

PAIRWISE SEQUENCE ALIGNMENT AT ARBITRARILY LARGE EVOLUTIONARY DISTANCE

BY BRANDON LEGRIED^{1,a} AND SEBASTIEN ROCHE^{2,b}

¹*School of Mathematics, Georgia Institute of Technology, ^ablegried3@gatech.edu*

²*Department of Mathematics, University of Wisconsin–Madison, ^broch@wisc.edu*

Ancestral sequence reconstruction is a key task in computational biology. It consists in inferring a molecular sequence at an ancestral species of a known phylogeny, given descendant sequences at the tip of the tree. In addition to its many biological applications, it has played a key role in elucidating the statistical performance of phylogeny estimation methods. Here we establish a formal connection to another important bioinformatics problem, multiple sequence alignment, where one attempts to best align a collection of molecular sequences under some mismatch penalty score by inserting gaps. Our result is counter-intuitive: we show that perfect pairwise sequence alignment with high probability is possible in principle *at arbitrary large evolutionary distances*—provided the phylogeny is known and dense enough. We use techniques from ancestral sequence reconstruction in the taxon-rich setting together with the probabilistic analysis of sequence evolution models involving insertions and deletions.

1. Introduction. Ancestral sequence reconstruction (ASR) is a key task in computational evolutionary biology [24]. It consists in inferring a molecular sequence at an ancestral species of a known phylogeny, given descendant sequences at the tip of the tree. Numerous approaches are available for this task. Some are based on statistical models of sequence evolution on a tree, while others rely on combinatorial optimization formulations [37, 41]. In addition to its many biological applications, ASR has played a key role in elucidating the statistical performance of phylogeny estimation methods [29, 30, 33, 34]. Here we establish a formal connection to sequence alignment.

Rigorous analyses of the accuracy of ASR methods have been performed mainly in two asymptotic settings. In phylogenies of arbitrarily large depth, an achievable goal is to infer a sequence that is correlated site-by-site with the true ancestral sequence [11, 22, 27, 38]. In the taxon-rich setting, on the other hand, where the depth of the phylogeny is bounded as the number of taxa increases, consistent estimators are known to exist [17, 35]. That is, under conditions on the branching of the phylogeny around its root, the correct inference of a single site in the ancestral sequence can be guaranteed as the number of leaves goes to infinity.

Most theoretical results in this area are derived under models of sequence evolution by single site substitutions. More complex models allowing for site insertions and deletions (indels) have also been considered [2, 12, 16]. The star case, also known as trace reconstruction, has been the subject of much recent interest [3, 5–7, 9, 10, 20, 21, 31]. See also [1, 8, 13, 26, 39] for rigorous analyses of indel models in other contexts, for example, distance-based phylogeny reconstruction.

Indel models are closely related to another important bioinformatics problem, multiple sequence alignment (MSA), in which one attempts to best align a collection of molecular sequences under some mismatch penalty score by inserting gaps. In practice, MSA is a hard

Received July 2022; revised July 2023.

MSC2020 subject classifications. Primary 60K35; secondary 92D15.

Key words and phrases. Reconstruction, branching processes.

problem, especially at large evolutionary distances [4, 36]. While statistical approaches based on indel models have also been developed [25], commonly used approaches involve progressively aligning the given sequences up a guide tree, in what is reminiscent of ASR procedures [32]. In fact, many trace reconstruction and ASR methods under indels involve partial local alignments of sequences.

In this paper, we combine insights from ASR in the taxon-rich setting together with the probabilistic analysis of indel models to prove the first (as far as we know) rigorous guarantee for sequence alignment under an indel model on a phylogenetic tree. Our result is somewhat counter-intuitive: we show that perfect pairwise sequence alignment with high probability is in principle possible *at arbitrary large evolutionary distances*—provided the phylogeny is known and dense enough. While such a condition may not be satisfied in real datasets, our analysis is a step towards a better theoretical understanding of MSA and its connections to ASR.

In a nutshell, we take advantage of the density of the phylogeny to estimate ancestral sequences with high probability along the path between two leaf sequences of interest, then reconstruct the history of mutations along the way. For the ASR step, we use a standard phylogenetic method known as parsimony, which seeks to use the smallest number of mutations possible to explain sequences at the leaves of a phylogeny. Rigorous analyses of parsimony are often challenging and have revealed the intricate, often unintuitive, behavior of the method [14, 18, 19, 23, 42]. In our taxon-rich setting, branching process results lead to rigorous guarantees on the ancestral reconstruction.

The rest of the paper is organized as follows. In Section 2, we state our main result after introducing some background. The alignment algorithm is presented in Section 3. The proof is comprised of two parts: the ancestral estimation step is analyzed in Section 4 and the alignment step is analyzed in Section 5.

2. Background and main result. In this section, we state our main result. First, we introduce the model of sequence evolution we use here as well as the multiple sequence alignment problem.

2.1. Definitions. We consider the TKF91 insertion-deletion (indel) sequence evolution model. Technically, we use a slight variant of the TKF91 model defined in [40], where we only allow an alphabet with two letters 0 and 1 to simplify the analysis and its presentation. Our results extend naturally to more general settings, such as to any finite alphabet (see Section 2.3).

DEFINITION 1 (TKF91 model: two-state version). Consider the following Markov process $\mathcal{I} = \{\mathcal{I}_t\}_{t \geq 0}$ on the space \mathcal{S} of binary digit sequences together with an *immortal link* “•”, that is,

$$\mathcal{S} := \text{“•”} \otimes \bigcup_{M \geq 1} \{0, 1\}^M,$$

where the notation above indicates that all sequences begin with the immortal link. Positions of a sequence, except for that of the immortal link, are called *sites* or *mortal links*. Let $(\eta, \lambda, \mu) \in (0, \infty)^3$ and $(\pi_0, \pi_1) \in [0, 1]^2$ with $\pi_0 + \pi_1 = 1$ be given parameters. The continuous-time dynamics are as follows: If the current state is the sequence $\vec{x} \in \mathcal{S}$, then the following events occur independently:

- *Substitution:* Each site is substituted independently at rate $\eta > 0$. When a substitution occurs, the corresponding digit is replaced by 0 and 1 with probabilities π_0 and π_1 , respectively.

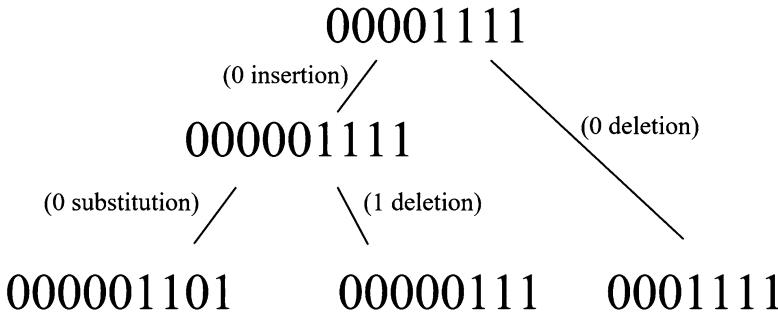


FIG. 1. *The evolution of a two-state TKF91 process on a tree T with three leaves.*

- *Deletion:* Each site is removed independently at rate μ .
- *Insertion:* Each site, as well as the immortal link, gives birth to a new digit independently at rate λ . When a birth occurs, the new site is added immediately to the right of its parent site. The newborn site has digit 0 and 1 with probabilities π_0 and π_1 , respectively.

We run this process on a rooted metric tree as follows. Consider a *rooted binary tree* $T = (V, E, \rho, \mathbf{t})$ with vertices V , edges E , root ρ , and edge lengths $\mathbf{t} = \{t_e\}_{e \in E}$ (in time units). It is standard in phylogenetics to use a binary tree, so that each interior node reflects the creation of a single new species. We restrict ourselves to ultrametric trees, that is, the sum of edge lengths from root to leaf is the same for every leaf. We refer to this common quantity as the *depth* of the tree and denote it by h . The rooted metric tree T is then indexed by all points along the edges of T . The root vertex has an initial sequence $\sigma_\rho \in \mathcal{S}$. With an initial sequence $\sigma_u \in \mathcal{S}$, the TKF91 process is recursively performed on each descending edge $e = (u, v)$ over the time interval $[0, t_e]$ to obtain another sequence $\sigma_v \in \mathcal{S}$. Processes running along descending edges of u are independent, conditioned on state σ_u at u . We refer to the full process as the *(two-state) TKF91 process on tree T* . An example of how the sequence evolves on the tree is given in Figure 1.

For any sequence $\sigma \in \mathcal{S}$, let $|\sigma|$ be the length of the sequence, and let $|\sigma|_0$ and $|\sigma|_1$ be the number of 0's and 1's in the sequence, respectively. The stationary distribution of the sequence length $|\sigma| = M$ is known [40] to be

$$(1) \quad \gamma_M = \left(1 - \frac{\lambda}{\mu}\right) \left(\frac{\lambda}{\mu}\right)^M, \quad M \in \mathbb{Z}_+,$$

provided $\mu > \lambda$. We assume that the root sequence σ_ρ follows its stationary distribution. That is, $|\sigma_\rho|$ is distributed according to γ_M and its sites are i.i.d. in $\{0, 1\}$ with respective probabilities π_0 and π_1 . Stationarity of σ_ρ implies stationarity of the TKF91 process throughout the tree. We assume from now on that $\mu > \lambda$ and that stationarity holds.

Some notation. Later on, we will need the following notation. For a sequence $\sigma \in \mathcal{S}$, let $\mathcal{S}_s(\sigma)$, $\mathcal{S}_d(\sigma)$, and $\mathcal{S}_i(\sigma)$ be the sequences that differ from σ respectively by a single substitution, a single deletion, and a single insertion. Observe that these sets are disjoint as the sequence lengths in each necessarily differ. Further, let $\mathcal{S}_1(\sigma) = \mathcal{S}_s(\sigma) \cup \mathcal{S}_d(\sigma) \cup \mathcal{S}_i(\sigma)$ be the sequences obtained by performing a single mutation on σ , and define

$$(2) \quad \lambda^*(\sigma) = \sum_{\tau \in \mathcal{S}_1(\sigma)} Q(\sigma, \tau) = \lambda(|\sigma| + 1) + \mu|\sigma| + \eta\pi_1|\sigma|_0 + \eta\pi_0|\sigma|_1$$

as the total rate under the TKF91 process of moving away from σ , where $Q(\sigma, \tau)$ is the rate at which the TKF91 process on an edge jumps from σ to τ . Formula (2) is derived formally in the [Appendix](#).

