LP), and a situation where only some segments of the sampling trajectory are preferential (UP). The combination of the two acronyms will be used in the plots, for example, B-U refers to bottleneck trajectory and uniform sampling. The first rows of Figures 2-3 depict the N_e (red) and $\lambda(t)$ (black) trajectories (up to a scaling constant and in log-scale) of the eight simulation scenarios considered. Exact specifics of the trajectories used are given in the supplementary material.

We consider three sample sizes ($n = 100$, $n = 300$, and $n = 500$) and 50 simulations for each combination of trajectories and sample size. For a fixed n , $\beta(t)$, and $N_e(t)$, we set $n_1 = 2$, and simulate $n - n_1$ sampling locations from a Poisson process with rate $\beta(t)N_e(t)$. This defines a vector of sampling times and $\mathbf{n} = (2, 1, \dots, 1)$. Conditionally on \mathbf{n} , \mathbf{s} , and $N_e(t)$, we sample coalescent times \mathbf{t} according to Palacios and Minin (2013, algor. 3). Both inhomogeneous Poisson processes are sampled using the Lewis-Shedler thinning algorithm (Lewis and Shedler 1979). Note that there is no need to sample tree topologies because the vector \mathbf{t} is a sufficient statistic for $N_e(t)$. Last, we approximate the posterior distribution of $N_e(t)$ conditionally on \mathbf{t} , \mathbf{s} , and \mathbf{n} for the three methods; and the posterior of $\beta(t)$ conditionally on \mathbf{s} and \mathbf{n} (only adaPref). The code for reproducing the simulation study is available at <https://github.com/lorenzocapp/adapref> as a R package.

Results. We first do a qualitative assessment of the three models considering a single simulated dataset from each simulation scenario (case $n = 500$). The four panels of Figure 2 second row depict the posterior medians and 95% credible regions of $N_e(t)$ for the constant-exponential trajectory (CE). Our adaPref results are depicted in black and gray scale, the parPref method in red, and the noPref in blue. Each column corresponds to a sampling protocol. All posterior medians of $N_e(t)$ are very similar and close to the truth (black dashed line in the same panel, red line in the panel above) except for parPref (red) in the last two scenarios. Indeed, the last two scenarios correspond to the cases in which the preferential sampling assumption of parPref is violated. ParPref and adaPref show similar credible region widths in the case of proportional preferential sampling (first column), however, adaPref consistently shows narrower credible regions across sampling protocols. Figure 3 shows the same type of comparisons for the bottleneck trajectory (B). The posterior median and credible intervals obtained with parPref are particularly off during the periods of no preferential sampling in the last simulation scenario (fourth column). In all other sampling scenarios, all methods have very similar posterior medians, however again, adaPref shows narrower credible regions while keeping high coverage across sampling protocols.

We now turn to the quantitative comparison of the accuracy of the $N_e(t)$ estimators obtained with the three different models: adaPref, noPref, and parPref considering the 1200 datasets

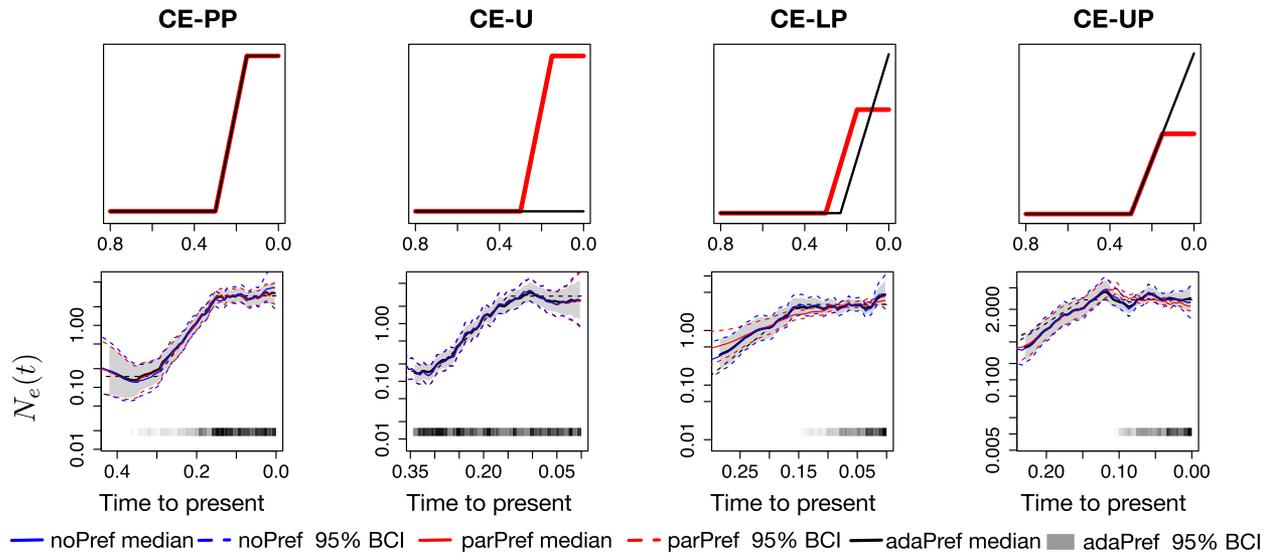


Figure 2. Simulations from constant-exponential trajectories and posterior inferences of $N_e(t)$. First row panels depict the simulated log-effective population size (red) and log-sampling intensity (black) trajectories (up to a constant). Second row panels depict posterior estimates for the four simulation scenarios from a single simulated genealogy picked at random with $n = 500$ tips. All models used GMRF of order 1 priors and posterior inference is approximated with INLA. The posterior medians of adaPref are depicted as solid black curves and the 95% Bayesian credible regions are depicted as shaded areas. Posterior medians of parPref and noPref are depicted, respectively, as solid blue and red curves, and the 95% Bayesian credible regions are depicted by the corresponding dashed curves. \mathbf{n} and \mathbf{s} are depicted by the heat maps at the bottom of the last four panels: the squares along the time axis depict the sampling times, while the intensity of the black color depicts the number of samples. The true trajectories are depicted as a black dashed curves.

