TRAILS: tree reconstruction of ancestry using incomplete lineage sorting

IKER RIVAS-GONZÁLEZ[1,*], MIKKEL H SCHIERUP[1],
JOHN WAKELEY[2], ASGER HOBOLTH[3]

[1]Bioinformatics Research Center (BiRC), Aarhus University

[2]Department of Organismic and Evolutionary Biology, Harvard University

[3]Department of Mathematics, Aarhus University

[*]Corresponding author (irg@birc.au.dk)

ABSTRACT

Genome-wide genealogies of multiple species carry detailed information about demographic and selection processes on individual branches of the phylogeny. Here, we introduce TRAILS, a hidden Markov model that accurately infers time-resolved population genetics parameters, such as ancestral effective population sizes and speciation times, for ancestral branches using a multi-species alignment of three species and an outgroup. TRAILS leverages the information contained in incomplete lineage sorting fragments by modelling genealogies along the genome as rooted three-leaved trees, each with a topology and two coalescent events happening in discretized time intervals within the phylogeny. Posterior decoding of the hidden Markov model can be used to infer the ancestral recombination graph for the alignment and details on demographic changes within a branch. Since TRAILS performs posterior decoding at the base-pair level, genome-wide scans based on the posterior probabilities can be devised to detect deviations from neutrality. Using TRAILS on a human-chimp-gorilla-orangutan alignment, we recover speciation parameters and extract information about the topology and coalescent times at high resolution.

INTRODUCTION

Orthologous sites in two or more sequences share a unique genealogical history, with coalescent events happening at certain time points in the past. In the absence of recombination, all sites along the sequences follow the same genealogy. In reality, however, ancestral recombination events might have decoupled consecutive sites, generating an array of segments with different yet correlated genealogies, collectively known as the ancestral recombination graph (ARG) [1, 2]. In principle, if inferred accurately, the ARG contains all available information about the demography of the samples, and it can be used to

estimate population parameters (such as the recombination rate and the ancestral effective population sizes), historical events (such as introgression and hybridization), and selective processes [3]. The ARG, however, is challenging to infer because the underlying genealogies along the genome alignment cannot be directly observed. Instead, inference of the genealogy along the genome relies on the site patterns of the accumulated mutations.

The ARG can also be formulated as a spatial process along the genomic alignment [4]. This process, however, contains a long-range correlation structure because if two recombination events happen flanking a genomic fragment, the fragment might be surrounded by the exact same genealogy. However, disregarding the fact that the process is non-Markovian in nature, the ARG can be approximated by a hidden Markov model (HMM), where the genealogy of a certain genomic position only depends on the genealogy of the previous position [5, 6]. It has been shown that this approach, commonly referred to as sequentially Markovian coalescent or SMC, is a good approximation of the true coalescent-with-recombination process [7]. Perhaps the simplest of such models is the pairwise sequentially Markovian coalescent (PSMC) [8], in which the ARG between two sequences (typically, the two copies of a diploid individual) is modelled. Here, the hidden states are coalescent events that happen in discretized time intervals, which correspond to two-leafed gene trees (fig. 1A). The transition probabilities between pairs of hidden states can be calculated using standard coalescent theory, parameterized by the recombination rate and the ancestral effective population sizes (N_e) in each time interval [8]. PSMC, and other SMCs, such as MSMC [9], MSMC2 [10], ASMC [11], and SMC++ [12], allow the use of standard HMM machinery to infer population parameters, and are thus also useful for inferring the most plausible coalescent times from the posterior decoding. However, SMC models are generally restricted to a single coalescent event between a pair of samples, which limits their usefulness. More recently, there have been new developments to model multiple samples explicitly. For example, ARGweaver [13], Relate [14], tsinfer+tsdate [15, 16] or ARG-Needle [17] use techniques such as resampling, threading and mathematical approximations to sequentially build the ARG [18].

These models are typically used to analyze samples from the same species to get within-species information about the ancestral process. Analyzing inter-species coalescent events adds another layer of complexity, since the coalescent events need to be contained within the underlying phylogeny or speciation tree [22, 23, 24]. Moreover, the models described above typically use the presence or absence of a certain mutation to construct haplotypes, but ignore or filter out instances where more than two alleles are observed. This infinite sites model poses a problem for inter-species analysis, because recurrent mutation is more likely to happen, generating instances of sites that have experienced more than a single mutation [25, 26].

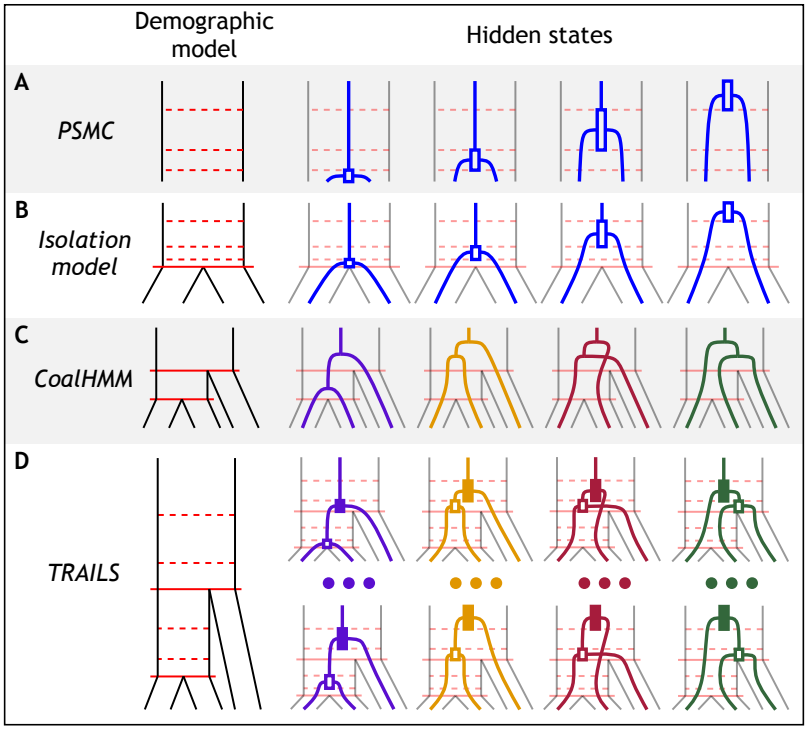

Figure 1: TRAILS is a HMM that reconstructs the time-resolved multi-species ARG for three genomes. TRAILS extends the isolation model (B) [19] to three species, by combining the time discretization of PSMC-like models (A) [8] with the topologically aware hidden states of CoalHMM (C) [20, 21]. The resulting hidden states of TRAILS (D) are three-leaved genealogies with two discretized coalescent events and one of four possible topologies. A full list of the 27 possible hidden states when $n_{\mathrm{AB}} = n_{\mathrm{ABC}} = 3$ can be consulted in fig. S11

Some other models have tried to extend these concepts for the analyses of multiple species. For example, the coalescent-with-isolation model [19] is conceptually similar to PSMC, but, backwards in time, the two analyzed samples are kept isolated until the speciation event, after which they can coalesce (fig. 1B). This model can be used to estimate the speciation time between the two samples and the N_e of the ancestral species, and an extension of it can be used to model isolation-with-migration [27]. These models, similar to SMCs, can output a posterior decoding of the coalescent times.

Beyond two samples, CoalHMM models the coalescent with recombination of three species [20, 21], where the hidden states are the four possible genealogies that might arise within the underlying species tree (fig. 1C). Two of the four genealogies differ from the species tree, which generate incongruencies that might pose a problem for standard phylogenetic reconstruction. Nevertheless, this phenomenon, commonly referred to as incomplete lineage sorting or ILS, is very informative about the demographic parameters of the underlying species tree, and CoalHMM can thus be used to estimate ancestral N_e and two speciation times. Moreover, CoalHMM uses a substitution model for mutations, so recurrent mutations are allowed. However, unlike SMCs, CoalHMM does not model coalescent events at discretized time intervals and, instead, coalescent times are modelled as single time points within an individual

branch. Because of this, some of the parameter estimates of CoalHMM are biased [21], and, although obtaining accurate estimates is still possible [28], the debiasing procedure involves costly coalescent simulations. Moreover, posterior decoding can only be performed on the topology of the gene trees, and not on the coalescent times.

