

The Tajima heterochronous n -coalescent: inference from heterochronously sampled molecular data

Lorenzo Cappello^{1*}, Amandine Véber^{2*}, Julia A. Palacios^{1,3*}

¹Department of Statistics, Stanford University

²CMAP, CNRS, École Polytechnique, I.P. Paris

³Department of Biomedical Data Science, Stanford Medicine

April 16, 2020

Abstract

The observed sequence variation at a locus informs about the evolutionary history of the sample and past population size dynamics. The standard Kingman coalescent model on genealogies – timed trees that represent the ancestry of the sample – is used in a generative model of molecular sequence variation to infer evolutionary parameters. However, the state space of Kingman’s genealogies grows superexponentially with sample size n , making inference computationally unfeasible already for small n . We introduce a new coalescent model called Tajima heterochronous n -coalescent with a substantially smaller cardinality of the genealogical space. This process allows to analyze samples collected at different times, a situation that in applications is both met (*e.g.* ancient DNA and RNA from rapidly evolving pathogens like viruses) and statistically desirable (variance reduction and parameter identifiability). We propose an algorithm to calculate the likelihood efficiently and present a Bayesian nonparametric procedure to infer the population size trajectory. We provide a new MCMC sampler to explore the space of Tajima’s genealogies and model parameters. We compare our procedure with state-of-the-art methodologies in simulations and applications. We use our method to re-examine the scientific question of how Beringian bison went extinct analyzing modern and ancient molecular sequences of bison in North America, and to reconstruct population size trajectory of SARS-CoV-2 from viral sequences collected in France and Germany.

Keywords: Bayesian nonparametric, Kingman n -coalescent, multi-resolution, ancient DNA, Gaussian process.

*The authors gratefully acknowledge partial funding from the France-Stanford Center for Interdisciplinary Studies. JAP acknowledges support from National Institutes of Health grant R01-GM-131404 and the Alfred P. Sloan Foundation. AV acknowledges partial funding from the chaire program Mathematical Modeling and Biodiversity (Ecole polytechnique, Museum National d’Histoire Naturelle, Veolia Environment, Foundation X).

1 Introduction

Statistical inference of evolutionary parameters from a sample of n DNA sequences accounts for the dependence among samples and models observed variation through two stochastic processes: an ancestral process of the sample represented by a genealogy \mathbf{g} , which depends on the *effective population size trajectory* $(N_e(t))_{t \geq 0}$, and a mutation process with a given set of parameters μ that, conditionally on \mathbf{g} , models the phenomena that have given rise to the sequences. However, state-of-the-art methodologies are not scalable to the amount of data available because the latent space of genealogies lives in a high dimensional space. In this paper, we tackle the problem of scalability from a modeling perspective: we propose a new ancestral process for heterochronous data that dramatically reduces the state space of genealogies. We complement this model with a new algorithm for fast likelihood calculations.

Inference for $(N_e(t))_{t \geq 0}$ has important applications in many fields, such as genetics, anthropology, and public health. In the absence of natural selection, the effective population size can be used to approximate census population size. While census population size estimates can be difficult to obtain due to high costs and challenging sampling designs, we can reconstruct past population sizes from observed signatures of genetic diversity in a sample of the population. For example, one can estimate the population size of a virus from genetic samples in a situation where census counts are believed to be inaccurate, as it is common during an epidemic. Inferring population size dynamics – timing of population events, growth and decline rates – rather than estimating census counts, may be of scientific interest. For example, Shapiro et al. (2004) reconstructed bison population dynamics, providing new insights into the extinction of Beringian bison. In this paper, we include two studies supporting the motivations highlighted above: in one study we analyze viral samples of SARS-CoV-2, the virus responsible for the coronavirus disease, and in a second study, we analyze ancient samples of bison in North America (dataset described by Froese et al. (2017)).

Both Bayesian and frequentist methods rely on the marginal likelihood computed by integrating over the space \mathcal{G}_n of tree topologies with n leaves, and over the space of branch lengths. That

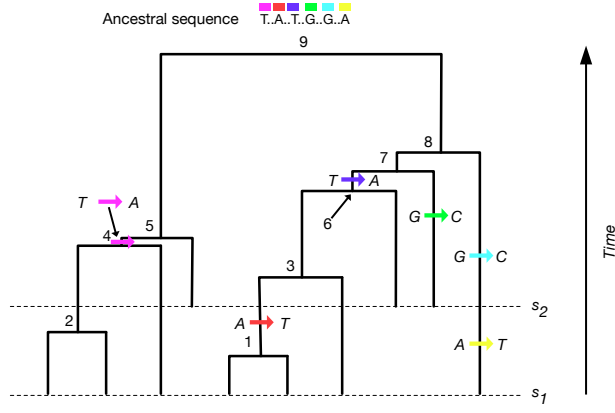


Figure 1: Coalescence and mutation. A genealogy of 10 individuals at a locus of 100 sites is depicted as a bifurcating tree. Six mutations (at different sites) along the branches of the tree give rise to the 10 sequences. The 94 sites that do not mutate are represented by black dots in the ancestral sequence. The nucleotides at the polymorphic sites are shown, and the colored arrows depict how ancestral sites are modified by mutation.

is:

$$P(\mathbf{Y} \mid (N_e(t))_{t \geq 0}, \mu) = \int_{\mathbf{g} \in \mathcal{G}_n \times \mathbb{R}_+^{n-1}} P(\mathbf{Y} \mid \mathbf{g}, (N_e(t))_{t \geq 0}, \mu) d\pi(\mathbf{g} \mid (N_e(t))_{t \geq 0}), \quad (1)$$

where the $n - 1$ random variables with values in \mathbb{R}_+ are the times between consecutive coalescence events in the genealogy, and $\pi(\cdot \mid (N_e(t))_{t \geq 0})$ denotes the probability distribution on $\mathcal{G}_n \times \mathbb{R}_+^{n-1}$ implied by the ancestral process as a function of the past population size trajectory $(N_e(t))_{t \geq 0}$. The genealogy \mathbf{g} is an auxiliary variable introduced to compute $P(\mathbf{Y} \mid \mathbf{g}, N_e(t), \mu)$ because direct calculation of the marginal $P(\mathbf{Y} \mid N_e(t), \mu)$ is intractable. The prevailing consensus in the literature is to compute (1) with π defined as the Kingman-coalescent prior law on leaf-labeled genealogies (formally introduced in Kingman (1982a,b)). However, the cardinality of \mathcal{G}_n grows superexponentially with n ($|\mathcal{G}_n| = n!(n - 1)!/2^{n-1}$), creating a computational bottleneck in the calculation of the integral (1).

An alternative to the Kingman n -coalescent is to use a lower resolution coalescent process, known as the Tajima n -coalescent (Tajima, 1983; Sainudiin et al., 2015; Palacios et al., 2019). While the state space of the Kingman n -coalescent is in bijection with the set of timed and *labeled* binary trees with n leaves, the state space of the Tajima n -coalescent is in bijection with the set of timed and *unlabeled* binary trees with n leaves. The cardinality of the space of the timed and unlabeled binary tree topologies with n leaves is given by the $(n - 1)$ -th Euler zigzag number (Disanto and Wiehe, 2013), which behaves like $2(2/\pi)^n (n - 1)!$ when n increases. While the cardinality of the space of Tajima trees still grows superexponentially in n , its rate of growth is drastically smaller than that of the space of Kingman trees. Palacios et al. (2019) show that

when all sequences are sampled at the same time, employing the Tajima coalescent allows fast inference of $(N_e(t))_{t \geq 0}$.

The main focus of this work is to develop a scalable model for sequences observed at different time points like those at the tips of the genealogy in Figure 1. The need to model heterochronous data is motivated by both applications and statistical reasons. In applications, viral samples (HIV, influenza) are routinely collected serially. Ancient DNA studies are another very active area of research in which the sampling design is intrinsically sequential. At least two statistical reasons are motivating this model. First, it usually leads to a decrease in the variance of the estimate (Felsenstein and Rodrigo, 1999). In coalescent-based inference, the smaller the number of extant lineages in a given time interval, the greater the variance of our estimate of the interval length, and consequently, the greater the variance of any estimate that depends on that length such as $(N_e(t))_{t \geq 0}$. By including heterochronous samples, we can increase the number of active lineages in the past, and thus obtain better estimates. Second, Parag and Pybus (2019) show that including heterochronous samples is necessary (in some cases) to make the parameters describing $(N_e(t))_{t \geq 0}$ identifiable.

The Tajima heterochronous n -coalescent fundamentally differs from the Tajima *isochronous* n -coalescent in that sequences sampled at different times are not exchangeable. The Tajima isochronous n -coalescent distinguishes between singletons and *vintaged* lineages, where a singleton lineage refers to a lineage that subtends a leaf in \mathbf{g} , and a vintaged lineage refers to a lineage that subtends an internal node in \mathbf{g} . Singletons are indistinguishable while vintages are labeled by the ranking of the coalescence event at which they were created. When dealing with heterochronous samples, singletons are instead implicitly labeled by their underlying sampling times. To account for this difference, we define a new Markov chain.

The present paper contains three main contributions: first, we propose a new lower-resolution continuous-time Markov chain (CTMC) which we call the Tajima *heterochronous* n -coalescent, that allows to model partially-labeled genealogies of heterochronous samples. Second, we introduce a new algorithm to compute the likelihood. Likelihood calculation in Palacios et al. (2019) called BESTT relies on a backtracking algorithm that is computationally unfeasible for sample

sizes greater than 35. Our new algorithm can accommodate up to 100 sequences. Lastly, we introduce new MCMC proposals for efficient exploration of the space of Tajima genealogies and do Bayesian nonparametric inference on $(N_e(t))_{t \geq 0}$.

The main challenge in employing Tajima genealogies for coalescent-based inference is that the sequence data can be allocated, *i.e.* mapped, to a given genealogy, in many possible ways. The allocation is necessary to compute the likelihood. To find all possible maps, Palacios et al. (2019) use a backtracking algorithm that proceeds in a bottom-up fashion: starting from the tips of the tree, the algorithm moves along the tree to the root checking for possible allocations of subsets of the data \mathbf{Y} to clades of the tree \mathbf{g} . Our proposal reverts this process and proceeds in a top to bottom fashion eliminating the backtracking step and reducing its computational complexity. While there is no exact analytical expression for the computational complexity of the backtracking algorithm, a loose upper bound of the backtracking algorithm is $\mathcal{O}(n!)$, which we bring down to $\mathcal{O}(n^2)$ with our new proposal. The lower bound is of the order of $\mathcal{O}(n)$ for both algorithms. The algorithm relies on a novel graphical representation of the data as a tree structure. We note that this tree is related to the definition of the directed acyclic graph (DAG) used in Palacios et al. (2019), with some important differences. The DAG depends on \mathbf{g} whereas the tree we introduced is solely a function of the data. This implies that we need to define it only once. Also, the DAG groups sequences differently since it does not incorporate sampling time information.

The rest of the paper proceeds as follows. In Section 2, we define the Tajima heterochronous n -coalescent. In Section 3, we introduce the mutation model we shall assume, describe the data, define the likelihood and the new algorithm to compute it. Section 4 describes the MCMC algorithm for posterior inference, and in Section 5, we present a comprehensive simulation study outlining how the model works and comparing our method to state-of-the-art alternatives. In Section 6, we analyze modern and ancient bison sequences described in Froese et al. (2017). In Section 7, we apply our method to SARS-CoV-2 viral sequences collected in France and in Germany. Section 8 concludes. An open-source implementation is given in R package `phylodyn`, which is available for download at <https://github.com/JuliaPalacios/phylodyn>.

2 The Tajima heterochronous n -coalescent

The Tajima heterochronous n -coalescent is an inhomogeneous continuous-time Markov chain that describes the ancestral relationships of a set of n individuals sampled, possibly at different times, from a very large population. The set of ancestral relationships of the sample is represented by a ranked genealogy as the one depicted in Figure 2. Every organism is dated and labeled according to the time in which the organism lived (if ancient, by radiocarbon date) or in which the living organism was sequenced. In this generalization of the Tajima coalescent, each pair of extant ancestral lineages merges into a single lineage at an instantaneous rate which depends on the current effective population size $N_e(t)$, and new lineages are added when one of the prescribed sampling times is reached. In this work, we do not model the stochasticity of sampling times but we condition on them as being fixed.

