

# Evolutionary Inference via the Poisson Indel Process\*

ALEXANDRE BOUCHARD-CÔTÉ<sup>1</sup>, AND MICHAEL I. JORDAN<sup>2</sup>

<sup>1</sup>*Department of Statistics, University of British Columbia, V6T 1Z2, Canada;*

<sup>2</sup>*Department of Statistics and Computer Science Division, University of California, Berkeley, 94720, USA*

## Abstract

We address the problem of the joint statistical inference of phylogenetic trees and multiple sequence alignments from unaligned molecular sequences. This problem is generally formulated in terms of string-valued evolutionary processes along the branches of a phylogenetic tree. The classical evolutionary process, the TKF91 model [62], is a continuous-time Markov chain model comprised of insertion, deletion and substitution events. Unfortunately this model gives rise to an intractable computational problem—the computation of the marginal likelihood under the TKF91 model is exponential in the number of taxa [39]. In this work, we present a new stochastic process, the Poisson Indel Process (PIP), in which the complexity of this computation is reduced to linear. The new model is closely related to the TKF91 model, differing only in its treatment of insertions, but the new model has a global characterization as a Poisson process on the phylogeny. Standard results for Poisson processes allow key computations to be decoupled, which yields the favorable computational profile of inference under the PIP model. We present illustrative experiments in which Bayesian inference under the PIP model is compared to separate inference of phylogenies and alignments.

---

\*Accepted for the Proceedings of the National Academy of Sciences.

# 1 Introduction

The field of phylogenetic inference is being transformed by the rapid growth in availability of molecular sequence data. There is an urgent need for inferential procedures that can cope with data from large numbers of taxa and that can provide inferences for ancestral states and evolutionary parameters over increasingly large timespans. Existing procedures are often not scalable along these dimensions and can be a bottleneck in analyses of modern molecular datasets.

A key issue that renders phylogenetic inference difficult is that sequence data are generally not aligned a priori, having undergone evolutionary processes that involve insertions and deletions. Consider Figure 1, which depicts an evolutionary tree in which each node is associated with a string of nucleotides, and where the string evolves via insertion, deletion and substitution processes along each branch of the tree. Even if we consider evolutionary models that are stochastically independent along the branches of the tree (conditioning on ancestral states), the inferential problem of inferring evolutionary paths (conditioning on observed data at the leaves of the tree) does not generally decouple into independent computations along the branches of the tree. Rather, alignment decisions made throughout the tree can influence the posterior distribution on alignments along any branch.

This issue has come to the fore in a line of research beginning in 1991 with a seminal paper by Thorne, Kishino and Felsenstein [62]. In the “TKF91 model,” a simple continuous-time Markov chain (CTMC) provides a string-valued stochastic process along each branch of an evolutionary tree. This makes it possible to define joint probabilities on trees and alignments, and thereby obtain likelihoods and posterior distributions for statistical inference. A further important development has been the realization that the TKF91 model can be represented as a hidden Markov model (HMM), and that generalizations to a broader class of string-valued stochastic processes with finite-dimensional marginals are therefore possible [21, 17, 61, 37, 18, 40, 36, 45]. This has the appeal that statistical inference under these processes (known as transducers) can be based on dynamic programming [2, 29, 57, 42]. Unfortunately, however, despite some analytic simplification that is feasible in restricted cases [60], the memory needed to represent the state space in these models is generally exponential in the number of leaves in the tree [12]. Moreover, even in the simple TKF91 model, there does not appear to be additional structure in the state space that allows for simplification of the

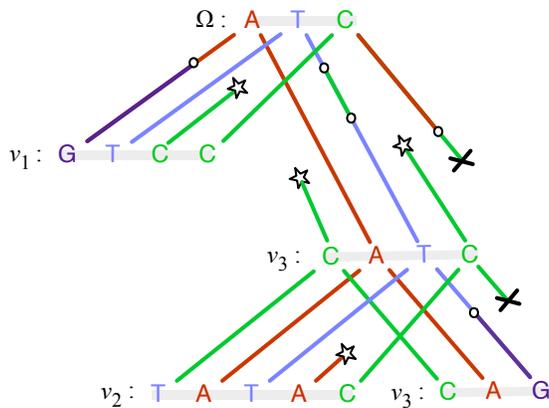


Figure 1: A depiction of the evolution of a set of strings of nucleotides along the branches of a tree with leaves  $\mathcal{L} = \{v_1, v_2, v_3\}$  and root  $\Omega$ , where each string is subject to insertion, deletion and substitution processes. Stars denote nucleotide insertion events, crosses, deletion events, and circles, substitution events.

dynamic program. Indeed, the running time of the most sophisticated algorithm for computing marginals [39] depends on the number of homology linearizations, which is exponential in sparse alignments [56].

As a consequence of this unfavorable computational complexity, there has been extensive work on approximations, specifically on approximations to the joint marginal probability of a tree and an alignment, obtained by integrating over the derivation [36, 65]. A difficulty, however, is that these marginal probabilities play a role in tree inference procedures as the numerators and denominators of acceptance probabilities for Markov chain Monte Carlo algorithms. Loss of accuracy in these values can have large, uncontrolled effects on the overall inference. A second approach is to consider joint models that are not obtained by marginalization of a joint continuous-time string-valued process. A range of combinatorial [54, 66, 31, 64, 58, 34] and probabilistic [16, 28, 48, 46, 47] models fall in this category. Although often inspired by continuous-time processes, obtaining a coherent and calibrated estimate of uncertainty in these models is difficult.

A third possible response to the computational complexity of joint inference of trees and alignments is to retreat to methods that treat these problems separately. In particular, as is often done in practice, one can ob-

tain a Multiple Sequence Alignment (MSA) via any method (often based on a heuristically chosen “guide tree”), and then infer a tree based on the fixed alignment. This latter inferential process is generally based on the assumption that the columns of the alignment are independent, in which case the problem decouples into a simple recursion on the tree (the “Felsenstein” or “sum-product” recursion [13]). Such an approach can introduce numerous artifacts, however, both in the inferred phylogeny [46, 47, 67], and in the inferred alignment [52, 43].

It is also possible to iterate the solution of the MSA problem and the tree inference problem [32, 33], which can be viewed as a heuristic methodology for attempting to perform joint inference. The drawbacks of these systems include a lack of theoretical understanding, the difficulty of getting calibrated confidence intervals, and over-alignment problems [56, 35].

Finally, other methods have focused on analyzing only pairs of sequences at a time [53, 51, 56, 6]. While this approach can considerably simplify computation [20, 9], it has the disadvantage that it is not based on an underlying joint posterior probability distribution.

In the current paper we present a new approach to the joint probabilistic inference of trees and alignments. Our approach is based on a model that is closely related to TKF91, altering only the insertion process while leaving the deletion and substitution processes intact. Surprisingly, this relatively small change has a major mathematical consequence—under the new model evolutionary paths have an equivalent global description as a Poisson process on the phylogenetic tree. We are then able to exploit standard results for Poisson processes (notably, Poisson thinning) to obtain significant computational leverage on the problem of computing the joint probability of a tree and an alignment. Indeed, under the new model this computation decouples in such a way that this joint probability can be obtained in linear time (linear in the number of taxa and the sequence length), rather than in exponential time as in TKF91.

Our new model has two descriptions: the first as a local continuous-time Markov process that is closely related to the TKF91 model, and the second as a global Poisson process. We treat the latter as the fundamental description and refer to the new process as the *Poisson Indel Process* (PIP). The new description not only sheds light on computational issues, but it also opens up new ways to extend evolutionary models, allowing, for example, models that incorporate structural constraints and slipped-strand mispairing phenomena.

Under the Poisson process representation, another interesting perspec-

tive on our process is to view it as a string-valued counterpart to stochastic Dollo models [1, 44], which are defined on finite state spaces. In particular, the general idea of the two steps generation process used in Section 3 has antecedents in the literature on probabilistic modeling of morphological or lexical characters, but the literature did not address the string-valued processes that are our focus here.

The remainder of the paper is organized as follows. Section 2 provides some basic background on the TKF91 model. In Section 3 presents the PIP model, in both its local and global formulations. Section 4 delves into the computational aspects of inference under the PIP model, describing the linear-time algorithm for computing the exact marginal density of an MSA and a tree. In Section 5 we present an empirical evaluation of the inference algorithm, and finally we present our conclusions in Section 6.

## 2 Background

We begin by giving a brief overview of the TKF91 model. Instead of following the standard treatment based on differential equations, we present a Doob-Gillespie view of the model [11, 14] that will be useful in our subsequent development.