Here we present TRAILS, an HMM that combines modelling the information-rich ILS signal in the style of CoalHMM and the time discretization of SMC-like models to infer unbiased estimates of the demographic parameters (ancestral N_e and speciation times), and to enable the posterior decoding of both topology and coalescent times. In TRAILS, the hidden states are three-leaved gene trees, each with a specified topology and two coalescent events that happen at discretized time intervals on an underlying speciation tree (fig. 1D, fig. S11). The genealogies are rooted by a fourth sample from an outgroup species. The transition probabilities between the hidden states of TRAILS are calculated using coalescent-with-recombination theory for one, two and three lineages that segregate within the branches of the phylogeny. We provide formulas in matrix notation to calculate these transition probabilities for a varying number of discretized time intervals (see Methods Summary for a short explanation, and Supporting Information for an in-depth description of the theory). The emitted states are sites in a four-way multiple genome alignment, containing the sequences of the three species and the outgroup. The transition and emission probabilities are parameterized by two ancestral N_e, speciation times, and the recombination rate. Keeping the mutation rate at a fixed value, TRAILS allows for the estimation of the other parameters by optimizing the HMM likelihood given the alignment. After fitting the HMM, TRAILS can perform posterior decoding of the hidden states, inferring a posterior probability of coalescent events through time within the speciation tree.

Here we derive the transition and emission probabilities, implement the model and demonstrate its use on simulated and real data. After optimizing the population parameters using TRAILS on a simulated dataset, we show that increasing the number of discrete coalescence intervals, especially deep in time, reduces the bias in the parameter estimation. We also show how the posterior decoding can accurately reconstruct the true ARG, by inferring the topology of gene trees and the time in which coalescent events occurred. We perform additional simulations to show that the posterior decoding of TRAILS can be used to detect selective sweeps that happened on ancestral branches of the phylogeny. Finally, we analyze a human-chimp-gorilla-orangutan alignment, inferring the demographic parameters of the underlying species tree and performing genome-wide posterior decoding at the base-pair level.

RESULTS

Parameter estimation

The transition and emission probabilities of coalescent hidden Markov models (HMMs) are parameterized by the demographic model, i.e., by the speciation times, ancestral effective population sizes (N_e) and recombination rate. This means that using standard HMM algorithms, these parameters can be optimized to obtain the model that best explains the observed data. Both in CoalHMM and in TRAILS, numerical optimization is performed on the log-likelihood calculated using the forward algorithm, given a four-way genome alignment of three focal species and an outgroup. The maximum likelihood estimates of the demographic model are then found using a bound-constrained search algorithm.

Previous work using coalescent HMMs has shown that the estimation of the demographic parameters is challenging. In CoalHMM, for example, the parameter estimates are highly biased [21], especially for the ancestral N_e and the recombination rate. It is possible to obtain close to unbiased estimates, but this requires a costly simulation procedure [28]. The source of the bias seems to be the restrictive state space of CoalHMM [21], which includes the topology of the genealogy but no information on when the coalescents happened within each branch of the tree (fig. 1C).

TRAILS overcomes this issue by extending the state space to include coalescent events that can happen in discretized time intervals. To demonstrate that TRAILS can perform unbiased parameter estimation, we generated five 10 Mb four-way alignments using `msprime` [29] by choosing a demographic model similar to the human-chimp-gorilla-orangutan speciation tree (see Methods for further details). The simulated sequence alignments were analyzed using TRAILS, estimating the times, ancestral N_e values and recombination rate depicted in fig. 2A. Parameters were estimated assuming different combinations of the number of intervals between speciation events ($n_{\mathrm{AB}} \in \{1, 3, 5\}$) and the number of intervals deep in time, in the common ancestor of all three species ($n_{\mathrm{ABC}} \in \{1, 3, 5\}$).

In the model where $n_{\mathrm{AB}} = n_{\mathrm{ABC}} = 1$, which is equivalent to the original CoalHMM model, parameter estimates deviate from their true values in the simulations, especially for t_2 and N_{AB} (fig. 2B). While increasing the number of intervals between speciation events (n_{AB}) only improves the accuracy of the parameter estimation slightly, a larger number of intervals deep in time (n_{ABC}) recovers estimates very close to the simulated demographic parameters. An exception to this is the recombination rate, which is underestimated, likely due to recombination events that are immediately followed by a coalescent event. For all the other parameters, the simulation results demonstrate that the source of the bias in the parameter estimation in CoalHMM can be alleviated with a more flexible model that includes coalescent

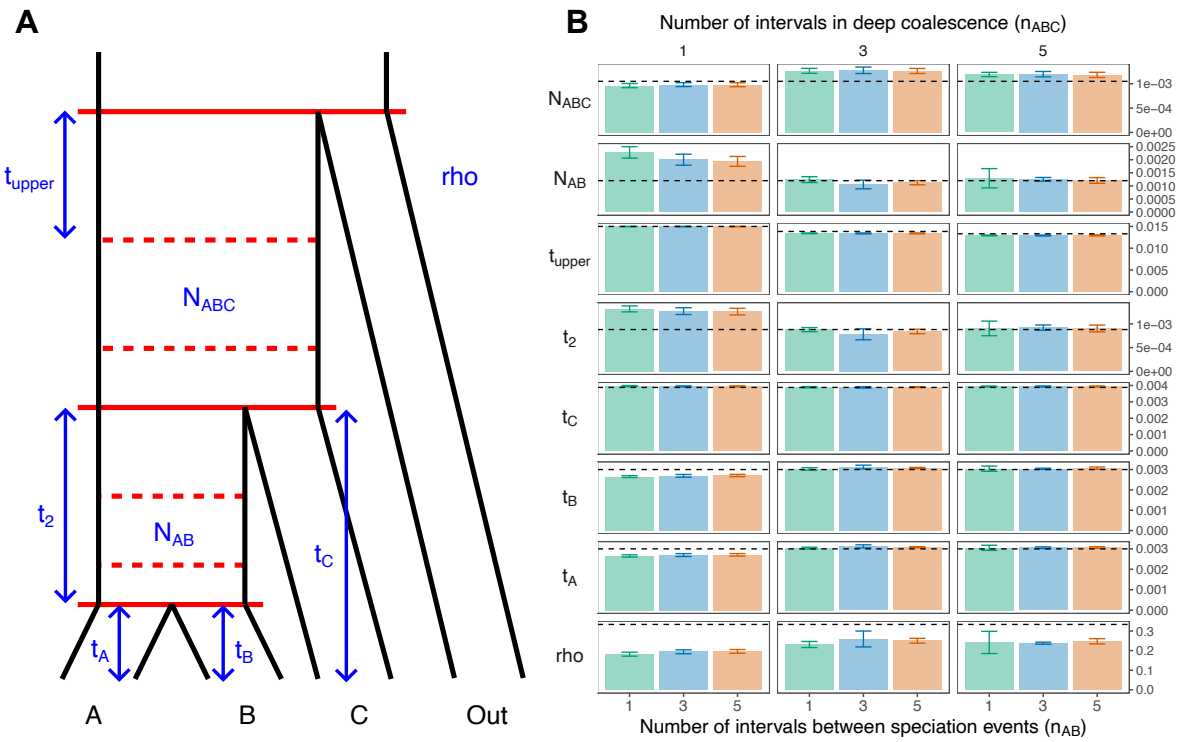

Figure 2: TRAILS performs unbiased estimation of the demographic parameters. (A) Diagram (not to scale) of the demographic model with all the optimized parameters in blue for the non-ultrametric case. In an ultrametric model, t_1 would correspond to the time from present to the shallowest speciation event, where $t_1 = t_A = t_B = t_C - t_2$. (B) Mean values and standard deviations for the parameters estimated from five simulated `msprime` genomes for a different combination of n_{AB} (bottom X axis) and n_{ABC} (top X axis). The parameter values are represented by horizontal lines. The time units are generations times mutation rate, the N_e units are effective population size times mutation rate, and the units of rho are crossovers per generation divided by the mutation rate.

times at discretized time intervals.

Posterior decoding of simulated data

Posterior decoding using the parameters estimated from the alignment can be performed using the transition and emission probabilities computed by TRAILS for a specific demographic model. In contrast to other coalescent-based HMMs, the resulting posterior probabilities are, however, hard to visualize, since each hidden state will have its own topology, and first (or more recent) and second (or more ancient) coalescent time intervals (fig. 1D, fig. S11). To overcome this, we summarize the posterior probabilities by grouping states that share certain features. For example, the posterior probabilities of all states that share the same topology can be summed. In a similar fashion, the posteriors of all states with the same first or second coalescent times can also be summed.