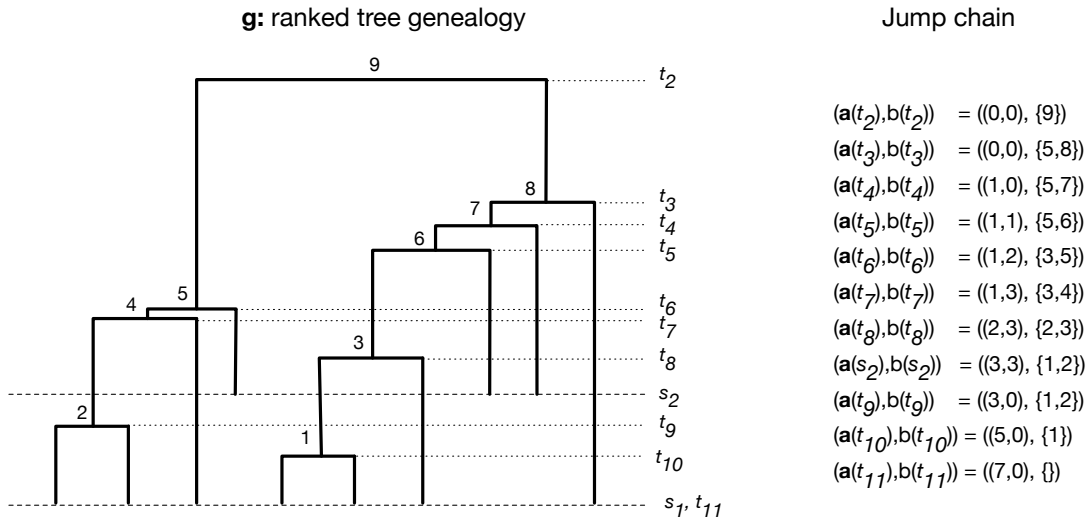


Figure 2: **Example of a Tajima heterochronous genealogy and its jump chain.** A realization of a Tajima heterochronous n -coalescent with $\mathbf{n} = (7, 3)$ and $\mathbf{s} = (s_1, s_2)$, represented as a ranked tree shape with coalescence and sampling times, denoted \mathbf{g} . The column to the right displays the corresponding jump chain (see the text for notation).

Let us introduce some notation. Let m be the number of sampling time points and n be the total number of samples. Let $\mathbf{n} = (n_1, \dots, n_m)$ denote the number of sequences collected at times $\mathbf{s} = (s_1, \dots, s_m)$, with $s_1 = 0$ denoting the present time, and $s_j > s_{j-1}$ for $j = 2, \dots, m$ (time goes from present towards the past). We refer to the sequences counted in n_i as “belonging

to sampling group s_i ”. Let $\mathbf{t} = (t_{n+1}, \dots, t_2)$ be the vector of coalescent times with $t_{n+1} = 0 < t_n < \dots < t_2$; these are the times when two lineages have a common ancestor. Note that the subscript in t_k does not indicate the current number of lineages, as it is often done in the coalescent literature, but it indicates the number of lineages that have yet to coalesce (some sequences may not have been sampled yet). We use the rank order of the coalescent events (bottom-up) to label the internal nodes of the genealogy. That is, the node corresponding to the coalescent event occurring at time t_n is labeled 1 (see t_{10} in Figure 2), the node corresponding to the coalescence event occurring at time t_{n-1} is labeled 2, *etc.* We refer to the internal node labels as *vintages* (*i.e.*, rankings).

The Tajima heterochronous n -coalescent is the process $(\mathbf{a}(t), b(t))_{t \geq 0}$ that keeps track of $\mathbf{a}(t)$, a vector of length m whose j -th position indicates the number of singletons (*i.e.*, lineages that have not been involved in a coalescence event) from sampling group s_j at time t , and $b(t)$ is the set of vintaged lineages at time t . The process starts at $t = 0$ in state $(\mathbf{a}(0) = (n_1, 0, \dots, 0), b(0) = \emptyset)$, jumps deterministically at every sampling time and jumps stochastically at every random coalescent time until it reaches the unique absorbing state $(\mathbf{a}(t_2) = (0, \dots, 0), b(t_2) = \{n - 1\})$ at time t_2 , when all n samples have a single most recent common ancestor at the root (Figure 2). At each sampling time s_i , the state of the Tajima coalescent jumps deterministically as follows:

$$(\mathbf{a}(s_i), b(s_i)) = (\mathbf{a}(s_i-) + n_i \mathbf{e}_i, b(s_i-)),$$

where $f(s_i-)$ denotes the left-limit of the function f at s_i and \mathbf{e}_i is the i -th unit vector.

Let us now turn to the embedded jump chain at coalescent times. At time t_i , two extant lineages coalesce to create a new lineage with vintage $n + 1 - i$. Four types of coalescence transitions are possible depending on which and how many sampling groups are involved: (1) two singletons of the same sampling group coalesce (up to m possible moves for the chain), (2) two singletons of different sampling groups coalesce (up to $m(m - 1)/2$ possible moves), (3) one singleton lineage and one vintaged lineage coalesce (up to m possible moves), or (4) two vintaged lineages coalesce (only one possibility because for vintages, the sampling information is irrelevant). Each pair coalesces with the same probability and the transition probabilities at

coalescent times are thus given by

$$P\left[(\mathbf{a}(t_i), b(t_i)) \mid (\mathbf{a}(t_i-), b(t_i-))\right] \quad (2)$$

$$= \begin{cases} \frac{\prod_{j=1}^m \binom{a_j(t_i-)}{a_j(t_i-) - a_j(t_i)}}{\binom{\sum_{j=1}^m a_j(t_i-) + |b(t_i-)|}{2}} & \text{if } (\mathbf{a}(t_i), b(t_i)) \prec (\mathbf{a}(t_i-), b(t_i-)) \\ 0 & \text{otherwise} \end{cases}$$

where $(\mathbf{a}(t_i), b(t_i)) \prec (\mathbf{a}(t_i-), b(t_i-))$ means that $(\mathbf{a}(t_i), b(t_i))$ can be obtained by merging two lineages of $(\mathbf{a}(t_i-), b(t_i-))$ and $|b|$ denotes the cardinality of the set b .

Observe that the quantity $\sum_{j=1}^m a_j(t_i-) + |b(t_i-)|$ appearing in (2) corresponds to the total number of extant lineages just before the event at t_i . Furthermore, since only two lineages coalesce at time t_i , at most two terms in the product appearing in the numerator of (2) are not equal to one. Finally, if $m = 1$, (2) degenerates into the transition probabilities of the Tajima isochronous n -coalescent; on the other hand if $m = n$, the process degenerates into the Kingman heterochronous n -coalescent since all singletons are uniquely labeled by their sampling times. Figure 2 shows a possible realization from the Tajima heterochronous n -coalescent.

To define the distribution of the holding times, we introduce the following notation. We denote the intervals that end with a coalescent event at t_k by $I_{0,k}$ and the intervals that end with a sampling time within the interval (t_{k+1}, t_k) as $I_{i,k}$ where $i \geq 1$ is an index tracking the sampling events in (t_{k+1}, t_k) . More specifically, for every $k \in \{2, \dots, n\}$, we define

$$I_{0,k} = [\max\{t_{k+1}, s_j\}, t_k), \quad \text{where the maximum is taken over all } s_j < t_k, \quad (3)$$

and for every $i \geq 1$ we set

$$I_{i,k} = [\max\{t_{k+1}, s_{j-i}\}, s_{j-i+1}) \text{ with the max taken over all } s_{j-i+1} > t_{k+1} \text{ and } s_j < t_k. \quad (4)$$

We also let $n_{i,k}$ denote the number of extant lineages during the time interval $I_{i,k}$. For example,

in Figure 2, in (t_9, t_8) we have $I_{0,8} = [s_2, t_8)$, $I_{1,8} = [t_9, s_2)$ and no $I_{i,8}$ for $i \geq 2$. The vector of coalescent times \mathbf{t} is a random vector whose density with respect to Lebesgue measure on \mathbb{R}_+^{n-1} can be factorized as the product of the conditional densities of t_{k-1} knowing t_k , which reads: for $k = 3, \dots, n+1$,

$$p(t_{k-1} \mid t_k, \mathbf{s}, \mathbf{n}, (N_e(t))_{t \geq 0}) = \frac{C_{0,k-1}}{N_e(t_{k-1})} \exp \left\{ - \int_{I_{0,k-1}} \frac{C_{0,k-1}}{N_e(t)} dt + \sum_{i=1}^m \int_{I_{i,k-1}} \frac{C_{i,k-1}}{N_e(t)} dt \right\}, \quad (5)$$

where $t_{n+1} = 0$ by convention, $C_{i,k} := \binom{n_{i,k}}{2}$, and the integral over $I_{i,k-1}$ is zero if there are less than i sampling times between t_k and t_{k-1} . The distribution of the holding times defined above corresponds to the same distribution of holding times in the heterochronous Kingman coalescent (Felsenstein and Rodrigo, 1999). Although the heterochronous Tajima coalescent takes value on a different state space, it remains true that every pair of extant lineages coalesces at equal rate.

Finally, given \mathbf{n} , \mathbf{s} and \mathbf{t} , a complete realization of the Tajima heterochronous n -coalescent chain can be uniquely identified with a partially labeled binary ranked tree shape g of $\mathbf{n} = (n_1, \dots, n_m)$ samples at (s_1, \dots, s_m) with its $n - 1$ coalescent transitions, so that

$$P(g \mid \mathbf{t}, \mathbf{s}, \mathbf{n}) = \prod_{i=2}^n P \left[(\mathbf{a}(t_i), b(t_i)) \mid (\mathbf{a}(t_i-), b(t_i-)) \right]. \quad (6)$$

Equation (6) gives the prior probability of the tree topology under the Tajima heterochronous n -coalescent. Putting together (5) and (6), we obtain a prior $\pi(\mathbf{g} \mid \mathbf{s}, \mathbf{n}, (N_e(t))_{t \geq 0})$

$$\pi(\mathbf{g} \mid \mathbf{s}, \mathbf{n}, (N_e(t))_{t \geq 0}) = P(g \mid \mathbf{t}, \mathbf{s}, \mathbf{n}) \prod_{k=3}^{n+1} p(t_{k-1} \mid t_k, \mathbf{s}, \mathbf{n}, (N_e(t))_{t \geq 0}), \quad (7)$$

which can be used in (1).

3 Data and Likelihood

3.1 Infinite Sites Model and the Perfect Phylogeny

We assume that the observed data \mathbf{Y} consists of n sequences at z polymorphic (mutating) sites at a non-recombining contiguous segment of DNA of organisms with low mutation rate. Under these assumptions, a widely studied mutation model is the *infinite sites model* (ISM) (Kimura, 1969; Watterson, 1975) with Poissonian mutation, which corresponds to throwing a Poisson point process of mutations on the branches of \mathbf{g} at rate μ such that every mutation occurs at a different site and no mutations are hidden by a second mutation affecting the same site.

An important consequence of the ISM is that \mathbf{Y} can be represented as an incidence matrix \mathbf{Y}_1 and a frequency counts matrix \mathbf{Y}_2 . \mathbf{Y}_1 is a $k \times z$ matrix with 0-1 entries, where 0 indicates the ancestral type and 1 indicates the mutant type; k is the number of unique sequences (or haplotypes) observed in the sample, and the columns correspond to polymorphic sites. \mathbf{Y}_2 is a $k \times m$ count matrix where the (i, j) th entry denotes how many haplotype i sequences belonging to group s_j are sampled. For example, the $n = 10$ sequences defined by the realizations of the ancestral and mutation processes depicted in Figure 1 can be summarized into \mathbf{Y}_1 and \mathbf{Y}_2 displayed in Figure 3(A).

\mathbf{Y}_1 and \mathbf{Y}_2 can alternatively be represented graphically as an *augmented perfect phylogeny* \mathbf{T} . This graphical representation of the data is exploited by our likelihood algorithm. The augmented perfect phylogeny representation is an extension of the *gene tree* or *perfect phylogeny* (Gusfield, 1991; Griffiths and Tavaré, 1994; Palacios et al., 2019) representation to the heterochronous case. The standard perfect phylogeny definition leaves out the information carried by \mathbf{Y}_2 . To our knowledge, Gusfield’s approach has never been generalized to the heterochronous case. In the augmented perfect phylogeny $\mathbf{T} = (\mathbf{V}, \mathbf{E})$, \mathbf{V} is the set of nodes of \mathbf{T} , and \mathbf{E} is the set of weighted edges. We define \mathbf{T} as follows:

1. Each haplotype labels at least one leaf in \mathbf{T} . If a haplotype is observed at k different sampling times, then k leaves in \mathbf{T} will be labeled by the same haplotype. The pair (haplotype label, sampling group) uniquely labels each leaf node.




Figure 3: Incidence matrix, frequency matrix and perfect phylogeny representation. Panel (A): data is summarized as an incidence matrix Y_1 (h denotes the haplotypes, l the segregating sites, the colors correspond to those depicted in Figure 1) and a matrix of frequencies Y_2 (s denotes the sampling group). Panel (B): T denotes the perfect phylogeny corresponding to Y_1 and Y_2 ; each of the 6 polymorphic sites labels exactly one edge. When an edge has multiple labels, the order of the labels is irrelevant. Each leaf node is labeled by a pair (haplotype, sampling time), with each haplotype possibly labeling more than one leaf nodes. Panel (C): $|E_i|$ corresponds to the number of mutations along the edge subtending node V_i in (B) and $|V_i|$ corresponds to the number of sequences descending from V_i in (B), see the text for details.

- Each of the z polymorphic sites labels exactly one edge. When multiple sites label the same edge, the order of the labels along the edge is arbitrary. Some external edges (edges subtending leaves) may not be labeled, indicating that they do not carry additional mutations to their parent node.
- For any pair (haplotype h_k , sampling group), the labels of the edges along the unique path from the root to the leaf h_k specify all the sites where h_k has the mutant type.

Figure 3(B) plots T corresponding to Y_1 and Y_2 displayed in Figure 3(A). Observe that T includes sampling information in the leaf labels. In the example, h_c labels two leaves because it is observed at times s_1 and s_2 . The corresponding edges E_3 and E_4 are unlabeled, *i.e.*, no mutations are allocated to those edges because the underlying nodes carry identical sequences (same haplotype). We “augment” Gusfield’s perfect phylogeny because the sampling information is crucial in the likelihood calculation.