Let us assume that at some point in time  $t$ , a sequence has length  $n$ . In the TKF91 process, the sequence stays unchanged for a random interval of time  $\Delta t$ , and after this interval, a single random mutation (substitution, insertion or deletion) alters the sequence. This is achieved by defining a total of  $3n + 1$  independent exponential random variables,  $n$  of which correspond to deletion of a single character,  $n$  of which correspond to the mutation of a single character and  $n + 1$  of which yield insertions after one of the  $n$  characters (including one “immortal” position at the leftmost position in the string). These  $3n + 1$  exponential random variables are simulated in parallel, and the value of the smallest of these random variables determines  $\Delta t$ . The index of the winner determines the nature of the event at time  $\Delta t$  (whether it is a substitution, deletion or insertion).

The random variables corresponding to a deletion have exponential rate  $\mu_{\text{TKF}}$  while those corresponding to an insertion have exponential rate  $\lambda_{\text{TKF}}$ . If the event is a mutation, a multinomial random variable with parameters obtained from the substitution rate matrix  $\theta$  is drawn to determine the new value of the character. Finally, if an insertion occurs, a multinomial ran-

dom variable is drawn to determine the value of the new character, with parameters generally taken from the stationary distribution of  $\theta$ .

This describes the evolution of a string of characters along a single edge of a phylogenetic tree. The extension to the entire phylogeny is straightforward; we simply visit the tree in preorder and apply the single-edge process to each edge. The distribution of the sequences at the root is generally assumed to be the stationary distribution of the single-edge process (conceptually, the distribution obtained along an infinitely-long edge).

Although the TKF91 model is reversible (and the PIP model as well, as we prove in Section 3.3), making the location of the root unidentifiable, it is useful to assume for simplicity that an arbitrary root has been picked, and we will make that assumption throughout. The likelihood is not affected by this arbitrary choice.

### 3 The Poisson Indel Process

In this section we introduce the *Poisson Indel Process* (PIP). This process has two descriptions, a local description which is closely related to the TKF91 model, and a global description as a Poisson process.

We require some additional notation (see Figure 2). A phylogeny  $\tau$  will be viewed as a continuous set of points, and its topology will be denoted by  $(\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V} \subset \tau$  is equal to the finite subset containing the branching points, the leaves  $\mathcal{L} \subset \mathcal{V}$  and the root  $\Omega$ , and where  $\mathcal{E}$  is the set of edges. Parent nodes will be denoted by  $\text{pa}(v)$ , for  $v \in \mathcal{V}$ , and the branch lengths by  $b(v)$ , which is the length of the edge from  $\text{pa}(v)$  to  $v$ . For any  $x \in \tau$  (whether  $x$  is a branch point in  $\mathcal{V}$ , or an intermediate point on an edge), we write  $\tau_x$  for the rooted phylogenetic subtree of  $\tau$  rooted at  $x$  (dropping all points in the original tree that are not descendants of  $x$ ). Finally, the set of characters (nucleotides or amino acids) will be denoted  $\Sigma$ .

#### 3.1 Local description

The stochastic process we propose has a local description that is very similar to the TKF91 process, the only change being that the insertion rate no longer depends on the sequence length. Therefore, instead of using  $3n+1$  competing exponential random variables to determine the next event as in the TKF91 model ( $n$  for substitutions,  $n+1$  for insertions, and  $n$  for deletions), we now

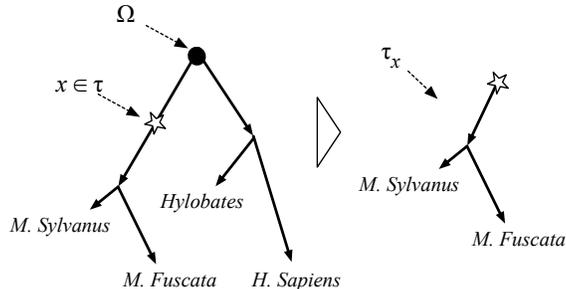


Figure 2: Notation used for describing the PIP. Given a phylogenetic tree  $\tau$  and a point  $x \in \tau$  on that tree,  $\tau_x$  is defined as the subtree rooted at  $x$ .

have  $2n + 1$  variables ( $n$  for substitutions, 1 for insertion, with rate  $\lambda$ , and  $n$  for deletion, each of rate  $\mu$ ). When an insertion occurs, its position is selected uniformly at random.<sup>1</sup> We assume that the process is initialized by sampling a Poisson-distributed number of characters, with parameter  $\lambda/\mu$ . Each character is sampled independently and identically according to the stationary distribution of  $\theta$ .

Note that if  $\lambda/\lambda_{\text{TKF}}$  is an integer, and the sequence has length  $(\lambda/\lambda_{\text{TKF}}) - 1$  at some point in time, the distribution over the time and type of the next mutation is the same as in TKF91, by using the fact that the minimum of exponential variables with  $\lambda_i$  is exponential, with rate equal to the sum of the  $\lambda_i$ . However, in general the distributions are different. We discuss some of the biological aspects of these differences in Section 6; for now we focus on the computational and statistical aspects of the PIP model.

### 3.2 Poisson process representation

We turn to a seemingly very different process for associating character strings with a phylogeny. This process consists of two steps, the first involving insertions and the second involving deletions and substitutions.

<sup>1</sup>More precisely, assume there is a real number  $r_i$  in the interval  $[0, 1]$  assigned to each character in the string in increasing order:  $0 < r_1 < r_2 < \dots < r_n < 1$ . When an insertion occurs, sample a new real number  $r'$  uniformly in the interval  $[0, 1]$  and insert the new character at the unique position (with probability one) such that an increasing sequence of real numbers  $0 < \dots < r' < \dots < 1$  is maintained.

In the first step, depicted in Figure 3A, a multiset of insertion points is sampled from a Poisson process defined on the phylogeny  $\tau$  [23]. The rate measure for this Poisson process has atomic mass at the root of the tree; hence the need for multisets rather than simple point sets. Except for the root, no other points on the tree have an atomic mass (in particular, and in contrast to population genetics models, the probability that evolutionary events occur at branching points is zero). We denote this multiset of insertion points by  $\mathbf{X}$ .

In the second step, we visit the insertion points one at a time. The order of the visits of the insertions is sampled uniformly at random,  $(X_1, X_2, \dots, X_I) \sim \text{Perm}(\mathbf{X})$ . An insertion visit consists of two substeps. First, we extract the directed subtree rooted at the insertion location  $X_i$ . Examples of these subtrees are shown in Figure 3B, left. Second, we simulate the fate of the inserted character along  $\tau_{X_i}$ . This is done via a *substitution-deletion* CTMC whose state space  $\Sigma_\varepsilon = \Sigma \cup \{\varepsilon\}$  consists of the basic alphabet  $\Sigma$  augmented with an empty string symbol  $\varepsilon$ . As shown in Figure 3B, right, the substitution-deletion CTMC yields paths along subtrees in which a single character either mutates or is deleted. The latter event, represented by  $\varepsilon$ , is an absorbing state.

We define a *homology path*  $H_i$  as the single-character history generated by a substitution-deletion CTMC along a phylogeny. If a point  $x \in \tau$  is a descendant of the insertion  $X_i$ ,  $H_i(x)$  is set to the state of the substitution-deletion CTMC at  $x$ . If  $x \in \tau$  is not a descendant of  $X_i$ , we set  $H_i(x)$  to the absorbing symbol  $\varepsilon$ . Thus, formally, a homology path  $H_i$  is a random map from any point on  $\tau$  to  $\Sigma_\varepsilon$ .

Given a set of homology paths for each inserted character index  $i$ , the sequence at any point on the tree,  $x \in \tau$ , is obtained as follows (see Figure 3C, right). First, we construct a list of all the values taken by  $H_i(x)$  at the given point:  $(H_1(x), H_2(x), \dots, H_I(x))$ . Second, we remove from the list any characters that are equal to the absorbing symbol  $\varepsilon$ . The string obtained thereby is denoted by  $Y(x)$ . The set of observed data comprises the values of  $Y$  at the leaves of the tree:  $\mathcal{Y} = \{(v, Y(v)) : v \in \mathcal{L}\}$ .

We can also construct an MSA  $M$  from a set of homology paths (see Figure 3C, left). From each homology path  $H_i$ , we extract the characters at the leaves, arranging these characters in a column. Delete any column in which all of the characters are the character  $\varepsilon$ . Arranging these columns in the order of the visits to the insertion points. The resulting matrix, whose entries range over the augmented alphabet  $\Sigma_\varepsilon$ , is the MSA  $M$ .