000001101	0000011 - 01
00000111	0000011 - - 1
000001101	00000110 - 1
000001111	00000111 - 1
0011111	-0001111 - 1

FIG. 2. *Left: A collection of sequences to align. Right: A possible multiple sequence alignment..*

2.2. *Multiple sequence alignment.* To compare sequences descending from a common ancestor through substitutions, insertions, and deletions, it is natural to attempt to align them as best as possible, that is, to construct a multiple sequence alignment.

DEFINITION 2 (Multiple sequence alignment). For any integer $n \geq 1$ and sequences $\sigma = (\sigma_{v_1}, \dots, \sigma_{v_m}) \in \mathcal{S}^m$ at points $v_1, \dots, v_m \in T$, a *multiple sequence alignment* (or pairwise alignment when $m = 2$) is a collection of sequences $\mathbf{a}(\sigma) = (a_1(\sigma), \dots, a_m(\sigma))$ whose entries come from $\{0, 1, -\}$ ($-$ is called a *gap*) such that:

- the lengths satisfy

$$|a_1(\sigma)| = |a_2(\sigma)| = \dots = |a_m(\sigma)| \geq \max\{|\sigma_{v_1}|, |\sigma_{v_2}|, \dots, |\sigma_{v_m}|\},$$

- no corresponding entries of $a_1(\sigma), \dots, a_m(\sigma)$ all equal $-$, and
- removing $-$ from $a_i(\sigma)$ yields σ_{v_i} for all $i \in \{1, 2, \dots, m\}$.

A multiple sequence alignment can be expressed as an $m \times |a_1(\sigma)|$ matrix where the rows are the sequence alignments and where no column consists of all gaps. If $m = 2$, the alignment is referred to as pairwise.

An example of a multiple sequence alignment is given in Figure 2. More generally, a multiple sequence alignment procedure may take as input further auxiliary information (beyond the sequences to be aligned), such as a tree or sequences at other points of the tree. Our alignment algorithm (see Section 3) will indeed use additional information.

Two sites, one from one sequence and the other from another sequence, are said to be *homologous* provided they descend from a common site in their most recent common ancestral sequence *only through substitutions* under the evolutionary process on the tree. A *true* multiple sequence alignment is one that places homologous sites in the same column and nonhomologous sites in different columns. We note however that certain homology relationships are unknowable a priori: for example, if in the course of evolution a 0 is inserted in a sequence next to another 0, which of them descends from the ancestral 0 is arbitrary. Here we take the convention that a repeated site is always inserted at the beginning of a run; and that similarly a repeated site is always deleted at the beginning of a run.

2.3. *Statement of main result.* The following theorem states that it is possible to construct with high probability a true pairwise alignment of the sequences at two arbitrary leaves v and w of a phylogeny as long as the maximal branch length is sufficiently small.

THEOREM 1 (Main result). *Fix $\eta, \mu, \lambda \in (0, \infty)$, the substitution, deletion, and insertion rates under the TKF91 model. There is a polynomial-time alignment procedure A such that for any tree depth $h > 0$ and any failure probability $\varepsilon > 0$, there exists a maximum branch length $t_{\max} := t_{\max}(h, \varepsilon) > 0$ such that the following property holds. For any rooted binary tree $T = (V, E, \rho, \mathbf{t})$ with vertices V , edges E , root ρ , and edge lengths $\mathbf{t} = \{t_e\}_{e \in E}$, assume*

that the leaves $\partial T = \{\ell_i\}_{i=1}^n$ are ordered from left to right in a planar realization of T , and let $v = \ell_1$ and $w = \ell_n$. Then the alignment procedure applied to the sequences $\sigma_{\ell_1}, \sigma_{\ell_2}, \dots, \sigma_{\ell_n}$ outputs a true pairwise alignment of σ_v and σ_w with probability at least $1 - \varepsilon$, provided that $t_e \leq t_{\max}$ for all edges $e \in E$.

Note that the tree depth h is arbitrary. The alignment procedure, which is described in Section 3, takes as input leaf sequences at the leaves of T as well as T itself. We show in Section 3 that the overall computational complexity of the algorithm is $O(n^2(\ell + \log n))$ where ℓ is the maximum sequence length at the leaves. Note however that we have not attempted to optimize the computational efficiency of the algorithm.

The maximum branch length, t_{\max} , depends only on h and ε . Whether the specific dependence in our proofs can be improved is a nontrivial problem. It is ultimately driven by our use of exact sequence reconstruction deep inside the tree. One could potentially get better quantitative bounds by allowing for errors in the reconstruction. However, little is known about approximate ancestral reconstruction under indel models [2, 16], even in the trace reconstruction setting [3, 5–7, 10] where the tree is a star and the evolution model is often simpler.

Extensions. While we have stated our main result for binary sequences to simplify the presentation, it holds unchanged for an arbitrary alphabet size. Firstly our algorithm, as described in Section 3, can be applied as is to sequences over an arbitrary alphabet. In particular, we state our ancestral reconstruction estimator in a form that does not rely on the size of the alphabet (see Definition 3). Second, the model we use, the TKF91 model, was in fact introduced over alphabets of any size in [40] (and applied specifically to nucleotide sequences, which have four states). There, insertions and substitutions are simply drawn from the stationary distribution over nucleotides. Under this more general model, all steps of the analysis remain unchanged. Indeed, the proofs rely on identifying substructures where none or one mutation has occurred (see Propositions 2 and 3).

While we assume above that the rate of substitution is the same throughout the tree, our proof still goes through if the parameter η is merely an upper bound on that rate across edges.

3. Alignment algorithm. In this section, we describe the alignment procedure of Theorem 1. We emphasize that this algorithm is not meant to be practical, but rather serve as a proof of our main result.

3.1. Overview of full alignment algorithm. We introduce the following alignment algorithm A which takes as input a rooted metric tree T , two distinguished leaves v and w , all leaf sequences, and a preprocessing parameter δ_1 . Here, we assume that the most recent common ancestor of v and w is the root. If it is not the root, then we consider only the subtree descending from the most recent common ancestor. We take δ_1 to satisfy $t_{\max} \leq \delta_1 \leq h$. The algorithm outputs a pairwise alignment for the sequences at v and w .

There is a unique path between v and w that we henceforth call the *backbone*. We let B be the number of *nonroot* vertices on the backbone. Then $v = x_1$ and $w = x_B$ and the other *nonroot* backbone vertices are in order x_2, \dots, x_{B-1} . For some parts of the algorithm and analysis, it will be convenient to use an alternative numbering of the backbone vertices—away from the root, numbering the left side, then numbering the right side. Specifically, let x_1, \dots, x_{B-} be the backbone vertices on the same side of the root as x_1 . Let x_{B-+1}, \dots, x_B be the backbone vertices on the same side of the root as x_B and let B^+ be the number of such vertices. Then we set

$$\tilde{x}_i^- := x_{B--(i-1)}, \quad i = 1, \dots, B^-$$

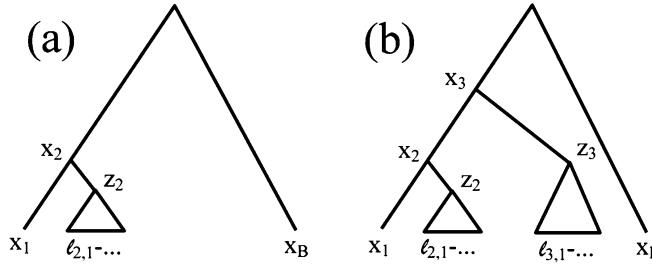


FIG. 3. (a) Tree T with leaf sequences $\sigma_1 = \sigma_{x_1}, \sigma_2 = \sigma_{\ell_{2,1}}, \dots, \sigma_{\ell_{2,i_2}}, \sigma_B = \sigma_{x_B}$. (b) Tree T with leaf sequences $\sigma_1 = \sigma_{x_1}, \sigma_2 = \sigma_{x_{2,1}}, \dots, \sigma_{\ell_{2,i_2}}, \sigma_3 = \sigma_{x_{3,1}}, \dots, \sigma_{\ell_{3,i_3}}, \sigma_B$.

and

$$\tilde{x}_i^+ := x_{B^-+i}, \quad i = 1, \dots, B^+.$$

Notice in particular that \tilde{x}_1^- and \tilde{x}_1^+ are the children of the root and $\tilde{x}_{B^-}^- = x_1$ and $\tilde{x}_{B^+}^+ = x_B$.

We now describe the main steps of the algorithm. Some details will be given in the following subsections. Figure 3 illustrates part of this algorithm at a high level.

We start with a preprocessing step. This step is used in the analysis of the algorithm. Indeed, it will be convenient for all edges on the backbone to be roughly of the same length δ_1 (up to a constant factor) and, furthermore, to have explicit control over this length separately from the overall bound t_{\max} on the edge lengths. This is because δ_1 and t_{\max} play different roles in the proof: t_{\max} serves to guarantee correctness of ancestral reconstruction on the subtrees hanging from the backbone; δ_1 on the other hand must be small enough to ensure that at most one mutation occurs along the edges of the backbone, but large enough to allow a union bound to be taken over the vertices of the backbone (see the proof of Proposition 3 in Section 5.2 where this argument is detailed and an explicit relationship between δ_1 and t_{\max} is derived).

