

MIT Open Access Articles

Evolutionary trees and the Ising model on the Bethe lattice: a proof of Steel's conjecture

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Daskalakis, Constantinos, Mossel, Elchanan and Roch, Sébastien. 2009. "Evolutionary trees and the Ising model on the Bethe lattice: a proof of Steel's conjecture."

As Published: <https://doi.org/10.1007/s00440-009-0246-2>

Publisher: Springer-Verlag

Persistent URL: <https://hdl.handle.net/1721.1/148126>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution



Evolutionary trees and the Ising model on the Bethe lattice: a proof of Steel's conjecture

Constantinos Daskalakis · Elchanan Mossel · Sébastien Roch

Received: 15 July 2008 / Revised: 14 August 2009 / Published online: 1 October 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract A major task of evolutionary biology is the reconstruction of phylogenetic trees from molecular data. The evolutionary model is given by a Markov chain on a tree. Given samples from the leaves of the Markov chain, the goal is to reconstruct the leaf-labelled tree. It is well known that in order to reconstruct a tree on n leaves, sample sequences of length $\Omega(\log n)$ are needed. It was conjectured by Steel that for the CFN/Ising evolutionary model, if the mutation probability on all edges of the tree is less than $p^* = (\sqrt{2} - 1)/2^{3/2}$, then the tree can be recovered from sequences of length $O(\log n)$. The value p^* is given by the transition point for the extremality of the free Gibbs measure for the Ising model on the binary tree. Steel's conjecture was proven by the second author in the special case where the tree is "balanced." The second author also proved that if all edges have mutation probability larger than p^* then

Partially supported by CIPRES (NSF ITR grant # NSF EF 03-31494). Supported by a Miller fellowship in Statistics and Computer Science, by a Sloan fellowship in Mathematics and by NSF grants DMS-0504245, DMS-0528488 DMS-0548249 (CAREER), ONR grant N0014-07-1-05-06 and ISF grant 1300/08.

C. Daskalakis
CSAIL, MIT, Cambridge, USA
e-mail: costis@csail.mit.edu

E. Mossel
Statistics and Computer Science, U.C. Berkeley, Berkeley, USA
e-mail: mossel@stat.berkeley.edu

E. Mossel
Weizmann Institute of Science, Rehovot, Israel

S. Roch (✉)
Department of Mathematics, UCLA, Los Angeles, USA
e-mail: roch@math.ucla.edu

the length needed is $n^{\Omega(1)}$. Here we show that Steel's conjecture holds true for general trees by giving a reconstruction algorithm that recovers the tree from $O(\log n)$ -length sequences when the mutation probabilities are discretized and less than p^* . Our proof and results demonstrate that extremality of the free Gibbs measure on the infinite binary tree, which has been studied before in probability, statistical physics and computer science, determines how distinguishable are Gibbs measures on finite binary trees.

Keywords Phylogenetics · CFN model · Ising model · Phase transitions · Reconstruction problem · Jukes–Cantor

Mathematics Subject Classification (2000) Primary 60K35 · 92D15; Secondary 60J85 · 82B26.

1 Introduction

In this paper we prove a central conjecture in mathematical phylogenetics [34]: We show that every phylogenetic tree with short, discretized edges on n leaves can be reconstructed from sequences of length $O(\log n)$, where by short we mean that the mutation probability on every edge is bounded above by the critical transition probability for the extremality of the Ising model on the infinite binary tree.

This result establishes that the extremality of the free Gibbs measure for the Ising model on the infinite binary tree, studied in probability, statistical physics and computer science, determines the sampling complexity of the phylogenetic problem, a central problem in evolutionary biology. We proceed with background on the phylogenetic problem, on the reconstruction problem and a statement of our results.

Phylogenetic background. Phylogenies are used in evolutionary biology to model the stochastic evolution of genetic data on the ancestral tree relating a group of species. The leaves of the tree correspond to (known) extant species. Internal nodes represent extinct species. In particular the root of the tree represents the most recent ancestor of all species in the tree. Following paths from the root to the leaves, each bifurcation indicates a speciation event whereby two new species are created from a parent. We refer the reader to [33] for an excellent introduction to phylogenetics. The underlying assumption is that genetic information evolves from the root to the leaves according to a Markov model on the tree. It is further assumed that this process is repeated independently a number of times denoted k . Thus each node of the tree is associated with a sequence of length k . The vector of the i th letter of all sequences at the leaves is called the i th *character*. One of the major tasks in molecular biology, the *reconstruction of phylogenetic trees*, is to infer the topology of the tree from the characters at the leaves.