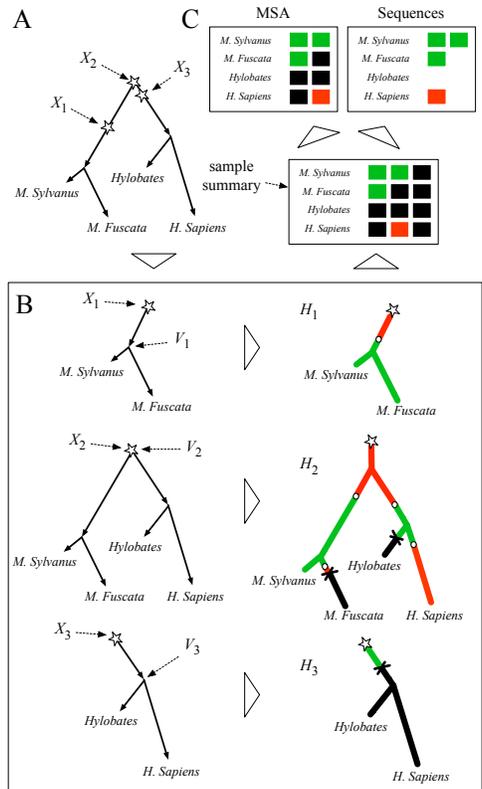


Figure 3: Example of a PIP sample. Here  $\Sigma$  has two symbols, represented by red and green squares, and the absorbing deletion symbol  $\varepsilon$  is represented in black. (A) A sample from a Poisson process on  $\tau$ . (B) Each sampled point corresponds to a rooted tree on which a CTMC path is sampled. (C) The alignments and sequences are obtained as a deterministic function of the first two steps.

For a given rooted phylogenetic tree  $\tau$ , we will denote by  $p_\tau(m)$  the marginal probability that this process generates an MSA  $m$ , integrating over all homology paths,  $p_\tau(m) = \mathbb{P}(M = m)$ . For joint inference, we make the phylogenetic tree  $T$  random, with a distribution specified by a prior with density  $p(\tau)$ .

### 3.3 Characterization

In this section we show that the local and the global descriptions of the PIP given in the previous two subsections are in fact alternative descriptions of the same string-valued stochastic process. In stating our theorem, we let  $\nu$  denote the rate measure characterizing the insertion process in the global description, and let  $Q$  and  $\pi$  denote the transition matrix and the initial distribution for the substitution-deletion CTMC.

**Theorem 1** *Let  $\tau$  be a phylogenetic tree with an arbitrary rooting, and let us denote the Lebesgue measure on  $\tau$  by the same symbol. For any insertion rate  $\lambda > 0$ , deletion rate  $\mu > 0$ , and reversible substitution rate matrix  $\theta$ , the local and global processes described in Sections 3.1 and 3.2 coincide if we set, for all  $\sigma, \sigma' \in \Sigma_\varepsilon$ :*

$$\nu(dx) = \lambda \left( \tau(dx) + \frac{1}{\mu} \delta_\Omega(dx) \right),$$

$$Q_{\sigma, \sigma'} = \begin{cases} -\sum_{\sigma'' \neq \sigma'} Q_{\sigma, \sigma''} & \text{if } \sigma = \sigma' \\ 0 & \text{if } \sigma = \varepsilon \\ \mu & \text{if } \sigma' = \varepsilon \\ \theta_{\sigma, \sigma'} & \text{o.w.,} \end{cases}$$

and set  $\pi$  to be the quasi-stationary distribution of  $Q$  [7].

The proof is given in the Appendix A. Note that in the case of interest here, where the rate of deletion does not depend on the character being deleted,  $\pi_\sigma$  is equal to the entry of the stationary distribution of  $\theta$  corresponding to  $\sigma$  when  $\sigma \neq \varepsilon$ , and zero otherwise. The following result establishes some basic properties of the PIP model. Its proof can be found in the Appendix A.

**Proposition 2** *For all  $\mu, \lambda > 0$  and reversible rate matrix  $\theta$ , the PIP model is reversible, with a stationary length distribution given by a Poisson distribution with mean  $\lambda/\mu$ .*

The Poisson stationary length distribution represents a modeling advantage of PIP over TKF91, which has a geometrically distributed stationary distribution. Based on a study of protein-length distributions for the three domains of life [69], the Poisson distribution has been suggested [38] as a more adequate length distribution.

From Proposition 2, we can also obtain an alternative reparameterization of the PIP model, in terms of asymptotic expected length  $\eta = \lambda/\mu$  and indel intensity  $\zeta = \lambda \cdot \mu$ .

## 4 Computational Aspects

We turn to a consideration of the computational consequences of the Poisson representation of the PIP model. We first consider how the Poisson process characterization allows us to compute the marginal likelihood,  $p_\tau(m)$ , in linear time; a significant improvement over methods based on the TKF91 model. We provide in Appendix D a brief discussion of the role that the marginal likelihood plays in inference.

To compute the marginal likelihood,  $p_\tau(m)$ , we first condition on the number of homology paths,  $|\mathbf{X}|$ . While the number of homology paths is random and unknown, we know that it can be no less than the number of columns  $|m|$  in the postulated alignment  $m$ . We need to consider an unknown and unbounded number of birth events with no observed offsprings in the MSA, but as they are exchangeable, they can be marginalized analytically. This is done as follows:

$$\begin{aligned} p_\tau(m) &= \mathbb{E}[\mathbb{P}(M = m | |\mathbf{X}|)] \\ &= \sum_{n=|m|}^{\infty} \mathbb{P}(|\mathbf{X}| = n) \cdot \binom{n}{|m|} \cdot (p(c_\emptyset))^{n-|m|} \prod_{c \in m} p(c), \end{aligned}$$

where the first factor captures the probability of sampling  $n$  homology paths, the second, the number of ways to pick the  $|m|$  observed homology paths (the columns, which contain at least one descendent character at the leaves) out of the  $n$  paths, the factor  $p(c) = \mathbb{P}(C = c)$  is the likelihood of a single MSA column  $c$ , and  $c_\emptyset$  is a column with an absorbing deletion symbol at every leaf  $v \in \mathcal{L}$ :  $c_\emptyset \equiv \varepsilon$  (in this section, we drop subscripts for column-specific random variables such as  $C$ ,  $H$  and  $X$  since they are exchangeable). Note that such simplification is not possible in the TKF91 model, because the rate

of insertion depends on the length of the internal sequences, and hence of the deletion events.

This expression can be simplified by introducing the function  $\varphi$  defined as follows for all  $z \in (0, 1)$ ,  $k \in \{1, 2, \dots\}$ :

$$\begin{aligned}\varphi(z, k) &= \frac{1}{k!} \|\nu\|^k \exp\{(z-1)\|\nu\|\}, \\ \|\nu\| &= \lambda \left( \|\tau\| + \frac{1}{\mu} \right),\end{aligned}$$

where  $\|\tau\|$  is the normalization of the measure  $\tau$ , i.e., the sum of all the branch lengths in the topology. We show in the Appendix B that this yields the simple formula:

$$p_\tau(m) = \varphi(p(c_\emptyset), |m|) \prod_{c \in m} p(c).$$

The next step is to compute the likelihood  $p(c)$  of each individual alignment column  $c$ . We do this by partitioning the computation into subcases depending on the location of the tree at which the insertion point  $X$  is located for column  $c$ . More precisely, we look at the most recent common ancestor  $V = v \in \mathcal{V}$  of the characters in  $c$  that are not equal to  $\varepsilon$  (see Figure 3B). If  $v \neq \Omega$ , this corresponds to the most recent endpoint of the edge  $e \in \mathcal{E}$  where the insertion occurred.

Computing the prior probability of the insertion location is greatly simplified by the fact that  $X|\mathbf{X} \sim \bar{\nu}$  (see [27], Chapter 2.4), where  $\bar{\nu} = \nu/\|\nu\|$  denotes the probability obtained by normalizing the measure  $\nu$ . We can therefore write:

$$\begin{aligned}\mathbb{P}(V = v) &= \begin{cases} \bar{\nu}(e \setminus \{\Omega\}) & \text{if } v \neq \Omega \\ \bar{\nu}(\{\Omega\}) & \text{o.w.} \end{cases} \\ &= \frac{1}{\|\tau\| + 1/\mu} \times \begin{cases} b(v) & \text{if } v \neq \Omega \\ 1/\mu & \text{o.w.} \end{cases}\end{aligned}$$

Finally, the column probabilities are computed as follows:

$$\begin{aligned}\mathbb{P}(C = c) &= \sum_{v \in \mathcal{V}} \mathbb{P}(V = v) \mathbb{P}(C = c | V = v) \\ &= \sum_{v \in \mathcal{V}} \mathbb{P}(V = v) f_v,\end{aligned}$$

where  $f_v$  is the output of a slight modification of Felsenstein’s peeling recursion [13] applied on the subtree rooted at  $v$  (the derivation for  $f_v$  can be found in the Appendix B). Since computing the peeling recursion for one column takes time  $O(|\mathcal{L}|)$ , we get a total running time of  $O(|\mathcal{L}| \cdot |m|)$ , where  $|\mathcal{L}|$  is the number of observed taxa, and  $|m|$  is the number of columns in the alignment.