In order to have a ground truth for comparison, the posterior decoding was performed on 100 kb of an alignment simulated using `msprime`, with a demographic model identical to that used in fig. 2B (see

Summary of Methods for a full description of the model). The resulting posteriors can capture the true topology and the second coalescent time quite accurately (fig. 3 A and B, respectively), while the first coalescent time (fig. 3C) is harder to estimate. Additionally, V1 segments (fig. 3A) are easily misclassified as V0 segments, since V0 and V1 only differ in branch lengths but not in topology, and, thus, the emitted site patterns are similar.

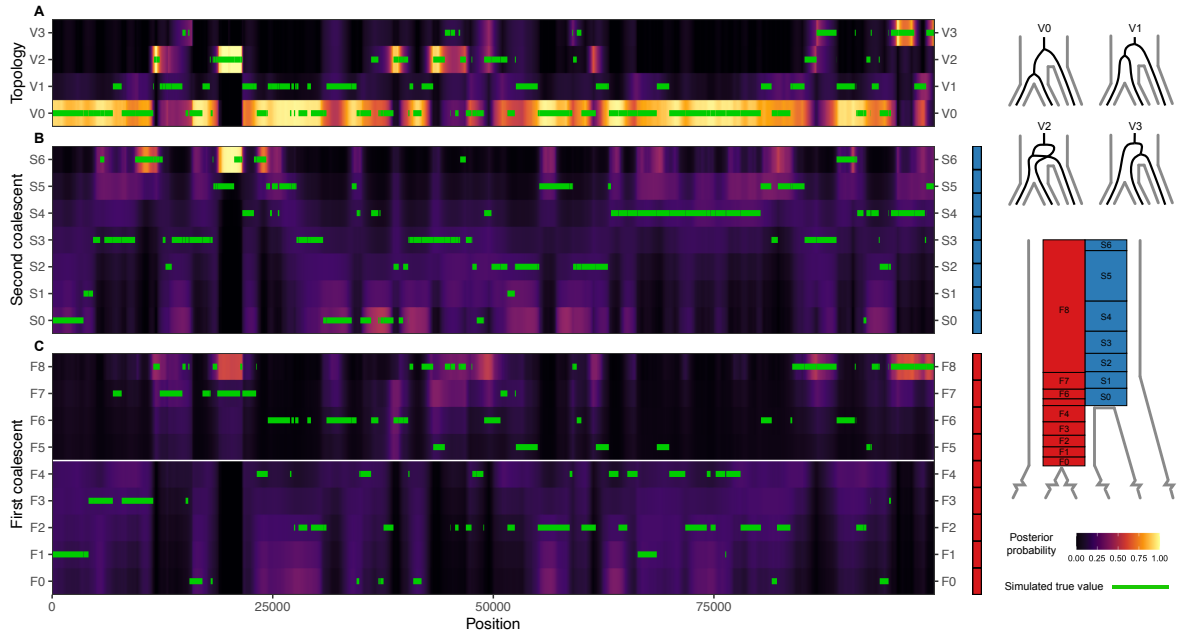

Figure 3: Posterior probability for the topology (A), second coalescent event (B) and first coalescent event (C) of a 100 kb `msprime` simulation. The true empirical topology and coalescent times are plotted as green lines.

Posterior decoding from simulated data with selection

To showcase how the posterior decoding could be useful to infer deviations from the standard coalescent, we simulated a 200 kb alignment using SLiM [30] containing a single positively selected variant ($2N_e s = 175$) in position 100 kb of the simulated alignment that arises at the first interval backwards in time, where the first two species merge in the speciation tree (interval S0 in fig. 4A). The site is strongly positively selected, but it lies well within the possible range of values for selection coefficients recorded in humans, with the lactase gene having up to $2N_e s = 1000$ in some human populations [31]. The demographic parameters were the same as those used for fig. 2 and fig. 3.

Both the true empirical values and the posterior decoding show that there is an overrepresentation of second coalescent events happening in interval S0 (fig. 4A), which is qualitatively different from the neutral case (fig. 3B). The positively selected variant confers a big advantage and is fixed rapidly in the population, with an expected fixation time of $\frac{2\ln(2\times N_e - 1)}{s} = \frac{2\times\ln(2\times 10{,}000)}{175/(2\times 10{,}000)} = 2{,}264$ generations or

56,591 years for an $N_e = 10,000$ and a generation time of $g = 25$ years, compared to the $4N_e = 40,000$ generations or 1 million years for a neutrally evolving site [32]. This means that whatever polymorphism existed at the selected locus is quickly purged from the population, and, with it, linked neutral variants are hitchhiking along, causing a selective sweep. As a result, there is an excess of coalescent events happening in interval S0, which can be discerned in the posterior decoding.

The signal observed in the posterior decoding can be summarized by averaging over the time interval with the maximum posterior probability, and comparing it to the theoretical expectation (fig. 4B). There is a clear excess of coalescent events estimated to happen at the interval where the beneficial mutation arises, which does not happen in the neutral case (fig. 3, fig. S16). This demonstrates that deviations from neutrality can be inferred using the posterior decoding of TRAILS.

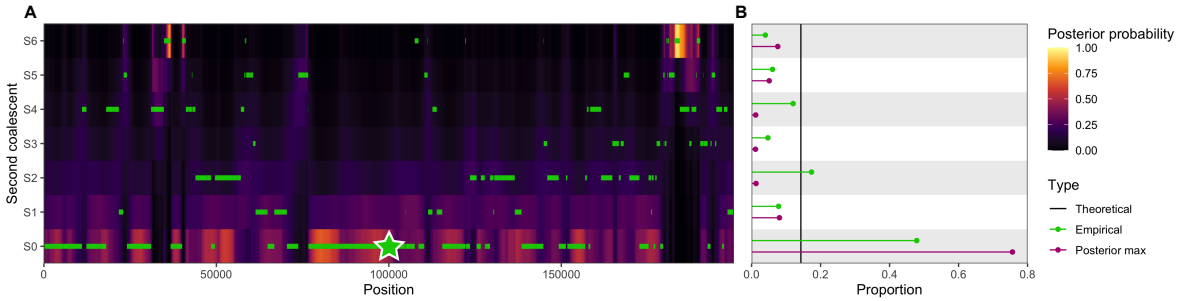

Figure 4: Posterior probability of the second coalescent time for a selective sweep. (A) Posterior probability of the second coalescent event for a simulated 200 kb region containing a positively selected variant at position 100 kb that arises in interval S0. The true simulated coalescent times are plotted as green horizontal lines. (B) Proportion of sites with maximum posterior probability for each interval of the second coalescent (purple), and the empirical true proportion of sites for each interval (green), compared to the theoretical expectation (black).

Parameter estimation from a HCGO alignment

ILS happens pervasively on the branches of the tree of life, spanning taxonomically diverse groups such as marsupials [33], birds [34, 35], fishes [36], plants [37], and mammals [38, 39], including primates [28, 40, 41]. For example, there is around 32% of ILS in the human-chimp ancestor, with 16% of the genome grouping human and gorilla, and another 16% grouping chimp and gorilla [28]. These estimates were obtained using CoalHMM [21], together with estimates for ancestral N_e and split times, which were debiased using simulations. Here, we apply TRAILS to a 50 Mb human-chimp-gorilla alignment from chromosome 1 with orangutan as an outgroup to infer population genetics parameters and to gain information about the coalescent times and the topology through posterior decoding.

Using MafFilter [42], the alignment was first preprocessed to extract the species of interest (human, chimp, gorilla and orangutan), to merge consecutive alignment blocks, and to filter out small blocks (see Methods for further details). Using `biopython` [43], 50 Mb were extracted from chr1, namely the region

from 25 Mb to 75 Mb. This region was used as the input for TRAILS, choosing the parameter values estimated for that branch in Rivas-González et al. [28] as starting values for the optimized parameters, setting $n_{\mathrm{AB}} = n_{\mathrm{ABC}} = 3$, and using the L-BFGS-B algorithm for model fitting [44, 45]. To get more accurate parameter estimates, the optimized parameters were used as starting parameters for a second TRAILS run, where $n_{\mathrm{AB}} = 3$ and $n_{\mathrm{ABC}} = 5$.