In this paper we will be mostly interested in two evolutionary models, the so-called Cavender–Farris–Neyman (CFN) [5, 10, 28] and Jukes–Cantor (JC) [15] models. In the CFN model the states at the nodes of the tree are 0 and 1 and their a priori probability at the root is uniform. To each edge e corresponds a *mutation probability* $p(e)$

which is the probability that the state changes along the edge e . Note that this model is identical to the free Gibbs measure of the Ising model on the tree. See [18]. In the JC model the states are A , C , G and T with a priori probability $1/4$ each. To each edge e corresponds a mutation probability $p(e)$ and it is assumed that every state transitions with probability $p(e)$ to each of the other states. This model is equivalent to the ferromagnetic Potts model on the tree.

Extremality and the reconstruction problem. A problem that is closely related to the *phylogenetic problem* is that of inferring the *ancestral state*, that is, the state at the root of the tree, given the states at the leaves. This problem was studied earlier in statistical physics, probability and computer science under the name of *reconstruction problem*, or *extremality of the free Gibbs measure*. See [12, 13, 32]. The reconstruction problem for the CFN model was analyzed in [1–3, 8, 14, 27]. In particular, the role of the reconstruction problem in the analysis of the mixing time of Glauber dynamics on trees was established in [2, 27].

Roughly speaking, the reconstruction problem is *solvable* when the correlation between the root and the leaves persists no matter how large the tree is. When it is *unsolvable*, the correlation decays to 0 for large trees. The results of [1–3, 8, 14, 27] show that for the CFN model, if for all e it holds that $p(e) \leq p_{\max} < p^*$, then the reconstruction problem is solvable, where

$$p^* = \frac{\sqrt{2} - 1}{\sqrt{8}} \approx 15\%.$$

If, on the other hand, for all e it holds that $p(e) \geq p_{\min} > p^*$ and the tree is balanced in the sense that all leaves are at the same distance from the root, then the reconstruction problem is *unsolvable*. Moreover in this case, the correlation between the root state and any function of the character states at the leaves decays as $n^{-\Omega(1)}$.

Our results. Steel [34] conjectured that when $0 < p_{\min} \leq p(e) \leq p_{\max} < p^*$ for all edges e , one can reconstruct with high probability the phylogenetic tree from $O(\log n)$ characters. Steel's insightful conjecture suggests that there are deep connections between the reconstruction problem and phylogenetic reconstruction.

This conjecture has been proven to hold for trees where all the leaves are at the same graph distance from the root—the so-called “balanced” case—in [22]. It is also shown there that the number of characters needed when $p(e) \geq p_{\min} > p^*$, for all e , is $n^{\Omega(1)}$. The second result intuitively follows from the fact that the topology of the part of the tree that is close to the root is essentially independent of the characters at the leaves if the number of characters is not at least $n^{\Omega(1)}$.

The basic intuition behind Steel's conjecture is that: since in the regime where $p(e) \leq p_{\max} < p^*$ there is no decay of the quality of reconstructed sequences, it should be as easy to reconstruct deep trees as it is to reconstruct shallow trees. In [9] (see also [23]) it is shown that “shallow” trees can be reconstructed from $O(\log n)$ characters if all mutation probabilities are bounded away from 0 and $1/2$ (the results of [9] also show that in this regime, any tree can be recovered from sequences of polynomial length). The same high-level reasoning has also yielded a complete proof that $O(\log n)$ characters suffice for a percolation-type mutation model when all edges are short [25]. See [25] for details.