## 5 Experiments

We implemented a system based on our model that performs joint Bayesian inference of phylogenies and alignments. We used this system to quantify the relative benefits of joint inference relative to separate inference under the PIP and TKF91 models; i.e., the benefits of inferring trees on accuracy of the inferred MSA and the benefits of inferring MSAs on the accuracy of the inferred tree.

We used synthetic data to assess the quality of the tree reconstructions produced by PIP, compared to the reconstructions of PhyML 2.4.4, a widely-used platform for phylogenetic tree inference [15]. We also compared the inferred MSAs to those produced by Clustal 2.0.12 [19], a popular MSA inference system.

While our implementation evaluated in this section is based on the Bayesian framework, we evaluate it using a frequentist methodology. More precisely, we use Bayes estimators (described in Appendix D) to obtain two point estimates from the posterior, one for the MSA, and one for the phylogeny. Each point estimate is compared to the true alignment and tree. It is therefore possible to compare the method to the well-known frequentist methods mentioned above.

In this study, we explored four types of potential improvements: (1) resampling trees and MSAs increasing the quality of inferred MSAs, compared to resampling only MSAs; (2) resampling trees and MSAs increasing the quality of inferred trees, compared to resampling only trees; (3) resampling trees increasing the quality of inferred trees, compared to trees inferred by PhyML, and fixing the MSA to the one produced by Clustal; (4) resampling MSAs increasing the quality of inferred MSAs, compared to MSAs inferred by Clustal, and fixing the tree to the one produced by PhyML. The results are shown in Table 1. These experiments were based on 100 replicas, each having 7 taxa at the leaves, a topology sampled from the uniform distribu-

Table 1: PIP results on simulated data

Exp.	Tree resampled?	No	Yes	No	Yes
	MSA resampled?	No	No	Yes	Yes
MSAs	Edge recall (SP)	<b>0.25</b>	-	0.22	0.24
	Edge Precision	0.22	-	0.56	<b>0.58</b>
	Edge F1	0.23	-	0.31	<b>0.32</b>
Trees	Partition Metric	0.24	0.22	-	<b>0.19</b>
	Robinson-Foulds	0.45	0.38	-	<b>0.33</b>

tion, branch lengths sampled from rate 2 exponential distributions, indels generated from the PIP with parameters  $\eta = 100, \zeta = 1$ , and nucleotides sampled from the Kimura two-parameter model (K2P) [26].

We measured the quality of MSA reconstructions using the F1 score, defined as the harmonic mean of the reconstructed alignment edge recall (called the sum-of-pairs score or developer’s score in the MSA literature [55]) and alignment edge precision (modeler’s score [68]). We measured the quality of tree reconstructions using the partition (symmetric clade difference) metric [5] and the weighted Robinson-Foulds metric [50]. Relative improvements were obtained by computing the absolute value of the quality difference (in terms of the F1 for alignments, and Robinson-Foulds distance for trees), divided by the initial value of the measure. We report relative improvements averaged over the 100 replicas.

We observed improvements of all four types. Comparing Edge F1 relative improvements to Robinson-Foulds relative improvements, the relative additional improvement of type (2) is larger (13%) than that of type (1) (3%). Overall (i.e., comparing the baselines to the joint system), the full improvements of both trees and MSAs are substantial: 43% Edge F1 improvement, and 27% Robinson-Foulds improvement. See Figure 4 for a summary of the relative improvements.

We also tested our system on data generated from the TKF91 model instead of the PIP model. We used the same tree distribution and number of replicas as in the previous experiments, and the same generating TKF91 parameters as [21]. We again observed improvements over the baselines, both

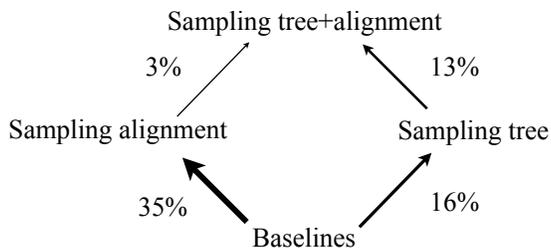


Figure 4: Relative improvements for enabling each component of the sampler. Arrows on the left are relative alignment improvements, and arrows on the right are relative tree improvements.

in terms of MSA and tree quality. For MSAs, the relative improvement over the baseline was actually larger on the TKF91-generated data than on the PIP-generated data (47% versus 43%, as measured by Edge F1 improvement over Clustal), and lower but still substantial for phylogenetic trees (13% versus 27%, as measured by Robinson-Foulds improvement over PhyML).

It should be noted that the MCMC kernels used in these experiments (described in the Appendix C) are based on simple Metropolis-Hastings proposals, and can therefore suffer from high rejection rates in large datasets. Fortunately, previous work in the statistical alignment literature has developed sophisticated MCMC kernels, some of which could be applied to inference in our model, for example [46]. Another potential direction would be to replace MCMC by a Sequential Monte Carlo posterior approximation [4].

It should also be emphasized that point indels is certainly not the exclusive driving force behind sequence evolution. In particular, “long indels” (atomic insertions and deletions of long segments, with a probability higher than the product of their point indels), are also prominent. As a consequence, any system purely based on point indels will have significant biases on biological data. In practice, these biases will introduce three undesirable artifacts: overestimation of the branch lengths; “gappy alignments,” where the reconstructed MSA has many scattered gaps instead of a few long ones; and the related “ragged end” problem, where the prefix and suffix of sequences are poorly aligned because observed sequences are often truncated in practice. In the next section, we propose ways to address these limitations.

## 6 Discussion

We have presented a novel string-valued evolutionary model that can be used for joint inference of phylogenies and multiple sequence alignments. As with its predecessor, the TKF91 model, the new model can be used to capture the homology of characters evolving on a phylogenetic tree under insertion, deletion and substitution events. Its advantage over TKF91 is that it permits a representation as a Poisson process on the tree. This representation has the consequence that the marginal likelihood of a tree and an alignment (marginalizing over ancestral states) can be computed in time linear in the number of taxa, rather than exponential as in the case of TKF91. Poisson representations have played an important role in pure substitution processes [22, 44, 1], but to the best of knowledge, this is the first time Poisson representations are used for indel inference.

Although the insertion process in TKF91 might be argued to be more realistic biologically than that of the PIP model, in that it allows the insertion rate to vary as the sequence length varies, in the common setting in which all of the sequences being aligned are of roughly similar lengths, this extra degree of freedom may be of limited value for inference. Indeed, in our experiments we saw that the PIP model can perform well even when data are generated from the TKF91 model. We might also note that there are biological processes in which insertions originate from a source that is extrinsic to the sequence (e.g., viruses or other genomic regions), in which case the constant-rate assumption of PIP may actually be preferred.

It is also important to acknowledge, however, that neither TKF91 nor PIP are accurate representations of biology. Their use in phylogenetic modeling reflects the hope that the statistical inferences they permit—most notably taking into account the effect of indels on the tree topology—will nonetheless be useful as data accrue. This hope is more likely to be realized in larger datasets, motivating our goal of obtaining a method that scales to larger sets of species. But both models should also be viewed as jumping-off points for further modeling that is more faithful to the biology while retaining the inferential power of the basic models. For example, there has been significant work on extending TKF91 to models that capture the “long indels” that arise biologically but are not captured by the basic model [63, 41, 40].

In this regard, we wish to note that the Poisson representation of the PIP model provides new avenues for extension that are not available within the TKF91 framework. In particular, the superposition property of Poisson

processes makes it possible to combine the PIP model with other models that follow a Poisson law. For example, if the location  $X'$  of long indels, slipped-strand mispairing [25] or other non-local changes follow a Poisson point process, the union  $U = X \cup X'$  of the non-local changes with the point indels  $X$  provided by a PIP will also be distributed according to a Poisson process. Moreover, the thinning property of Poisson processes provides a principled approach to inference for such superpositions. Indeed, an MCMC sampler for the superposition model can be constructed as follows: first, we can exploit the decomposition to analytically marginalize  $X$  (using the algorithm presented in this paper). Second, the other terms of the superposition and the sequences at these point in time can be represented explicitly as auxiliary variables. Since we have an efficient algorithm for computing the marginal likelihood, the auxiliary variables can be resampled easily. Note that designing an irreducible sampler without marginalizing  $X$  would be difficult: integrating out  $X$  creates a bridge of positive probability between any pair of patterns of non-local changes.