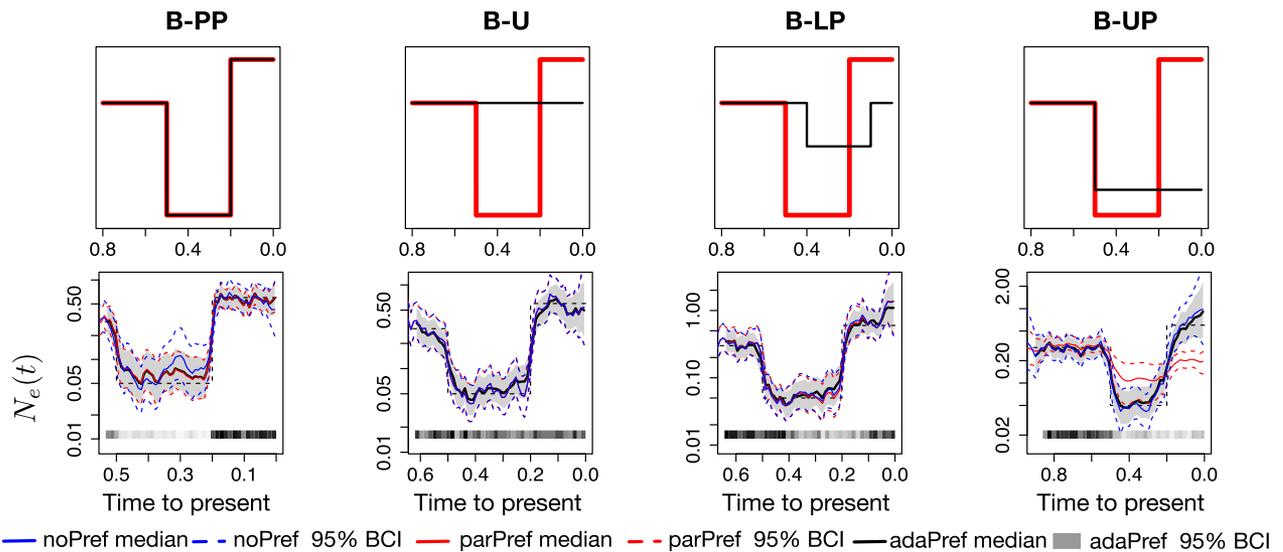


Figure 3. Simulations from bottleneck trajectories and posterior inferences of $N_e(t)$. First row panels depict the simulated log-effective population size (red) and log-sampling intensity (black) trajectories (up to a constant). Second row panels depict posterior estimates for the four simulation scenarios and from a single simulated genealogy picked at random with $n = 500$ tips. All models used GMRF of order 1 priors and posterior inference is approximated with INLA. The posterior medians of adaPref are depicted as solid black curves and the 95% Bayesian credible regions are depicted by shaded areas. Posterior medians of parPref and noPref are depicted respectively as solid blue and red curves, and the 95% Bayesian credible regions are depicted by the corresponding dashed curves. \mathbf{n} and \mathbf{s} are depicted by the heat maps at the bottom of the last four panels: the squares along the time axis depict the sampling times, while the intensity of the black color depicts the number of samples. The true trajectories are depicted as a black dashed curves.

sampled. We use the three metrics described above: ENV as a measure of coverage, RWD as a measure of credible region width, and DEV as a measure of bias. Table 1 reports the average values of the three metrics relative to the value obtained by the adaPref method (grouped by sample size, model, and simulation scenario). A value higher than one means a better performance in terms of ENV, and a worst performance in terms of RWD and DEV. Figure 4 top two rows plot the ENV, RWD, and DEV summary statistics obtained from the CE (first row) and the B (second row) trajectories. In the top two rows, we are not group-

ing anymore by sample size: a boxplot includes the datasets of all the three sample sizes. Figure 4 last row panels depict the boxplots of the three statistics considered, now grouping simulations by sample size (now, a boxplot includes the datasets of all simulation scenarios).

Table 1 suggests that adaPref has generally a better performance in terms of DEV and RWD. The adaPref model (red boxplots) has the best mean performance in six out of the eight scenarios in terms of both RWD and DEV (B-PP, B-U, all CE scenarios). In the last two scenarios (B-LP, B-UP), the

Table 1. Simulation: Summary statistics of adaPref posterior inference of $N_e(t)$ grouped by simulation study and sample size.

model	method	$n = 100$			$n = 300$			$n = 500$		
		ENV*	DEV*	RWD*	ENV*	DEV*	RWD*	ENV*	DEV*	RWD*
B-LP	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.09	1.08	2.02	1.05	1.21	3.02	1.04	1.08	1.33
	parPref	1.08	1.04	1.85	1.05	1.19	2.92	1.04	1.07	1.30
B-PP	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.01	1.26	1.77	1.00	1.20	1.92	1.00	1.10	1.20
	parPref	1.00	1.04	1.10	0.99	1.04	1.15	0.99	0.99	0.92
B-U	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.06	1.18	2.36	1.04	1.12	1.49	1.03	1.19	1.68
	parPref	1.06	1.15	2.28	1.04	1.11	1.47	1.03	1.19	1.67
B-UP	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.04	1.31	3.41	1.05	1.13	1.92	1.04	1.07	1.45
	parPref	1.03	1.27	2.74	0.80	1.32	0.86	0.76	1.40	0.76
CE-LP	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.01	0.99	1.49	1.01	1.00	1.11	1.01	0.99	1.10
	parPref	0.90	1.11	0.66	0.84	1.33	0.86	0.59	2.33	1.02
CE-PP	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.00	1.11	1.36	1.00	1.15	1.36	1.00	1.19	1.27
	parPref	1.00	1.07	1.16	1.00	1.00	0.99	1.00	1.00	0.99
CE-U	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.01	1.20	1.70	1.01	1.12	1.22	1.00	1.12	1.23
	parPref	1.01	1.15	1.56	1.01	1.11	1.21	1.00	1.12	1.22
CE-UP	adaPref	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	noPref	1.01	0.97	1.10	1.01	0.99	1.07	1.01	0.98	1.07
	parPref	0.98	1.01	0.90	0.87	1.25	0.89	0.74	1.60	0.87