Here we prove Steel’s conjecture for general trees under the assumption that the mutation probabilities are discretized. We show that, if $0 < p_{\min} \leq p(e) \leq p_{\max} < p^*$ for all edges e of the tree, then the tree can be reconstructed from $c(p_{\min}, p_{\max})(\log n + \log 1/\delta)$ characters with error probability at most δ . The discretization assumption amounts to assuming that all edge lengths are a multiple of a small constant Δ . (See below for a formal definition of “edge length.”) This result further implies that sequences of logarithmic length suffice to reconstruct phylogenetic trees in the Jukes–Cantor model, when all the edges are sufficiently short.

Compared to [22], our main technical contribution is the design and analysis of a tree-building procedure that uses only “local information” to build parts of the tree, while also maintaining “disjointness” between the different reconstructed subtrees. The disjointness property is crucial in order to maintain the conditional independence properties of the Gibbs distribution on the original tree—for the purpose of performing estimation of the states at internal nodes of the tree. Note that in the balanced case of [22] this property can be achieved in a straightforward manner by building the tree “level-by-level.”

1.1 Formal definitions and main results

Trees and path metrics. Let T be a tree. Write $\mathcal{V}(T)$ for the nodes of T , $\mathcal{E}(T)$ for the edges of T and $\mathcal{L}(T)$ for the leaves of T . If the tree is rooted, then we denote by $\rho(T)$ the root of T . Unless stated otherwise, all trees are assumed to be *binary* (all internal degrees are 3) and it is further assumed that $\mathcal{L}(T)$ is labelled.

Let T be a tree equipped with a length function on its edges, $d : \mathcal{E}(T) \rightarrow \mathbb{R}_+$. The function d will also denote the induced path metric on $\mathcal{V}(T)$: $d(v, w) = \sum\{d(e) : e \in \text{path}_T(v, w)\}$, for all $v, w \in \mathcal{V}(T)$, where $\text{path}_T(x, y)$ is the path (sequence of edges) connecting x to y in T .

We will further assume below that the length of all edges is bounded between f and g for all $e \in \mathcal{E}(T)$. In other words, for all $e \in \mathcal{E}(T)$, $f \leq d(e) \leq g$.

Markov model of evolution. The evolutionary process is determined by a rooted tree $T = (V, E)$ equipped with a path metric d and a *rate matrix* Q . We will be mostly interested in the case where $Q = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$ corresponding to the CFN model and in the case where Q is a 4×4 matrix given by $Q_{i,j} = 1 - 4 \cdot \mathbb{1}\{i = j\}$ corresponding to the Jukes–Cantor model. To edge e of length $d(e)$ we associate the transition matrix $M^e = \exp(d(e)Q)$.

In the evolutionary model on the tree T rooted at ρ each vertex iteratively chooses its state from the state at its parent by an application of the Markov transition rule M^e , where e is the edge connecting it to its parent. We assume that all edges in E are directed away from the root. Thus the probability distribution on the tree is the probability distribution on $\{0, 1\}^V$ ($\{A, C, G, T\}^V$) given by

$$\bar{\mu}[\sigma] = \pi(\sigma(\rho)) \prod_{(x \rightarrow y) \in E} M_{\sigma(x), \sigma(y)}^{(x \rightarrow y)}$$

where π is given by the uniform distribution at the root, so that $\pi(0) = \pi(1) = 1/2$ for the CFN model and $\pi(A) = \pi(C) = \pi(G) = \pi(T) = 1/4$ for the JC model. We let the measure μ denote the marginal of $\bar{\mu}$ on the set of leaves which we identify with $[n] = \{1, \dots, n\}$. Thus

$$\mu(\sigma) = \sum \{\bar{\mu}(\tau) : \forall i \in [n], \tau(i) = \sigma(i)\}.$$

The measure μ defines the probability distribution at the leaves of the tree.

We note that both for the CFN model and for the JC model, the transition matrices M^e admit a simple alternative representation. For the CFN model, with probability $p(e) = (1 - \exp(-2d(e)))/2$, there is a transition and, otherwise, there is no transition. Similarly for the JC model with probability $p(e) = (1 - \exp(-4d(e)))/4$ each of the three possible transitions occur. In particular, defining

$$g^* = \frac{\ln 2}{4}, \tag{1}$$

we may formulate the result on the reconstruction problem for the phase transition of the CFN model as follows: “If $d(e) \leq g < g^*$ for all e then the reconstruction problem is solvable.”