Under the parameterization of the process used in this paper, the model assumes both an equal deletion rate for all characters, and a uniform probability over inserted characters. It is worth noting that our inference algorithm can be modified to handle models relaxing both assumptions, by replacing the calculation of  $\beta(v)$  in Appendix B by a quasi-stationary distribution calculation [7]. It would be interesting to use this idea to investigate what non-uniformities are present in biological indel data.

Finally, another avenue to improve PIP models is to make the insertion rate mean measure more realistic: instead of being uniform across the tree, it could be modeled using a prior distribution, hence forming a Cox process [8]. This would be most useful when the sequences under study have large length or indel intensity variations across sites and branches [59].

## Acknowledgments

We would like to thank Bastien Boussau, Ian Holmes, Michael Newton and Marc Suchard for their comments and suggestions. This work was partially supported by a grant from the Office of Naval Research under Contract Number N00014-11-1-0688, by grant K22 HG00056 from the National Institutes of Health, and by grant SciDAC BER KP110201 from the Department of Energy.

## References

- [1] A. Alekseyenko, C. Lee, and M. A. Suchard. Wagner and Dollo: a stochastic duet by composing two parsimonious solos. *Systematic Biology*, 57 (5):772–784, 2008.
- [2] L. Allison, C. S. Wallace, and C. N Yee. Finite-state models in the alignment of macromolecules. *J. Mol. Evol.*, 35:77–89, 1992.
- [3] A. Bouchard-Côté, M. I. Jordan, and D. Klein. Efficient inference in phylogenetic InDel trees. *In Proceedings of Advances in Neural Information Processing Systems*, 21:177–184, 2009.
- [4] A Bouchard-Côté, S. Sankararaman, and M.I. Jordan. Phylogenetic Inference via Sequential Monte Carlo. *Systematic Biology*, 61:579–593, 2012.
- [5] M. Bourque. *Arbres de Steiner et reseaux dont certains sommets sont á localisation variable*. PhD thesis, Université de Montréal, 1978.
- [6] R. K. Bradley, A. Roberts, M. Smoot, S. Juvekar, J. Do, C. Dewey, I. Holmes, and L. Pachter. Fast statistical alignment. *PLoS Comput Biol.*, 5(5):e1000392, 2009.
- [7] M. Buiculescu. Quasi-stationary distributions for continuous-time Markov processes with a denumerable set of states. *Rev. Roum. Math. Pures et Appl.*, XVII:1013–1023, 1972.
- [8] D. R. Cox. Some statistical methods connected with series of events. *J. R. Statist. Soc. Ser. B*, 17(2):129–164, 1955.
- [9] F. W. Crawford and M. A. Suchard. Transition probabilities for general birth-death processes with applications in ecology, genetics and evolution. *Journal of Mathematical Biology*, (In Press), 2012.
- [10] B. Delyon. Convergence of the simulated annealing algorithm. Technical report, Massachusetts Institute of Technology, 1988.
- [11] J. L. Doob. Markoff chains: Denumerable case. *Transactions of the American Mathematical Society*, 58 (3):455–473, 1945.

- [12] M. Dreyer, J. R. Smith, and J. Eisner. Latent-variable modeling of string transductions with finite-state methods. *In Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 13:1080–1089, 2008.
- [13] J. Felsenstein. Evolutionary trees from DNA sequences: A maximum likelihood approach. *Journal of Molecular Evolution*, 17:368–376, 1981.
- [14] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81 (25):2340–2361, 1977.
- [15] S. Guindon and O. Gascuel. A simple, fast, and accurate algorithm to estimate large phylogenies by maximum likelihood. *Systematic Biology*, 52(5):696–704, 2004.
- [16] J. Hein. A unified approach to phylogenies and alignments. *Methods in Enzymology*, 183:625–944, 1990.
- [17] J. Hein. An algorithm for statistical alignment of sequences related by a binary tree. *Pac. Symp. Biocomp.*, 6:179–190, 2001.
- [18] J. Hein, J. Jensen, and C. Pedersen. Recursions for statistical multiple alignment. *Proc Natl Acad Sci U S A*, 100(25):14960–5, 2003.
- [19] D. G. Higgins and P. M. Sharp. CLUSTAL: a package for performing multiple sequence alignment on a microcomputer. *Gene*, 73:237–244, 1988.
- [20] I. Holmes. A probabilistic model for the evolution of RNA structure. *BMC Bioinformatics*, 5:166, 2004.
- [21] I. Holmes and W. J. Bruno. Evolutionary HMM: A Bayesian approach to multiple alignment. *Bioinformatics*, 17:803–820, 2001.
- [22] J. P. Huelsenbeck, B. Larget, and D. L. Swofford. A compound poisson process for relaxing the molecular clock. *Genetics*, 154:1879–1892, 2000.
- [23] J. P. Huelsenbeck and R. Nielsen. Effect of non-independent substitution on phylogenetic accuracy. *Systematic Biology*, 48:317–328, 1999.
- [24] J. Jensen and J. Hein. Gibbs sampler for statistical multiple alignment. Technical report, Dept of Theor Stat, U Aarhus, 2002.

- [25] S. A. Kelchner. The evolution of non-coding chloroplast DNA and its application in plant systematics. *Annals of the Missouri Botanical Garden*, 87:482–498, 2000.
- [26] M. Kimura. A simple method for estimating evolutionary rate of base substitutions through comparative studies of nucleotide sequences. *J. Mol. Evol.*, 16:111–120, 1980.
- [27] J. F. C. Kingman. *Poisson Processes*. Oxford Studies in Probabilities, 1993.
- [28] B. Knudsen and M. Miyamoto. Sequence alignments and pair hidden Markov models using evolutionary history. *J Mol Biol.*, 333:453–460, 2003.
- [29] A. Krogh, M. Brown, I. S. Mian, K. Sjolander, and D. Haussler. Hidden Markov models in computational biology: Applications to protein modeling. *J. Mol. Biol.*, 235:1501–1531, 1994.
- [30] C. Lakner, P. van der Mark, J. P. Huelsenbeck, B. Larget, and F. Ronquist. Efficiency of Markov Chain Monte Carlo Tree Proposals in Bayesian Phylogenetics. *Systematic Biology*, 57(1):86–103, 2008.
- [31] G. Lancia and R. Ravi. GESTALT: Genomic Steiner alignments. *Notes in Computer Science*, 1645:101–114, 1999.
- [32] K. Liu, S. Raghavan, S. Nelesen, C. R. Linder, and T. Warnow. Rapid and accurate large-scale coestimation of sequence alignments and phylogenetic trees. *Science*, 324(5934):1561–1564, 2009.
- [33] K. Liu, T. J. Warnow, M. T. Holder, S. M. Nelesen, J. Yu, A. P. Stamatakis, and C. R. Linder. SATe-II: very fast and accurate simultaneous estimation of multiple sequence alignments and phylogenetic trees. *Syst Biol*, 61(1):90–106, 2012.
- [34] A. Löytynoja, A. J. Vilella, and N. Goldman. Accurate extension of multiple sequence alignments using a phylogeny-aware graph algorithm. *Bioinformatics*, 28:1684–1691, 2012.
- [35] G. Lunter, A. Drummond, I. Miklós, and J. Hein. Statistical alignment: Recent progress, new applications, and challenges. In R. Nielsen, editor,

*Statistical methods in Molecular Evolution*, Series in Statistics in Health and Medicine. Springer Verlag, 2004.

- [36] G. Lunter, I. Miklós, A. Drummond, J. Jensen, and J. Hein. Bayesian coestimation of phylogeny and sequence alignment. *BMC Bioinformatics*, 6(83), 2005.
- [37] D. Metzler, R. Fleissner, A. Wakolbinger, and von Haeseler A. Assessing variability by joint sampling of alignments and mutation rates. *J Mol Evol*, 53(6):660–669, 2001.
- [38] I. Miklós. Algorithm for statistical alignment of sequences derived from a Poisson sequence length distribution. *Disc. Appl. Math.*, 127:79–84, 2003.
- [39] I. Miklós, A. Drummond, G. Lunter, and J. Hein. *Algorithms in Bioinformatics*, chapter Bayesian Phylogenetic Inference under a Statistical Insertion-Deletion Model. Springer, 2003.
- [40] I. Miklós, G. A. Lunter, and I. Holmes. A long indel model for evolutionary sequence alignment. *Mol Biol Evol.*, 21 (3):529–540, 2004.
- [41] I. Miklós and Z. Toroczka. An improved model for statistical alignment. *Lecture Notes in Computer Science*, 2149:1–10, 2001.
- [42] M. Mohri. *Handbook of Weighted Automata*, chapter 6, pages 213–254. Monographs in Theoretical Computer Science. Springer, 2009.
- [43] S. Nelesen, K. Liu, D. Zhao, C. R. Linder, and T. Warnow. The effect of the guide tree on multiple sequence alignments and subsequent phylogenetic analyses. *Pac Symp Biocomput*, 13:25–36, 2008.
- [44] G. Nicholls and R. Gray. *Phylogenetic methods and the prehistory languages*, chapter Quantifying uncertainty in a stochastic Dollo model of vocabulary evolution, pages 161–172. McDonald Institute for Archaeological Research, 2006.
- [45] A. Novák, I. Miklós, R. Lyngsoe, and J. Hein. StatAlign: An extendable software package for joint Bayesian estimation of alignments and evolutionary trees. *Bioinformatics*, 24:2403–2404, 2008.