NOTES: The bar on top of the criterion means we are considering the average value across the 50 estimated statistics for each simulation trajectory based on $N_e(t)$ posterior, *rel* refers that the value is relative to the result of adaPref. adaPref (our method), parPref (Karcher et al. 2016), and noPref (Palacios and Minin 2012). A value higher than one means a better performance in terms of ENV, and a worst performance in terms of RWD and DEV.

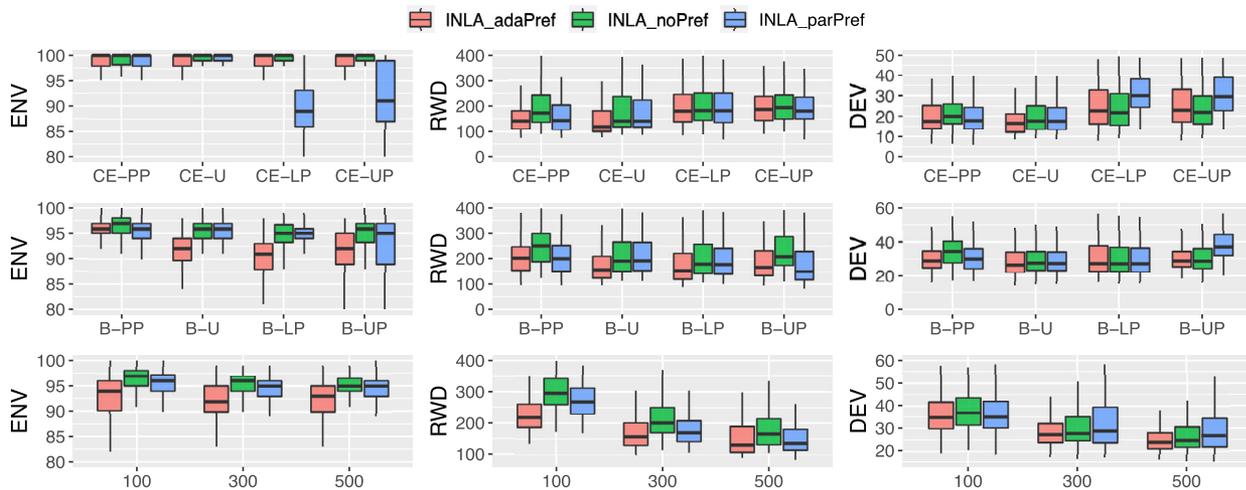


Figure 4. Summary statistics of parPref, noPref and adaPref inference of $N_e(t)$ approximated with INLA. In the first two rows, each box refers to one method (color in the legend) and depicts the distribution of the 150 estimated statistics for each simulation trajectory (50 datasets for each sample size, 150 in total) based on $N_e(t)$ posterior; ENV, first column; RWD, second column; DEV, third column. In the third row, the grouping is done according to the sample size: each box is based on 400 simulated genealogies (50 genealogies for each of the eight trajectories). In the legend, INLA_adaPref is our method, INLA_noPref is the method in Palacios and Minin (2012), INLA_parPref is the method of Karcher et al. (2016). All models used GMRF of order 1 priors.

parPref model achieved the lowest mean RWD and noPref the lowest mean DEV. Surprisingly, the adaPref model outperforms parPref also in the CE-PP and B-PP scenarios, where the parametric assumptions are met. noPref is the model that has a slightly better performance in terms of coverage (ENV). However, this is achieved with much wider credible regions. We see how the adaPref coverage is never too different from the noPref ones, while the credible regions are substantially narrower. We judge this as a better performance given that one can always achieve perfect coverage with very wide credible intervals.

The adaPref model is more heavily parameterized and one may be led to think that the performance of the adaPref estima-

tor is affected by the sample size. However, there is no detectable sample size effect: the relative performance of the estimators is roughly similar as n increases. The adaPref estimator is the best performing according to DEV and RWD averaging over all the simulation scenarios jointly. See Table 1 and the last row of Figure 4.

5. SARS-CoV-2 in Los Angeles and Santa Clara counties

SARS-CoV-2 is the virus responsible for the coronavirus disease pandemic in 2019–2020. Molecular surveillance of