The resulting estimates are displayed in fig. 5A. Assuming a mutation rate of $\mu = 1.25 \times 10^{-8}$ per site per generation and a generation time of $g = 25$ years, the speciation time estimates are in good agreement with previously inferred values, with the HC split dating 5.51-5.81 million years ago ($\sim$4-7 MYA from literature [21, 28, 46, 47, 48]), the HCG split at 10.40-11.70 MYA ($\sim$8-12 from literature [21, 28, 47, 48]), and the HCGO split at 18.55-19.85 MYA ($\sim$10-20 from literature [40, 47, 28]). Moreover, ancestral N_e inferred for the HC ancestor (167,400) and for the HCG ancestor (101,290) are consistent with previous estimates using CoalHMM (177,368 and 106,702, respectively [28]). Using the estimates for t_2 (in generations) and N_{ABC}, we get a probability of ILS equal to

$$\mathrm{ILS} = \frac{2}{3}\exp(-t_2/(2N_{\mathrm{ABC}})) = \frac{2}{3}\exp(-195,000/(2 \times 167,400)) = 37\%,$$

so our estimates suggest the ((human, gorilla), chimp) topology in 18.5% of the genome and the ((chimp, gorilla), human) topology in 18.5% of the genome. Finally, the recombination rate was estimated to be $\rho = 1.19 \times 10^{-8}$ per site per generation, which matches the rate estimated for present-day humans [49].

TRAILS allows for the independent estimation of each individual branch length, which is useful for non-ultrametric trees. Figure 5A shows that the branch leading to chimps is longer than that leading to humans by around 5.9%, and the gorilla branch is longer than both the human (12.6%) and the chimp (9.1%) branches (calculated from the second speciation event to present). This deviation from the molecular clock is well supported by previous studies [50], and is likely because of different branches accumulating a different number of mutations per year, either due to an acceleration or deceleration of the mutational process, changes in the average time of reproduction, or a combination of these [51, 52, 53].

In summary, we have demonstrated that TRAILS is able to infer demographic parameters that are in agreement with estimates from the literature. More importantly, it does so without the need for any post-processing or corrections, avoiding the use of fossil calibrations [47] or debiasing procedures [21, 28].

Posterior decoding of the HCGO alignment

Posterior decoding was then performed using the optimized parameters and setting $n_{\mathrm{AB}} = n_{\mathrm{ABC}} = 5$. In order to get an understanding of the genome-wide variation of ILS and coalescent times, the resulting posterior probabilities were summarized in 100 kb windows along the chr1 region in three different ways. First, the mean posterior probability was calculated for each of the four possible topologies, by first summing the posteriors of all hidden states sharing the same topology for each site, and then averaging over all sites on the 100 kb window. The resulting probabilities were then used to calculate a proxy for ILS, by summing the probabilities of observing the ILS topologies (V2 and V3 in fig. 3). Second, using a similar procedure, the mean posterior probability was calculated for each of the six possible intervals for the first coalescent event, by first summing the posteriors of all hidden states sharing the same first coalescent interval for each site, and then averaging over all sites in the 100 kb window. As a proxy for the first coalescent time, integers from 1 to 6 were assigned to each interval in chronological order backwards in time, and a weighted mean of those integers was computed, where the weights were the mean posterior probabilities per window. Third, the same quantity as for the first coalescent was computed for the second coalescent.

After filtering outliers smaller than the 1st percentile and larger than the 99th percentile, these proxies were plotted as heatmaps (fig. 5B). The first coalescent and the ILS proportion show a very strong correlation ($\rho = 0.979$), likely reflecting that V2 and V3 can only happen in the common ancestor of all three species, so when ILS is present, coalescent times are generally deeper (and vice versa). This signal is also captured, although more weakly, by the correlation between the second coalescent and the ILS proportion ($\rho = 0.483$). This can be explained by knowing that, conditional on ILS, the second coalescent follows a convolution of exponentials of rates 3 and 1 [54], while, conditional on V0 (i.e., conditional on the first coalescent happening between speciation events), the second coalescent simply follows an exponential or rate 1. Thus, if more ILS is present in a certain window, then, on average, the second coalescent will tend to happen deeper in time.

Figure 5B also shows how, at the 100 kb level, the genome displays spatial covariation in the amount of ILS and the time to coalescence that exceeds stochastic effects of a neutral coalescent process. This is in line with previous results [28], where ILS proportions are affected by genomic features such as gene density, recombination rate, and the effects of linked selection.

A zoomed-in region of around 41 kb is shown in fig. 5C, which shows a long fragment of the ((chimp, gorilla), human) topology (V3). This fragment is unusually long, spanning 6,800 bp, and it is highly implausible following the demographic model inferred by TRAILS. Thanks to the posterior decoding of

the coalescent times performed by TRAILS, we can observe that, for this fragment, the first coalescent event backwards in time happens close to the second speciation time (in interval F5), while the second coalescent event happens in the deepest time interval (S4).

One explanation for such a long V3 fragment is that it might be influenced by selection, which would maintain the alternative topology uninterrupted for a long period of time. Another explanation could be that this fragment is introgressed, especially given that the first coalescent event is shallow and that the fragment is long. For comparison, another region in chromosome 1, which also shows an excess of V3 topology, has a much more variable distribution of coalescent times and it is more fragmented (fig. S17). Such detailed information about the timing of coalescent events is only possible thanks to the time discretization of TRAILS, and these details would have been missed, for example, in the posterior decoding of CoalHMM (recall fig. 1C).

The V3 fragment in fig. 5C overlaps with the last exon of the *LDLRAD1* gene, which codes for a lipoprotein receptor (UniProt: Q5T700). This showcases how TRAILS can be used to infer the topology and coalescent times of protein-coding genes at the base-pair level across millions of years of evolution, potentially solving cases of phenotypic hemiplasy [33], and illuminating selective processes affecting the sorting of lineages [28].

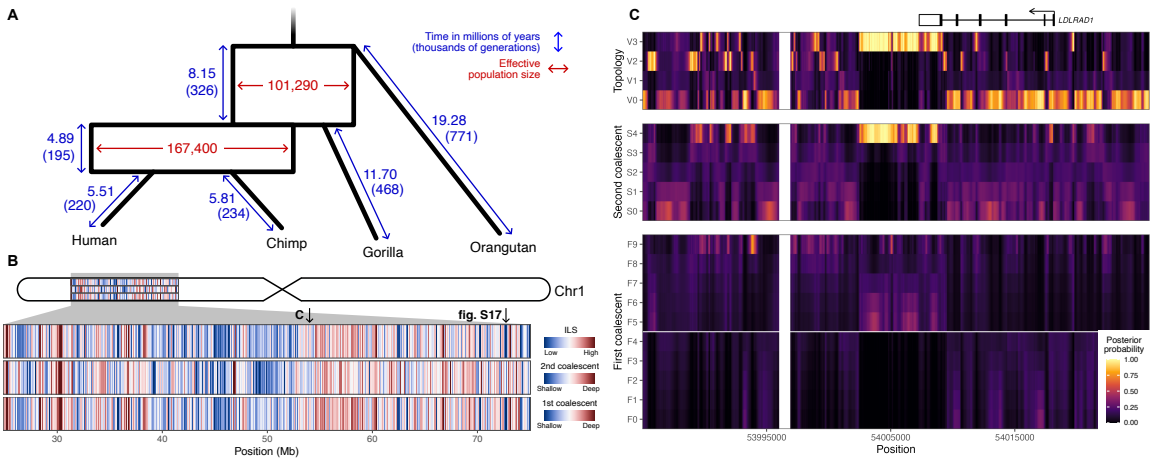

Figure 5: TRAILS output for 50 Mb of chromosome 1 of the HCGO alignment. (A) Estimates for the speciation times (blue) and ancestral N_e (red) of the speciation process, optimized using TRAILS and assuming a mutation rate of $\mu = 1.25 \times 10^{-8}$ per site per generation. To convert time from generations to millions of years, a generation time of $g = 25$ years per generation was used. (B) Genome-wide variation of ILS, and first and second coalescent times. (C) Posterior decoding of the topology, and first and second coalescent events for a zoomed-in region in chromosome 1. The *LDLRAD1* gene is plotted on top, where exons are represented as boxes, coding regions as filled boxes, and introns as horizontal lines.