- [46] B. D. Redelings and M. A. Suchard. Joint Bayesian estimation of alignment and phylogeny. *Syst. Biol.*, 54(3):401–418, 2005.
- [47] B. D. Redelings and M. A. Suchard. Incorporating indel information into phylogeny estimation for rapidly emerging pathogens. *BMC Evolutionary Biology*, 7(40), 2007.
- [48] E. Rivas. Evolutionary models for insertions and deletions in a probabilistic modeling framework. *BMC Bioinformatics*, 6(1):63, 2005.
- [49] C. P. Robert. *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer, 2001.
- [50] D.F. Robinson and L.R. Foulds. Comparison of weighted labelled trees. *Lecture Notes in Mathematics*, 748:119–126, 1979.
- [51] S. Roch. Toward extracting all phylogenetic information from matrices of evolutionary distances. *Science*, 327(5971):1376–1379, 2010.
- [52] U. Roshan, D. R. Livesay, and S. Chikkagoudar. Improving progressive alignment for phylogeny reconstruction using parsimonious guide-trees. *In Proceedings of the IEEE Symposium on Bioinformatics and Bioengineering*, 6:159–164, 2006.
- [53] N. Saitou and M. Nei. The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Mol Biol Evol.*, 4(4):406-25, 1987.
- [54] D. Sankoff. Minimal mutation trees of sequences. *SIAM Journal on Applied Mathematics*, 28:35–42, 1975.
- [55] J. Sauder, J. Arthur, and R. Dunbrack. Large-scale comparison of protein sequence alignment algorithms with structure alignments. *Proteins*, 40:6–22, 2000.
- [56] A. Schwartz and L. Pachter. Multiple alignment by sequence annealing. *Bioinformatics*, 23:e24–e29, 2006.
- [57] D. B. Searls and K. P. Murphy. Automata-theoretic models of mutation and alignment. In *Proc Int Conf Intell Syst Mol Biol.*, 1995.
- [58] S. Snir and L. Pachter. Tracing the most parsimonious indel history. *Journal of Computational Biology*, 18(8):967–986, 2011.

- [59] S. Sniret and L. Pachter. Phylogenetic profiling of insertions and deletions in vertebrate genomes. *Lecture Notes in Computer Science*, 3909:265–280, 2006.
- [60] Y. S. Song. A sufficient condition for reducing recursions in hidden Markov models. *Bulletin of Mathematical Biology*, 68:361–384, 2006.
- [61] M. Steel and J. Hein. Applying the Thorne-Kishino-Felsenstein model to sequence evolution on a star-shaped tree. *Appl. Math. Let.*, 14:679–684, 2001.
- [62] J. L. Thorne, H. Kishino, and J. Felsenstein. An evolutionary model for maximum likelihood alignment of DNA sequences. *Journal of Molecular Evolution*, 33:114–124, 1991.
- [63] J. L. Thorne, H. Kishino, and J. Felsenstein. Inching toward reality: an improved likelihood model of sequence evolution. *J. of Mol. Evol.*, 34:3–16, 1992.
- [64] A. Varon, L. S. Vinh, and W. C. Wheeler. POY version 4: phylogenetic analysis using dynamic homologies. *Cladistics*, 26:72–85, 2010.
- [65] O. Westesson, G. Lunter, B. Paten, and I. Holmes. Accurate reconstruction of insertion-deletion histories by statistical phylogenetics. *PLoS One*, 7(4):e34572, 2012.
- [66] W. C. Wheeler and D. S. Gladstein. MALIGN: A multiple sequence alignment program. *J Hered*, 85 (5):417–418, 1994.
- [67] K. M. Wong, M. A. Suchard, and J. P. Huelsenbeck. Alignment uncertainty and genomic analysis. *Science*, 319(5862):473–476, 2008.
- [68] M. Zachariah, G. Crooks, S. Holbrook, and S. Brenner. A generalized affine gap model significantly improves protein sequence alignment accuracy. *Proteins*, 58:329–38, 2005.
- [69] J. Zhang. Protein length distributions for the three domains of life. *Trends Genet.*, 16:107–109, 2000.

## A Proofs for the Main PIP Properties

In this section, we prove Theorem 1 and Proposition 2. We begin by stating and proving two lemmas.

**Lemma 3** *Let  $U \sim \text{Unif}(0, t)$  and  $W \sim \text{Exp}(\mu)$  be independent for fixed  $t, \mu > 0$ . Then*

$$\mathbb{P}(W + U > t) = \frac{1 - \exp(-t\mu)}{t\mu}.$$

**Proof** By conditioning:

$$\begin{aligned} \mathbb{P}(W + U > t) &= \mathbb{E}[\mathbb{P}(W + U > t|U)] \\ &= \int_0^t \frac{\exp(-x\mu)}{t} dx \\ &= \frac{1 - \exp(-t\mu)}{t\mu}. \end{aligned}$$

□

**Lemma 4** *Let  $\tau_0$  denote a degenerate topology consisting of a root  $\Omega$  connected to a single leaf  $v_0$  by an edge of length  $t$ . Let  $H_i$  be a homology path as defined in the main paper, with  $\tau = \tau_0$ . For all  $x \in \tau$ , define  $I(x) = \{i : H_i(x) \neq \varepsilon, 1 \leq i \leq I\}$  and:*

$$\begin{aligned} N &= |I(\Omega)| \\ N' &= |I(v_0)|. \end{aligned}$$

*Then  $N \sim \text{Poi}(\lambda/\mu)$  implies  $N' \sim \text{Poi}(\lambda/\mu)$ .*

**Proof** To prove the result, we decompose  $N$  and  $N'$  as follows (see Figure 5):

$$\begin{aligned} N_1 &= |I(\Omega) \setminus I(v_0)| \\ N_2 &= |I(\Omega) \cap I(v_0)| \\ N_3 &= |I(v_0) \setminus I(\Omega)| \\ N_4 &= |I \setminus I(\Omega) \setminus I(v_0)| \\ N &= N_1 + N_2 \\ N' &= N_2 + N_3. \end{aligned}$$

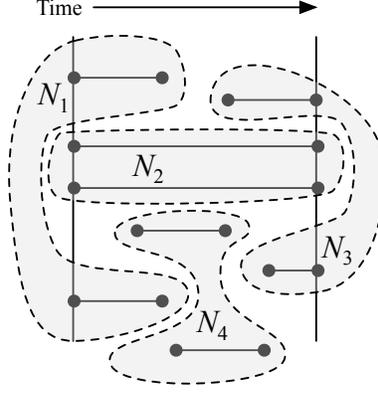


Figure 5: Notation used in the appendix. The horizontal lines denote the times where each character is present in the sequence. The vertical line on the left denotes the sequence at  $\Omega$ , and the vertical line on the right, the sequence at  $v_0$ . The sites are decomposed depending on whether they are present at each of two points  $\Omega, v_0$  in  $\tau_0$ .

By the Coloring Theorem [27],

$$N_2 \sim \text{Poi}(\nu(\{\Omega\})\mathbb{P}(W > t)),$$

where  $W$  is a rate  $\mu$  exponential random variable, and  $\nu$  is as in the condition of Theorem 1. Therefore  $N_2 \sim \text{Poi}(\lambda \exp(-t\mu)/\mu)$ . Similarly,

$$N_3 \sim \text{Poi}(\nu(\tau \setminus \{\Omega\})\mathbb{P}(W + U > t)),$$

where  $U \sim \text{Unif}(0, t)$ , and therefore from Lemma 3,  $N_3 \sim \text{Poi}(\lambda(1 - \exp(-t\mu))/\mu)$ . It follows that:

$$\begin{aligned} N' &= N_2 + N_3 \\ &\sim \text{Poi}\left(\frac{\lambda}{\mu}e^{-\mu} + \frac{\lambda}{\mu}(1 - e^{-\mu})\right) \\ &= \text{Poi}\left(\frac{\lambda}{\mu}\right), \end{aligned}$$

which concludes the proof of the lemma.  $\square$

We can now prove Theorem 1:

**Proof** In order to establish the equivalence, it is enough to show that for all edges  $e = (v \rightarrow v')$  in the tree, the following two properties hold:

1. The distribution of the string length at the ancestral endpoint,  $|Y(v)|$ , is identical in the local and global descriptions: a Poisson distribution with rate  $\lambda/\mu$ .
2. The distribution of the number and locations of mutations that fall on  $e \setminus \{v, v'\}$  are also identical in the local and global descriptions.