- *Preprocessing: backbone sparsification.* We first construct a subtree T' by pruning some backbone vertices and their descendants. Initialize $T' := T$. Then, for $o = -, +$ and for $k = 1, \dots, B^o - 1$,

1. Check whether the vertex \tilde{x}_k^o is a vertex in the tree T' . If not, do nothing.
2. If \tilde{x}_k^o is in the tree T' , find the minimal $\ell \geq 1$ such that the distance (accounting for edge lengths) between \tilde{x}_k^o and $\tilde{x}_{k+\ell}^o$ is at least δ_1 . Observe that, by assumption, the distance between \tilde{x}_k^o and $\tilde{x}_{k+\ell}^o$ is necessarily at most $\delta_1 + t_{\max} \leq 2\delta_1$.
3. Remove the vertices $\tilde{x}_{k+1}^o, \dots, \tilde{x}_{k+\ell-1}^o$, except $\tilde{x}_{B^o}^o$, and all of their off-backbone descendants from the tree T' (if they exist).

The result is a tree where the distance between consecutive vertices on the backbone is in $[\delta_1, 2\delta_1]$ (by the observation in Item 2), with the possible exception of the children of the root and the last pair on each side of the root all of whose distances are in $(0, 2\delta_1]$. To simplify the notation, we reassign T to be this new rooted metric tree and we reassign x_1, x_2, \dots, x_B to be the backbone vertices on this tree (with an updated value for B and updated alternative numbering \tilde{x}_k^o for $o = -, +$ and $k = 1, \dots, B^o$). Because the initial tree is binary, it has at most $O(n)$ internal vertices. In particular the number of vertices on the path between v and w is $O(n)$. Hence the preprocessing step has $O(n)$ computational complexity.

We then proceed with the alignment algorithm, which consists of two main steps both proceeding along the backbone:

1. *Ancestral estimation:* We infer the ancestral sequences at the backbone vertices as follows. For $k = 2, \dots, B - 1$:

(a) For the child vertex z_k of x_k that is off the backbone, infer the sequence $\hat{\sigma}_{z_k}$ at z_k using the Fitch method [15] (described below in Section 3.2) applied to the subtree rooted at z_k .
 (b) Set $\hat{\sigma}_{x_k}$ equal to $\hat{\sigma}_{z_k}$.

2. *Recursive alignment:* Now that the sequences at the nonroot backbone vertices $\{x_k\}_{k=1}^B$ have been estimated, we construct a multiple sequence alignment sequentially, starting from x_1 , going to x_2 , and ending at x_{B-1} and x_B . This stepwise alignment procedure is described in Section 3.3 below. If the inferred sequences of successive backbone vertices are not at most one mutation apart, then we terminate the algorithm with no output. Else, a pairwise sequence alignment is produced for vertices $v = x_1$ and $w = x_B$.

We will show in Proposition 2 below that, with high probability, $\hat{\sigma}_{z_k} = \sigma_{z_k}$ for all $k = 2, \dots, B-1$. We will then show in Proposition 3 below that the above stepwise alignment outputs a true pairwise alignment with high probability.

3.2. *Ancestral sequence reconstruction.* We briefly describe below the ancestral sequence reconstruction subroutine. Note that we use the Fitch method for the convenience of its analysis, but other methods could also be used. Though we focus on the $\{0, 1\}$ alphabet, this method can be performed on sequences over any alphabet.

DEFINITION 3 (Fitch estimator). Let $T = (V, E)$ be a finite binary rooted tree with root z and leaf set $\partial T \subset V$ with given leaf sequences $(\sigma_\ell)_{\ell \in \partial T}$. For any leaf vertex ℓ , define $\hat{S}_\ell \subset \mathcal{S}$ to be the subset $\hat{S}_\ell = \{\sigma_\ell\}$. For each nonleaf vertex v with children v_1 and v_2 , define $\hat{S}_v \subset \mathcal{S}$ recursively to be

$$\hat{S}_v = \begin{cases} \hat{S}_{v_1} \cap \hat{S}_{v_2} & \text{if } \hat{S}_{v_1} \cap \hat{S}_{v_2} \neq \emptyset, \\ \hat{S}_{v_1} \cup \hat{S}_{v_2} & \text{otherwise.} \end{cases}$$

Then define the *Fitch estimator* $\hat{\sigma}_z$ of σ_z to be a uniformly chosen member of \hat{S}_z .

An analysis of this method in our setting is provided in Proposition 2 below.

Let ℓ be the maximum sequence length at the leaves. Identifying all unique sequences at the leaves, of which there are at most n , has $O(n^2\ell)$ computational complexity. For each backbone vertex, the ancestral reconstruction step then has $O(n^2)$ computational complexity. The number of leaves n is at least exponential in $1/t_{\max} \geq 1/\delta_1$. Hence the number of backbone vertices is $O(\log n)$. Hence the ancestral reconstruction step can be accomplished in time $O(n^2\ell + n^2 \log n)$.

3.3. *Stepwise alignment.* In this section, we describe the stepwise alignment subroutine. It is based on the assumption that along the backbone (of the pruned tree):

- (i) the sequences have been correctly inferred; and
- (ii) consecutive sequences differ by at most one mutation.

We establish these facts in Propositions 2 and 3 below. In these circumstances, we show that homologous sites can be traced (up to the convention we described earlier). We will construct a sequence of alignments $\mathbf{a}^2, \mathbf{a}^3, \dots$. We first describe the alignment of two sequences, then the alignment of alignments, and so on.

Given two sequences $\hat{\sigma}, \hat{\tau}$ satisfying the assumptions (i) and (ii) above, we construct an alignment $\mathbf{a}^2(\hat{\sigma}, \hat{\tau})$. For $\ell \in \{1, 2\}$, we let $a_\ell^2(\hat{\sigma}, \hat{\tau})$ denote the ℓ th sequence (or ℓ th row) of the alignment constructed from the sequences $\hat{\sigma}, \hat{\tau}$. There are three possible cases:

(A) If $\hat{\sigma} = \hat{\tau}$, then a true alignment is obtained by setting $a_1^2(\hat{\sigma}, \hat{\tau}) = \hat{\sigma}$ and $a_2^2(\hat{\sigma}, \hat{\tau}) = \hat{\tau}$, corresponding to no mutation.

(B) If $|\hat{\sigma}| = |\hat{\tau}|$ but $\hat{\sigma}$ and $\hat{\tau}$ agree on all sites except one, then a true alignment is obtained by setting $a_1^2(\hat{\sigma}, \hat{\tau}) = \hat{\sigma}$ and $a_2^2(\hat{\sigma}, \hat{\tau}) = \hat{\tau}$, corresponding to exactly one substitution between the sequences.

(C) If $|\hat{\tau}| = |\hat{\sigma}| + 1$ and there exists $j \in \{1, 2, \dots, |\hat{\sigma}|\}$ and $\hat{\tau}_{\text{ins}} \in \{0, 1\}$ such that

$$\hat{\tau}_i = \begin{cases} \hat{\sigma}_i & i < j, \\ \hat{\tau}_{\text{ins}} & i = j, \\ \hat{\sigma}_{i-1} & i > j. \end{cases}$$

As we discussed before, the location of the indel cannot be determined from the sequences alone. For example, if $\hat{\sigma}$ and $\hat{\tau}$ are separated by an indel so that they are given by

$$\hat{\sigma} = (0, 1, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0),$$

$$\hat{\tau} = (0, 1, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0),$$

we cannot tell which site gave birth to the new 0 to obtain $\hat{\tau}$ (assuming that the evolutionary process transformed $\hat{\sigma}$ into $\hat{\tau}$). So we assume by convention that j is the minimal choice possible. Then a true alignment is obtained by setting $a_2^2(\hat{\sigma}, \hat{\tau}) = \hat{\tau}$ and for $i = 1, \dots, |\hat{\sigma}| + 1$

$$a_1^2(\hat{\sigma}, \hat{\tau})_i = \begin{cases} \hat{\sigma}_i & i < j, \\ - & i = j, \\ \hat{\sigma}_{i-1} & i > j, \end{cases}$$

corresponding to a single site $\hat{\tau}_{\text{ins}}$ being inserted into the sequence $\hat{\tau}$ to the left of the j th site to obtain $\hat{\sigma}$. Similarly, if instead $|\hat{\sigma}| = |\hat{\tau}| + 1$ (in which case a deletion has occurred), we interchange the roles of $\hat{\sigma}$ and $\hat{\tau}$ and use the same convention.

In fact, we will need to align alignments along the backbone, rather than sequences. Suppose we have sequences $\hat{\sigma}_{x_1}, \hat{\sigma}_{x_2}, \dots, \hat{\sigma}_{x_B}$ and successive pairs

$$\{\hat{\sigma}_{x_1}, \hat{\sigma}_{x_2}\}, \{\hat{\sigma}_{x_2}, \hat{\sigma}_{x_3}\}, \dots, \{\hat{\sigma}_{x_{B-1}}, \hat{\sigma}_{x_B}\}$$

each satisfy exactly one of the cases (A), (B), or (C). (We terminate without output if the assumptions do not hold.) Then we recursively construct a multiple sequence alignment as follows. To simplify the notation, we let $\hat{\sigma}_{1:k} = (\hat{\sigma}_{x_1}, \dots, \hat{\sigma}_{x_k})$. For any $\ell \leq k$, we let $a_\ell^k(\hat{\sigma}_{1:k})$ denote the ℓ th sequence (or ℓ th row) of the alignment constructed from the given sequences $\hat{\sigma}_{1:k}$.

- Given $\hat{\sigma}_{x_1}$ and $\hat{\sigma}_{x_2}$, let $a_1^2(\hat{\sigma}_{1:2})$ and $a_2^2(\hat{\sigma}_{1:2})$ be the pairwise alignment constructed above.
- For $k = 3, \dots, B$:

(a) We are given a multiple alignment $a_1^{k-1}(\hat{\sigma}_{1:k-1}), \dots, a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})$ of the sequences $\hat{\sigma}_{x_1}, \dots, \hat{\sigma}_{x_{k-1}}$, and a new sequence $\hat{\sigma}_{x_k}$ that is at most one mutation away from $\hat{\sigma}_{x_{k-1}}$.