Discussion

Coalescent-based approaches for analyzing genomic data are essential tools for understanding the ancestral history of species. Here, we have introduced TRAILS, an HMM that models the topology and the two coalescent events for gene genealogies within a phylogeny of three species. TRAILS can accurately infer population genetic parameters (ancestral N_e, speciation times and recombination rate). From the posterior decoding, the three-species ARG can be inferred at the base-pair level, providing insight into the ancestral history of the species at high resolution. Deviations from neutrality can be detected by summarizing the posterior decoding in windows and running genomic scans to find excess of coalescents happening at certain time intervals. Moreover, more coarse-grained summaries of the posterior decoding spanning several kilobases could be used to infer the genome-wide variation of ILS or coalescent times, potentially revealing correlations with other genomic features such as variation in the recombination rate or selection [28, 55].

As demonstrated here, the posterior decoding from TRAILS is a powerful way to infer details of the ARG in the context of speciation, together with departures from the neutral expectation. Recurrent selective sweeps that have happened during the speciation process are hypothesized to be drivers of speciation, and to greatly influence the genealogical landscape of present-day genomes. For example, the human X chromosome contains long haplotypes shared across all non-African populations [56], spanning large genomic regions that are both lacking Neanderthal introgression [57], and showing very low rates of ILS in the human-chimp ancestor [58]. This suggests that the X chromosome has a unique evolutionary history which is greatly affected by gene flow (or lack thereof), and that these low-diversity regions might be related to genetic incompatibilities that arose during the speciation of ancestral hominids. TRAILS can help locate these ancient sweeps and infer when they occurred, potentially illuminating when and how genomes were affected by selection during the speciation process.

Using posterior decoding, regions that show unusually high levels of an alternative topology with very shallow coalescents can also be detected, which could indicate ancient introgression or hybridization events happening between ancestral branches of the species tree. Such ancient introgression events have been reported to be pervasive among some branches in the primate species tree [47], although they can be difficult to distinguish from ILS [59] unless explicitly modelled. TRAILS could be extended to model introgression more directly by including additional hidden states representing introgressed genomic fragments. These would have exceedingly short coalescent times compared to the deep coalescent ILS states [28], and TRAILS provides the mathematical framework to distinguish between these two cases.

TRAILS could also be extended to accommodate variation in N_e along individual ancestral branches

in the species tree, conceptually very similar to what is done in PSMC analyses from a single extant genome [8]. Modelling variation in N_e can elucidate how speciation events might have happened. For example, population sizes that are maintained more constant during the speciation event might indicate a cleaner split, while increased ancestral N_e just prior to the estimated time of speciation (here equalled to the total cessation of gene flow) might point to a period of elevated population structure and a prolonged species separation with migration [27]. Modelling changes in the demography around speciation events might also help us detect and characterize instances of complex speciation, as proposed, for example, by Patterson et al. [46].

Finally, with more efficient algorithms to compute the transition probability matrix, future extensions of TRAILS could be devised to analyze more than three species, thus allowing for the inference of the speciation tree and the multi-species ARG for more taxa. This can help us resolve more complex patterns of ILS involving more than three species, which include phenomena such as anomaly zones [35, 60, 61].

METHODS SUMMARY

The transition probabilities between the hidden states of TRAILS can be calculated from a series of interconnected continuous-time Markov chains (CTMCs) that model the coalescent with recombination of two contiguous nucleotides for one, two or three sequences. The CTMCs are parameterized by the ancestral N_e, speciation times and recombination rate. The transitions for TRAILS are subsequently calculated by conditioning the CTMCs on the topology and coalescent times of the gene trees at those two sites, binning coalescent events into discretized time intervals along the speciation process. Additionally, the emission probabilities for each hidden state are calculated from a CTMC of the mutational process by choosing a certain substitution model. In this section, we provide a summary of the model, and the full explanation can be consulted in the Supporting Information.

Continuous-time Markov chains for the ancestral process

The coalescent with recombination between two sites can be approximated as a continuous-time Markov chain (CTMC). For one sequence, the left and the right sites can be either linked or unlinked, so there are only two possible states for the CTMC. Two linked sites become unlinked when a recombination event happens between them, which happens with a rate of ρ_1. On the other hand, two unlinked sites become linked when a coalescent event happens between them, with a rate of γ_1. These two transitions

can be gathered in a transition rate matrix

$$\boldsymbol{Q_1} = \begin{pmatrix} -\gamma_1 & \gamma_1 \\ \rho_1 & -\rho_1 \end{pmatrix}. \tag{1}$$

From this transition rate matrix, we can calculate the probability matrix P_{A} as $\exp(t\boldsymbol{Q_1})$, which gives the probability of the sites being unlinked or linked at a certain time t given that the chain starts in the unlinked state (first row) or unlinked state (second row).

When two sequences are involved, the state-space of the CTMC becomes more complex. Apart from the coalescent and the recombination events described above, sites can also coalesce irreversibly backwards in time with rate γ_2, which happens when two left (or two right) sites from two different sequences find common ancestry. The resulting rate matrix (fig. S5) for the coalescent with recombination with two sequences then corresponds to a CTMC with 15 states (fig. S4), which was originally described by Simonsen and Churchill [62]. Due to these irreversible coalescent events, the rate matrix has a block-like structure, and it contains sets of states in which sequences can freely recombine and coalesce until a irreversible coalescent event occurs. Ultimately, both the left and the right sites will have irreversibly coalesced, reaching one of two absorbing states. Note that the matrix is quite sparse, since most of the transitions are not allowed.

Following a similar reasoning, the coalescent with recombination for three sequences can also be modelled as a CTMC. In this case, both the left and the right site will eventually undergo two irreversible coalescent events, which can potentially happen in any order between the three sequences. This creates 203 possible states (fig. S7), the transitions of which can also be gathered in a rate matrix (fig. S8). This matrix also has a block-like structure, and, given enough time, states will transition into one of the two absorbing states.

If all three sequences belonged to the same species, a three-sequence CTMC would be sufficient to model the coalescent with recombination. However, the sequences belong to three different species, so the speciation process has to be overlaid on top. Subsequently, the coalescent with recombination along the speciation tree is modelled as a series of interconnected CTMCs.

Because sequences are sampled in present time, the left and the right sites are fully linked at time 0, meaning that the starting probability vector for the one-sequence CTMCs is $(0, 1)$. Backwards in time, each of the sequences will remain isolated for a certain period of time in which the two sites can recombine and coalesce freely. The sequences for species A and B will remain isolated until the first

speciation event at time t_A and t_B, respectively. Then, the final probabilities of the one-sequence CTMCs for A and B are merged to create the initial probabilities for the two-sequence CTMC. After a certain time, where the two sequences are allowed to coalesce and recombine, the final probabilities for the two-sequence CTMC and the final probabilities for the one-sequence CTMC of species C will be merged, thus creating the starting probabilities for the three-sequence CTMC. Finally, given enough time, all sequences for both the right and the left site will eventually coalesce into one of the two absorbing states of the last CTMC.

Transition probabilities of the HMM

The hidden states of TRAILS include a topology and two coalescent events that can happen within discretized time intervals. The breakpoints of the time intervals can thus be used to divide the CTMCs. Now, instead of calculating the end probabilities of each of the CTMCs, we must calculate the end probabilities of each of the discretized time intervals, conditioned on the path taken by the left and the right sites.

Emission probabilities of the HMM

For each hidden state, the emission probabilities are calculated using the Jukes-Cantor mutational model [63]. Instead of calculating the emitted nucleotides for the three species only, TRAILS also includes the nucleotides emitted by an outgroup, which provides essential information about the ancestral state in each site. This additional species must have a sufficient enough divergence with the rest of the species such that there is no ILS between them.