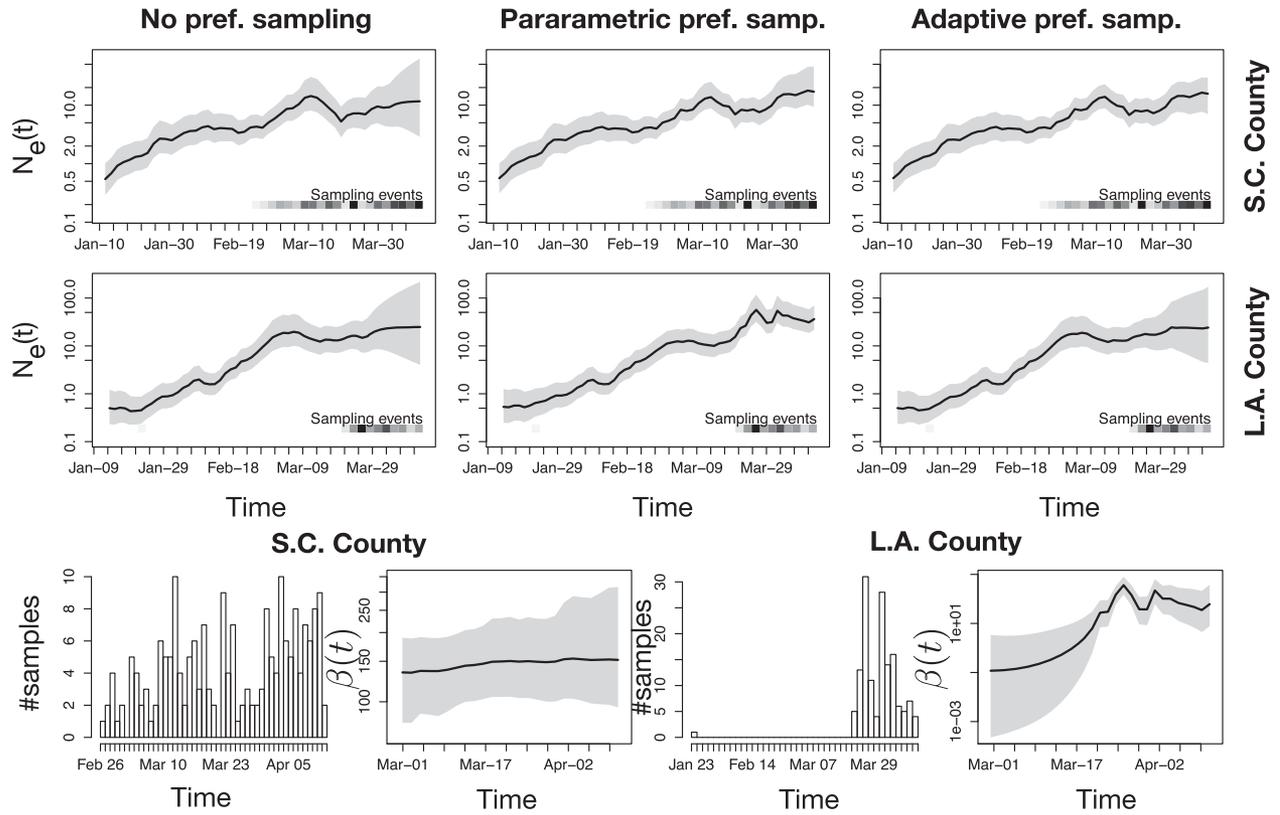


Figure 5. Phylodynamic inference with noPref, parPref and adaPref models from SARS-CoV-2 genealogies inferred from GISAID data obtained from Los Angeles and Santa Clara counties. First two rows panels depict $N_e(t)$ posterior distributions inferred using three possible priors: the noPref (Palacios and Minin 2012), the parPref Karcher et al. (2016), and our adaPref prior. The last row first and third panels depict how many samples are collected each day. The second and fourth panels depict $\beta(t)$ posterior distribution (only under the adaPref prior). The posterior medians are depicted as solid black lines and the 95% Bayesian credible regions are depicted by shaded areas. Sampling times are also depicted by the heat maps at the bottom of the top two rows panels: the squares along the time axis depicts the sampling time, while the intensity of the black color depicts the number of samples.

SARS-CoV-2 complements traditional surveillance methods based on case count data and provides a unique opportunity to retrospectively learn past disease dynamics. Researchers and public health officials agree that a successful response to an outbreak involves the prompt collection, sequencing, and sharing of molecular samples (Gire et al. 2014; Polonsky et al. 2019; WHO 2021). The benefit of a rapid response has been evident in the COVID-19 pandemic; for example, in the Netherlands (Munnink et al. 2020) and in the UK (Meredith et al. 2020). However, the number of sequences collected and shared varied substantially country by country, as can be seen in the GISAID EpiCov database (Shu and McCauley 2017). In addition to a timely response, deciding when, where and how many sequences should be sampled has been receiving attention (Parag and Pybus 2019). The heterogeneity in sampling designs and speed of implementations of genomic surveillance has motivated our work. In this section, we provide a case study to highlight the usefulness of our procedure.

Here, we use the adaPref model for estimating the viral genetic diversity trajectory $N_e(t)$, from currently available viral molecular sequences in GISAID obtained from infected individuals. We analyzed viral whole-genome sequences collected in California in Santa Clara (S.C.) and Los Angeles (L.A.) counties. The GISAID reference numbers of the sequences included in this study, together with data access acknowledgments, are included in the supplementary material.

We downloaded all molecular sequences available on June 27, 2020. The datasets consist of 195 and 134 sequences from S.C and L.A. counties, respectively, with collection dates ranging from mid-February, 2020 to April 13 of 2020. We included only high coverage sequences with more than 25,000 base pairs. Sampling frequency information is depicted in the first and third panels of the last row of Figure 5. We note that the sampling effort varied in the two counties: most of the L.A. samples are concentrated in late March-mid April, while samples have been collected throughout late February-mid April in S.C. county. We consider this case study interesting because the two counties exhibit a different response to the outbreak: S.C. public health officials and researchers promptly responded collecting sequences, while in L.A. country most sequences were collected at the peak of the first outbreak.

The two estimated genealogies employed in the analysis are the maximum clade credibility trees of the posterior distributions obtained with BEAST2 (Bouckaert et al. 2019). The MCMC parameters are 20×10^6 iterations, thinning every 1000 and burnin of 10×10^6 iterations. We selected the following priors: Extended Bayesian Skyline prior on $N_e(t)$ (Heled and Drummond 2008), HKY mutation model with empirically estimated base frequencies (Hasegawa, Kishino and Yano 1985), and uniform prior on the mutation rate with support constrained between 9×10^{-4} and 1.1×10^{-3} substitutions per site per year. The support of the uniform prior was centered around

1×10^{-3} mutations per site per year, an estimate obtained by regressing the Hamming distances of the sequences to the ancestral reference sequence (GenBank MN908947; Wu et al. 2020) on the time difference between the sampling times and the reference sampling time.