(b) The sequences $\hat{\sigma}_{x_{k-1}}$ and $\hat{\sigma}_{x_k}$ satisfy one of the three cases (A), (B), or (C) by assumption, so their alignment $a_{k-1}^k(\hat{\sigma}_{1:k})$ and $a_k^k(\hat{\sigma}_{1:k})$ (within the larger multiple sequence alignment) will differ by at most one entry similar to the sequence case above. To describe the alignment, it will be convenient to imagine that the tree is rooted at x_1 , and that the evolutionary process transforms $\hat{\sigma}_{x_1}$ into $\hat{\sigma}_{x_2}$, and so on up to $\hat{\sigma}_{x_B}$. Indeed, observe that the direction of time simply turns insertions into deletions and vice versa, and that it plays no role in the alignment procedure. The full alignment is defined as follows:

- If $\hat{\sigma}_{x_{k-1}} = \hat{\sigma}_{x_k}$, then set $a_k^k(\hat{\sigma}_{1:k})$ to be equal to $a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})$ and $a_i^k(\hat{\sigma}_{1:k})$ to be equal to $a_i^{k-1}(\hat{\sigma}_{1:k-1})$ for all $i < k$.

- If $\hat{\sigma}_{x_{k-1}}$ and $\hat{\sigma}_{x_k}$ have equal length and disagree at a single segregating site, set $a_i^k(\hat{\sigma}_{1:k})$ to $a_i^{k-1}(\hat{\sigma}_{1:k})$ for all $i \leq k-1$. Each entry of $a_k^k(\hat{\sigma}_{1:k})$ is set to the corresponding entry of $a_{k-1}^k(\hat{\sigma}_{1:k})$, except for the segregating site. If the latter occurs at position i within $a_{k-1}^k(\hat{\sigma}_{1:k})$, then we set $a_k^k(\hat{\sigma}_{1:k})_i$ to $a_{k-1}^k(\hat{\sigma}_{1:k})_i + 1 \pmod{2}$.
- If $\hat{\sigma}_{x_k}$ has one more site than $\hat{\sigma}_{x_{k-1}}$, then an insertion has occurred and the inserted site in $\hat{\sigma}_{x_k}$ cannot be ancestral to any site in $\hat{\sigma}_{x_1}, \dots, \hat{\sigma}_{x_{k-1}}$. This is because the insertion corresponds to the creation of a new column in the alignment. For the same reason, the inserted site in $\hat{\sigma}_{x_k}$ must correspond to a gap in all previous sequences. More specifically, if the site $\hat{\sigma}_{\text{ins}}$ is inserted to the left of position $j^* \in \{1, \dots, |a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})|\}$ in the $(k-1)$ st sequence in the previously constructed alignment $a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})$ (where j^* is the minimal such choice) then set

$$a_k^k(\hat{\sigma}_{1:k})_i = \begin{cases} a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})_i & 1 \leq i < j^*, \\ \hat{\sigma}_{\text{ins}} & i = j^*, \\ a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})_{i-1} & j^* < i \leq |a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})| + 1 \end{cases}$$

and for all $\ell \leq k-1$

$$a_\ell^k(\hat{\sigma}_{1:k})_i = \begin{cases} a_\ell^{k-1}(\hat{\sigma}_{1:k-1})_i & 1 \leq i < j^*, \\ - & i = j^*, \\ a_\ell^{k-1}(\hat{\sigma}_{1:k-1})_{i-1} & j^* < i \leq |a_{k-1}^{k-1}(\hat{\sigma}_{1:k-1})| + 1. \end{cases}$$

- The case where $\hat{\sigma}_{x_k}$ has one fewer site than $\hat{\sigma}_{x_{k-1}}$ is handled similarly. This time we include a gap in the k th sequence of the alignment, while all other rows remain unchanged from the previous multiple alignment.

3. Output the pairwise alignment $(a_1^B(\hat{\sigma}_{1:B}), a_B^B(\hat{\sigma}_{1:B}))$ after removing all columns with only gaps.

We provide an example of repeated use of steps 1 and 2 in the stepwise alignment in Figure 4.

Because each loop through the alignment algorithm, of which there are at most $O(\log n)$, introduces at most one gap in the current alignment of alignments, the overall computational complexity of this phase is $O((\log n \vee \ell) \log^2 n)$.

3.4. Theoretical guarantee. We establish the two claims below in the next sections.

PROPOSITION 1 (Correctness of alignment). *Let T be the output of the preprocessing step and let x_1, \dots, x_B be the resulting backbone vertices. Then the alignment algorithm produces a true pairwise alignment of σ_{x_1} and σ_{x_B} provided that:*

1. (Correctness of ancestral estimation) For $k = 2, \dots, B-1$, $\hat{\sigma}_{x_k} = \sigma_{x_k}$.
2. (One-mutation condition) Successive pairs of true backbone sequences

$$\{\sigma_{x_1}, \sigma_{x_2}\}, \{\sigma_{x_2}, \sigma_{x_3}\}, \dots, \{\sigma_{x_{B-1}}, \sigma_{x_B}\},$$

are at most one mutation away.

4. Correctness of ancestral estimation. In this section, we analyze the ancestral sequence estimation step. The analysis proceeds by coupling the TKF91 process with a percolation process. Roughly, we say that an edge is *open* if the sequence does not change along it under the TKF91 process. We will show that, provided the edge lengths are short enough, the open cluster of the root forms a fairly “dense” subtree with high probability. The latter property will lead to a correct reconstruction by the Fitch method.

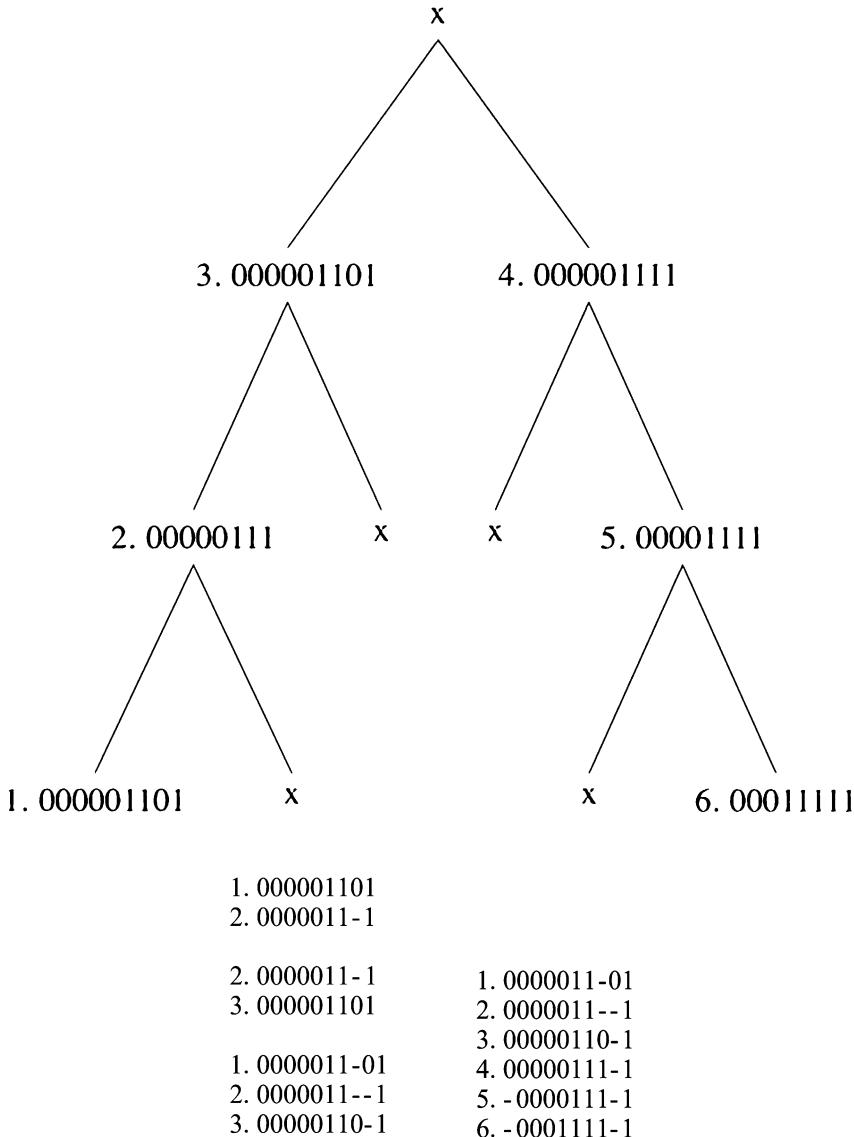


FIG. 4. Top: A rooted binary tree whose backbone graph-distance neighbors are all exactly one mutation apart. Bottom Left: The alignment of sequences 1,2,3 along the backbone. In the first step, a deletion is observed going from sequence 1 into sequence 2. So in the alignment algorithm, a gap in sequence 2 is written to correspond to the 0 in site 8 of sequence 1. In the second step, an insertion is observed going from sequence 2 into sequence 3, so an additional gap must be written in sequences 1 and 2. Bottom right: The alignment of all six backbone sequences.

4.1. *The percolation process.* Consider again the backbone vertices $\{x_k\}_{k=1}^B$ and the off-backbone child vertices $\{z_k\}_{k=2}^{B-1}$. For each $k = 2, \dots, B-1$ separately, we couple the sequence evolution process on the subtree descending from z_k with a simpler percolation process.

We will need some notation. We denote by T_k the subtree of T (after preprocessing) rooted at z_k . For two sequences σ, τ , we let $P_t(\sigma, \tau)$ be the probability under the TKF91 model on an edge that, started at σ , the state is τ after time t . Similarly, we let $\tilde{P}_t(\sigma, \tau)$ be the same probability *conditioned on not being at state σ at time t* .