Parameterization

The transition and emission probabilities are parameterized by the speciation times (t_A, t_B, t_C, t_2, t_{upper}), the effective population sizes (N_{AB}, N_{ABC}), and the recombination rate (fig. 2A). Implicitly, TRAILS is also parameterized by the mutation rate, but this cannot be jointly inferred with the rest of the parameters because the parameter values can be scaled by any factor and still produce the same coalescent model [54, 64]. Instead, the mutation rate in the model is fixed to 1, and all other parameters are rescaled appropriately. The resulting units for the speciation times are number of generations multiplied by the mutation rate, and, similarly, the effective population sizes are number of individuals times the mutation rate. Accordingly, the recombination rate is divided by the mutation rate, so the optimized parameter is the ratio between the recombination and the mutation rate. After estimating the parameters, parameter values with more interpretable units can be obtained by choosing an appropriate

mutation rate.

TRAILS allows for two different parameterizations, namely the *ultrametric* model and the *non-ultrametric* model. In the ultrametric model, all sequences are sampled at time 0, so the molecular clock is assumed ($t_A = t_B = t_C - t_2 = t_1$). Instead, in the non-ultrametric model, each sequence (A, B and C) is allowed to be sampled at a different time. This is useful to model deviations from the molecular clock, for example, when the number of generations for each of the species from present time until the speciation event is different, or the mutation rate varies between the species.

Simulations

Parameter estimation

Simulations to validate the model were performed in `msprime` [29]. The underlying demographic model follows a speciation tree with four species, namely A, B, C and D. The time from present to the first speciation event was set to $t_1 = 200,000$ generations. The (haploid) ancestral N_e for the time between speciation events was set to $N_{AB} = 80,000$. In order to keep an ILS proportion of 32%, the time between the first and the second speciation events was set to $t_2 = -N_{AB}\log(\frac{3}{2} \times 0.32) = 25,501$ generations. The (haploid) ancestral N_e earlier than the second speciation event was set to $N_{ABC} = 70,000$, and the time between the second speciation event and the speciation event with the outgroup was set to $t_3 = 1,000,000$ generations. The recombination rate was set to $\rho = 0.5 \times 10^{-8}$ per site per generation. The tree was kept ultrametric in number of generations, meaning that all species were sampled at generation 0, so $t_A = t_B = t_1$, $t_C = t_1 + t_2$, and $t_D = t_1 + t_2 + t_3$. Mutations were then added on top of the simulated genealogies according to the Jukes-Cantor model [63] with a mutation rate of $\mu = 1.5 \times 10^{-8}$ per generation per site.

In order to investigate how the number of intervals in the AB-ancestor (n_{AB}) and the ABC-ancestor (n_{ABC}) affect the parameter estimation, five 10 Mb alignments were simulated using `msprime`, and then TRAILS was run to estimate the demographic parameters for each combination of $n_{AB} \in \{1, 3, 5\}$ and $n_{ABC} \in \{1, 3, 5\}$, using the bound-constrained Nelder-Mead algorithm. The starting parameters of the optimization were randomly drawn from a normal distribution centered on the true value and with a standard deviation of the true value divided by 5.

Posterior probability

The demographic model described above was also used to generate a 100 kb alignment to perform posterior decoding with the true parameters. In TRAILS, the default way of dividing up the coalescent

space into discretized time intervals is by taking quantiles of a truncated exponential of rate 1 (measured in units of N_{AB}) for the time between speciation events, while the time previous to the earliest speciation event is divided following the quantiles of an exponential of rate 1 (measured in units of N_{ABC}). This scheme is appropriate for computing the posterior decoding of the first coalescent event when it happens between speciation events, since it is expected to follow exactly a truncated exponential of rate 1 according to the standard coalescent. However, if the first coalescent happens deep in time, then the coalescent times will follow an exponential with rate 3, since there are 3 lineages present. This means that most coalescent events happen very fast, and there would be an overrepresentation of coalescents in the first interval if the default cutpoint scheme is used.

TRAILS is, however, not restricted to a specific discretization, and it can compute the transition and emission probabilities, and, thus, perform posterior decoding, for user-specified intervals. For the first coalescent in deep time, posterior decoding was performed using cutpoints from the quantiles of an exponential with rate 3, with $n_{\mathrm{ABC}} = 7$ and the true parameters from the `msprime` simulation.

Moreover, the second coalescent event will follow a mixture of an exponential with a rate of 1 (for the V0 states) and a convolution of two exponentials with rates 3 and 1 (for the deep coalescent states) [54], which will happen with probability $1 - \exp(-t_2/N_{\mathrm{AB}})$ and $\exp(-t_2/N_{\mathrm{AB}})$, respectively. Here, t_2 is the time between speciation events in number of generations, and N_{AB} is the effective population population size. The second coalescent time can thus be represented as a phase-type distribution [65], with sub-intensity matrix $\boldsymbol{S}$ and initial probability vector $\boldsymbol{\pi}$, such that

$$\boldsymbol{S} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -3 & 3 \\ 0 & 0 & -1 \end{pmatrix}$$

and

$$\boldsymbol{\pi} = \left(1 - \exp\left(-\frac{t_2}{N_{\mathrm{AB}}}\right), \exp\left(-\frac{t_2}{N_{\mathrm{AB}}}\right), 0\right).$$

Therefore, posterior decoding for the second coalescent was performed using the quantiles of this phase-type distribution using `PhaseTypeR` [66], with $n_{\mathrm{AB}} = 5$, $n_{\mathrm{ABC}} = 7$, and the true parameters used to generate the `msprime` simulation.

Selection

Using the same demographic model, a 200 kb alignment was simulated using SLiM [30], assuming a single positively selected variant in the middle of the region, with population-scaled selection parameter

$2N_e s = 175$. Since SLiM is a forward simulator and runs much slower than forward simulators such as `msprime`, all the demographic parameters of the model were rescaled by a factor of 200 in order to increase computational speed. Posterior decoding was performed on the resulting alignment simulated using the same discretization scheme as described above.

Real data

The chromosome 1 multiz alignment of 30 mammalian species (27 primates) was downloaded from the UCSC Genome Browser database in MAF format. Using MafFilter [42], the species of interest were filtered (human, chimp, gorilla and orangutan), syntenic blocks separated by 200 nucleotides or less were merged using human as a reference, and blocks smaller than 2,000 bp were filtered out. The resulting filtered MAF was used as input for TRAILS, using the parameters estimated in Rivas-González et al. [28] as starting values. The optimization was performed using a bound-constrained version of the L-BFGS-B algorithm implemented in `numpy` [44, 45, 67], by setting $n_{\mathrm{AB}} = n_{\mathrm{ABC}} = 3$, and using the L-BFGS-B algorithm for model fitting. To get a more accurate parameter estimation, the optimized estimates were used as starting values for a second TRAILS run where $n_{\mathrm{ABC}} = 5$, optimized using a bound-constrained Nelder-Mead algorithm [68, 69].

Code availability

TRAILS can be downloaded and installed from pip (`https://pypi.org/project/trails-rivasiker/`), and the source code can be browsed at `https://github.com/rivasiker/trails`. The code for reproducing the figures can be found at `https://github.com/rivasiker/trails_paper`. Simulations and analysis of real data can be re-run using a gwf workflow (`https://gwf.app`).

Author contributions

I.R.-G. performed conceptualization, methodology, software development, validation, formal analysis, investigation, writing, and visualization. M.H.S. performed conceptualization, writing (review and editing), and supervision. J.W. performed conceptualization, methodology, writing (review and editing), and supervision. A.H. performed conceptualization, methodology, writing (review and editing), and supervision.

Acknowledgements

We gratefully acknowledge Julien Dutheil and Nick Patterson for useful discussions on the implementation of the model. We also thank GenomeDK for providing the computational resources for performing the analyses. This work was supported by grants NNF18OC0031004 from the Novo Nordisk Foundation and 6108-00385 from the Independent Research Fund Denmark, Natural Sciences.

References

[1] R. C. Griffiths. The two-locus ancestral graph. *Lecture Notes-Monograph Series*, 18:100–117, 1991. ISSN 07492170. URL `http://www.jstor.org/stable/4355649`.

[2] Robert C Griffiths and Paul Marjoram. An ancestral recombination graph. *Institute for Mathematics and its Applications*, 87:257–270, 1997.

[3] Melissa Hubisz and Adam Siepel. *Inference of Ancestral Recombination Graphs Using ARGweaver*, pages 231–266. Springer US, New York, NY, 2020.

[4] Carsten Wiuf and Jotun Hein. Recombination as a point process along sequences. *Theoretical population biology*, 55(3):248–259, 1999.

[5] Gilean AT McVean and Niall J Cardin. Approximating the coalescent with recombination. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1459):1387–1393, 2005.