Phylogenetic reconstruction problem. We will be interested in reconstructing phylogenies in the regime where the reconstruction problem is solvable. The objective is to reconstruct the underlying tree whose internal nodes are unknown from the collection of sequences at the leaves. Since for both the CFN model and the JC model, the distribution $\bar{\mu}[\sigma]$ described above is independent of the location of the root we can only aim to reconstruct the underlying unrooted topology. Let \mathcal{T} represent the set of all *binary topologies* (that is, unrooted undirected binary trees) and $\mathcal{M}_{f,g}^{CFN}$ the family of CFN transition matrices, as described above, which correspond to distances d satisfying:

$$0 < f \leq d \leq g < g^*,$$

where g^* is given by (1) and f is an arbitrary positive constant. Let $\mathcal{T} \otimes \mathcal{M}_{f,g}^{CFN}$ denote the set of all unrooted phylogenies, where the underlying topology is in \mathcal{T} and all transition matrices on the edges are in $\mathcal{M}_{f,g}^{CFN}$. Rooting $T \in \mathcal{T} \otimes \mathcal{M}_{f,g}^{CFN}$ at an arbitrary node, let μ_T be the measure at the leaves of T as described above. It is well known, e.g. [6,9], that different elements in $\mathcal{T} \otimes \mathcal{M}_{f,g}^{CFN}$ correspond to different measures; therefore we will identify measures with their corresponding elements of $\mathcal{T} \otimes \mathcal{M}_{f,g}^{CFN}$. We are interested in finding a (efficiently computable) map Ψ such that $\Psi(\sigma_\partial^1, \dots, \sigma_\partial^k) \in \mathcal{T}$, where $\sigma_\partial = (\sigma_\partial^i)_{i=1}^k$ are k characters at the leaves of the tree. Moreover, we require that for every distribution $\mu_T \in \mathcal{T} \otimes \mathcal{M}_{f,g}^{CFN}$, if the characters $\sigma_\partial^1, \dots, \sigma_\partial^k$ are generated independently from μ_T , then with high probability $\Psi(\sigma_\partial^1, \dots, \sigma_\partial^k) = T$. The problem of finding an efficiently computable map

Ψ is called the *phylogenetic reconstruction problem* for the CFN model. The phylogenetic reconstruction problem for the JC model is defined similarly. In [9], it is shown that there exists a polynomial-time algorithm that reconstructs the topology from $k = \text{poly}(n, \log 1/\delta)$ characters, with probability of error δ . Our results are the following. We first define a subset of $\mathcal{T} \otimes \mathcal{M}_{f,g}^{\text{CFN}}$. In words, the Δ -Branch Model (Δ -BM) is a subset of $\mathcal{T} \otimes \mathcal{M}_{f,g}^{\text{CFN}}$ where the edge lengths $d(e), e \in E$, are discretized. This extra assumption is made for technical reasons. See Sect. 7.3.

Definition 1.1 (Δ -Branch Model) Let $\Delta > 0$. We denote by $\mathcal{Y}[\Delta]$ the subset of $\mathcal{T} \otimes \mathcal{M}_{f,g}^{\text{CFN}}$ where all $d(e)$'s are multiples of Δ . We call $\mathcal{Y}[\Delta]$ the Δ -Branch Model (Δ -BM).

Theorem 1.2 (Main Result) Consider the Δ -Branch Model above with: $0 < f \leq g < g^*$ and $\Delta > 0$. Then there exists a polynomial-time algorithm that reconstructs the topology of the tree from $k = c(f, g, \Delta)(\log n + \log 1/\delta)$ characters with error probability at most δ . In particular,

$$c(f, g, \Delta) = \frac{c(g)}{\min\{\Delta^2, f^2\}}.$$

Moreover, the value g^* given by (1) is tight.