It will be convenient to work on an infinite tree. Specifically, let \overline{T}_k be the completion of T_k into an infinite binary tree where new edges have length 0. We now describe the coupling:

- *The percolation process on \overline{T}_k .* We condition on the sequence σ_{z_k} at z_k , the root of T_k . For an edge e on \overline{T}_k , let t_e be its length and set

$$\zeta_e^{(k)} = P_{t_e}(\sigma_{z_k}, \sigma_{z_k}),$$

that is, the probability that the sequence does not change along edge e if started at σ_{z_k} .

We then perform percolation on \overline{T}_k with probabilities $\zeta_e^{(k)}$: for each edge e , it is open independently with probability $\zeta_e^{(k)}$. Let \mathcal{C}_k be the resulting open cluster including z_k (i.e., all vertices of \overline{T}_k that can be reached from z_k using only open edges).

- *The joint process on \overline{T}_k .* For each vertex v in \mathcal{C}_k , set $\sigma_v = \sigma_{z_k}$. For each descendant w of a vertex $v \in \mathcal{C}_k$ that is not itself in \mathcal{C}_k , assign a sequence to w taken from the conditional distribution $P_{t_e}(\sigma_{z_k}, \cdot)$ where $e = (v, w)$. For each remaining vertex, we run the TKF91 process recursively from the states already assigned. Note that the edges of length 0 added in the completion of T_k simply entail copying the sequences at the leaves of T_k to all their descendants.

We will be interested in the properties of the cluster \mathcal{C}_k . Let $\overline{T}_k^\mathcal{O}$ be the subtree of \overline{T}_k made of all the vertices in \mathcal{C}_k and the edges connecting them. For any $\ell, b \in \mathbb{Z}_+$, a rooted tree T' is said [28] to be a (ℓ, b) -diluted tree if: for all $i \in \mathbb{Z}_+$, each of the vertices of T' at graph distance $i\ell$ from the root has at least b descendants at graph distance $i\ell + \ell$ from the root. From the following lemma adapted from [28], $\overline{T}_k^\mathcal{O}$ is $(2, 3)$ -diluted with arbitrarily high probability provided edge lengths are short enough. Figure 5 depicts an infinite 2-diluted 3-regular tree. Later in the proof, we will need to condition on the length of σ_{z_k} being less than a threshold \bar{L} . Let $\mathbf{P}^{\bar{L}}$ be the probability measure of the joint process where σ_{z_k} is drawn from the stationary distribution of the TKF91 process conditioned on $|\sigma_{z_k}| \leq \bar{L}$. We first record a simple observation.

LEMMA 1 (Staying probability). *Fix $\bar{L} < +\infty$. For any sequence σ such that $|\sigma| \leq \bar{L}$ and any $t > 0$, we have*

$$P_t(\sigma, \sigma) \geq 1 - t(\bar{L} + 1)(\mu + \lambda + \eta).$$

PROOF. Indeed $P_t(\sigma, \sigma)$ is lower bounded by the probability that no mutation occurs up to time t which, by (2), is at least

$$P_t(\sigma, \sigma) \geq \exp(-(\bar{L} + 1)[\mu + \lambda + \eta]t) \geq 1 - (\bar{L} + 1)[\mu + \lambda + \eta]t,$$

as claimed. \square

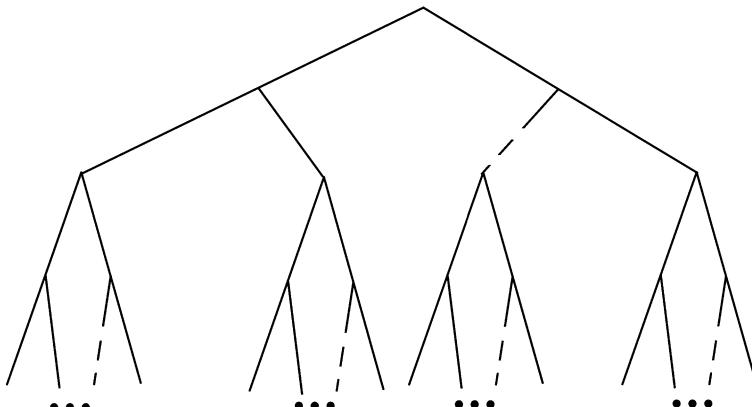


FIG. 5. *An open 2-diluted 3-regular subtree of the infinite binary rooted tree. Solid lines not descending from any dashed line indicate no mutation. Dashed lines indicate a mutation may have occurred. For every vertex in the even generations, at least three of its grandchildren share the same trait.*

LEMMA 2 (Existence of an open diluted tree). *For any $\bar{L} \in \mathbb{Z}_+$ and $\delta_a > 0$, there is $t_{\max} > 0$ small enough that, if $t_e \leq t_{\max}$ for all e ,*

$$\mathbf{P}^{\bar{L}}[\bar{T}_k^{\mathcal{O}} \text{ is } (2, 3)\text{-diluted}] \geq 1 - \delta_a.$$

PROOF. By Lemma 1, for any sequence σ_{z_k} such that $|\sigma_{z_k}| \leq \bar{L}$, we have that

$$\zeta_e^{(k)} = P_{t_e}(\sigma_{z_k}, \sigma_{z_k}) \geq 1 - (\bar{L} + 1)[\mu + \lambda + \eta]t_{\max},$$

where we recall that $t_e \leq t_{\max}$ by assumption (and that of course includes the added edges of length 0). Hence $\zeta_e^{(k)}$ can be made arbitrarily close to 1 (uniformly in e) by taking t_{\max} small enough (as a function of \bar{L}). The result then follows directly from [28], Lemma 8, (which can be extended in a straightforward manner to the case where percolation probabilities vary across edges but are uniformly bounded). \square

4.2. *Analyzing the Fitch estimator.* Next, we analyze the Fitch estimator in the event that \bar{T}_k contains an open $(2, 3)$ -diluted subtree.

For any $D \in \mathbb{Z}_+$, let $\bar{T}_{k,D}$ be the truncation of \bar{T}_k at level D , that is, the finite tree obtained by removing all vertices of \bar{T}_k at graph distance greater than D from its root. Let β_k be the smallest positive integer such that $\bar{\bar{T}}_k := \bar{T}_{k,2\beta_k}$ contains all of T_k . Importantly, we make the following observation about the Fitch estimator.

LEMMA 3 (Fitch estimator on the completion). *The Fitch estimator applied to the leaves of $\bar{\bar{T}}_k$ produces the same ancestral sequence estimate as the Fitch estimator applied to the leaves of T_k .*

PROOF. All leaves $\bar{\ell}$ of $\bar{\bar{T}}_k$ descending from a leaf ℓ of T_k satisfy $\sigma_{\bar{\ell}} = \sigma_{\ell}$, so by definition of the Fitch estimator $\hat{S}_{\ell} = \sigma_{\ell}$. The claim follows. \square

Let $\bar{\bar{T}}_k^{\mathcal{O}}$ be the truncation of $\bar{T}_k^{\mathcal{O}}$ at level $2\beta_k$.

LEMMA 4 (Fitch estimator in the presence of an open diluted tree). *If $\bar{\bar{T}}_k^{\mathcal{O}}$ is $(2, 3)$ -diluted, then the Fitch estimator $\hat{\sigma}_{z_k}$ over the tree $\bar{\bar{T}}_k$ equals the true sequence σ_{z_k} at z_k .*

PROOF. We prove this claim by induction on β_k . We start with the $\beta_k = 1$ case. Then $\bar{\bar{T}}_k$ consists of z_k , two children z_k^1 and z_k^2 , and the grandchildren $z_k^{1,1}, z_k^{1,2}, z_k^{2,1}, z_k^{2,2}$. If all four grandchildren belong to $\bar{\bar{T}}_k^{\mathcal{O}}$, then we are done. The other case, without loss of generality, is $z_k^{2,1} \notin \bar{\bar{T}}_k^{\mathcal{O}}$. Then $\sigma_{z_k^{2,1}} \neq \sigma_{z_k^{2,2}} = \sigma_{z_k}$, so the Fitch method gives $\hat{S}_{z_k^2} = \hat{S}_{z_k^{2,1}} \cup \hat{S}_{z_k^{2,2}} = \{\sigma_{z_k^{2,1}}, \sigma_{z_k}\}$. Since $\sigma_{z_k^{1,1}} = \sigma_{z_k^{1,2}} = \sigma_{z_k}$, we have $\hat{S}_{z_k^1} = \{\sigma_{z_k}\}$. Continuing on, we have $\hat{S}_{z_k} = \hat{S}_{z_k^1} \cap \hat{S}_{z_k^2} = \{\sigma_{z_k}\}$. Since \hat{S}_{z_k} contains only the state σ_{z_k} , the Fitch method is guaranteed to return σ_{z_k} .

Now, we assume the r th case holds for $r \geq 1$ and we show that the $(r+1)$ st case holds as well. As before, consider the four grandchildren of z_k and the same cases. If all four grandchildren belong to $\bar{\bar{T}}_k^{\mathcal{O}}$, then they are each the root of a subtree of $2r$ levels with root state equal to σ_{z_k} . The induction assumption implies that the Fitch method returns σ_{z_k} as estimates for $\sigma_{z_k^{i,j}}$, $i, j \in \{1, 2\}$. The Fitch method then returns $\hat{S}_{z_k} = \{\sigma_{z_k}\}$, as required. For

the other case when $z_k^{2,1} \notin \overline{\overline{T}}_k$, we know only that $\hat{S}_{z_k^{2,1}}$ is an arbitrary set of sequences. If $\sigma_{z_k} \in \hat{S}_{z_k^{2,1}}$, then $\hat{S}_{z_k^2} = \hat{S}_{z_k^{2,1}} \cap \hat{S}_{z_k^{2,2}} = \{\sigma_{z_k}\}$, and we are done. Else, we have $\hat{S}_{z_k^2} = \hat{S}_{z_k^{2,1}} \cup \hat{S}_{z_k^{2,2}}$, where $\sigma_{z_k} \in \hat{S}_{z_k^{2,2}}$ so that $\hat{S}_{z_k} = \hat{S}_{z_k^1} \cap \hat{S}_{z_k^2} = \{\sigma_{z_k}\}$, as required. This completes the proof for the $(r+1)$ st case, and hence of the lemma. \square

Combining Lemmas 2, 3, and 4, we get the following.