T implicitly carries some quantitative information that can be quickly summarized. We denote the number of observed sequences subtended by an internal node V by $|V|$. If V is a

leaf node, $|V|$ denotes the frequency of the haplotype h observed at the corresponding sampling time s . Similarly, we denote the number of mutation labels assigned to an edge E by $|E|$. If no mutations are assigned to E , then $|E| = 0$. For parsimony, the edge that connects node V_i to its parent node is denoted by E_i . See Figure 3(C) for an example.

Gusfield (1991) gives an algorithm to construct the “traditional perfect phylogeny” \mathbf{T} in linear time. Constructing \mathbf{T} from \mathbf{T}' is straightforward since all we need is to incorporate the sampling information and add leaf nodes if a haplotype is observed at multiple sampling times. If we drew \mathbf{T}' from the data in Figure 3, it would not have node V_4 , but only a single node V_3 labeled by haplotype h_C . A description of the algorithm can be found in the supplementary material.

3.2 Likelihood

The crucial step needed to compute the likelihood of a Tajima genealogy \mathbf{g} is to sum over all possible allocations of mutations to its branches. This can be efficiently done by exploiting the augmented perfect phylogeny representation of the data \mathbf{T} and by first mapping nodes of \mathbf{T} to subtrees of \mathbf{g} . We stress that the need for an allocation step arises only when working with Tajima genealogies. In Kingman’s coalescent, tree leaves are labeled by the sequences to which they correspond, and so there is a unique possible allocation. In Tajima’s coalescent, leaves are unlabeled, creating potential symmetries in the tree, and so we have to scan all the possible ways in which the observed sequences may be allocated to \mathbf{g} .

3.2.1 Allocations

Let \mathbf{a} denote a possible mapping of nodes of \mathbf{T} to subtrees of \mathbf{g} . \mathbf{a} is encoded as a vector of length $n - 1$, where the i -th entry gives the node in \mathbf{T} which is mapped to the subtree with vintage i , \mathbf{g}_i (including the branch that subtends vintage i). Our algorithm first maps all *non-singleton* nodes \mathbf{V} of \mathbf{T} to subtrees of \mathbf{g} , that is, only nodes such that $|V| > 1$ are entries of \mathbf{a} . Singleton nodes in \mathbf{T} ($V \in \mathbf{V}$ such that $|V| = 1$) are treated separately and are initially excluded from the allocation step. For example, Figure 4 shows a possible vector \mathbf{a} whose entries are the non-singleton nodes




Figure 4: **A possible allocation of non-singleton nodes of \mathbf{V} to subtrees of \mathbf{g} .** For a given allocation \mathbf{a} (bottom figure), we display how subtrees in \mathbf{g} (identified by the vintage tag at their root – black number in the top figure) are allocated to the nodes of \mathbf{T} . Each color depicts an allocation of a subtree to a node: V_5 (red), V_0 (blue), V_1 (green) and V_2 (yellow).

V_0 , V_1 , V_2 , and V_5 of \mathbf{T} of Figure 3. We note that nodes can appear more than once in \mathbf{a} , meaning that they can be mapped to more than one subtree. On the other hand, a single node V_i is not necessarily mapped to all the vintages, leaves and internal branches of \mathbf{g}_j ; different nodes may be mapped to some subtrees of \mathbf{g}_j (including external branches), leading to a situation where V_i is mapped to only a subset of the vintages and branches constituting \mathbf{g}_j . For example, in Figure 4, V_1 is mapped to \mathbf{g}_6 and \mathbf{g}_3 , but V_2 is mapped to \mathbf{g}_1 , a subtree of both \mathbf{g}_6 and \mathbf{g}_3 ; hence V_1 is only mapped to the green part of \mathbf{g}_6 and \mathbf{g}_3 as depicted in the Figure.

The precise mapping of nodes in \mathbf{T} to subtrees of \mathbf{g} described below is needed to allocate mutations in \mathbf{T} to branches of \mathbf{g} . We will explain the allocation of mutations on \mathbf{g} for a given \mathbf{a} in the next subsection.

We now define an algorithm to efficiently find all possible mappings \mathbf{a} for a given \mathbf{g} . We encode the set of all possible \mathbf{a} , as an $\#\mathbf{a} \times (n - 1)$ matrix \mathbf{A} , where each row is a possible \mathbf{a} ($n - 1$ columns) and the number of rows $\#\mathbf{a}$ is equal to the number of possible allocations. To generate \mathbf{A} , the algorithm proceeds recursively from top to bottom in \mathbf{g} , by sweeping through subtrees in \mathbf{g} and matching them to nodes in \mathbf{T} according to parent-offspring relationships and number of descendants in both \mathbf{T} and \mathbf{g} . To be more precise, the algorithm is initialized by setting the $1 \times (n - 1)$ \mathbf{A} matrix to $\mathbf{A} = (V_0, \dots, V_0)$, i.e., V_0 is mapped to all subtrees in \mathbf{g} . The

algorithm proceeds iteratively, adding and removing rows from A , iterating over an index i going from $n - 2$ to 1. The first step is to define $A(i)$, the set of node allocations in the i -th column of \mathbf{A} . Then for all $V \in A(i)$, the algorithm iterates through the following steps: define T_V as the set of child nodes of V that have $|\mathbf{g}_i|$ descendants. If the number of child nodes of V is at least 3, V is also included in T_V . If $T_V = \emptyset$, for example if V is a leaf node, the algorithm does nothing. If $|T_V| = 1$, the algorithm replaces V by the element of T_V in the columns I of \mathbf{A} corresponding to all subtrees of \mathbf{g}_i . If $|T_V| > 1$, the matrix \mathbf{A} is augmented by stacking $|T_V| - 1$ copies of $\mathbf{A}_V(I)$, the submatrix of \mathbf{A} obtained by extracting all the row vectors whose I -th elements are V . The original submatrix $\mathbf{A}_V(I)$ is referred to as $\mathbf{A}_V^{(1)}(I)$, and $\mathbf{A}_V^{(2)}(I), \dots, \mathbf{A}_V^{(|T_V|)}(I)$ denote its copies. Lastly, the algorithm replaces V by the first element of T_V in $\mathbf{A}_V^{(1)}(I)$, by the second element of T_V in $\mathbf{A}_V^{(2)}(I)$ and so on, until the last element of T_V is substituted in $\mathbf{A}_V^{(|T_V|)}(I)$.

The simple rule described above is fast to compute but it leads to incorrect allocations because nodes may be mapped a redundant number of times. For example, it is easy to see that implementing the algorithm above, we could define an allocation \mathbf{a} where node V_2 is allocated to all subtrees of size two; however, V_2 should be allocated at most once. This issue can be avoided by noting that internal nodes in \mathbf{V} should appear in each \mathbf{a} a number of times equal to their number of child nodes minus one, while leaf nodes, say $V' \in \mathbf{V}$, should appear $|V'| - 1$ times. Hence, we complete each iteration by eliminating rows of \mathbf{A} where this rule is violated. A second elimination rule is needed to account for the constraints imposed by the sampling time information: rows are eliminated when their assignments involve nodes labeled by a sampling time s' “matched” to subtrees of \mathbf{g} that have leaf branches terminating at a different sampling time. Algorithm 1 in the Appendix summarizes the above description.

Figure 5 gives examples of possible allocations of \mathbf{T} to two different genealogies \mathbf{g} and \mathbf{g}' . The second genealogy \mathbf{g}' differs from \mathbf{g} by the order of the coalescent events 3 and 6 which are inverted. \mathbf{g} and \mathbf{g}' share the common allocation $\mathbf{a}_1 = (V_2, V_5, V_1, V_5, V_0, V_1, V_0, V_0, V_0)$; however, \mathbf{g} has a second possible allocation $\mathbf{a}_2 = (V_5, V_2, V_5, V_1, V_1, V_0, V_0, V_0, V_0)$ that it is not compatible with \mathbf{g}' . This difference is due to the fact that V_5 has three descendants belonging to sampling group s_1 , while \mathbf{g} has two subtrees with 3 leaves sampled at s_1 , and \mathbf{g}' has only one. We note that

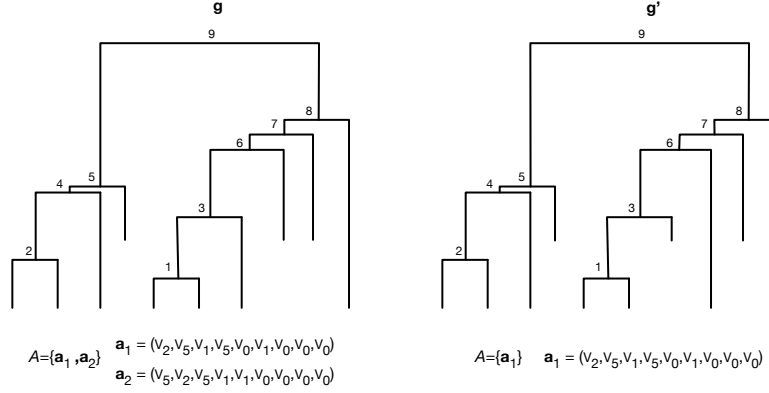


Figure 5: **Example of allocations for two distinct genealogies.** Two possible samples from the Tajima heterochronous n -coalescent. Below we list all the possible allocations of nodes of \mathbf{T} to \mathbf{g} and \mathbf{g}' . The two genealogies differ solely by the inversion of the coalescent events 3 and 6. This change gives rise to differences in the possible allocations: for example, V_5 can be mapped to the subtree defined by node 3 in \mathbf{g} but not in \mathbf{g}' .

singleton nodes also need to be allocated, both in \mathbf{a}_1 and \mathbf{a}_2 . We will elaborate on this point in the next subsection.

3.2.2 Likelihood Calculations

To calculate the likelihood, we assume the ISM of mutations and that mutations occur according to a Poisson point process with rate μ on the branches of \mathbf{g} , where μ is the total mutation rate. To compute the likelihood we need to map mutations in \mathbf{T} to branches of \mathbf{g} and this is done for each mapping \mathbf{a}_i of non-singleton nodes of \mathbf{T} to subtrees of \mathbf{g} . For every V in \mathbf{T} such that $|V| > 1$, we define \mathbf{E}_V as the set formed by the edges in \mathbf{T} that subtend singleton children of V and, with the exception of $V = V_0$, \mathbf{E}_V in addition includes the edge that subtends V . For the example in Figure 3(B), $\mathbf{E}_{V_1} = \{E_1, E_3, E_4\}$. Let \mathbf{V}^* be the set of all $V \in \mathbf{V}$ such that $|V| > 1$, then the likelihood function is defined as

$$\begin{aligned}
 P(\mathbf{Y} \mid \mathbf{g}, N_e, \mu) &= \sum_{i=1}^{\#\mathbf{a}} P(\mathbf{Y}, \mathbf{a}_i \mid \mathbf{g}, N_e, \mu) \\
 &= \sum_{i=1}^{\#\mathbf{a}} \prod_{V \in \mathbf{V}^*} P(V, \mathbf{E}_V, \mathbf{a}_i \mid \mathbf{g}, N_e, \mu),
 \end{aligned} \tag{8}$$

where we recall that $\#\mathbf{a}$ is the number of possible allocations, we have written $N_e = (N_e(t))_{t \geq 0}$, and $P(V, \mathbf{E}_V, \mathbf{a}_i \mid \mathbf{g}, N_e, \mu)$ is the probability of observing the mutations of the \mathbf{E}_V edges along the corresponding branches of \mathbf{g} defined by the mapping \mathbf{a}_i as follows.

If V has no singleton child nodes, then $\mathbf{E}_V = \{E\}$ and

$$P(V, \{E\}, \mathbf{a}_i \mid \mathbf{g}, N_e, \mu) \propto (\mu l)^{|E|} e^{-\mu \mathcal{T}}, \quad (9)$$

where l is the length of the branch in \mathbf{g} that subtends \mathbf{g}_j , j is the largest index such that $\mathbf{a}_{i,j} = V$, and \mathcal{T} denotes the length of the subtree in \mathbf{g} to which V is mapped in \mathbf{a}_i (as described in Subsection 3.2.1). For example, considering V_2 in Figure 4, we have $\mathcal{T}_2 = 2t_n + (t_{n-2} - t_n)$ and $l = (t_{n-2} - t_n)$ is the length of the branch connecting vintage 1 to vintage 3.

If node V has singleton child nodes,

$$P(V, \{E, E_{ch_1}, \dots, E_{ch_k}\}, \mathbf{a}_i \mid \mathbf{g}, N_e, \mu) \propto (\mu l)^{|E|} e^{-\mu \mathcal{T}} \sum_{\mathbf{R} \in \Pi(\mathbf{E}_V)} \prod_{j=1}^k (\mu l_{R_j})^{|E_{ch_j}|}, \quad (10)$$

where the first term on the r.h.s is defined exactly as the quantity on the r.h.s. of (9), while the second term corresponds to the probability of all possible different matchings between R_1, \dots, R_k , the first k indexes such that $\mathbf{a}_{i,R_j} = V$, and $|E_{ch_1}|, |E_{ch_2}|, \dots, |E_{ch_k}|$, the k numbers of mutations observed on the edges $E_{ch_1}, \dots, E_{ch_k}$ leading to the child nodes of V . In this expression, $\Pi(\mathbf{E}_V)$ is the set of all possible such matchings \mathbf{R} .