We will enumerate the edges in the tree in preorder, using induction to establish these two hypotheses on this list of edges.

In the base case, hypothesis 1 is satisfied by construction: the local description is initialized with a  $\text{Poi}(\lambda/\mu)$ -distributed number of characters, and in the global description, the intensity measure  $\nu$  of the Poisson process  $\mathbf{X}$  assigns a point mass  $\lambda/\mu$  to  $v = \Omega$ .

To establish hypothesis 1 in the inductive case, let  $e' = (v'' \rightarrow v)$  denote the parent edge. By hypothesis 1 on  $e'$ ,  $|Y(v'')| \sim \text{Poi}(\lambda/\mu)$ , therefore by Lemma 4 and hypothesis 2 on  $e'$ , hypothesis 1 is satisfied on  $e$  as well.

To establish hypothesis 2, it is enough to show that for all  $x \in e \setminus \{v, v'\}$  the waiting time for each type of mutation given  $Y(x)$  is exponential, with rates:

- (a)  $\lambda$  for insertion,
- (b)  $\mu \cdot |Y(x)|$  for deletion, and
- (c)  $\sum_{\sigma \neq \varepsilon} \theta_{\sigma, \sigma'} |Y(x)|_{\sigma}$  for substitutions to  $\sigma' \neq \varepsilon$ , where  $|s|_{\sigma}$  denotes the number of characters of type  $\sigma \in \Sigma$  in the string  $s \in \Sigma^*$ .

Item (a) follows from the Poisson Interval Theorem [27]. Items (b) and (c) follow from the standard Doob-Gillespie characterization of CTMCs: if  $X_t$  is a CTMC with rate matrix  $Q = (q_{i,j})$  and  $Z_{i,j}$  are independent exponential random variables with rate  $q_{i,j}$ , then

$$(\Delta, J)|(X_0 = i) \stackrel{d}{=} (\min_{j \neq i} Z_{i,j}, \operatorname{argmin}_{j \neq i} Z_{i,j}),$$

where  $\Delta = \inf\{t : X_t \neq i\}$ ,  $J = X_{\Delta}$ . □

We now turn to Proposition 2 and establish reversibility.

**Proof** Let  $h(n_1, n_2, n_3, n_4) = \mathbb{P}(N_i = n_i, i \in \{1, 2, 3, 4\})$ . Using reversibility of  $\theta$ , it is enough to show that  $h$  is invariant under the permutation (1 3); i.e.,  $h(n_1, n_2, n_3, n_4) = h(n_3, n_2, n_1, n_4)$ .

We have that  $h(n_1, n_2, n_3, n_4)$  is equal to:

$$\begin{aligned}
& \mathbb{P}\left(N_i = n_i, \sum_i N_i = \sum_i n_i, N_1 + N_2 = n_1 + n_2, N_3 + N_4 = n_3 + n_4\right) \\
&= \mathbb{P}\left(\sum_i N_i = \sum_i n_i\right) \times \\
& \quad \mathbb{P}\left(N_1 + N_2 = n_1 + n_2, N_3 + N_4 = n_3 + n_4 \mid \sum_i N_i = \sum_i n_i\right) \times \\
& \quad \mathbb{P}(N_1 = n_1, N_2 = n_2 \mid N_1 + N_2 = n_1 + n_2) \times \\
& \quad \mathbb{P}(N_3 = n_3, N_4 = n_4 \mid N_3 + N_4 = n_3 + n_4) \\
&= f_1(n_1 + n_2 + n_3 + n_4) \times \\
& \quad \left(\frac{1/\mu}{1/\mu + t}\right)^{n_1 + n_2} \left(\frac{t}{1/\mu + t}\right)^{n_3 + n_4} \times \\
& \quad (1 - e^{-\mu t})^{n_1} f_2(n_2) \times \\
& \quad \left(\frac{1 - e^{-\mu t}}{t\mu}\right)^{n_3} f_3(n_4),
\end{aligned}$$

where only the dependencies of the functions  $f_1, f_2$  and  $f_3$  is important in this argument, not their exact form. By inspection, it is clear that  $h$  is invariant under the permutation (1 3).  $\square$

## B Proofs for the Likelihood Computation

First, we show how the function  $\varphi$ , defined in the main paper, simplifies the computation of  $p_\tau(m)$ :

$$\begin{aligned}
p_\tau(m) &= \mathbb{E}[\mathbb{P}(M = m | \mathbf{X})] \\
&= \sum_{n=|m|}^{\infty} \mathbb{P}(|\mathbf{X}| = n) \cdot \binom{n}{|m|} \cdot (p(c_\emptyset))^{n-|m|} \prod_{c \in m} p(c) \\
&= \frac{e^{\|\nu\|} \prod_{c \in m} p(c)}{|m|! (p(c_\emptyset))^{|m|}} \sum_{n=|m|}^{\infty} \frac{(\|\nu\| p(c_\emptyset))^n}{(n - |m|)!} \\
&= \frac{e^{\|\nu\|} (\|\nu\| p(c_\emptyset))^{|m|} \prod_{c \in m} p(c)}{|m|! (p(c_\emptyset))^{|m|}} \sum_{k=0}^{\infty} \frac{(\|\nu\| p(c_\emptyset))^k}{k!} \\
&= \frac{e^{\|\nu\|} (\|\nu\| p(c_\emptyset))^{|m|} \prod_{c \in m} p(c)}{|m|! (p(c_\emptyset))^{|m|}} \exp(\|\nu\| p(c_\emptyset)) \\
&= \varphi(p(c_\emptyset), |m|) \prod_{c \in m} p(c).
\end{aligned}$$

Next, we show how to compute  $f_v = \mathbb{P}(C = c | V = v)$  for all  $v \in \mathcal{V}$ . The recursions for  $f_v$  are similar to those found in stochastic Dollo models [1]. Note first that  $f_v$  can be zero for some vertices. To see where and why, consider the subset of leaves  $S$  that have an extant nucleotide in the current column  $c$ ,  $S = \{v \in \mathcal{L} : H(v) \neq \varepsilon\}$ . Then  $f_v$  will be non-zero only for the vertices ancestral to all the leaves in  $S$ . Let us call this set of vertices  $A$  (see Figure 6).

To compute  $f_v$  on the remaining vertices, we introduce an intermediate variable,  $\tilde{f}_v = \mathbb{P}(C = c | V = v, H(v) \neq \varepsilon)$ . This variable can be computed using the standard Felsenstein peeling recursion (dynamic programming) as follows:

$$\tilde{f}_v(\sigma) = \begin{cases} \mathbf{1}(c(v) = \sigma) & \text{if } v \in \mathcal{L} \\ \sum_{\sigma' \in \Sigma_\varepsilon} \exp(b(v)Q)_{\sigma, \sigma'} \prod_{w \in \text{child}(v)} \tilde{f}_w(\sigma') & \text{o.w.} \end{cases} \quad (1)$$

$$\tilde{f}_v = \sum_{\sigma \in \Sigma} \pi_\sigma \tilde{f}_v(\sigma). \quad (2)$$

From Lemma 3, we have an expression for the survival probability at  $v$

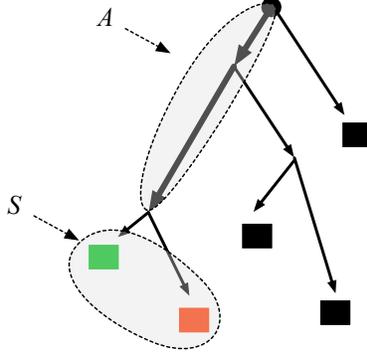


Figure 6: Given a set  $S$  of leaves  $v$  with  $H(v) \neq \varepsilon$ , we define the set  $A$  of vertices with nonzero modified Felsenstein peeling weight to be those ancestral to the leaves in  $S$ . In this example,  $A$  contains three vertices.

given an insertion on the edge  $(\text{pa}(v) \rightarrow v)$ :

$$\begin{aligned} \beta(v) &= \mathbb{P}(H(v) \neq \varepsilon | V = v) \\ &= \frac{1}{b(v)} \frac{1}{\mu} (1 - e^{-\mu b(v)}). \end{aligned} \quad (3)$$

Finally, for  $c \neq c_\emptyset$ , we have:

$$\begin{aligned} f_v &= \mathbb{P}(C = c | V = v) \\ &= \mathbb{E}[\mathbb{P}(C = c | V = v, H(v))] \\ &= \begin{cases} \tilde{f}_v & \text{if } v = \Omega \\ \mathbf{1}[v \in A] \beta(v) \tilde{f}_v & \text{o.w.,} \end{cases} \end{aligned} \quad (4)$$

and for  $c = c_\emptyset$ :

$$f_v = \begin{cases} \tilde{f}_v & \text{if } v = \Omega \\ 1 + \beta(v)(\tilde{f}_v - 1) & \text{o.w.} \end{cases} \quad (5)$$

## C Proposal distributions

To perform full joint inference over trees and alignments using Markov chain Monte Carlo, several objects need to be resampled: the tree topology, the branch lengths, the MSA, and the parameters.