PROPOSITION 2 (Correctness of ancestral estimation off the backbone). *For any $\bar{L} \in \mathbb{Z}_+$ and $\delta_a > 0$, there is $t_{\max} > 0$ small enough that, under $\mathbf{P}^{\bar{L}}$, the Fitch estimator on T_k returns the correct ancestral state $\hat{\sigma}_{z_k} = \sigma_{z_k}$ with probability at least $1 - \delta_a$.*

5. One-mutation condition. In this section, we establish the one-mutation condition required by Proposition 1 and use it to finish the proof of the main result.

5.1. A bound on the transition probabilities. We will need a bound on the probability that at most one mutation occurs on an edge along the backbone. Because the state space of the sequence process is infinite, the rates are unbounded and we state the next bound explicitly in terms of the length of the sequence at the start of the edge. Later on, we will use the fact that the length is stationary to control it.

LEMMA 5 (At most one mutation). *Fix $\bar{L} < +\infty$. For any sequence σ such that $|\sigma| \leq \bar{L}$ and any $t > 0$, we have*

$$P_t(\sigma, Y_\sigma) \geq 1 - \{t(\bar{L} + 2)[\mu + \lambda + \eta]\}^2,$$

where $Y_\sigma = \{\sigma\} \cup \mathcal{S}_1(\sigma)$ are the sequences at most one mutation away from σ .

PROOF. For a TKF91 process on an edge started at σ , let $X_s \in \mathcal{S}$ be the sequence observed at time $s \in [0, t]$ and T_i be the time of the i th jump from one state to another state. Then

$$P_t(\sigma, Y_\sigma) \geq \mathbf{P}_\sigma[T_2 > t],$$

as the event on the right-hand side guarantees a single jump, which in turn guarantees that $X_t \in Y_\sigma$. Here \mathbf{P}_σ indicates that the edge process is started at σ . Letting

$$f_{T_1|\sigma}(s) = \lambda^*(\sigma) \exp(-s\lambda^*(\sigma)), \quad F_{T_1|\sigma}(s) = 1 - \exp(-s\lambda^*(\sigma)),$$

be the probability density function and cumulative distribution function of the time of the first jump started at σ , we get by the strong Markov property

$$\begin{aligned} \mathbf{P}_\sigma[T_2 \leq t] &= \int_0^t f_{T_1|\sigma}(s) \sum_{\tau \in \mathcal{S}_1(\sigma)} \frac{Q(\sigma, \tau)}{\lambda^*(\sigma)} F_{T_1|\tau}(t-s) \, ds \\ &\leq \int_0^t \lambda^*(\sigma) \exp(-s\lambda^*(\sigma)) \max_{\tau \in \mathcal{S}_1(\sigma)} \{1 - \exp(-(t-s)\lambda^*(\tau))\} \, ds. \end{aligned}$$

Under the assumption that $|\sigma| \leq \bar{L}$, it holds that $|\tau| \leq \bar{L} + 1$ for any $\tau \in \mathcal{S}_1(\sigma)$, and hence $\max\{\lambda^*(\sigma), \lambda^*(\tau)\} \leq (\bar{L} + 2)[\mu + \lambda + \eta]$. Continuing on, the last line in the previous display is

$$\begin{aligned} &\leq \{1 - \exp(-t(\bar{L} + 2)[\mu + \lambda + \eta])\} \int_0^t \lambda^*(\sigma) \exp(-s\lambda^*(\sigma)) \, ds \\ &= \{1 - \exp(-t(\bar{L} + 2)[\mu + \lambda + \eta])\} \{1 - \exp(-t\lambda^*(\sigma))\} \\ &\leq \{t(\bar{L} + 2)[\mu + \lambda + \eta]\}^2, \end{aligned}$$

establishing the claim. \square

5.2. *Union bound over the backbone.* We define a number of events whose joint occurrence guarantees the success of our alignment procedure:

- (*One-mutation condition*) For $o = -, +$ and $k = 1, \dots, B^o - 1$, let F_k^o be the event that the sequences at \tilde{x}_k^o and the backbone child vertex of \tilde{x}_k^o (i.e., \tilde{x}_{k+1}^o) satisfy constraints (A), (B), or (C) from Section 3.3.
- (*Ancestral reconstruction*) For $o = -, +$ and $k = 1, \dots, B^o - 1$, let G_k^o be the event that there is no mutation between the sequences at \tilde{x}_k^o and its off-backbone child vertex \tilde{z}_k^o and that $\sigma_{\tilde{z}_k^o}$ is correctly reconstructed by applying the Fitch method on the subtree rooted at \tilde{z}_k^o .
- (*Root segment*) For $o = -, +$, let H^o be the event that the sequences at the root and at \tilde{x}_1^o are identical.

The following proposition provides a requirement on the maximum branch length t_{\max} for all the above events to occur simultaneously. Define the bad event

$$\mathcal{B} = (H^-)^c \cup (H^+)^c \cup \left\{ \bigcup_{o=-,+} \bigcup_{k=1}^{B^o-1} (F_k^o)^c \cup (G_k^o)^c \right\}.$$

Recall that the preprocessing procedure has a parameter δ_1 .

PROPOSITION 3 (Union bound over the backbone). *Fix a tree height $h > 0$. For any $0 < \delta_1 < h$, there is a t_{\max} small enough that*

$$\mathbf{P}[\mathcal{B}] \leq Ch\delta_1 \log^2(\delta_1^{-1}),$$

where C is a constant depending only on λ, μ, η .

PROOF. We take a union bound over the events making up \mathcal{B} .

Controlling the lengths. For each event, we first apply the law of total probability to control for the length of the starting sequence as follows. Suppose that sequence τ is stationary, which we denote by $\tau \sim \Pi$. Using the stationary distribution for the length (i.e., (1)), we have

$$\mathbf{P}_{\tau \sim \Pi}[|\tau| > \bar{L}] = \sum_{M=\bar{L}+1}^{\infty} \left(1 - \frac{\lambda}{\mu}\right) \left(\frac{\lambda}{\mu}\right)^M = \left(\frac{\lambda}{\mu}\right)^{\bar{L}+1}.$$

The expression on the right is made less than δ_1^2 by choosing

$$(3) \quad \bar{L} = \left\lceil \frac{\log(\delta_1^2)}{\log(\lambda/\mu)} \right\rceil \leq C' \log(\delta_1^{-1}),$$

for a constant $C' > 0$ depending only on μ, λ , where recall that $\mu > \lambda$. Then for any event \mathcal{E} which depends on τ , we can write

$$(4) \quad \begin{aligned} \mathbf{P}[\mathcal{E}] &= \mathbf{P}[\mathcal{E} \mid |\tau| \leq \bar{L}] \mathbf{P}[|\tau| \leq \bar{L}] + \mathbf{P}[\mathcal{E} \mid |\tau| > \bar{L}] \mathbf{P}[|\tau| > \bar{L}] \\ &\leq \mathbf{P}[\mathcal{E} \mid |\tau| \leq \bar{L}] + \mathbf{P}[|\tau| > \bar{L}] \\ &\leq \mathbf{P}[\mathcal{E} \mid |\tau| \leq \bar{L}] + \delta_1^2, \end{aligned}$$

for the choice of \bar{L} above.

Events H^o . For $o = -, +$, we use Lemma 1 to bound the probability of $(H^o)^c$. By construction, \tilde{x}_1^o is a child of the root, so the edge length between the root and \tilde{x}_1^o is at most t_{\max} . Hence, using Lemma 1 and (4) with $\tau := \sigma_\rho$ and $\mathcal{E} := (H^o)^c$, we get

$$(5) \quad \mathbf{P}[(H^o)^c] \leq t_{\max}(\bar{L} + 1)(\mu + \lambda + \eta) + \delta_1^2.$$

Events G_k^o . For $o = -, +$ and $k = 1, \dots, B^o - 1$, we use Lemma 1 together with Proposition 2 to bound the probability of $(G_k^o)^c$. Here we take $\tau := \sigma_{\tilde{x}_k^o}$ and $\mathcal{E} := (G^o)^c$. By assumption, the edge length between \tilde{x}_k^o and its off-backbone child \tilde{z}_k^o is at most t_{\max} . Further, for any fixed failure probability $\delta_a > 0$ and length threshold \bar{L} , the maximum branch length t_{\max} can be taken small enough for Proposition 2 to hold. By (4), we get

$$\begin{aligned} \mathbf{P}[(G_k^o)^c] &\leq \mathbf{P}[(G_k^o)^c \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] + \delta_1^2 \\ &\leq \mathbf{P}[\{\sigma_{\tilde{x}_k^o} \neq \sigma_{\tilde{z}_k^o}\} \cup \{\{\sigma_{\tilde{x}_k^o} = \sigma_{\tilde{z}_k^o}\} \cap \{\hat{\sigma}_{\tilde{z}_k^o} \neq \sigma_{\tilde{z}_k^o}\}\} \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] + \delta_1^2 \\ &\leq \mathbf{P}[\{\sigma_{\tilde{x}_k^o} \neq \sigma_{\tilde{z}_k^o}\} \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] \\ &\quad + \mathbf{P}[\{\{\sigma_{\tilde{x}_k^o} = \sigma_{\tilde{z}_k^o}\} \cap \{\hat{\sigma}_{\tilde{z}_k^o} \neq \sigma_{\tilde{z}_k^o}\}\} \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] + \delta_1^2. \end{aligned}$$