Before defining $\Pi(\mathbf{E}_V)$ more precisely, we make two observations. First, not all matchings are possible since not all leaf branches terminate at the same time (heterochronous sampling). Second, it is enough to consider the allocations that contribute to distinct likelihood values, *i.e.* allocations for which the underlying samples are “distinguishable” in the sense that they have a different number of mutations.

We define $\Pi(\mathbf{E}_V)$ as the set of all possible “distinct matchings of number of observed singleton mutations to singleton branches”, that is, allocations which lead to a distinct likelihood values. To construct $\Pi(\mathbf{E}_V)$, we first partition the singleton edges $E_{ch_1}, \dots, E_{ch_k}$ according to the sampling times of the corresponding nodes $V_{ch_1}, \dots, V_{ch_k}$. Let k_i be the number of nodes

in $\{V_{ch_1}, \dots, V_{ch_k}\}$ with sampling time s_i , *i.e.*, the size of each subset of the partition. We then further partition these subsets by grouping together the edges carrying the same number of mutations (defined as $|E_{ch_1}|, \dots, |E_{ch_k}|$). For each given sampling time s_j , let $k_j^{(1)}, \dots, k_j^{(m_j)}$ denote the cardinalities of the m_j sub-subsets obtained by this procedure, so that $k_j = \sum_{h=1}^{m_j} k_j^{(h)}$. The cardinality of $\Pi(\mathbf{E}_V)$ is then

$$|\Pi(\mathbf{E}_V)| = \prod_{j=1}^m \frac{k_j!}{k_j^{(1)}! \dots k_j^{(m_j)}!}, \quad (11)$$

where the product in (11) is the number of permutations with repetition of the different edges that are compatible with the data in terms of sampling times and numbers of mutations carried. Note that Equation (11) is not the same as Equation (6) in Palacios et al. (2019) because here we account for the different sampling groups. It degenerates into Equation (6) in Palacios et al. (2019) in the isochronous case.

Lastly, we note that knowing a priori the full matrix \mathbf{A} allows to compute efficiently the likelihood (8) via a sum-product algorithm. Indeed, for each $V \in \mathbf{V}^*$ there may be several rows \mathbf{a} of \mathbf{A} such that $P(V, \mathbf{E}_V, \mathbf{a} \mid \mathbf{g}, N_e, \mu)$ is the same, due to the fact that V is mapped to the same subtree in all these allocations. For such a V , one could compute the likelihood corresponding to these r allocations $\mathbf{a}'_1, \dots, \mathbf{a}'_r$ in the following way:

$$\begin{aligned} & \sum_{i=1}^r \prod_{V \in \mathbf{V}^*} P(V, \mathbf{E}_V, \mathbf{a}'_i \mid \mathbf{g}, N_e, \mu) \\ &= P(V, \mathbf{E}_V, \mathbf{a}'_1 \mid \mathbf{g}, N_e, \mu) \sum_{i=1}^r \prod_{V' \in \mathbf{V}^* \setminus \{V\}} P(V', \mathbf{E}_{V'}, \mathbf{a}'_i \mid \mathbf{g}, N_e, \mu). \end{aligned} \quad (12)$$

The exact sum-product formulation of (8) is specific to the observed \mathbf{Y} and \mathbf{A} .

4 Bayesian Model and MCMC inference

In Section 2 we have introduced a new prior for genealogies and in Section 3, we have expounded how to compute the likelihood of heterochronous data \mathbf{Y} generated by a Poisson process of mu-

tations superimposed on this new genealogy. We finally need to specify a prior distribution on $(N_e(t))_{t \geq 0}$ to complete our Bayesian model. In this paper, we follow Palacios and Minin (2013), who place a Gaussian process (GP) prior on $(\log(N_e(t)))_{t \geq 0}$ (the logarithm is used to ensure that $N_e(t) \geq 0$ for all t). We thus have:

$$\begin{aligned}
\mathbf{Y} \mid \mathbf{g}, \mu, (N_e(t))_{t \geq 0}, \mathbf{n}, \mathbf{s} &\sim \text{Poisson process} \\
\mathbf{g} \mid (N_e(t))_{t \geq 0}, \mathbf{s}, \mathbf{n} &\sim \text{Tajima heterochronous } n\text{-coalescent} \\
(\log(N_e(t)))_{t \geq 0} \mid \tau &\sim \text{GP}(0, C(\tau)) \\
\tau &\sim \text{Gamma}(\alpha, \beta)
\end{aligned} \tag{13}$$

where $C(\tau)$ is the covariance function of the Gaussian process. As in Palacios and Minin (2013), for computational convenience we use Brownian motion with covariance elements

$$\text{Cov}(\log(N_e(t)), \log(N_e(t'))) = \tau \min(t, t')$$

for any $t, t' > 0$ as our GP prior. From (13), the posterior distribution can be written as

$$\pi((\log(N_e(t)))_{t \geq 0}, \tau, \mathbf{g} \mid \mathbf{Y}, \mu) \propto P(\mathbf{Y} \mid \mathbf{g}, (\log(N_e(t)))_{t \geq 0}, \mu) \pi(\mathbf{g} \mid (\log(N_e(t)))_{t \geq 0}) \pi((\log(N_e(t)))_{t \geq 0} \mid \tau) \pi(\tau), \tag{14}$$

which we approximate via MCMC methods. Full conditionals are not available, and so we use Metropolis-within-Gibbs updates. At each MCMC iteration, we jointly update $((\log(N_e(t)))_{t \geq 0}, \tau)$ via a Split Hamiltonian Monte Carlo (HMC) (Shahbaba et al., 2014) suitably adapted to phylogenetics inference by Lan et al. (2015); then we update the topology g and \mathbf{t} . We propose two Metropolis steps to update g and \mathbf{t} . The latter may also be combined in a single step. The transitions for g and \mathbf{t} are tailored to the Tajima n -coalescent genealogies. To update g , we employ the scheme in Palacios et al. (2019). To update \mathbf{t} , we propose a new sampler, which allows to propose branch lengths that account for the observed sampling times constraints, an issue specific to heterochronous samples under the ISM assumption and detailed in the next subsection.

4.1 Constraints imposed by the ISM hypothesis

Under the ISM hypothesis, mutations partition the observed sequences into two sets: the sequences that carry the mutations and the sequences that do not. This recursive partitioning of the sequences is graphically represented by \mathbf{T} . As a consequence, not all topologies g and not all vectors \mathbf{t} are compatible with the data, i.e. have posterior probability or density greater than 0. The combinatorial constraints imposed by the ISM on the space of topologies are discussed in detail in Cappello and Palacios (2020).

The constraints on \mathbf{t} solely arise in the heterochronous case. First note that the definition of the Tajima heterochronous n -coalescent implies that there can be at most $n_1 - 1$ coalescence events before s_2 , at most $n_1 + n_2 - 1$ events before s_3 , and so on. Moreover, if there are shared mutations between some (but not all) samples with different sampling times, the maximum number of coalescent events between the involved sampling times is further restricted. In the example of Figure 3(A), there is a shared mutation l_3 between 3 samples with sampling time s_1 and a sample with sampling time $s_2 > s_1$. Out of the 7 samples obtained at time s_1 , the 3 samples that share the l_3 mutation could coalesce first some time between s_1 and s_2 (at most 2 coalescent events among the 4 sequences descending from node V_1), but they need to coalesce with the sample at time s_2 in node V_1 before they coalesce with the other 4 samples collected at time s_1 (those can coalesce at most 3 times between s_1 and s_2). Therefore, there are at most 5 coalescent events before s_2 .

To encode the constraints imposed by the sampling information, we define a vector \mathbf{c} of length m , where the i th entry denotes the maximum number of coalescent events that can happen (strictly) before time s_i for given \mathbf{Y} , \mathbf{s} and \mathbf{n} . Trivially $c_1=0$ because there are no samples. Let us stress that \mathbf{c} does not define the number of coalescent events in a given interval (a quantity determined by \mathbf{t}), but it is the maximum number of coalescent events that can happen before each sampling time to ensure compatibility with the data. In the example of Figure 3(A), we have $\mathbf{c} = (0, 5)$. Note that c_2 is 5 and not $n_1 - 1 = 6$. In the online supplementary material, we provide a greedy search algorithm to define \mathbf{c} .

4.2 Coalescent times updates

Let $\Delta \mathbf{t} := (t_n - t_{n+1}, \dots, t_2 - t_3)$ be the vector of intercoalescence times, and $(\Delta t_i)_{i \in I}$ the subvector of elements of $\Delta \mathbf{t}$ at positions $I \subseteq \{1, \dots, n-1\}$. The proposal is generated in three steps. First, we uniformly sample the number of intercoalescent times proposal moves – *i.e.*, the cardinality of I , then we uniformly choose which times to modify – *i.e.*, we define I , and lastly, we sample the proposals $(\Delta t_i)'_{i \in I}$. The first two steps balance between fast exploration of the coalescent times state-space and a high acceptance probability – few changes are expected to lead to higher acceptance rates while many changes are expected to lead to faster exploration of the state space. In our implementation we limit the maximum possible number of intercoalescent times moves to a fixed number $Z \ll n-1$. Lastly, we sample new states $(\Delta t_i)'$, for $i \in I$ from a truncated normal with mean Δt_i and standard deviation $\sigma \Delta t_i$. The left tail is truncated by a parameter lo_i , and the right tail is left unbounded. Three reasons motivate this choice: it has positive support, it can be centered and scaled around the current Δt_i using a single parameter σ , and we can set the lower bound lo_i to ensure that only compatible times \mathbf{t}' are proposed. To set the values of lo_i , we rely on \mathbf{c} , the vector that specifies the maximum number of coalescent events possible before each sampling time. We note that the elements of \mathbf{c} can be used to index coalescent times. In particular, t_{n-c_i} denotes the time of the $(c_i + 1)$ th coalescent event. For example in Figure 2, $t_{n-c_1} = t_{10}$ is the first coalescent event, and $t_{n-c_2} = t_5$ is the sixth coalescent event. Given \mathbf{c} , lo_i is set to

$$lo_i = \max_{j=1, \dots, m} \{0, \{[s_j - (t_{n-c_j} - \Delta t_i)] \mathbb{1}(i \leq c_j + 1)\}\}, \quad (15)$$

where $\mathbb{1}(i \leq c_j + 1)$ is an indicator function. Equation (15) ensures the proposal $t'_{n-c_j} \geq s_j$ for all j . Indeed, note that $t_{n-c_j} = \sum_{k=1}^{c_j+1} \Delta t_k$. Hence, if $(t_{n-c_j} - \Delta t_i) - s_j > 0$ for any given j such that $i \leq c_j + 1$, then the proposed value of $(\Delta t_i)'$ could be zero and still t'_{n-c_j} would be a compatible time. In this case, we do not need to impose any restriction on the lower bound of the truncated normal. On the other hand, if the vector considered in (15) has one or more positive values, the proposed value $(\Delta t_i)'$ should be large enough to ensure that for all sampling times

s_j , there will never be more than c_j coalescent events before s_j . In other words, we truncate the proposal distribution support to ensure the compatibility of \mathbf{t}' . We discuss how to set Z and σ in Section 5.

The transition density of coalescent times is given by

$$k(\mathbf{t}, \mathbf{t}') = \frac{1}{Z} \binom{n-1}{|I|}^{-1} \prod_{i \in I} \text{Truncated N}(\Delta t_i, \sigma \Delta t_i, lo_i, \infty), \quad (16)$$

with $\text{Truncated N}(\Delta t_i, \sigma \Delta t_i, lo_i, \infty)$ denoting a truncated normal density function with mean Δt_i , standard deviation $\sigma \Delta t_i$, lower bound lo_i and upper bound ∞ .

5 Simulations

We explore the ability of our procedure to reconstruct $(N_e(t))_{t \geq 0}$ in simulation across a range of demographic scenarios which capture realistic and challenging population size trajectories encountered in applications. The code for simulations and inference is implemented in R package `phylodyn`, which is available for download at <https://github.com/JuliaPalacios/phylodyn>.

Simulation setup. Given \mathbf{n} , \mathbf{s} , and $(N_e(t))_{t \geq 0}$, we simulate genealogies under the Tajima heterochronous n -coalescent (Section 2). Given a realized \mathbf{g} and fixed μ , we draw M mutations from a Poisson distribution with parameter μL (L is the length of the tree \mathbf{g} : the sum of all branch lengths of \mathbf{g}) and place them independently and uniformly at random along the branches of the timed genealogy. \mathbf{Y}_1 , \mathbf{Y}_2 and \mathbf{T} are then constructed as described in Section 3.1. We simulate genealogies with three population scenarios:

1. A bottleneck (“bottleneck”):

$$N_e(t) = \begin{cases} 3 & \text{if } t \in [0, 0.1), \\ 0.1 & \text{if } t \in [0.1, 0.3), \\ 2 & \text{if } t \in [0.3, \infty). \end{cases} \quad (17)$$

2. An instantaneous drop (“drop”):

$$N_e(t) = \begin{cases} 0.5 & \text{if } t \in [0, 0.5), \\ 2 & \text{if } t \in [0.5, \infty). \end{cases} \quad (18)$$

3. Two periods of constant population size with an exponential growth in between (“exp”):

$$N_e(t) = \begin{cases} 10 & \text{if } t \in [0, 0.1), \\ 10 \exp(2 - 20t) & \text{if } t \in [0.1, 0.25), \\ 0.5 & \text{if } t \in [0.25, \infty). \end{cases} \quad (19)$$

For each scenario, we generated genealogies with three numbers of leaves ($n = 14, 35, 70$) and different \mathbf{n}, \mathbf{s} as summarized in Table 1. The mutation parameter is varied to analyze the effect of the number of segregating sites on the quality of the estimation.