[6] Paul Marjoram and Jeff D Wall. Fast "coalescent" simulation. *BMC genetics*, 7:1–9, 2006.

[7] Peter R Wilton, Shai Carmi, and Asger Hobolth. The SMC′ is a highly accurate approximation to the ancestral recombination graph. *Genetics*, 200(1):343–355, 2015.

[8] Heng Li and Richard Durbin. Inference of human population history from individual whole-genome sequences. *Nature*, 475(7357):493–496, 2011.

[9] Stephan Schiffels and Richard Durbin. Inferring human population size and separation history from multiple genome sequences. *Nature genetics*, 46(8):919–925, 2014.

[10] Anna-Sapfo Malaspinas, Michael C Westaway, Craig Muller, Vitor C Sousa, Oscar Lao, Isabel Alves, Anders Bergström, Georgios Athanasiadis, Jade Y Cheng, Jacob E Crawford, et al. A genomic history of aboriginal australia. *Nature*, 538(7624):207–214, 2016.

[11] Pier Francesco Palamara, Jonathan Terhorst, Yun S Song, and Alkes L Price. High-throughput inference of pairwise coalescence times identifies signals of selection and enriched disease heritability. *Nature Genetics*, 50(9):1311–1317, 2018.

[12] Jonathan Terhorst, John A Kamm, and Yun S Song. Robust and scalable inference of population history from hundreds of unphased whole genomes. *Nature genetics*, 49(2):303–309, 2017.

[13] Matthew D Rasmussen, Melissa J Hubisz, Ilan Gronau, and Adam Siepel. Genome-wide inference of ancestral recombination graphs. *PLoS genetics*, 10(5):e1004342, 2014.

[14] Leo Speidel, Marie Forest, Sinan Shi, and Simon R Myers. A method for genome-wide genealogy estimation for thousands of samples. *Nature genetics*, 51(9):1321–1329, 2019.

[15] Jerome Kelleher, Yan Wong, Anthony W Wohns, Chaimaa Fadil, Patrick K Albers, and Gil McVean. Inferring whole-genome histories in large population datasets. *Nature genetics*, 51(9):1330–1338, 2019.

[16] Anthony Wilder Wohns, Yan Wong, Ben Jeffery, Ali Akbari, Swapan Mallick, Ron Pinhasi, Nick Patterson, David Reich, Jerome Kelleher, and Gil McVean. A unified genealogy of modern and ancient genomes. *Science*, 375(6583):eabi8264, 2022.

[17] Brian C Zhang, Arjun Biddanda, Árni Freyr Gunnarsson, Fergus Cooper, and Pier Francesco Palamara. Biobank-scale inference of ancestral recombination graphs enables genealogical analysis of complex traits. *Nature Genetics*, pages 1–9, 2023.

[18] Débora YC Brandt, Xinzhu Wei, Yun Deng, Andrew H Vaughn, and Rasmus Nielsen. Evaluation of methods for estimating coalescence times using ancestral recombination graphs. *Genetics*, 221(1): iyac044, 2022.

[19] Thomas Mailund, Julien Y Dutheil, Asger Hobolth, Gerton Lunter, and Mikkel H Schierup. Estimating divergence time and ancestral effective population size of Bornean and Sumatran orangutan subspecies using a coalescent hidden Markov model. *PLoS genetics*, 7(3):e1001319, 2011.

[20] Asger Hobolth, Ole F Christensen, Thomas Mailund, and Mikkel H Schierup. Genomic relationships and speciation times of human, chimpanzee, and gorilla inferred from a coalescent hidden Markov model. *PLoS genetics*, 3(2):e7, 2007.

[21] Julien Y Dutheil, Ganesh Ganapathy, Asger Hobolth, Thomas Mailund, Marcy K Uyenoyama, and Mikkel H Schierup. Ancestral population genomics: the coalescent hidden Markov model approach. *Genetics*, 183(1):259–274, 2009.

[22] James H Degnan and Noah A Rosenberg. Gene tree discordance, phylogenetic inference and the multispecies coalescent. *Trends in Ecology & Evolution*, 24(6):332–340, 2009.

[23] Bruce Rannala, A Leache, S Edwards, and Ziheng Yang. The multispecies coalescent model and species tree inference. Self Published, 2020.

[24] Siavash Mirarab, Luay Nakhleh, and Tandy Warnow. Multispecies coalescent: theory and applications in phylogenetics. *Annual Review of Ecology, Evolution, and Systematics*, 52:247–268, 2021.

[25] Colm O'hUigin, Yoko Satta, Naoyuki Takahata, and Jan Klein. Contribution of homoplasy and of ancestral polymorphism to the evolution of genes in anthropoid primates. *Molecular Biology and Evolution*, 19(9):1501–1513, 2002.

[26] David B Wake, Marvalee H Wake, and Chelsea D Specht. Homoplasy: from detecting pattern to determining process and mechanism of evolution. *science*, 331(6020):1032–1035, 2011.

[27] Thomas Mailund, Anders E Halager, Michael Westergaard, Julien Y Dutheil, Kasper Munch, Lars N Andersen, Gerton Lunter, Kay Prüfer, Aylwyn Scally, Asger Hobolth, et al. A new isolation with migration model along complete genomes infers very different divergence processes among closely related great ape species. *PLoS genetics*, 8(12):e1003125, 2012.

[28] Iker Rivas-González, Marjolaine Rousselle, Fang Li, Long Zhou, Julien Y. Dutheil, Kasper Munch, Yong Shao, Dongdong Wu, Mikkel H. Schierup, and Guojie Zhang. Pervasive incomplete lineage sorting illuminates speciation and selection in primates. *Science*, 380(6648):eabn4409, 2023.

[29] Franz Baumdicker, Gertjan Bisschop, Daniel Goldstein, Graham Gower, Aaron P Ragsdale, Georgia Tsambos, Sha Zhu, Bjarki Eldon, E Castedo Ellerman, Jared G Galloway, Ariella L Gladstein, Gregor Gorjanc, Bing Guo, Ben Jeffery, Warren W Kretzschumar, Konrad Lohse, Michael Matschiner, Dominic Nelson, Nathaniel S Pope, Consuelo D Quinto-Cortés, Murillo F Rodrigues, Kumar Saunack, Thibaut Sellinger, Kevin Thornton, Hugo van Kemenade, Anthony W Wohns, Yan Wong, Simon Gravel, Andrew D Kern, Jere Koskela, Peter L Ralph, and Jerome Kelleher. Efficient ancestry and mutation simulation with msprime 1.0. *Genetics*, 220(3):iyab229, 2022.

[30] Benjamin C Haller and Philipp W Messer. Slim 4: Multispecies eco-evolutionary modeling. *The American Naturalist*, 201(5):E000–E000, 2023.

[31] Laure Ségurel and Céline Bon. On the evolution of lactase persistence in humans. *Annual review of genomics and human genetics*, 18:297–319, 2017.

[32] Sarah P Otto and Michael C Whitlock. Fixation probabilities and times. *Encyclopedia of Life Sciences*, pages 1–5, 2013.

[33] Shaohong Feng, Ming Bai, Iker Rivas-González, Cai Li, Shiping Liu, Yijie Tong, Haidong Yang, Guangji Chen, Duo Xie, Karen E Sears, et al. Incomplete lineage sorting and phenotypic evolution in marsupials. *Cell*, 185(10):1646–1660, 2022.

[34] Alexander Suh, Linnéa Smeds, and Hans Ellegren. The dynamics of incomplete lineage sorting across the ancient adaptive radiation of neoavian birds. *PLoS biology*, 13(8):e1002224, 2015.

[35] Alison Cloutier, Timothy B Sackton, Phil Grayson, Michele Clamp, Allan J Baker, and Scott V Edwards. Whole-genome analyses resolve the phylogeny of flightless birds (Palaeognathae) in the presence of an empirical anomaly zone. *Systematic Biology*, 68(6):937–955, 2019.

[36] Fernando Alda, Victor A Tagliacollo, Maxwell J Bernt, Brandon T Waltz, William B Ludt, Brant C Faircloth, Michael E Alfaro, James S Albert, and Prosanta Chakrabarty. Resolving deep nodes in an ancient radiation of neotropical fishes in the presence of conflicting signals from incomplete lineage sorting. *Systematic biology*, 68(4):573–593, 2019.