Given the two estimated genealogies, we approximate posterior marginal distributions of $N_e(t)$ through the INLA approximations of the noPref model (Palacios and Minin 2012), the parPref model (Karcher et al. 2016), and our adaPref model. In the first two rows of Figure 5, we show the estimates of effective population size trajectories with the noPref model (first column), the parPref model (second column), and the adaPref model (third column). Results for S.C. county correspond to the first row and for L.A. county to the second row. Sampling intensity posteriors (computed only through the adaPref model) are given in the third row of Figure 5 in the second and third panels.

The median posteriors of $N_e(t)$ obtained with the noPref and the adaPref models in L.A. county have almost identical trajectories, while the one with the parPref model has a more pronounced maximum later on (around April 1). In the S.C. county dataset, the median posterior estimates of $N_e(t)$ obtained with the parPref and adaPref models are in this case almost identical, with the estimate obtained with noPref not recovering a steep growth at the end of March. The split behavior of the adaPref posterior, once matching with the noPref posterior and once with parPref posterior, can be explained by looking at the posterior of $\beta(t)$: in the S.C. dataset, $\beta(t)$ median posterior is practically flat, a situation consistent with the parametric assumption of the parPref model, while the time-varying $\beta(t)$ accounts for the fact that sampling in L.A. is concentrated in a short time frame.

The average width of the credible regions ($RW = \frac{1}{K} \sum_{i=1}^K |\hat{N}_{97.5}(v_i) - \hat{N}_{2.5}(v_i)|$) differ across methods and datasets. In S.C. county, RWD is 8.5 for noPref, 6.6 for parPref, and 4.7 for adaPref inferences. In L.A. county, the RWD is 23.9 for noPref, 13.6 for parPref, and 20.8 for adaPref inferences. We get a general confirmation that preferential sampling estimators lead to narrower credible regions.

We consider the “split performance” of adaPref (in S.C. similar to parPref and in L.A. similar to noPref) a positive feature of our model. The timely response to the outbreak in S.C. county was followed by a sampling collection that spanned the early stages of the pandemic. Here, adaPref detects an association between the sampling and the coalescent processes. In L.A. county, there is one sequence collected on January 23rd, no sequences collected for about 2 months, followed by an intense sampling effort. The preferential sampling assumption “forces $N_e(t)$ down” in February and early March because there are no samples. This is evident in the result of parPref. On the other hand, adaPref does not detect a strong association between the two processes; hence, the estimates are closer to noPref. To understand which model is correct in this setting, we would need to have access to the sampling protocols of L.A. county to know why there are no samples in those two months.

A final remark. The estimates of $N_e(t)$ presented here are representative of genetic diversity over time and do not directly translate to the number of infections. The coalescent

we employed ignores recombination, population structure, and selection. Also, we note that observed nucleotide substitutions may be caused by sequencing errors and these are being ignored in our study.

6. Discussion

We have introduced an adaptive preferential sampling model to estimate the effective population size $N_e(t)$ of a coalescent process accounting for a situation in which sampling dates are stochastically dependent on the effective population size. We model sampling dates as an inhomogeneous Poisson process with rate $\beta(t)N_e(t)$, where $\beta(t)$ is a time-varying coefficient that modulates how this dependence varies over time. We assume that both $N_e(t)$ and $\beta(t)$ are continuous functions and model them in a Bayesian framework with Markov random field priors. This methodology allows us to account for preferential sampling while making minimal assumptions on the dependence between the sampling process and the genealogical process. We term the model proposed adaptive preferential sampling.

The adaptive preferential sampling model allows for a situation in which the sampling protocol changes over time but no detailed knowledge on the way samples are collected is available. In particular, the local adaptivity of the Horseshoe Markov random field prior allows also to model abrupt changes in the sampling protocol.

We show through simulation studies that the estimates obtained through the adaptive preferential sampling are more accurate than some of the available alternatives, leading to smaller absolute deviations from the true trajectories and narrower credible regions. The performance is competitive also in a broad set of scenarios in which the parametric assumptions of the alternative methods are met. We provide an application to SARS-CoV-2 monitoring and show the “adaptive nature” of our methodology: in one scenario the estimate was comparable to that of the model without preferential sampling. In a second one, the estimate was matched that of the parametric preferential sampling methodology.

The most direct extension for future work is to include genealogical uncertainty, which is being ignored in the present work. While Kingman heterochronous n -coalescent is the standard coalescent model choice to include genealogical uncertainty, recent works have proposed to infer $N_e(t)$ employing lower resolution coalescent models (Sainudiin, Stadler, and Véber 2015; Cappello, Veber, and Palacios 2020). Our adaptive preferential sampling framework can be paired with any of the ancestral processes.

Another natural extension to the proposed adaptive preferential framework is to incorporate covariates into $\lambda(t)$, the sampling rate, as it is done in Karcher et al. (2020), to include auxiliary information about the sampling protocols available to the modeler. For example, it is easy to imagine that one may have direct control over the sampling protocol. The resulting rate of the sampling process would be $\lambda(t) = \beta(t)N_e(t) + \beta'X(t)$, where X is a vector of covariates and β' the corresponding linear coefficients.

In this article, we model jointly the coalescent process and a sampling process depending on $N_e(t)$. An interesting direction

of future work is to model jointly the coalescent process with other processes that depend on $N_e(t)$, such as the total number of infected individuals in an epidemic (Volz et al. 2009). The adaptive framework introduced in this article seems to be suitable to such an extension, given that we make limited assumptions on the dependence between the two processes.

Acknowledgments

We thank an associate editor and two anonymous reviewers for constructive comments. Computing resources were provided by Stanford University research computing center (Sherlock cluster).