Table 1: **Summary of parameter values used in simulations.** List of parameters $\mathbf{n}, \mathbf{s}, \mu$ and demographic scenarios $(N_e(t))_{t \geq 0}$ used to simulate data. We report the realized number M of mutations for each of the 9 data sets (# mutations).

| | | $n = 14$ | $n = 35$ | $n = 70$ |
|------------|--------------|-------------|---------------------|--|
| Bottleneck | \mathbf{n} | (5,5,4) | (10,10,10,5) | (10,10,10,10,5,10,5,5,5) |
| | \mathbf{s} | (0,11,32) | (0,0.045,0.11,0.32) | (0,0.045,0.075,0.11,0.2,0.25,0.31,0.35,0.45) |
| | μ | 15 | 30 | 18 |
| | # mutations | 122 | 186 | 252 |
| Drop | \mathbf{n} | (8,3,3) | (10,10,10,5) | (15,10,10,15,10,5,5) |
| | \mathbf{s} | (0,0.4,0.6) | (0,0.2,0.4,0.6) | (0,0.1,0.2,0.4,0.47,0.6,0.8) |
| | μ | 12 | 12 | 12 |
| | # mutations | 121 | 127 | 190 |
| Exp | \mathbf{n} | (14) | (20,5,5,5) | (20,15,10,10,10,5) |
| | \mathbf{s} | (0) | (0,0.11,0.16,0.255) | (0,0.05,0.07,0.11,0.21,0.26) |
| | μ | 15 | 22 | 22 |
| | # mutations | 66 | 174 | 254 |

We empirically assess the accuracy of our estimates with three commonly used criteria. The first one is the sum of relative errors (SRE).

$$SRE = \sum_{i=1}^k \frac{|\hat{N}_e(v_i) - N_e(v_i)|}{N_e(v_i)},$$

where (v_1, \dots, v_k) is a regular grid of k time points, $\hat{N}_e(v_i)$ is the posterior median of N_e at

time v_i and $N_e(v_i)$ is the value of the true trajectory at time v_i . The second criterion is the mean relative width, defined by

$$MRW = \frac{1}{k} \sum_{i=1}^k \frac{|\hat{N}_{97.5}(v_i) - \hat{N}_{2.5}(v_i)|}{N(v_i)},$$

where $\hat{N}_{97.5}(v_i)$ and $\hat{N}_{2.5}(v_i)$ are respectively the 97.5% and 2.5% quantiles of the posterior distribution of $N(v_i)$. Lastly, we consider the envelope measure defined by

$$ENV = \frac{1}{k} \sum_{i=1}^k \mathbf{1}_{\{\hat{N}_{2.5}(v_i) \leq N_e(v_i) \leq \hat{N}_{97.5}(v_i)\}},$$

which measures the proportion of the curve that is covered by the 95% credible region. In this simulation study we fix $k = 100$, $v_1 = 0$ and $v_k = .8 t_2$.

MCMC tuning parameters. The posterior approximation is sensitive to both the initial values of $(\mathbf{g}, (N_e(t))_{t \geq 0})$ and the MCMC parameters. We initialize \mathbf{g} with the serial UPGMA (Drummond and Rodrigo, 2000). In addition to the usual MCMC parameters such as chain length, burnin and thinning, there are three parameters specific to our method: the HMC step size ϵ , the maximum number of intercoalescent times proposals (Z), and the standard deviation σ that parametrizes the transition kernel $k(\mathbf{t}, \mathbf{t}')$. While all three parameters contribute to the mixing of the Markov chain and acceptance rates, in our experience, ϵ and σ are the most influential. In settings similar to the ones analyzed here (time scale, type of trajectory patterns, and mutation rate), parameter values $\epsilon \in [0.03, 0.09]$, $Z \in \{1, 2, 3\}$, and $\sigma \in [0.01, 0.03]$ lead to a similar mixing of the Markov chain and accuracy (w.r.t the metrics considered). We based these guidelines on extensive simulation studies on the 9 datasets considered (Table 1), which we believe to be representative of a broad set of settings encountered in applications. In our simulations, we set $\epsilon = 0.07$, $Z = 2$, and $\sigma = 0.02$ for the “bottleneck” and “drop” trajectories, and we set $\epsilon = 0.08$, $Z = 2$, and $\sigma = 0.02$ for the “exp” trajectories.

Inference is carried out with 3×10^5 iterations for $n = 14$, 4×10^5 iterations for $n = 35$ and 5×10^5 iterations for $n = 70$. Posterior distributions are approximated after a burn-in period of

5×10^4 iterations and after thinning every 20 iterations.

Comparison to other methods. To our knowledge, no other method simultaneously deals with heterochronous data, assumes the ISM, samples Tajima genealogies and does Bayesian nonparametric inference. All available methodologies rely on the Kingman coalescent coupled with finite sites mutation models, with the Jukes-Cantor model (Jukes and Cantor, 1969) being the closest to the ISM. Hence, it is not fully possible to isolate the impact of using Tajima’s genealogies in lieu of Kingman’s. Nevertheless, we include some alternative estimates for completeness. We compare our results to two popular methodologies implemented in BEAST (Drummond et al., 2012): the Bayesian Skyline (SKY) (Drummond et al., 2005) and the Gaussian Markov Random Field Skyride (GMRF) (Minin et al., 2008). For the SKY and GMRF, we use the Jukes Cantor mutation model (Jukes and Cantor, 1969), and approximate posterior distributions with 10^7 iterations after a burn-in period of 10^6 iterations and after thinning every 10^3 iterations. We also compare our results to an oracle estimator that infers $(N_e(t))_{t \geq 0}$ from the “true” \mathbf{g} . The oracle estimation is obtained using the method of Palacios and Minin (2012) with the same Gaussian process prior on $(N_e(t))_{t \geq 0}$. Note that the goal of the comparison is not to determine whether our method is superior, rather see if the performance of Tajima-based inference is in line with the results obtained through two popular Kingman-based methods in some challenging population scenarios.

Results. The results of the nine curves estimated with our method are plotted in Figure 6. The supplementary material includes the plots for SKY and GMRF. True trajectories are depicted as dashed lines, posterior medians as black lines and 95% credible regions as gray shaded areas. Table 2 summarizes SRE, MRW, and ENV for the 9 simulated data sets achieved with our method (“Tajima”), SKY, GMRF, and “Oracle”. SKY estimates for “Exp” $n = 14$ are not included because we could not obtain convergent runs. Accuracy increases with sample size: credible regions shrink substantially in all three scenarios. As n increases, posterior medians track more closely the true trajectories. It is well known in the literature that abrupt population size changes are the most difficult to recover. The “drop” and “bottleneck” scenarios are less accurate for $n = 14$, as exhibited by the wider credible region. We recover the bottleneck (panel first row and

first column), but we do not recover the instantaneous drop (panel first row and third column).

Table 2: Simulation: performance comparison between Tajima and Oracle models. We compute three statistics - envelope (ENV), sum of relative errors (SRE) and mean relative width (MRW) - for three population trajectories (Bottle, Exp, Drop) and three population sizes ($n = 14, 35, 70$). Tajima refers to the estimation of $(N_e(t))_{t \geq 0}$ through our model, SKY refers to Drummond et al. (2005), GMRF to Minin et al. (2008). Oracle refers to the method of Palacios and Minin (2012) (known \mathbf{g}). Bold depicts the method with the best performance (excluding the “oracle”). SKY “Exp” $n = 14$ results are not included because we could not obtain convergent runs.

| | n | %ENV | | | | SRE | | | | MRW | | | |
|--------|----|--------|------------|------------|------------|--------|--------------|--------------|--------------|-----------|-------------|-------------|--------------|
| | | Oracle | Tajima | SKY | GMRF | Oracle | Tajima | SKY | GMRF | Oracle | Tajima | SKY | GMRF |
| Bottle | 14 | 100 | 98 | 97 | 92 | 450.21 | 127.47 | 75.14 | 94.27 | 137943.26 | 6279.04 | 7.01 | 20.28 |
| | 35 | 100 | 97 | 93 | 92 | 121.27 | 53.96 | 58.89 | 56.29 | 192.46 | 33.74 | 2.25 | 6.25 |
| | 70 | 97 | 92 | 99 | 91 | 111.01 | 109.02 | 24.86 | 45.83 | 34.76 | 24.53 | 2.33 | 3.15 |
| Exp | 14 | 100 | 100 | - | 100 | 34.36 | 54.11 | - | 49.81 | 19.4 | 39.28 | - | 11.32 |
| | 35 | 100 | 91 | 100 | 100 | 29.16 | 69.34 | 50.45 | 26.06 | 10.12 | 14.59 | 8.32 | 3.43 |
| | 70 | 100 | 100 | 100 | 100 | 29.42 | 48.5 | 49.24 | 29.8 | 4.06 | 4.14 | 3.96 | 2.76 |
| Drop | 14 | 100 | 100 | 98 | 99 | 25.65 | 32.58 | 105.93 | 85.55 | 19.38 | 5.47 | 8.06 | 14.21 |
| | 35 | 99 | 97 | 100 | 100 | 21.53 | 17.26 | 30.9 | 21.87 | 3.22 | 3.39 | 3.32 | 4.41 |
| | 70 | 98 | 95 | 100 | 99 | 17.48 | 31.24 | 24.36 | 28.48 | 2.57 | 2.27 | 2.48 | 2.41 |

Table 5 largely confirms the analysis of Figure 6. Recall that SKY and GMRF rely on Kingman coalescent rather than Tajima coalescent; SKY and GMRF methods assume a different mutation model, whereas Oracle relies on knowing the true \mathbf{g} rather than computing its posterior. First, no method unequivocally outperforms the others. The oracle methodology is the method with the best overall performance more frequently. Surprisingly, the advantage of knowing \mathbf{g} is not as big as one would expect. Both SKY and GMRF have much narrower credible regions for the bottleneck trajectory. On the other hand, Tajima has the best overall performance in the “drop” trajectory (low SRE and MRW). Note that 100% ENV is not always an indicator of accuracy because it can be achieved with a very wide credible region.

Lastly, note that no method outperforms the others. This is consistent with theoretical expectations as we are comparing two resolutions of the same ancestral process. Reassuringly, Tajima-based estimates are competitive with Kingman-based estimates. The current simulation study cannot single out the benefit of employing Tajima vs Kingman topologies because no available implementations rely on an identical mutation model and MCMC scheme. This analysis is out of the current scope and will be the subject of future research.

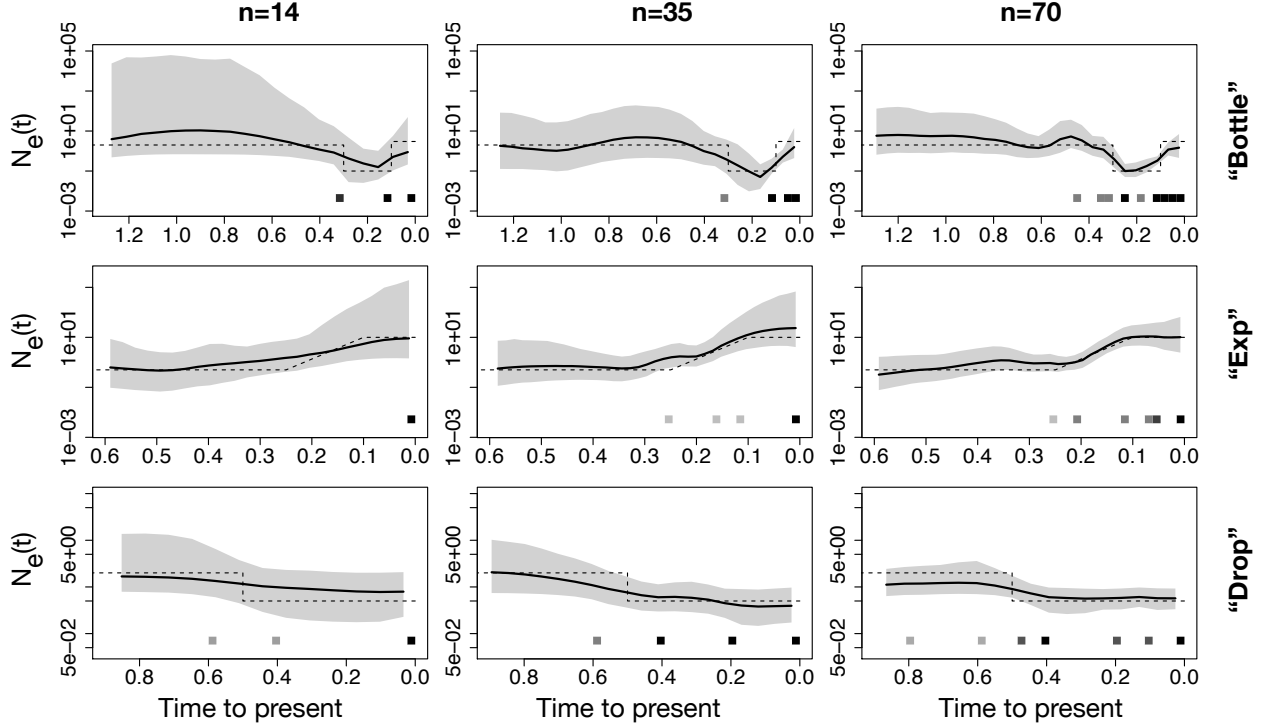


Figure 6: **Simulation: effective population size posterior medians from different trajectories and sample sizes.** $(N_e(t))_{t \geq 0}$ posterior distribution from simulated data with three population size trajectories (rows) - bottleneck (“Bottle”), exponential growth (“Exp”) and instantaneous fall (“Drop”) - different sample sizes (columns) - $n = 14$, $n = 35$ and $n = 70$. Posterior medians are depicted as solid black lines and 95% Bayesian credible intervals are depicted by shaded areas. **n** and **s** are depicted by the heat maps at the bottom of each panel: the squares along the time axis depicts the sampling time, while the intensity of the black color depicts the number of samples. More details are given in Table 1.