[37] Y Zhou, L Duvaux, G Ren, L Zhang, O Savolainen, and J Liu. Importance of incomplete lineage sorting and introgression in the origin of shared genetic variation between two closely related pines with overlapping distributions. *Heredity*, 118(3):211–220, 2017.

[38] Kun Wang, Johannes A Lenstra, Liang Liu, Quanjun Hu, Tao Ma, Qiang Qiu, and Jianquan Liu. Incomplete lineage sorting rather than hybridization explains the inconsistent phylogeny of the wisent. *Communications biology*, 1(1):169, 2018.

[39] Celine Scornavacca and Nicolas Galtier. Incomplete lineage sorting in mammalian phylogenomics. *Systematic biology*, 66(1):112–120, 2017.

[40] Asger Hobolth, Julien Y Dutheil, John Hawks, Mikkel H Schierup, and Thomas Mailund. Incomplete lineage sorting patterns among human, chimpanzee, and orangutan suggest recent orangutan speciation and widespread selection. *Genome research*, 21(3):349–356, 2011.

[41] Thomas Mailund, Kasper Munch, and Mikkel Heide Schierup. Lineage sorting in apes. *Annual review of genetics*, 48:519–535, 2014.

[42] Julien Y Dutheil, Sylvain Gaillard, and Eva H Stukenbrock. MafFilter: a highly flexible and extensible multiple genome alignment files processor. *BMC genomics*, 15:1–10, 2014.

[43] Peter JA Cock, Tiago Antao, Jeffrey T Chang, Brad A Chapman, Cymon J Cox, Andrew Dalke, Iddo Friedberg, Thomas Hamelryck, Frank Kauff, Bartek Wilczynski, et al. Biopython: freely available python tools for computational molecular biology and bioinformatics. *Bioinformatics*, 25 (11):1422–1423, 2009.

[44] Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on scientific computing*, 16(5):1190–1208, 1995.

[45] Ciyou Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on mathematical software (TOMS)*, 23(4):550–560, 1997.

[46] Nick Patterson, Daniel J Richter, Sante Gnerre, Eric S Lander, and David Reich. Genetic evidence for complex speciation of humans and chimpanzees. *Nature*, 441(7097):1103–1108, 2006.

[47] Dan Vanderpool, Bui Quang Minh, Robert Lanfear, Daniel Hughes, Shwetha Murali, R Alan Harris, Muthuswamy Raveendran, Donna M Muzny, Mark S Hibbins, Robert J Williamson, et al. Primate phylogenomics uncovers multiple rapid radiations and ancient interspecific introgression. *PLoS biology*, 18(12):e3000954, 2020.

[48] Aylwyn Scally, Julien Y Dutheil, LaDeana W Hillier, Gregory E Jordan, Ian Goodhead, Javier Herrero, Asger Hobolth, Tuuli Lappalainen, Thomas Mailund, Tomas Marques-Bonet, et al. Insights into hominid evolution from the gorilla genome sequence. *Nature*, 483(7388):169–175, 2012.

[49] Bjarni V Halldorsson, Gunnar Palsson, Olafur A Stefansson, Hakon Jonsson, Marteinn T Hardarson, Hannes P Eggertsson, Bjarni Gunnarsson, Asmundur Oddsson, Gisli H Halldorsson, Florian Zink, et al. Characterizing mutagenic effects of recombination through a sequence-level genetic map. *Science*, 363(6425):eaau1043, 2019.

[50] Priya Moorjani, Carlos Eduardo G Amorim, Peter F Arndt, and Molly Przeworski. Variation in the molecular clock of primates. *Proceedings of the National Academy of Sciences*, 113(38):10607–10612, 2016.

[51] Søren Besenbacher, Christina Hvilsom, Tomas Marques-Bonet, Thomas Mailund, and Mikkel Heide Schierup. Direct estimation of mutations in great apes reconciles phylogenetic dating. *Nature Ecology & Evolution*, 3(2):286–292, 2019.

[52] Gregg WC Thomas, Richard J Wang, Arthi Puri, R Alan Harris, Muthuswamy Raveendran, Daniel ST Hughes, Shwetha C Murali, Lawrence E Williams, Harsha Doddapaneni, Donna M

Muzny, et al. Reproductive longevity predicts mutation rates in primates. *Current Biology*, 28(19): 3193–3197, 2018.

[53] Lindell Bromham. The genome as a life-history character: why rate of molecular evolution varies between mammal species. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366 (1577):2503–2513, 2011.

[54] John Wakeley. *Coalescent Theory: An Introduction*. Roberts & Company Publishers, 1 edition, June 2008.

[55] James B Pease and Matthew W Hahn. More accurate phylogenies inferred from low-recombination regions in the presence of incomplete lineage sorting. *Evolution*, 67(8):2376–2384, 2013.

[56] Laurits Skov, Moises Coll Macia, Elise Anne Lucotte, Maria Izabel Alves Cavassim, David Castellano, Mikkel Heide Schierup, and Kasper Munch. Extraordinary selection on the human X chromosome associated with archaic admixture. *Cell Genomics*, 3(3), 2023.

[57] Sriram Sankararaman, Swapan Mallick, Michael Dannemann, Kay Prüfer, Janet Kelso, Svante Pääbo, Nick Patterson, and David Reich. The genomic landscape of Neanderthal ancestry in present-day humans. *Nature*, 507(7492):354–357, 2014.

[58] Julien Y Dutheil, Kasper Munch, Kiwoong Nam, Thomas Mailund, and Mikkel H Schierup. Strong selective sweeps on the X chromosome in the human-chimpanzee ancestor explain its low divergence. *PLoS Genetics*, 11(8):e1005451, 2015.

[59] Mark S Hibbins and Matthew W Hahn. Phylogenomic approaches to detecting and characterizing introgression. *Genetics*, 220(2):iyab173, 2022.

[60] James H Degnan and Noah A Rosenberg. Discordance of species trees with their most likely gene trees. *PLoS genetics*, 2(5):e68, 2006.

[61] Fábio K Mendes and Matthew W Hahn. Why concatenation fails near the anomaly zone. *Systematic biology*, 67(1):158–169, 2018.

[62] Katy L Simonsen and Gary A Churchill. A markov chain model of coalescence with recombination. *Theoretical population biology*, 52(1):43–59, 1997.

[63] Thomas H Jukes, Charles R Cantor, et al. Evolution of protein molecules. *Mammalian protein metabolism*, 3:21–132, 1969.

[64] Jotun Hein, Mikkel Schierup, and Carsten Wiuf. *Gene genealogies, variation and evolution: a primer in coalescent theory*. Oxford University Press, USA, 2004.

[65] Asger Hobolth, Arno Siri-Jegousse, and Mogens Bladt. Phase-type distributions in population genetics. *Theoretical population biology*, 127:16–32, 2019.

[66] Iker Rivas-González, Lars Nørvang Andersen, and Asger Hobolth. PhaseTypeR: an R package for phase-type distributions in population genetics. *Journal of Open Source Software*, 8(82):5054, 2023.

[67] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020.

[68] John A Nelder and Roger Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.

[69] Fuchang Gao and Lixing Han. Implementing the Nelder-Mead simplex algorithm with adaptive parameters. *Computational Optimization and Applications*, 51(1):259–277, 2012.

[70] Charles Van Loan. Computing integrals involving the matrix exponential. *IEEE transactions on automatic control*, 23(3):395–404, 1978.

[71] Wolfram Research, Inc. Mathematica, Version 13.2, 2023. URL `https://www.wolfram.com/mathematica`.

[72] Joseph Felsenstein. Evolutionary trees from DNA sequences: a maximum likelihood approach. *Journal of molecular evolution*, 17:368–376, 1981.

[73] Michael JD Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The computer journal*, 7(2):155–162, 1964.

[74] Stephen G Nash. Newton-type minimization via the Lanczos method. *SIAM Journal on Numerical Analysis*, 21(4):770–788, 1984.

[75] Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics*, 37(6):1554–1563, 1966.

[76] Lawrence R Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

[77] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. Ray: A distributed framework for emerging AI applications. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 561–577, 2018.

[78] The pandas development team. pandas-dev/pandas: Pandas, April 2023. URL `https://doi.org/10.5281/zenodo.7857418`.

[79] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[80] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6, 2015.