Funding

Julia A. Palacios acknowledges support from National Institutes of Health grant R01-GM-131404 and the Alfred P. Sloan Foundation.

References

- Bouckaert, R., Vaughan, T. G., Barido-Sottani, J., Duchêne, S., Fourment, M., Gavryushkina, A., Heled, J., Jones, G., Kühnert, D., De Maio, N., Matschiner, M., Mendes, F. K., Müller, N. F., Ogilvie, H. A., du Plessis, L., Poppinga, A., Rambaut, A., Rasmussen, D., Siveroni, I., Suchard, M. A., Wu, C.-H., Xie, D., Zhang, C., Stadler, T., Drummond, A. J. (2019), “BEAST 2.5: An Advanced Software Platform for Bayesian Evolutionary Analysis,” *PLoS Computational Biology*, 15, e1006650. [9]
- Cappello, L., Veber, A. and Palacios, J. A. (2020), “The Tajima Heterochronous n-Coalescent: Inference From Heterochronously Sampled Molecular Data,” arXiv:2004.06826. [10]
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017), “Stan: A Probabilistic Programming Language,” *Journal of Statistical Software*, 76, arXiv:2004.06826 [5]
- Carvalho, C. M., Polson, N. G. and Scott, J. G. (2010), “The Horseshoe Estimator for Sparse Signals,” *Biometrika*, 97, 465–480. [3]
- Diggle, P. J., Menezes, R., and Su, T.-I. (2010), “Geostatistical Inference Under Preferential Sampling,” *Journal of the Royal Statistical Society, Series C*, 59, 191–232. [1,4]
- Drummond, A. J., Rambaut, A., Shapiro, B., and Pybus, O. G. (2005), “Bayesian Coalescent Inference of Past Population Dynamics From Molecular Sequences,” *Molecular Biology and Evolution*, 22, 1185–1192. [2]
- Faria, N. R., da Silva Azevedo, R. d. S., Kraemer, M. U., Souza, R., Cunha, M. S., Hill, S. C., Thézé, J., Bonsall, M. B., Bowden, T. A., Rissanen, I., Rocco, I. M., Nogueira, J. S., Maeda, A. Y., Vasami, F. G. D. S., Macedo, F. L. L., Suzuki, A., Rodrigues, S. G., Cruz, A. C. R., Nunes, B. T., Medeiros, D. B. A., Rodrigues, D. S. G., Queiroz, A. L. N., da Silva, E. V. P., Henriques, D. F., da Rosa, E. S. T., de Oliveira, C. S., Martins, L. C., Vasconcelos, H. B., Casseb, L. M. N., Simith, D. B., Messina, J. P., Abade, L., Lourenço, J., Alcantara, L. C. J., de Lima, M. M., Giovanetti, M., Hay, S. I., de Oliveira, R. S., Lemos, P. D. S., de Oliveira, L. F., de Lima, C. P. S., da Silva, S. P., de Vasconcelos, J. M., Franco, L., Cardoso, J. F., Vianez-Júnior, J. L. D. S. G., Mir, D., Bello, G., Delatorre, E., Khan, K., Creatore, M., Coelho, G. E., de Oliveira, W. K., Tesh, R., Pybus, O. G., Nunes, M. R. T., Vasconcelos, P. F. C. (2016), “Zika Virus in the Americas: Early Epidemiological and Genetic Findings,” *Science*, 352, 345–349. [1]
- Faulkner, J. R., Magee, A. F., Shapiro, B. and Minin, V. N. (2020), “Horseshoe-Based Bayesian Nonparametric Estimation of Effective Population Size Trajectories,” *Biometrics*, 76, 677–690. [2,3,5]
- Felsenstein, J., and Rodrigo, A. G. (1999), “Coalescent Approaches to HIV Population Genetics,” in *The Evolution of HIV*, eds. Keith A. Crandall, Baltimore, MD: Johns Hopkins University Press, pp. 233–272. [1,2,3]
- Frost, S. D., and Volz, E. M. (2010), “Viral Phylodynamics and the Search for an Effective Number of Infections,” *Philosophical Transactions of the Royal Society, Series B*, 365, 1879–1890. [1]
- Gire, S. K., Goba, A., Andersen, K. G., Sealfon, R. S., Park, D. J., Kanneh, L., Jalloh, S., Momoh, M., Fullah, M., Dudas, G., Wohl, S., Moses, L. M., Yozwiak, N. L., Winnicki, S., Matranga, C. B., Malboeuf, C. M., Qu, J., Gladden, A. D., Schaffner, S. F., Yang, X., Jiang, P. P., Nekoui, M., Colubri, A., Coomber, M. R., Fonnies, M., Moigboi, A., Gbokie, M., Kamara, F. K., Tucker, V., Konuwa, E., Saffa, S., Sellu, J., Jalloh, A. A., Kovoma, A., Koninga, J., Mustapha, I., Kargbo, K., Foday, M., Yillah, M., Kanneh, F., Robert, W., Massally, J. L., Chapman, S. B., Bochicchio, J., Murphy, C., Nusbaum, C., Young, S., Birren, B. W., Grant, D. S., Scheffelin, J. S., Lander, E. S., Hapfi, C., Gevao, S. M., Gnirke, A., Rambaut, A., Garry, R. F., Khan, S. H., Sabeti, P. C. (2014), “Genomic Surveillance Elucidates Ebola Virus Origin and Transmission During the 2014 Outbreak,” *Science*, 345, 1369–1372. [9]