6 North American Bison data

Recent advances in molecular and sequencing technologies allow recovering genetic material from ancient specimens (Pääbo et al., 2004). In this section, we analyze modern and ancient bison sequences. These mammals offer a case study of a population experiencing a population growth followed by a decline. It was a long-standing question whether the drop was instigated by human intervention or by environmental changes. Shapiro et al. (2004) first reconstructed the genetic history of Beringian bison. Their estimate for the start of the decline supports the environmental hypothesis, in particular, they suggest that the decline may be due to environmental events preceding the last glacial maxima (LGM). This data-set has been the subject of extensive




Figure 7: **Bison study: data (Froese et al., 2017).** Perfect phylogeny T of 38 bison sequences selected from Froese et al. (2017) data-set. Node labels depict the number of sequences subtending that node. The 91 mutations are allocated along the edges of T (all of them are single digits). Sampling information are not written in this Figure. The two vectors n and s are represented by the columns to the right of T . Sampling times are obtained by radiocarbon dating. The scale is number of years before present.

research in the past decade.

We analyze new bison data recently described by Froese et al. (2017). We fit our coalescent model to these sequences and estimate population size dynamics. To our knowledge, there is no phylodynamics analysis of this data set in the literature. Two motivations underlie this study: first, Shapiro et al. (2004) sequences include 602 base pairs from the mitochondrial control region, while Froese et al. (2017) provide the full mitochondrial genome (16322 base pairs after alignment); second, we are interested in testing whether the previously published overwhelming evidence in favor of the environmentally induced population decline is confirmed by this new data. Froese et al. (2017) data comprises 50 sequences (14 modern and 36 ancient). DNA was extracted from bison specimens from Canada (28, three locations), USA (9, two locations); Siberia (7, three locations), and unknown locations (5). It includes sequences of 37 *Bison priscus* (extinct ancient bison), 1 *Bison latifrons* (extinct ancient bison), 11 *Bison bison* (modern bison), and 4 *Bos grunniens* (control group).

We selected 38 out of 50 sequences. We removed the control group sequences and the Siberian sequences to analyze samples from a single population (Froese et al. (2017) (Figure 1) suggested population structure). We removed the *Bison latifrons* sequence because it has 3803 ambiguities *i.e.*, sites in a sequence that cannot be unambiguously assigned to a unique nucleotide




Figure 8: **Bison in North America: effective population size “expected trajectory” and posterior median estimates from Froese et al. (2017) data set.** The first panel depicts a sketch of a “consensus population trajectory” obtained from the phylodynamics study of the data of Shapiro et al. (2004) in Faulkner et al. (2020). The second and third panels display estimated posterior medians of $(N_e(t))_{t \geq 0}$ (as black curves) obtained from $n = 38$ ancient and modern sequences from North America specimens in Froese et al. (2017) data. The second panel corresponds to our method and the third panel to GMRF. The posterior medians are depicted as solid black curves and the 95% Bayesian credible regions are depicted by shaded areas. **n** and **s** are depicted by the heat maps at the bottom of the last two panels: the squares along the time axis depicts the sampling time, while the intensity of the black color depicts the number of samples. More details are given in Figure 7.

basis at sites where all the other samples have valid entries. Out of the 94 observed polymorphic sites, we retain 91 sites compatible with the ISM assumption. To encode data in the 0 – 1 incidence matrix representation \mathbf{Y}_1 , we use the root of the UPGMA tree reconstructed using R function `upgma` (`phanghorn`) as the ancestral state. Figure 7 displays the perfect phylogeny \mathbf{T} and the vectors **s** and **n**.

For our inference procedure, we set $\epsilon = 0.09$, $Z_1 = 2$, $\sigma = 0.02$, and approximated the posterior distribution with 1.5×10^6 iterations after a burn-in of 8×10^5 and after thinning every 200 iterations. As a comparison, we ran GMRF on BEAST and approximated the posterior distribution with 1×10^7 iterations after a burn-in of 1×10^6 and after thinning every 1000 iteration. We used the default values for all GMRF hyperparameters. We initialized both methods with the same genealogy (serial UPGMA). To compute the likelihood, we used the BEAST mutation rate estimate per site per year of 2.52×10^{-8} .

The first panel of Figure 8 plots a summary of the effective population size pattern recovered by a recent analysis of Shapiro et al. (2004) data by Faulkner et al. (2020). While the precise timings and the details of the trajectory differ from method to method, the broad patterns are

consistent. The population peak is estimated to be between 41.6 and 47.3 kya. The timing of the start of the decline is the main feature of interest. We plot the posterior medians (black lines) of $(N_e(t))_{t \geq 0}$ along with the 95% credible regions (gray area) obtained from posterior samples by sampling Tajima’s trees (“Tajima”, second panel) and Kingman’s trees (“GMRF”, third panel).

Both our method and GMRF recover the pattern described in the first panel. We detect the population decline only up to about 60kya ago, afterward the median trajectory is quite flat while the credible regions are wide. This can be explained by the fact that we have no samples from 42kya to 128.5kya. On the other hand, GMRF detects more clearly the population decline. The GMRF median time estimate of the population peak is 29.6 kya, while the median time estimate for our method is 29.7 kya. Thus, the estimates of the main event of interest, the population decline, are practically identical. The difference between the estimates obtained analyzing 2017 data differ substantially from the estimates of a population peak between 41.6 and 47.3 kya obtained analyzing the 2004 data.

The LGM in the Northern hemisphere reached its peak between 26.5 and 19 kya (Clark et al., 2009). Hence, the analysis of the 2017 data still supports the hypothesis of a decline that initiated before the LGM. However, our estimates suggest an initial decline much closer to the LGM peak than the analysis of the 2004 data. Human arrival in North America via the Berigian bridge route should have happened around 14 – 16 kya (Llamas et al., 2016). Therefore, despite the mismatch of the timing, the human-induced decline hypothesis has little evidence also according to our analysis of this new dataset.

7 SARS-CoV-2

SARS-CoV-2 is the virus causing the pandemic of novel coronavirus disease in 2019-2020 and it is of interest to explore the utility of viral molecular sequences for surveillance during the outbreak of the epidemic. Here, we analyze 123 whole genome sequences collected in France, and 32 sequences collected in Germany that were made publicly available in the GISAID EpiCov database (Shu and McCauley, 2017). We note that our estimates may not reflect the whole countries effective population sizes but simply local effective population trajectories of the locations

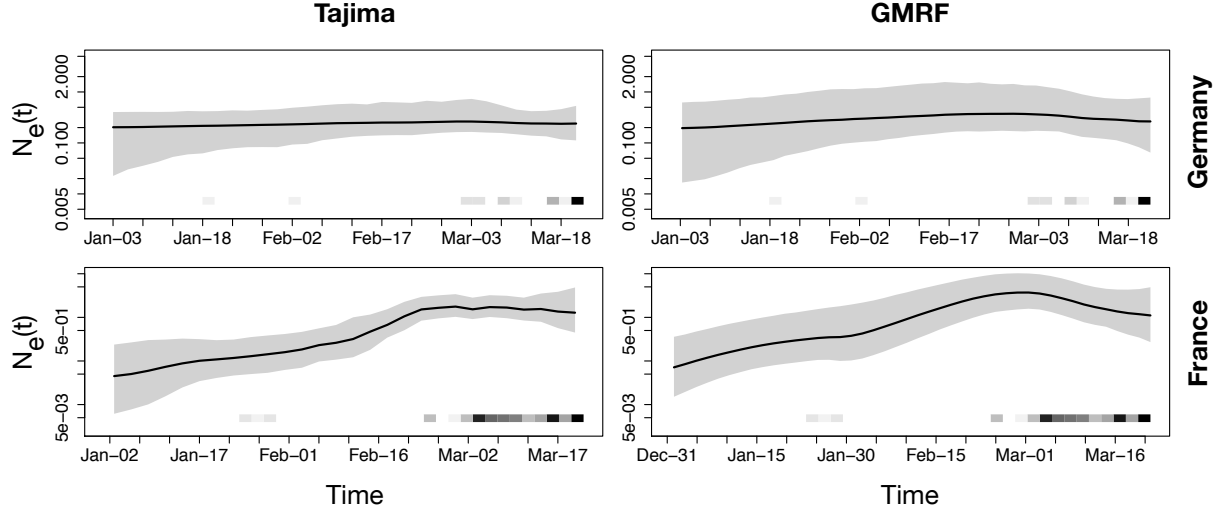


Figure 9: **2019-2020 SARS-CoV-2: $N_e(t)$ posterior median estimates from SARS-CoV-2 GISAID data sets from France and Germany.** Black curves in the first row panels depict estimated posterior medians of $(N_e(t))_{t \geq 0}$ obtained from $n = 32$ viral samples from Germany. Black curves in the second row panels depict estimated posterior medians of $(N_e(t))_{t \geq 0}$ obtained from $n = 123$ viral samples from France. The left column corresponds to our method's results and the second column to GMRf results. The posterior medians are depicted as solid black lines and the 95% Bayesian credible regions are depicted by shaded areas. \mathbf{n} and \mathbf{s} are depicted by the heat maps at the bottom of the last two panels: the squares along the time axis depicts the sampling time, while the intensity of the black color depicts the number of samples.

in which our samples were obtained. We only analyzed high coverage sequences with more than 25000 base pairs and performed multiple sequence alignment with Mafft (Katoh and Standley, 2013). To encode nucleotide data as binary sequences \mathbf{Y}_1 , we used the GenBank MN908947 (Wu et al., 2020) sequence as ancestral reference and eliminated sites that were not present in the ancestral sequence. The numbers of variable sites observed are 137 and 45 for France and Germany respectively. The observed patterns of mutations in both datasets are compatible with the ISM (no site was further removed). The Gisaidd reference numbers of the sequences included in this study and data access acknowledgment are included in the supplementary material. We note that observed differences may be caused by sequencing errors and these are being ignored in our study. The heat maps included in each panel of Figure 9 show the sampling frequency information. In the French dataset, 109 out of 123 samples were collected in March (at least one sample every day from 03/01/20 to 03/22/20), 9 in February (spread over 5 different dates), 5

in January (spread over 3 days, oldest sample dated 01/23/20). In the German dataset, 25 out of 32 samples were collected in March (spread over 7 different dates and 03/16/20 last sampling day), 6 in February (spread over 4 dates), 1 in January (oldest sample 01/28/20). We include in each dataset the reference sequence.

For our inference procedure, we set $\epsilon = 0.11$, $Z_1 = 2$, $\sigma = 0.02$, and approximate the posterior distribution with 1.4×10^6 iterations after a burn-in of 8×10^5 and after thinning every 100 iterations. For comparison, we ran GMRF on BEAST assuming the HKY mutation model (Hasegawa M, 1985) as proposed in previous studies (Scire et al., 2020) and approximate the posterior distribution with 5×10^7 iterations after a burn-in of 5×10^6 and after thinning every 1000 iteration. We used the default values for all GMRF hyperparameters. We initialized both methods with the serial UPGMA genealogy (Drummond and Rodrigo, 2000). BEAST estimates a mutation rate of 5.99×10^{-4} mutations per site per year in the French dataset, and 7.41×10^{-4} mutations per site per year in the German dataset. Our estimate follows the method discussed by Rambaut et al. (2016) obtained by regressing the Hamming distance of the sequences to the ancestral reference sequence on time difference between the sampling times and the reference sampling time. We estimated a mutation rate of 1.07×10^{-3} mutations per site per year in the French dataset, and 8.54×10^{-4} mutations per site per year in the German dataset.

We show the estimates of effective population size with our method in the first column of Figure 9 and with BEAST in the second column. Results for Germany correspond to the first row and for France to the second row. Both analyses of the French dataset exhibit exponential growth from mid-December of 2019 to the end of February (Tajima estimate of median population peak is 2020/02/29, GMRF estimate is 2020/03/1). Following the exponential growth, both methods suggest a decline. Both analyses of the German dataset recover nearly constant trajectories, possibly due to sampling time concentration in mid-march and spatial sampling concentration in Duesseldorf (see online supplementary material for details).

A final remark. Our estimates should be interpreted as estimates of genetic diversity over time and not as number of infections. Our model ignores recombination, population structure and selection. Viruses tend to exhibit antigenic drifts, selective sweeps, and tend to cluster spa-

tially following migration events (Rambaut et al., 2008). All these aspects may hinder the use of coalescent-based models to analyze viral population size dynamics. Indeed, the scientific knowledge on this virus is still limited and the validity of our model assumptions to SARS-CoV-2 is an active area of research.