We use the Markov property to bound the second term as follows:

$$\begin{aligned} \mathbf{P}[\{\{\sigma_{\tilde{x}_k^o} = \sigma_{\tilde{z}_k^o}\} \cap \{\hat{\sigma}_{\tilde{z}_k^o} \neq \sigma_{\tilde{z}_k^o}\}\} \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] \\ &= \mathbf{P}[\sigma_{\tilde{x}_k^o} = \sigma_{\tilde{z}_k^o} \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] \mathbf{P}[\hat{\sigma}_{\tilde{z}_k^o} \neq \sigma_{\tilde{z}_k^o} \mid \sigma_{\tilde{x}_k^o} = \sigma_{\tilde{z}_k^o}, \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] \\ &= \mathbf{P}[\sigma_{\tilde{x}_k^o} = \sigma_{\tilde{z}_k^o} \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] \mathbf{P}[\hat{\sigma}_{\tilde{z}_k^o} \neq \sigma_{\tilde{z}_k^o} \mid \mid \sigma_{\tilde{z}_k^o} \mid \leq \bar{L}] \\ &\leq \mathbf{P}[\hat{\sigma}_{\tilde{z}_k^o} \neq \sigma_{\tilde{z}_k^o} \mid \mid \sigma_{\tilde{z}_k^o} \mid \leq \bar{L}]. \end{aligned}$$

Plugging this back above and using Lemma 1 and Proposition 2 gives

$$(6) \quad \mathbf{P}[(G_k^o)^c] \leq t_{\max}(\bar{L} + 1)(\mu + \lambda + \eta) + \delta_a + \delta_1^2.$$

Events F_k^o . For $o = -, +$ and $k = 1, \dots, B^o - 1$, we use Lemma 5 to bound the probability of $(F_k^o)^c$. Here we take $\tau := \sigma_{\tilde{x}_k^o}$ and $\mathcal{E} := (F_k^o)^c$. By construction (i.e., by the backbone sparsification preprocessing step), the edge length between \tilde{x}_k^o and its backbone child \tilde{x}_{k+1}^o is at most $2\delta_1$. By (4) and Lemma 5, we get

$$(7) \quad \begin{aligned} \mathbf{P}[(F_k^o)^c] &\leq \mathbf{P}[(F_k^o)^c \mid \mid \sigma_{\tilde{x}_k^o} \mid \leq \bar{L}] + \delta_1^2 \\ &\leq \{2\delta_1(\bar{L} + 2)[\mu + \lambda + \eta]\}^2 + \delta_1^2. \end{aligned}$$

Union bound. Taking a union bound over all events above gives

$$\begin{aligned} \mathbf{P}[\mathcal{B}] &\leq 2[t_{\max}(\bar{L} + 1)(\mu + \lambda + \eta) + \delta_1^2] \\ &\quad + \sum_{o=-,+} \sum_{k=1}^{B^o-1} [t_{\max}(\bar{L} + 1)(\mu + \lambda + \eta) + \delta_a + \delta_1^2] \\ &\quad + \sum_{o=-,+} \sum_{k=1}^{B^o-1} [\{2\delta_1(\bar{L} + 2)[\mu + \lambda + \eta]\}^2 + \delta_1^2] \end{aligned}$$

by (5), (6), and (7). We make all terms in square brackets of order $\delta_1^2 \log \delta_1^{-1}$ by choosing $\delta_a := \delta_1^2$ and then choosing $0 < t_{\max} \leq \delta_1^2$ small enough for Proposition 2 to hold. Then we get, using (3),

$$\begin{aligned} \mathbf{P}[\mathcal{B}] &\leq 2[\delta_1^2(C' \log(\delta_1^{-1}) + 1)(\mu + \lambda + \eta) + \delta_1^2] \\ &\quad + 2(B^o - 1)[\delta_1^2(C' \log(\delta_1^{-1}) + 1)(\mu + \lambda + \eta) + 2\delta_1^2] \\ &\quad + 2(B^o - 1)[\{2\delta_1(C' \log(\delta_1^{-1}) + 2)[\mu + \lambda + \eta]\}^2 + \delta_1^2]. \end{aligned}$$

Because the tree has height h and each backbone edge has length at least δ_1 (after preprocessing), with the exception of the first and last one on each side of the root, we must have

$(B^o - 2)\delta_1 \leq h$, or after rearranging $B^o \leq h/\delta_1 + 2$. Employing this bound and simplifying gives finally

$$\mathbf{P}[\mathcal{B}] \leq Ch\delta_1 \log^2(\delta_1^{-1}),$$

for a constant C depending only on μ, λ, η , as claimed. \square

5.3. *Proof of the theorem.* We are now ready to finish the proof of the main result.

PROOF OF THEOREM 1. For a fixed failure probability ε , we first choose δ_1 small enough (as a function of h, μ, λ, η) such that $Ch\delta_1 \log^2(\delta_1^{-1}) \leq \varepsilon$. We then choose t_{\max} small enough (again as a function of h, μ, λ, η) that Proposition 3 implies $\mathbf{P}[\mathcal{B}] \leq \varepsilon$. Proposition 1 then completes the proof of the theorem. \square

APPENDIX: FURTHER LEMMAS

We formally justify equation (2).

LEMMA 6 (Total rate). *The rate matrix Q satisfies*

$$\sum_{\tau \in \mathcal{S}_1(\sigma)} Q(\sigma, \tau) = \lambda(|\sigma| + 1) + \mu|\sigma| + \eta\pi_1|\sigma|_0 + \eta\pi_0|\sigma|_1.$$

PROOF. The transition rate matrix Q entries are as follows. For the transition from σ to τ , where τ differs from σ only by a single substitution, we have

$$Q(\sigma, \tau) = \eta\pi_i,$$

when label $i \in \{0, 1\}$ is the product of the single substitution. The sum over all τ in $\mathcal{S}_s(\sigma)$, the sequences that differ by σ by a single substitution, is

$$(8) \quad \sum_{\tau \in \mathcal{S}_s(\sigma)} Q(\sigma, \tau) = \eta\pi_1|\sigma|_0 + \eta\pi_0|\sigma|_1.$$

Next, for the transitions from σ to τ in $\mathcal{S}_d(\sigma)$, the sequences obtained by deleting a single digit of σ , we have

$$Q(\sigma, \tau) = s_{\sigma, \tau}^{(D)}\mu,$$

where $s_{\sigma, \tau}^{(D)} \geq 1$ is the length of the repeated segment of letters where the deletion occurs. Here, we observe that the position of the deletion within the segment is not needed. This is because the letters are identical within the segment, so the remaining letters are identical when one is deleted. For example, if a zero is deleted from a string of 10 zeros, then $s_{\sigma, \tau}^{(D)} = 10$. Since only one deletion out of $|\sigma|$ can occur, the sum is

$$(9) \quad \sum_{\tau \in \mathcal{S}_d(\sigma)} Q(\sigma, \tau) = |\sigma|\mu.$$

Finally, for the transitions from σ to τ in $\mathcal{S}_i(\sigma)$, the sequences obtained by inserting a digit into σ , we have

$$Q(\sigma, \tau) = (s_{\sigma, \tau}^{(I)} + 1)\lambda\pi_i, \quad i \in \{0, 1\},$$

where i is the digit inserted and $s_{\sigma, \tau}^{(I)} \geq 0$ is the size of the repeated segment of digit i where the site is being inserted. Noting that only one insertion can occur and the immortal link might insert a site, the sum is

$$(10) \quad \sum_{\tau \in \mathcal{S}_i(\sigma)} Q(\sigma, \tau) = (|\sigma| + 1)\lambda.$$

The claim follows from combining (8), (9), and (10). \square

Acknowledgments. SR is grateful to Alexandre Bouchard-Côté (UBC) for insightful discussions in the early stages of this project.

Funding. SR was supported by NSF Grants DMS-1614242, DMS-1902892, DMS-1916378 and DMS-2023239 (TRIPODS Phase II), as well as a Simons Fellowship and a Vilas Associates Award.

BL was supported by NSF Grants DMS-1614242, CCF-1740707 (TRIPODS), DMS-1902892 and a Vilas Associates Award (to SR). BL was also supported by NSF Grant DMS-1646108 (University of Michigan, Ann Arbor, Department of Statistics, RTG).