- Grenfell, B. T., Pybus, O. G., Gog, J. R., Wood, J. L., Daly, J. M., Mumford, J. A. and Holmes, E. C. (2004), “Unifying the Epidemiological and Evolutionary Dynamics of Pathogens,” *Science*, 303, 327–332. [1]
- Hadfield, J., Megill, C., Bell, S. M., Huddleston, J., Potter, B., Callender, C., Sagulenko, P., Bedford, T., and Neher, R. A. (2018), “Nextstrain: Real-Time Tracking of Pathogen Evolution,” *Bioinformatics*, 34, 4121–4123. [1]
- Hasegawa, M., Kishino, H., and Yano, T. (1985), “Dating of the Human-Ape Splitting by a Molecular Clock of Mitochondrial DNA,” *Journal of Molecular Evolution*, 2, 160–164. [9]
- Heled, J., and Drummond, A. J. (2008), “Bayesian Inference of Population Size History From Multiple Loci,” *BMC Evolutionary Biology*, 8, 289. [9]
- Hudson, R. R. (1990), “Gene Genealogies and the Coalescent Process,” *Oxford Surveys in Evolutionary Biology*, 7, 1–44. [2]
- Huff, C. D., Xing, J., Rogers, A. R., Witherspoon, D., and Jorde, L. B. (2010), “Mobile Elements Reveal Small Population Size in the Ancient Ancestors of Homo Sapiens,” *Proceedings of the National Academy of Sciences*, 107, 2147–2152. [1]
- Karcher, M. D., Palacios, J. A., Bedford, T., Suchard, M. A., and Minin, V. N. (2016), “Quantifying and Mitigating the Effect of Preferential Sampling on Phylodynamic Inference,” *PLoS Computational Biology*, 12, e1004789. [1,2,4,5,6,8,9,10]
- Karcher, M. D., Palacios, J. A., Lan, S., and Minin, V. N. (2017), “PhyloDYN: An R Package for Phylodynamic Simulation and Inference,” *Molecular Ecology Resources*, 17, 96–100. [6]
- Karcher, M. D., Suchard, M. A., Dudas, G., and Minin, V. N. (2020), “Estimating Effective Population Size Changes From Preferentially Sampled Genetic Sequences,” *PLoS Computational Biology*, 16, [4,10]
- Kim, S.-J., Koh, K., Boyd, S., and Gorinevsky, D. (2009), “ l_1 Trend Filtering,” *SIAM Review*, 51, 339–360. [4]
- Kingman, J. F. (1982a), “On the Genealogy of Large Populations,” *Journal of Applied Probability*, 19, 27–43. [1,2]
- (1982b), “The Coalescent,” *Stochastic Processes and Their Applications*, 13, 235–248. [1]
- Lan, S., Palacios, J. A., Karcher, M., Minin, V. N., and Shahbaba, B. (2015), “An Efficient Bayesian Inference Framework for Coalescent-Based Nonparametric Phylodynamics,” *Bioinformatics*, 31, 3282–3289. [5]
- Lewis, P. W., and Shedler, G. S. (1979), “Simulation of Nonhomogeneous Poisson Processes by Thinning,” *Naval Research Logistics Quarterly*, 26, 403–413. [6]
- Lorenzen, E. D., Nogués-Bravo, D., Orlando, L., Weinstock, J., Binladen, J., Marske, K. A., Ugan, A., Borregaard, M. K., Gilbert, M. T. P., Nielsen, R., Ho, S. Y., Goebel, T., Graf, K.E., Byers, D., Stenderup, J. T., Rasmussen, M., Campos, P. F., Leonard, J. A., Koepfli, K. P., Froese, D., Zazula, G., Stafford, T. W. Jr, Aaris-Sørensen, K., Batra, P., Haywood, A. M., Singarayer, J. S., Valdes, P. J., Boeskorov, G., Burns, J. A., Davydov, S. P., Haile, J., Jenkins, D. L., Kosintsev, P., Kuznetsova, T., Lai, X., Martin, L. D., McDonald, H. G., Mol, D., Meldgaard, M., Munch, K., Stephan, E., Sablin M., Sommer RS, Sipko T, Scott, E., Suchard, M. A., Tikhonov, A., Willerslev, R., Wayne, R.K., Cooper, A., Hofreiter, M., Sher, A., Shapiro, B., Rahbek, C., Willerslev, E. (2011), “Species-Specific Responses of Late Quaternary Megafauna to Climate and Humans,” *Nature*, 479, 359–364. [1]
- Meredith, L. W., Hamilton, W. L., Warne, B., Houldcroft, C. J., Hosmillo, M., Jahun, A. S., Curran, M. D., Parmar, S., Caller, L. G., Caddy, S. L., Khokhar, F. A., Yakovleva, A., Hall, G., Feltwell, T., Forrest, S., Sridhar, S., Weekes, M. P., Baker, S., Brown, N., Moore, E., Popay, A., Roddick,