8 Discussion

We have introduced a new methodology for Bayesian nonparametric inference of population size trajectory from heterochronous DNA sequences collected at a single non-recombining locus. The main focus of this work is scalability. In this respect, we developed a fast alternative to the Kingman’s coalescent that can be used for nonparametric inference of serially sampled sequences. We also developed a fast algorithm to compute the likelihood of Tajima genealogies, which is in itself a relevant contribution to the literature.

We applied our method to a recent data set including modern and ancient bison sequences. There has been a lot of interest in determining whether the decline in the bison population was human-induced or climate-induced. Genetic evidence supported the environmental hypothesis, estimating the population peak to be approximately 45kya. Our analysis reconstructed a similar population size pattern. However, we estimated the peak to be about 29.5 kya. These analyses confirm that the population decline started sometimes before the LGM. We believe that this brings further genetic evidence to the environmentally induced population decline hypothesis.

This paper makes important steps in the direction of a more scalable coalescent-based inference. However, the Tajima heterochronous n -coalescent has some limitations which need to be addressed. An obvious one is that we do not model population structure. In the ancient bison application, we removed the Beringian sequences, keeping only North-American ones, because population structure violated the assumption of the standard coalescent, and consequently of any of its “resolutions”. We have also stressed the importance of this feature in the analysis of viral data.

In addition, throughout the paper, we assumed the infinite site model of mutation. This prevented us from analyzing the original bison data of Shapiro et al. (2004), as well as many other

data that violate the ISM assumption. Given the promise of Tajima based inference shown here, incorporating other mutation models seems to be an interesting avenue of research.

Appendix

Algorithm 1 Description of the algorithm to define the allocation matrix

Inputs: T, s

Output: A

1. Initialize $A = (V_0, \dots, V_0)$
 2. **For** $i = n - 2$ to 1 **do**
 - (a) Define $A(i)$ unique nodes in the i th column of A
 - (b) **For all** $V \in A(i)$ **do**
 - i. Define T_V , set of (non-singleton) child nodes of V having $|g_i|$ descendants
 - ii. Include V in T_V if it has more than two child nodes
 - iii. Define I , set of vintages corresponding to all subtrees of g_i
 - iv. **If** $|T_V| = 0$: do nothing
Else if $|T_V| = 1$: set column $A_V(\cdot, I)$ equal to T_V
Else if $|T_V| > 1$: copy $A_V(\cdot, I)$ $|T_V| - 1$ times, attach the copies to A and set each copy equal to one element of T_V
 - v. Eliminate rows in A where V appears too frequently (rule in the paper)
 - vi. Eliminate rows not compatible with s and t
 3. Return A .
-

References

- Cappello, L. and Palacios, J. A. (2020), ‘Sequential importance sampling for multi-resolution Kingman-Tajima coalescent counting’, Annals of Applied Statistics **in press**.
- Clark, P. U., Dyke, A. S., Shakun, J. D., Carlson, A. E., Clark, J., Wohlfarth, B., Mitrovica, J. X., Hostetler, S. W. and McCabe, A. M. (2009), ‘The last glacial maximum’, Science **325**(5941), 710–714.
- Disanto, F. and Wiehe, T. (2013), ‘Exact enumeration of cherries and pitchforks in ranked trees under the coalescent model’, Mathematical biosciences **242**(2), 195–200.

- 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**(5), 1185–1192.
- Drummond, A. and Rodrigo, A. G. (2000), ‘Reconstructing genealogies of serial samples under the assumption of a molecular clock using serial-sample UPGMA’, Molecular Biology and Evolution **17**(12), 1807–1815.
- Drummond, A., Suchard, M., Xie, D. and Rambaut, A. (2012), ‘Bayesian phylogenetics with BEAUti and the BEAST 1.7’, Molecular Biology and Evolution **29**(8), 1969–1973.
- Faulkner, J. R., Magee, A. F., Shapiro, B. and Minin, V. N. (2020), ‘Horseshoe-based Bayesian nonparametric estimation of effective population size trajectories’, Biometrics **in press**.
- Felsenstein, J. and Rodrigo, A. G. (1999), Coalescent approaches to HIV population genetics, in ‘The Evolution of HIV’, Johns Hopkins University Press, pp. 233–272.
- Froese, D., Stiller, M., Heintzman, P. D., Reyes, A. V., Zazula, G. D., Soares, A. E., Meyer, M., Hall, E., Jensen, B. J., Arnold, L. J. et al. (2017), ‘Fossil and genomic evidence constrains the timing of bison arrival in North America’, Proceedings of the National Academy of Sciences **114**(13), 3457–3462.
- Griffiths, R. C. and Tavaré, S. (1994), ‘Sampling theory for neutral alleles in a varying environment’, Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences **344**(1310), 403–410.
- Gusfield, D. (1991), ‘Efficient algorithms for inferring evolutionary trees’, Networks **21**(1), 19–28.
- Hasegawa M, Kishino H, Y. T. (1985), ‘Dating of the human-ape splitting by a molecular clock of mitochondrial DNA’, Journal of Molecular Evolution **2**, 160–164.
- Jukes, T. H. and Cantor, C. R. (1969), ‘Evolution of protein molecules’, Mammalian protein metabolism **3**(21), 132.
- Katoh, K. and Standley, D. M. (2013), ‘Mafft multiple sequence alignment software version 7: improvements in performance and usability’, Molecular biology and evolution **30**(4), 772–780.
- Kimura, M. (1969), ‘The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations’, Genetics **61**(4), 893.
- Kingman, J. F. (1982a), ‘On the genealogy of large populations’, Journal of Applied Probability **19**(A), 27–43.
- Kingman, J. F. C. (1982b), ‘The coalescent’, Stochastic processes and their applications **13**(3), 235–248.

- 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**(20), 3282–3289.
- Llamas, B., Fehren-Schmitz, L., Valverde, G., Soubrier, J., Mallick, S., Rohland, N., Norderfelt, S., Valdiosera, C., Richards, S. M., Rohrlach, A. et al. (2016), ‘Ancient mitochondrial DNA provides high-resolution time scale of the peopling of the americas’, Science advances **2**(4), e1501385.
- 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**(7), 1459–1471.
- Pääbo, S., Poinar, H., Serre, D., Jaenicke-Després, V., Hebler, J., Rohland, N., Kuch, M., Krause, J., Vigilant, L. and Hofreiter, M. (2004), ‘Genetic analyses from ancient DNA’, Annu. Rev. Genet. **38**, 645–679.
- 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, AUAI Press, Arlington, Virginia, United States, pp. 726–735.
- Palacios, J. A. and Minin, V. N. (2013), ‘Gaussian process-based Bayesian nonparametric inference of population size trajectories from gene genealogies’, Biometrics **69**(1), 8–18.
- Palacios, J. A., Véber, A., Cappello, L., Wang, Z., Wakeley, J. and Ramachandran, S. (2019), ‘Bayesian estimation of population size changes by sampling Tajimas trees’, Genetics **213**(2), 967–986.
- Parag, K. V. and Pybus, O. G. (2019), ‘Robust design for coalescent model inference’, Systematic biology **68**(5), 730–743.
- Rambaut, A., Lam, T. T., Max Carvalho, L. and Pybus, O. G. (2016), ‘Exploring the temporal structure of heterochronous sequences using tempest (formerly path-o-gen)’, Virus evolution **2**(1), vew007.
- 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**(7195), 615–619.
- 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**(6), 1207–1247.
- Scire, J., Vaughan, T. G. and Stadler, T. (2020), ‘Phylogenetic analyses based on 93 genomes’.

- Shahbaba, B., Lan, S., Johnson, W. O. and Neal, R. M. (2014), ‘Split Hamiltonian Monte Carlo’, Statistics and Computing **24**(3), 339–349.
- 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. et al. (2004), ‘Rise and fall of the beringian steppe bison’, Science **306**(5701), 1561–1565.
- Shu, Y. and McCauley, J. (2017), ‘Gisaid: Global initiative on sharing all influenza data—from vision to reality’, Eurosurveillance **22**(13).
- Tajima, F. (1983), ‘Evolutionary relationship of dna sequences in finite populations’, Genetics **105**(2), 437–460.
- Watterson, G. (1975), ‘On the number of segregating sites in genetical models without recombination’, Theoretical Population Biology **7**(2), 256–276.
- 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**(7798), 265–269.

SUPPLEMENTARY MATERIAL

Algorithm Description: Augmented Perfect Phylogeny. The algorithm below uses Gusfield’s perfect phylogeny as an input, duplicates nodes corresponding to haplotypes that are sampled more than once, and returns the augmented perfect phylogeny \mathbf{T} .

Algorithm 2 Define \mathbf{T}

Inputs: \mathbf{T}' (Gusfield, 1991), \mathbf{s} , \mathbf{Y}_2

Output: \mathbf{T}

1. **For** $i = 1$ to k **do**
 - If** h_i is observed at multiple sampling times (from \mathbf{Y}_2):
[let w.l.o.g. r be the number of sampling groups in which h_i is observed, and s_{i_1}, \dots, s_{i_r} the corresponding sampling times]
 - (a) Take the leaf node V' in \mathbf{T}' labeled by h_i (each haplotype labels a unique node in Gusfield \mathbf{T}')
 - (b) **If** $|E'| = 0$: make $r - 1$ copies of V ($r - 1$ nodes with edges connecting them to the same parent of V with no edge labels). Then label each of these nodes uniquely with a pair $(h_i, s_{i_1}), \dots, (h_i, s_{i_r})$
Else if $|E'| \geq 1$: create r new nodes with unlabeled edges connecting them to V' . Then label each of these nodes in a unique way with a pair $(h_i, s_{i_1}), \dots, (h_i, s_{i_r})$
 - Else if** h_i is observed at a single sampling time (from \mathbf{Y}_2):
 - (a) Identify V' in \mathbf{T}' labeled h_i
 - (b) Label V' with a pair $(h_i, \text{its corresponding sampling time})$
 2. **Return** \mathbf{T} .
-

Algorithm Description: Computing \mathbf{c} . We compute \mathbf{c} through greedy search. The idea is simple: it is not possible to build a compatible topology g conditionally on an incompatible vector \mathbf{t} . We initially assume that ISM does not impose any constraints on \mathbf{t} , check if we can build a compatible topology, if we are, it will mean that indeed ISM does not impose constraints, if not, it will mean that we need to add some constraints. We continue iteratively until we manage to sample a compatible g . To do this process, we consider one sampling group at a time. We define a vector **add** of length m whose i th entry is the number of coalescent events that happens before s_i . Note that if we are interested in sampling g (ignoring branch information), **add** is the only time information we need. We can sample compatible g ’s through a simple extension of an Algorithm 2 in Cappello and Palacios

(2020). We refer to that paper for details.

Algorithm 3 Define **c**

Inputs: **T**, **s**

Output: **c**

1. Initialize **c** = $(n_1 - 1, n_1 + n_2 - 1, \dots, \sum_{i=1}^{m-1} n_i - 1, n - 1)$
 2. **For** $i = 1$ to $m - 1$ **do**
 - (a) Set **add** = $(0, \dots, add_i = 0, add_{i+1} = \sum_{i=1}^i n_i - 1, \dots, add_m = \sum_{i=1}^i n_i - 1)$
 - (b) Given **add**, try to sample a compatible topology g
 - (c) **If** g compatible: set $c_i = add_{i+1}$
Else if g not compatible: set $add_{i+1} = a_{i+1} - 1, \dots, add_m = add_m - 1$ and return to (b)
 3. Return **c**.
-

Simulation study: Plots of the estimates obtained from BEAST of the methods GMRF and SKY for the examples discussed in Section 5.

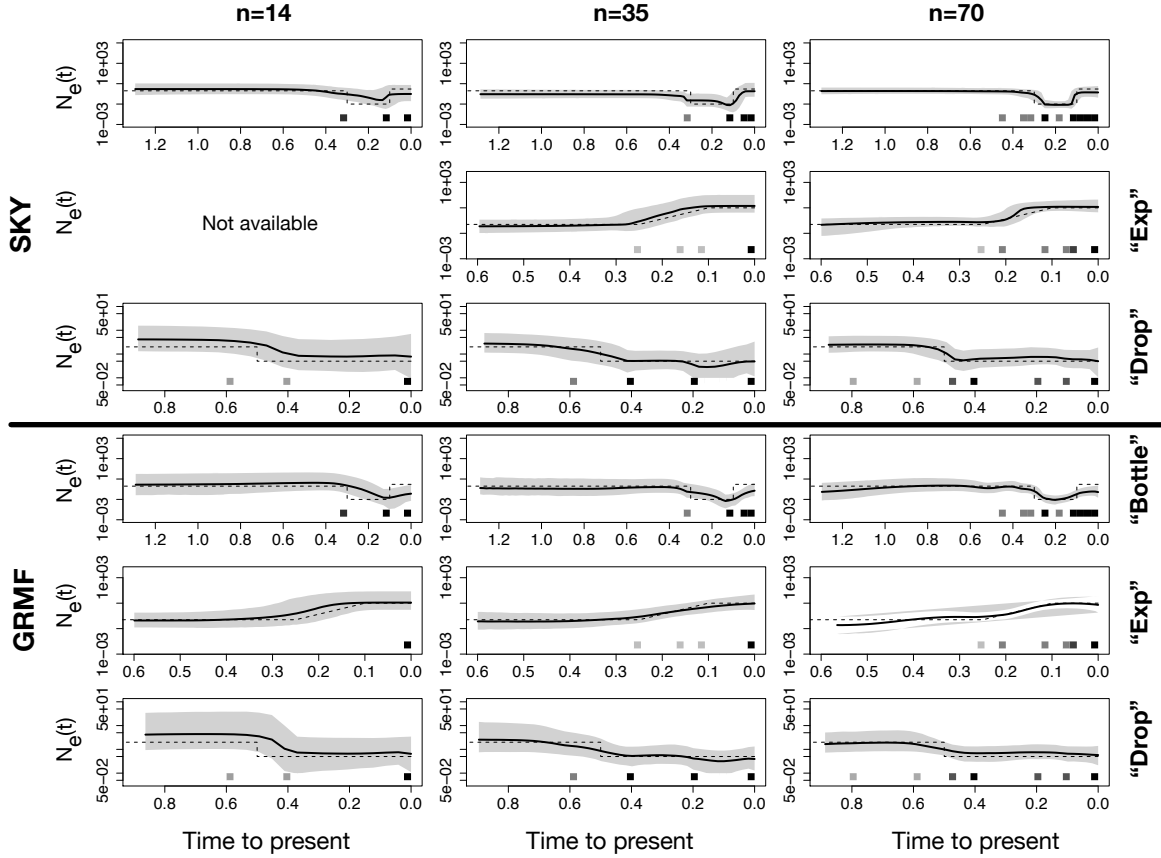


Figure 10: **Simulation: effective population size posterior medians from different trajectories and sample sizes.** $(N_e(t))_{t \geq 0}$ posterior distribution from simulated data of under three population size trajectories (rows) - bottleneck (“Bottle”), exponential growth (“Exp”) and instantaneous fall (“Drop”) - different sample sizes (columns) - $n = 14$, $n = 35$ and $n = 70$. Posterior medians are depicted as solid black lines and 95% Bayesian credible intervals are depicted by shaded areas. **n** and **s** are depicted by the heat maps at the bottom of each panel: the squares along the time axis depicts the sampling time, while the intensity of the black color depicts the number of samples. Top three rows panels depict estimates obtained through SKY, bottom three rows depicts estimates obtained through GMRF. More details are given in Table 1.