REFERENCES

- [1] ALLMAN, E. S., RHODES, J. A. and SULLIVANT, S. (2017). Statistically consistent k -mer methods for phylogenetic tree reconstruction. *J. Comput. Biol.* **24** 153–171. [MR3607847](#) <https://doi.org/10.1089/cmb.2015.0216>
- [2] ANDONI, A., DASKALAKIS, C., HASSIDIM, A. and ROCH, S. (2012). Global alignment of molecular sequences via ancestral state reconstruction. *Stochastic Process. Appl.* **122** 3852–3874. [MR2971717](#) <https://doi.org/10.1016/j.spa.2012.08.004>
- [3] CHAKRABORTY, D., DAS, D. and KRAUTHGAMER, R. (2021). Approximate trace reconstruction via median string (in average-case). In *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science* (M. and C. , eds.) *LIPICS. Leibniz Int. Proc. Inform.* **213** Art. No. 11, 23. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern. [MR4459293](#)
- [4] CHANG, G. S., HONG, Y., KO, K. D., BHARDWAJ, G., HOLMES, E. C., PATTERSON, R. L. and VAN ROSSUM, D. B. (2008). Phylogenetic profiles reveal evolutionary relationships within the “twilight zone” of sequence similarity. *Proc. Natl. Acad. Sci. USA* **105** 13474–13479. [https://doi.org/10.1073/pnas.0803860105](#)
- [5] CHASE, Z. and PERES, Y. Approximate trace reconstruction of random strings from a constant number of traces [math]. Available at [arXiv:2107.06454](#).
- [6] CHEN, X., DE, A., LEE, C. H., SERVEDIO, R. A. and SINHA, S. (2022). Near-optimal average-case approximate trace reconstruction from few traces. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* 779–821. SIAM, Philadelphia, PA. [MR4415071](#) <https://doi.org/10.1137/1.9781611977073.34>
- [7] CHEN, X., DE, A., LEE, C. H., SERVEDIO, R. A. and SINHA, S. (2023). Approximate trace reconstruction from a single trace. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* 605–637. SIAM, Philadelphia, PA. [MR4538159](#) <https://doi.org/10.1137/1.9781611977554.ch27>
- [8] DASKALAKIS, C. and ROCH, S. (2013). Alignment-free phylogenetic reconstruction: Sample complexity via a branching process analysis. *Ann. Appl. Probab.* **23** 693–721. [MR3059273](#) <https://doi.org/10.1214/12-AAP852>
- [9] DAVIES, S., RÁCZ, M. Z. and RASHTCHIAN, C. (2021). Reconstructing trees from traces. *Ann. Appl. Probab.* **31** 2772–2810. [MR4350973](#) <https://doi.org/10.1214/21-aap1662>
- [10] DAVIES, S., RÁCZ, M. Z., SCHIFFER, B. G. and RASHTCHIAN, C. (2021). Approximate trace reconstruction: Algorithms. *Abstr. Pap.—IEEE Int. Symp. Inf. Theory* 2525–2530. [https://doi.org/10.1109/ISIT45174.2021.9517926](#)
- [11] EVANS, W., KENYON, C., PERES, Y. and SCHULMAN, L. J. (2000). Broadcasting on trees and the Ising model. *Ann. Appl. Probab.* **10** 410–433. [MR1768240](#) <https://doi.org/10.1214/aoap/1019487349>
- [12] FAN, W.-T. and ROCH, S. (2020). Statistically consistent and computationally efficient inference of ancestral DNA sequences in the TKF91 model under dense taxon sampling. *Bull. Math. Biol.* **82** Paper No. 21, 32. [MR4056716](#) <https://doi.org/10.1007/s11538-020-00693-3>
- [13] FAN, W.-T. L., LEGRIED, B. and ROCH, S. (2020). Impossibility of consistent distance estimation from sequence lengths under the TKF91 model. *Bull. Math. Biol.* **82** Paper No. 123, 12. [MR4150869](#) <https://doi.org/10.1007/s11538-020-00801-3>
- [14] FISCHER, M. and THATTE, B. D. (2009). Maximum parsimony on subsets of taxa. *J. Theoret. Biol.* **260** 290–293. [MR2973084](#) <https://doi.org/10.1016/j.jtbi.2009.06.010>
- [15] FITCH, W. M. (1971). Toward defining the course of evolution: Minimum change for a specific tree topology. *Syst. Zool.* **20** 406–416.
- [16] GANESH, A. and ZHANG, Q. (2019). Optimal sequence length requirements for phylogenetic tree reconstruction with indels. In *STOC’19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* 721–732. ACM, New York. [MR4003378](#) <https://doi.org/10.1145/3313276.3316345>

- [17] GASCUEL, O. and STEEL, M. (2010). Inferring ancestral sequences in taxon-rich phylogenies. *Math. Biosci.* **227** 125–135. [MR2732449](https://doi.org/10.1016/j.mbs.2010.07.002) <https://doi.org/10.1016/j.mbs.2010.07.002>
- [18] HERBST, L. and FISCHER, M. (2017). Ancestral sequence reconstruction with maximum parsimony. *Bull. Math. Biol.* **79** 2865–2886. [MR3736430](https://doi.org/10.1007/s11538-017-0354-6) <https://doi.org/10.1007/s11538-017-0354-6>
- [19] HERBST, L. and FISCHER, M. (2018). On the accuracy of ancestral sequence reconstruction for ultrametric trees with parsimony. *Bull. Math. Biol.* **80** 864–879. [MR3778384](https://doi.org/10.1007/s11538-018-0407-5) <https://doi.org/10.1007/s11538-018-0407-5>
- [20] HOLDEN, N. and LYONS, R. (2020). Lower bounds for trace reconstruction. *Ann. Appl. Probab.* **30** 503–525. [MR4108114](https://doi.org/10.1214/19-AAP1506) <https://doi.org/10.1214/19-AAP1506>
- [21] HOLENSTEIN, T., MITZENMACHER, M., PANIGRAHY, R. and WIEDER, U. (2008). Trace reconstruction with constant deletion probability and related results. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms* 389–398. ACM, New York. [MR2487606](https://doi.org/10.1137/1.3053488)
- [22] IOFFE, D. (1996). On the extremality of the disordered state for the Ising model on the Bethe lattice. *Lett. Math. Phys.* **37** 137–143. [MR1391195](https://doi.org/10.1007/BF00416016) <https://doi.org/10.1007/BF00416016>
- [23] LI, G., STEEL, M. and ZHANG, L. (2008). More taxa are not necessarily better for the reconstruction of ancestral character states. *Syst. Biol.* **57** 647–653. <https://doi.org/10.1080/10635150802203898>
- [24] LIBERLES, D. A., ed. (2007). *Ancestral Sequence Reconstruction* Oxford Univ. Press, Oxford. <https://doi.org/10.1093/acprof:oso/9780199299188.001.0001>
- [25] LUNTER, G., DRUMMOND, A. J., MIKLÓS, I. and HEIN, J. (2005). Statistical alignment: Recent progress, new applications, and challenges. In *Statistical Methods in Molecular Evolution. Stat. Biol. Health* 375–405. Springer, New York. [MR2161838](https://doi.org/10.1007/0-387-27733-1_14) https://doi.org/10.1007/0-387-27733-1_14
- [26] MITROPHANOV, A. Y. and BORODOVSKY, M. (2007). Convergence rate estimation for the TKF91 model of biological sequence length evolution. *Math. Biosci.* **209** 470–485. [MR2359560](https://doi.org/10.1016/j.mbs.2007.02.011) <https://doi.org/10.1016/j.mbs.2007.02.011>
- [27] MOSSEL, E. (1998). Recursive reconstruction on periodic trees. *Random Structures Algorithms* **13** 81–97. [MR1634344](https://doi.org/10.1002/(SICI)1098-2418(199808)13:1<81::AID-RSA5>3.3.CO;2-K) [https://doi.org/10.1002/\(SICI\)1098-2418\(199808\)13:1<81::AID-RSA5>3.3.CO;2-K](https://doi.org/10.1002/(SICI)1098-2418(199808)13:1<81::AID-RSA5>3.3.CO;2-K)
- [28] MOSSEL, E. (2001). Reconstruction on trees: Beating the second eigenvalue. *Ann. Appl. Probab.* **11** 285–300. [MR1825467](https://doi.org/10.1214/aoap/998926994) <https://doi.org/10.1214/aoap/998926994>
- [29] MOSSEL, E. (2003). On the impossibility of reconstructing ancestral data and phylogenies. *J. Comput. Biol.* **10** 669–676. Publisher: Mary Ann Liebert, Inc., publishers. <https://doi.org/10.1089/106652703322539015>
- [30] MOSSEL, E. (2004). Phase transitions in phylogeny. *Trans. Amer. Math. Soc.* **356** 2379–2404. [MR2048522](https://doi.org/10.1090/S0002-9947-03-03382-8) <https://doi.org/10.1090/S0002-9947-03-03382-8>
- [31] NAZAROV, F. and PERES, Y. (2017). Trace reconstruction with $\exp(O(n^{1/3}))$ samples. In *STOC’17—Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing* 1042–1046. ACM, New York. [MR3678249](https://doi.org/10.1145/3055387.3055429)
- [32] RANWEZ, V. and CHANTRET, N. N. (2020). Strengths and limits of multiple sequence alignment and filtering methods. In *Phylogenetics in the Genomic Era* (C. Scornavacca, F. Delsuc and N. Galtier, eds.) 2.2:1–2.2:36. No commercial publisher. Authors open access book.
- [33] ROCH, S. (2010). Toward extracting all phylogenetic information from matrices of evolutionary distances. *Science* **327** 1376–1379. [MR2650508](https://doi.org/10.1126/science.1182300) <https://doi.org/10.1126/science.1182300>
- [34] ROCH, S. and SLY, A. (2017). Phase transition in the sample complexity of likelihood-based phylogeny inference. *Probab. Theory Related Fields* **169** 3–62. [MR3704765](https://doi.org/10.1007/s00440-017-0793-x) <https://doi.org/10.1007/s00440-017-0793-x>
- [35] ROCH, S. and WANG, K.-C. (2021). Sufficient condition for root reconstruction by parsimony on binary trees with general weights. *Electron. Commun. Probab.* **26** Paper No. 55, 13. [MR4319989](https://doi.org/10.1214/21-ecp423) <https://doi.org/10.1214/21-ecp423>
- [36] ROST, B. (1999). Twilight zone of protein sequence alignments. *Protein Eng. Des. Sel.* **12** 85–94. <https://doi.org/10.1093/protein/12.2.85>
- [37] SEMPLE, C. and STEEL, M. (2003). *Phylogenetics. Oxford Lecture Series in Mathematics and Its Applications* **24**. Oxford Univ. Press, Oxford. [MR2060009](https://doi.org/10.1093/0198509643)
- [38] STEEL, M. and CHARLESTON, M. (1995). Five surprising properties of parsimoniously colored trees. *Bull. Math. Biol.* **57** 367–375. <https://doi.org/10.1007/BF02460622>
- [39] THATTE, B. D. (2006). Invertibility of the TKF model of sequence evolution. *Math. Biosci.* **200** 58–75. [MR2211928](https://doi.org/10.1016/j.mbs.2005.12.025) <https://doi.org/10.1016/j.mbs.2005.12.025>
- [40] THORNE, J. L., KISHINO, H. and FELSENSTEIN, J. (1991). An evolutionary model for maximum likelihood alignment of DNA sequences. *J. Mol. Evol.* **33** 114–124. <https://doi.org/10.1007/BF02193625>

- [41] YANG, Z. (2014). *Molecular Evolution: A Statistical Approach*. Oxford Univ. Press, Oxford.
<https://doi.org/10.1093/acprof:oso/9780199602605.001.0001>
- [42] ZHANG, L., SHEN, J., YANG, J. and LI, G. (2010). Analyzing the Fitch method for reconstructing ancestral states of ultrametric phylogenetic trees. *Bull. Math. Biol.* **72** 1760–1782. MR2728004
<https://doi.org/10.1007/s11538-010-9505-8>