For trees and branch lengths, we use standard proposal mechanisms [30]. Our MSA proposal is inspired by the proposal of [36], avoiding the mixing problems of auxiliary variables [21, 24, 3]. Our proposal distribution consists of two steps. First, we partition the leaves into two sets  $A, B$ . Given a current MSA  $m_0$ , the support of the proposal is the set  $S$  of MSAs  $m$  satisfying the following constraints:

1. If  $e$  has both endpoints in  $A$  (or both in  $B$ ), then  $e \in m \iff e \in m_0$ .
2. If  $e, e'$  have both endpoints in  $A$  (or both in  $B$ ), then  $e \prec_m e' \iff e \prec_{m_0} e'$ .

The notation  $\prec_m$  is based on the concept of posets over the columns (and edges) of an MSA [56].

We propose an element  $m^* \in S$  with probability proportional to  $\prod_{c \in m^*} p(c)$ . The set  $S$  has exponential size, but can be sampled efficiently using standard pairwise alignment dynamic programming. A Metropolis-Hastings ratio is then computed to correct for  $\varphi$ . Note that the proposal induces an irreducible chain: one possible outcome of the move is to remove all links between two groups of sequences. The chain can therefore move to the empty MSA and then construct any MSA incrementally.

For the parameters, we used multiplicative proposals in the  $(\lambda, \mu)$  parameterization [30].

## D Computational Aspects

In this section, we provide a brief discussion of the role that the marginal likelihood plays in both frequentist and Bayesian inference methods.

### D.1 Maximum likelihood

In the case of maximum likelihood, the overall inference problem involves optimizing over the marginal likelihood:

$$\sup_{\tau \in \mathcal{T}(\mathcal{L}), m \in \mathcal{M}(y)} \log p_\tau(m),$$

where  $\tau$  ranges over phylogenies on the leaves  $\mathcal{L}$ , and  $m$  ranges over the alignments consistent with the observed sequences  $y$ . This optimization problem

can be approached using simulated annealing, where a candidate phylogeny and MSA pair  $(\tau', m')$  is proposed at each step  $i$ , and is accepted (meaning that it replaces the previous candidate  $(\tau, m)$ ) according to a sequence of acceptance functions  $f^{(i)}(p, p')$  depending only on the marginal probabilities  $p = p_\tau(m), p' = p_{\tau'}(m')$ . Provided  $\lim_{i \rightarrow \infty} f^{(i)}(p, p') = \mathbf{1}[p' > p]$  sufficiently slowly, this algorithm converges to the maximum likelihood phylogeny and MSA [10].

## D.2 Bayes estimators

In order to define a Bayes estimator, one typically specifies a decision space  $D$  (for example the space of MSAs, or the space of multifurcating tree topologies, or both), a projection into this space,  $(\tau, m) \mapsto \rho(\tau, m) \in D$ , and a loss function  $l : D \rightarrow [0, \infty)$  on  $D$  (for example, for tree topologies, the symmetric clade difference, or partition metric [5]; and for alignments, 1 – the edge recall or Sum-of-Pairs (SP) score [49]).

Given these objects, the optimal decision in the Bayesian framework (also known as the consensus tree or alignment), is obtained by minimizing over  $d \in D$  the risk  $\mathbb{E}[l(d, \rho(T, M)) | \mathcal{Y}]$ . This expectation is intractable, so it is usually approximated with the empirical distribution of the output  $(\tau^{(i)}, m^{(i)})$  of an Markov chain Monte Carlo (MCMC) algorithm. Producing MCMC samples boils down to computing acceptance ratios of the form:

$$\frac{p(\tau')p_{\tau'}(m')}{p(\tau)p_\tau(m)} \cdot \frac{q_{(\tau', m')}(\tau, m)}{q_{(\tau, m)}(\tau', m')},$$

for some proposal having density  $q$  with respect to a shared reference measure on  $\mathcal{T}(\mathcal{L}) \times \mathcal{M}(y)$ . We thus see that for both maximum likelihood and joint Bayesian inference of the MSA and phylogeny the key problem is that of computing the marginal likelihood  $p_\tau(m)$ .

## E Pseudocode and Example

In this section, we summarize the likelihood computation. We also give a concrete numerical example to illustrate the calculation.

1. Inputs:

- (a) PIP parameter values  $(\lambda, \mu)$ , substitution matrix  $\theta$  over  $\Sigma$ .

*Example:*  $(\lambda, \mu) = (2.0, 1.0), \Sigma = \{a\}$

- (b) Rooted phylogenetic tree  $\tau$

*Example:*  $\tau = ((v_2 : 1.0, v_3 : 1.0)v_0 : 1.0, v_4 : 2.0)v_1;$

- (c) Multiple sequence alignment  $m$

*Example:*  $m =$

```
v_2|-a
v_3|aa
v_4|a-
```

2. Computing modified Felsenstein recursion:

- (a) For each site, compute  $\tilde{f}_v(\sigma)$  in post-order using Equation (1), and from each  $\tilde{f}_v(\sigma)$ , compute  $f_v$  using Equation (2)

*Example:*

*for site 1,*  $(\tilde{f}_{v_2}, \tilde{f}_{v_3}, \tilde{f}_{v_0}, \tilde{f}_{v_4}, \tilde{f}_{v_1}) = (0.0, 1.0, 0.23, 1.0, 0.012);$

*for site 2,*  $(\tilde{f}_{v_2}, \tilde{f}_{v_3}, \tilde{f}_{v_0}, \tilde{f}_{v_4}, \tilde{f}_{v_1}) = (1.0, 1.0, 0.14, 0.0, 0.043);$

- (b) Do the same for an artificial site or column  $c_\emptyset$  where all leaves have a gap

*Example:*

*for site 3,*  $(\tilde{f}_{v_2}, \tilde{f}_{v_3}, \tilde{f}_{v_0}, \tilde{f}_{v_4}, \tilde{f}_{v_1}) = (0.0, 0.0, 0.40, 0.0, 0.67);$

3. For each node  $v$  in the tree, compute the survival probability  $\beta(v)$  using Equation (3) (setting it to 1 at the root for convenience)

*Example:*

$(\beta(v_2), \beta(v_3), \beta(v_0), \beta(v_4), \beta(v_1)) = (0.63, 0.63, 0.63, 0.43, 1.0)$

4. For each site, compute the set of nodes  $A$  ancestral to all extant characters, as described in the caption of Figure 6

*Example:*

*for site 1,*  $A = \{v_1\}$

*for site 2,*  $A = \{v_0, v_1\}$

5. Computing  $f_v$ :

- (a) For each site, compute  $f_v$  using Equation (4)

*Example:*

*for site 1,*  $(f_{v_2}, f_{v_3}, f_{v_0}, f_{v_4}, f_{v_1}) = (0.0, 0.0, 0.0, 0.0, 0.012);$

*for site 2,*  $(f_{v_2}, f_{v_3}, f_{v_0}, f_{v_4}, f_{v_1}) = (0.0, 0.0, 0.086, 0.0, 0.043);$

(b) For  $c_\emptyset$ , use Equation (5)

*Example:*

for site 3,  $(f_{v_2}, f_{v_3}, f_{v_0}, f_{v_4}, f_{v_1}) = (0.37, 0.37, 0.62, 0.57, 0.67)$ ;

6. For each node  $v$  in the tree, compute  $\iota_v = \mathbb{P}(V = v)$  as shown in Section 3 of the main paper

*Example:*

$(\iota(v_2), \iota(v_3), \iota(v_0), \iota(v_4), \iota(v_1)) = (0.17, 0.17, 0.17, 0.33, 0.17)$

7. Compute  $p_\tau(m)$  from the  $\iota_v$ 's,  $f_v$ 's as shown in Section 3 of the main paper

*Example:*  $\log p_\tau(m) = -11$