SARS-CoV-2 Molecular Data Description: Data set used in the study in Section 7. We acknowledge the following sequence submitting laboratories to Gisaidd.org:

- Charité Universitätsmedizin Berlin, Institute of Virology. Victor M Corman, Julia Schneider, Talitha Veith, Barbara Mühlemann, Markus Antwerpen, Christian Drosten,

Roman Wölfel.

- Bundeswehr Institute of Microbiology. Mathias C Walter, Markus H Antwerpen and Roman Wölfel.
- Center of Medical Microbiology, Virology, and Hospital Hygiene, University of Dues-seldorf. Ortwin Adams, Marcel Andree, Alexander Diltthey, Torsten Feldt, Sandra Hauka, Torsten Houwaart, Bjrn-Erik Jensen, Detlef Kindgen-Milles, Malte Kohns Vasconcelos, Klaus Pfeffer, Tina Senff, Daniel Strelow, Jrg Timm, Andreas Walker, Tobias Wienemann.
- CNR Virus des Infections Respiratoires - France SUD. Antonin Bal, Gregory Destras, Gwendolyne Burfin, Solenne Brun, Carine Moustaud, Raphaelle Lamy, Alexandre Gaymard, Maude Bouscambert-Duchamp, Florence Morfin-Sherpa, Martine Valette, Bruno Lina, Laurence Josset.
- National Reference Center for Viruses of Respiratory Infections, Institut Pasteur, Paris. Mélanie Albert, Marion Barbet, Sylvie Behillil, Méline Bizard, Angela Brise-barre, Flora Donati, Fabiana Gambaro, Etienne Simon-Lorière, Vincent Enouf, Maud Vanpeene, Sylvie van der Werf, Lèa Pilorge.
- Laboratoire Virpath, CIRI U111, UCBL1, INSERM, CNRS, ENS Lyon. Olivier Terrier, Aurlien Traversier, Julien Fouret, Yazdan Yazdanpanah, Xavier Lescure, Alexan-dre Gaymard, Bruno Lina, Manuel Rosa-Calatrava.

It follows a description of all sequence sampling locations and dates.

| gisaid_epi_isl | date | country | division |
|----------------|------------|---------|------------------------|
| EPI_ISL_412912 | 2020-02-25 | Germany | Baden-Wuerttemberg |
| EPI_ISL_406862 | 2020-01-28 | Germany | Bavaria |
| EPI_ISL_414520 | 2020-03-02 | Germany | Bavaria |
| EPI_ISL_414521 | 2020-03-02 | Germany | Bavaria |
| EPI_ISL_413488 | 2020-02-28 | Germany | North Rhine Westphalia |
| EPI_ISL_414497 | 2020-02-25 | Germany | North Rhine Westphalia |
| EPI_ISL_414499 | 2020-02-26 | Germany | North Rhine Westphalia |
| EPI_ISL_414505 | 2020-02-27 | Germany | North Rhine Westphalia |
| EPI_ISL_414509 | 2020-02-28 | Germany | North Rhine Westphalia |
| EPI_ISL_417457 | 2020-03-10 | Germany | Duesseldorf |
| EPI_ISL_417458 | 2020-03-11 | Germany | Duesseldorf |
| EPI_ISL_417459 | 2020-03-11 | Germany | Duesseldorf |
| EPI_ISL_417460 | 2020-03-11 | Germany | Duesseldorf |
| EPI_ISL_417461 | 2020-03-11 | Germany | Duesseldorf |
| EPI_ISL_417462 | 2020-03-11 | Germany | Duesseldorf |
| EPI_ISL_417463 | 2020-03-13 | Germany | Duesseldorf |
| EPI_ISL_417464 | 2020-03-14 | Germany | Duesseldorf |
| EPI_ISL_417465 | 2020-03-14 | Germany | Duesseldorf |
| EPI_ISL_417466 | 2020-03-14 | Germany | Duesseldorf |
| EPI_ISL_417467 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_417468 | 2020-03-16 | Germany | Duesseldorf |
| EPI_ISL_419541 | 2020-03-14 | Germany | Duesseldorf |
| EPI_ISL_419542 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_419543 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_419544 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_419545 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_419546 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_419548 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_419549 | 2020-03-15 | Germany | Duesseldorf |
| EPI_ISL_419550 | 2020-03-16 | Germany | Duesseldorf |
| EPI_ISL_419551 | 2020-03-16 | Germany | Duesseldorf |
| EPI_ISL_419552 | 2020-03-16 | Germany | Duesseldorf |
| EPI_ISL_402125 | 2019-12-26 | China | Hubei |

| gisaid_epi_isl | date | country | division |
|----------------|------------|---------|---------------------|
| EPI_ISL_418412 | 2020-03-15 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418413 | 2020-03-15 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418414 | 2020-03-15 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418416 | 2020-03-16 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418417 | 2020-03-16 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418418 | 2020-03-16 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418419 | 2020-03-16 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418420 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418422 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418423 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418424 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418425 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418426 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418427 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418428 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419168 | 2020-03-17 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418429 | 2020-03-18 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418430 | 2020-03-18 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418431 | 2020-03-18 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418432 | 2020-03-18 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419169 | 2020-03-21 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419170 | 2020-03-21 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419171 | 2020-03-21 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419172 | 2020-03-21 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419173 | 2020-03-21 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419174 | 2020-03-20 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419175 | 2020-03-21 | France | Auvergne-Rhne-Alpes |

| gisaid_epi_isl | date | country | division |
|----------------|------------|---------|-----------------------|
| EPI_ISL_419176 | 2020-03-21 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419177 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419178 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419179 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419180 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419181 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419182 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419183 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419184 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419185 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419186 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419187 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_419188 | 2020-03-22 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418219 | 2020-02-26 | France | Bretagne |
| EPI_ISL_416502 | 2020-02-26 | France | Bretagne |
| EPI_ISL_416503 | 2020-03-01 | France | Bretagne |
| EPI_ISL_416504 | 2020-03-02 | France | Bretagne |
| EPI_ISL_416505 | 2020-03-02 | France | Bretagne |
| EPI_ISL_416506 | 2020-03-03 | France | Bretagne |
| EPI_ISL_416507 | 2020-03-05 | France | Bretagne |
| EPI_ISL_416508 | 2020-03-06 | France | Bretagne |
| EPI_ISL_416509 | 2020-03-06 | France | Bretagne |
| EPI_ISL_416510 | 2020-03-06 | France | Bretagne |
| EPI_ISL_416511 | 2020-03-07 | France | Bretagne |
| EPI_ISL_416512 | 2020-03-07 | France | Bretagne |
| EPI_ISL_416513 | 2020-03-07 | France | Bretagne |
| EPI_ISL_415651 | 2020-03-05 | France | Bourgogne-France-Comt |
| EPI_ISL_415652 | 2020-03-05 | France | Bourgogne-France-Comt |
| EPI_ISL_416757 | 2020-03-07 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_417340 | 2020-03-07 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_418222 | 2020-03-04 | France | Centre-Val de Loire |
| EPI_ISL_416752 | 2020-03-04 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416751 | 2020-03-05 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_414623 | 2020-02-25 | France | Grand Est |
| EPI_ISL_414631 | 2020-03-04 | France | Grand Est |
| EPI_ISL_414632 | 2020-03-04 | France | Grand Est |

| gisaid_epi_isl | date | country | division |
|----------------|------------|---------|-----------------|
| EPI_ISL_418218 | 2020-02-21 | France | Hauts de France |
| EPI_ISL_418220 | 2020-02-28 | France | Hauts de France |
| EPI_ISL_414626 | 2020-02-29 | France | Hauts de France |
| EPI_ISL_414627 | 2020-03-02 | France | Hauts de France |
| EPI_ISL_414630 | 2020-03-03 | France | Hauts de France |
| EPI_ISL_414635 | 2020-03-04 | France | Hauts de France |
| EPI_ISL_414637 | 2020-03-04 | France | Hauts de France |
| EPI_ISL_414638 | 2020-03-04 | France | Hauts de France |
| EPI_ISL_415649 | 2020-03-05 | France | Hauts de France |
| EPI_ISL_418223 | 2020-03-05 | France | Hauts de France |
| EPI_ISL_418224 | 2020-03-08 | France | Hauts de France |
| EPI_ISL_418225 | 2020-03-08 | France | Hauts de France |
| EPI_ISL_415654 | 2020-03-09 | France | Hauts de France |
| EPI_ISL_416493 | 2020-03-08 | France | Hauts de France |
| EPI_ISL_416495 | 2020-03-10 | France | Hauts de France |
| EPI_ISL_416496 | 2020-03-10 | France | Hauts de France |
| EPI_ISL_416497 | 2020-03-10 | France | Hauts de France |
| EPI_ISL_418226 | 2020-03-09 | France | Hauts de France |
| EPI_ISL_418227 | 2020-03-12 | France | Hauts de France |
| EPI_ISL_418228 | 2020-03-12 | France | Hauts de France |
| EPI_ISL_418231 | 2020-03-15 | France | Hauts de France |
| EPI_ISL_418236 | 2020-03-16 | France | Hauts de France |
| EPI_ISL_418237 | 2020-03-16 | France | Hauts de France |
| EPI_ISL_418238 | 2020-03-16 | France | Hauts de France |
| EPI_ISL_418239 | 2020-03-16 | France | Hauts de France |
| EPI_ISL_406596 | 2020-01-23 | France | Ile de France |
| EPI_ISL_406597 | 2020-01-23 | France | Ile de France |
| EPI_ISL_411219 | 2020-01-28 | France | Ile de France |
| EPI_ISL_408430 | 2020-01-29 | France | Ile de France |
| EPI_ISL_408431 | 2020-01-29 | France | Ile de France |

| gisaid_epi_isl | date | country | division |
|----------------|------------|---------|---------------------|
| EPI_ISL_415650 | 2020-03-02 | France | Ile de France |
| EPI_ISL_416498 | 2020-03-11 | France | Ile de France |
| EPI_ISL_416499 | 2020-03-11 | France | Ile de France |
| EPI_ISL_416501 | 2020-03-10 | France | Ile de France |
| EPI_ISL_418229 | 2020-03-12 | France | Ile de France |
| EPI_ISL_418230 | 2020-03-13 | France | Ile de France |
| EPI_ISL_418232 | 2020-03-15 | France | Ile de France |
| EPI_ISL_418233 | 2020-03-15 | France | Ile de France |
| EPI_ISL_418234 | 2020-03-14 | France | Ile de France |
| EPI_ISL_418235 | 2020-03-16 | France | Ile de France |
| EPI_ISL_418240 | 2020-03-16 | France | Ile de France |
| EPI_ISL_417333 | 2020-03-04 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_417334 | 2020-03-04 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416753 | 2020-03-06 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416754 | 2020-03-06 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416756 | 2020-03-06 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_417337 | 2020-03-07 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_417336 | 2020-03-06 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_417339 | 2020-03-08 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416758 | 2020-03-08 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416747 | 2020-03-04 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416748 | 2020-03-04 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416750 | 2020-03-06 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_417338 | 2020-03-07 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_414624 | 2020-02-26 | France | Normandie |
| EPI_ISL_416494 | 2020-03-04 | France | Normandie |
| EPI_ISL_414625 | 2020-02-26 | France | Pays de la Loire |
| EPI_ISL_416745 | 2020-03-10 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416746 | 2020-03-03 | France | Auvergne-Rhne-Alpes |
| EPI_ISL_416749 | 2020-03-04 | France | Auvergne-Rhne-Alpes |