- I., Reacher, M., Gouliouris, T., Peacock, S. J., Dougan, G., Török, M. E., Goodfellow, I. (2020), “Rapid Implementation of SARS-CoV-2 Sequencing to Investigate Cases of Health-Care Associated Covid-19: A Prospective Genomic Surveillance Study,” *The Lancet Infectious Diseases*, 20, 1263–1272. [9]
- Minin, V. N., Bloomquist, E. W., and Suchard, M. A. (2008), “Smooth Skyride Through a Rough Skyline: Bayesian Coalescent-Based Inference of Population Dynamics,” *Molecular Biology and Evolution*, 25, 1459–1471. [2]
- Munnink, B. B. O., Nieuwenhuijse, D. F., Stein, M., O’Toole, Á., Haverkate, M., Mollers, M., Kamga, S. K., Schapendonk, C., Pronk, M., Lexmond, P., van der Linden, A., Bestebroer, T., Chestakova, I., Overmars, R. J., van Nieuwkoop, S., Molenkamp, R., van der Eijk, A. A., GeurtsvanKessel, C., Vennema, H., Meijer, A., Rambaut, A., van Dissel, J., Sikkema, R. S., Timen, A., Koopmans, M. (2020), “Rapid SARS-CoV-2 Whole-Genome Sequencing and Analysis for Informed Public Health Decision-Making in the Netherlands,” *Nature Medicine*, 26, 1405–1410. [9]
- Palacios, J. A., and Minin, V. N. (2012), “Integrated Nested Laplace Approximation for Bayesian Nonparametric Phylodynamics,” in *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI’12)*, Arlington, VA: AUAI Press, pp. 726–735. [2,3,4,5,6,8,9,10]
- Palacios, J. A., and Minin, V. N. (2013), “Gaussian Process-Based Bayesian Nonparametric Inference of Population Size Trajectories From Gene Genealogies,” *Biometrics*, 69, 8–18. [3,4,6]
- Parag, K. V., du Plessis, L., and Pybus, O. G. (2020), “Jointly Inferring the Dynamics of Population Size and Sampling Intensity From Molecular Sequences,” *Molecular Biology and Evolution*, 37, 2414–2429. [2,4,5]
- Parag, K. V., and Pybus, O. G. (2019), “Robust Design for Coalescent Model Inference,” *Systematic Biology*, 68, 730–743. [9]
- Polonsky, J. A., Baidjoe, A., Kamvar, Z. N., Cori, A., Durski, K., Edmunds, W. J., Eggo, R. M., Funk, S., Kaiser, L., Keating, P., de Waroux O. L. P., Marks, M., Moraga, P., Morgan, O., Nouvellet, P., Ratnayake, R., Roberts, C. H., Whitworth, J., Jombart, T. (2019), “Outbreak Analytics: A Developing Data Science for Informing the Response to Emerging Pathogens,” *Philosophical Transactions of the Royal Society, Series B*, 374, 20180276. [9]
- Pybus, O. G., Rambaut, A., and Harvey, P. H. (2000), “An Integrated Framework for the Inference of Viral Population History From Reconstructed Genealogies,” *Genetics*, 155, 1429–1437. [2]
- Rambaut, A., Pybus, O. G., Nelson, M. I., Viboud, C., Taubenberger, J. K., and Holmes, E. C. (2008), “The Genomic and Epidemiological Dynamics of Human Influenza A Virus,” *Nature*, 453, 615–619. [1]
- Rothenberg, T. J. (1971), “Identification in Parametric Models,” *Econometrica: Journal of the Econometric Society*, 39, 577–591. [4]
- Rue, H., and Held, L. (2005), *Gaussian Markov Random Fields: Theory and Applications*, London: Chapman and Hall–CRC Press. [3]
- Rue, H., Martino, S., and Chopin, N. (2009), “Approximate Bayesian Inference for Latent Gaussian Models by Using Integrated Nested Laplace Approximations,” *Journal of the Royal Statistical Society, Series B*, 71, 319–392. [5,6]
- Sainudiin, R., Stadler, T., and Véber, A. (2015), “Finding the Best Resolution for the Kingman–Tajima Coalescent: Theory and Applications,” *Journal of Mathematical Biology*, 70, 1207–1247. [10]
- Shapiro, B., Drummond, A. J., Rambaut, A., Wilson, M. C., Matheus, P. E., Sher, A. V., Pybus, O. G., Gilbert, M. T. P., Barnes, I., Binladen, J., Willerslev, E., Hansen, A. J., Baryshnikov, G. F., Burns, J. A., Davydov, S., Driver, J. C., Froese, D. G., Harington, C. R., Keddie, G., Kosintsev, P., Kunz, M. L., Martin, L. D., Stephenson, R. O., Storer, J., Tedford, R., Zimov, S., and Cooper, A. (2004), “Rise and Fall of the Beringian Steppe Bison,” *Science*, 306, 1561–1565. [1]
- Shu, Y., and McCauley, J. (2017), “GISAID: Global Initiative on Sharing All Influenza Data—From Vision to Reality,” *Eurosurveillance*, 22, 30494. [1,9]
- Slatkin, M., and Hudson, R. (1991), “Pairwise Comparisons of Mitochondrial DNA Sequences in Stable and Exponentially Growing Populations,” *Genetics*, 129, 555–562. [2]
- Stadler, T. (2010), “Sampling-Through-Time in Birth–Death Trees,” *Journal of Theoretical Biology*, 267, 396–404. [2]
- Stadler, T., Kühnert, D., Bonhoeffer, S., and Drummond, A. J. (2013), “Birth–Death Skyline Plot Reveals Temporal Changes of Epidemic Spread in HIV and Hepatitis C Virus (HCV),” *Proceedings of the National Academy of Sciences*, 110, pp. 228–233. [2]
- Strimmer, K., and Pybus, O. G. (2001), “Exploring the Demographic History of DNA Sequences Using the Generalized Skyline Plot,” *Molecular Biology and Evolution*, 18, 2298–2305. [2]
- Volz, E. M., and Frost, S. D. (2014), “Sampling Through Time and Phylodynamic Inference With Coalescent and Birth–Death Models,” *Journal of The Royal Society Interface*, 11, 20140945. [1,2,4]
- Volz, E. M., Pond, S. L. K., Ward, M. J., Brown, A. J. L., and Frost, S. D. (2009), “Phylodynamics of Infectious Disease Epidemics,” *Genetics*, 183, 1421–1430. [1,11]
- Wakeley, J. (2009), *Coalescent Theory: An Introduction*, Greenwood Village, Colorado: Roberts and Co. [2]
- WHO (2021), “Genomic Sequencing of SARS-CoV-2: A Guide to Implementation for Maximum Impact on Public Health, 8 January 2021”. [9]
- Wu, F., Zhao, S., Yu, B., Chen, Y.-M., Wang, W., Song, Z.-G., Hu, Y., Tao, Z.-W., Tian, J.-H., Pei, Y.-Y. et al. (2020), “A New Coronavirus Associated With Human Respiratory Disease in China,” *Nature*, 579, 265–269. [10]