Corollary 1.3 (Jukes–Cantor Model) Consider the JC model on binary trees where all edges satisfy

$$0 < f \leq g < g_{\text{JC}}^*, \text{ where } g_{\text{JC}}^* := g^*/2.$$

Under the Δ -BM, there exists a polynomial-time algorithm that reconstructs the topology of the tree from $c'(f, g, \Delta)(\log n + \log 1/\delta)$ characters with error probability at most δ and

$$c'(f, g, \Delta) = \frac{c'(g)}{\min\{\Delta^2, f^2\}}.$$

The value g_{JC}^* corresponds to the so-called Kesten–Stigum bound [17] on the reconstruction threshold, which has been conjectured to be tight for the JC model [19,31].

Theorem 1.2 and Corollary 1.3 extend also to cases where the data at the leaves is given with an arbitrary level of noise. For this *robust phylogenetic reconstruction problem* both values g^* and g_{JC}^* are tight. See [16].

The results stated here were first reported without proof in [7]. Note that in [7] the results were stated without the discretization assumption which is in fact needed for the final step of the proof. This is further explained in Sect. 7.3.

1.2 Organization of the paper

Roughly speaking, our reconstruction algorithm has two main components. First, the *statistical component* consists in

1. [Ancestral reconstruction] the reconstruction of sequences at internal nodes;
2. [Distance estimation] the estimation of distances between the nodes.

The former—detailed in Sect. 2—is borrowed from [22] where Steel’s conjecture is proved for the special case of balanced trees. In general trees, however, distance estimation is complicated by nontrivial correlations between reconstruction biases. We deal with these issues in Sect. 3.

Second, the *combinatorial component* of the algorithm—which uses quartet-based ideas from phylogenetics and is significantly more involved than [22]—is detailed in Sects. 4 and 5. A full description of the algorithm as well as an example of its execution can also be found in Sect. 5. Proof of the correctness of the algorithm is provided in Sects. 6 and 7.

1.3 Notation

Throughout we fix $0 < f \leq g < g' < g^*$, $\Delta > 0$, and $\gamma > 3$. By definition of g^* , we have $2e^{-4g} > 1$ and $2e^{-4g'} > 1$. We let $\theta = e^{-2g}$ and $\theta' = e^{-2g'}$.

2 Reconstruction of ancestral sequences

Ancestral reconstruction. In this section we state the results of [22] on ancestral reconstruction using recursive majority and we briefly explain how these results are used in the reconstruction algorithm. Following [22], we use state values ± 1 instead of $0/1$. Furthermore, we use the parameter $\theta(e) = 1 - 2p(e) = e^{-2d(e)}$. Note that $\theta(e)$ measures the correlation between the states at the endpoints of e . Because the CFN model is ferromagnetic, we have $0 \leq \theta(e) \leq 1$. In terms of θ we have reconstruction solvability whenever $\theta(e) \geq \theta > \theta_*$, for all edges e , where the value θ_* satisfies $2\theta_*^2 = 1$.

For the CFN model both the majority algorithm [13] and the recursive majority algorithm [20]¹ are effective in reconstructing the root value. Note that for other models, in general, most simple reconstruction algorithms are not effective all the way to the reconstruction threshold [16, 21, 24]. However, as noted in [22], there is an important difference between the two reconstruction algorithms when different edges e have different values of $\theta(e)$. Suppose that $\theta(e) > \theta' > \theta^*$ for all edges e . Then, the recursive majority function is effective in reconstructing the root value with probability bounded away from $1/2$ (as a function of θ'). On the other hand, it is easy to construct examples where the majority function reconstructs the root with probability tending to $1/2$ as the tree size increases.

¹ See below for a definition of these estimators.

The difference between the two algorithms can be roughly stated as follows. When different edges have different θ -values, different “parts” of the tree should have different weights in calculating the reconstructed value at the root. Indeed, when all θ -values are known, an appropriate weighted majority function can estimate the root value correctly with probability bounded away from $1/2$ [24]. However, when the θ -values are unknown, using uniform weights may result in an arbitrarily inaccurate estimator.

Recursive majority, on the other hand, does not require knowledge of the θ -values to be applied successfully. This essentially follows from the fact that the majority function is “noise-reducing” in the following sense. Suppose $\theta' > \theta^*$. Then, as we shall see shortly, there exists an integer ℓ and noise level $q < 1/2$ such that majority on the ℓ -level binary tree has the following property: if all leaf values are given with stochastic noise at most q , then the majority of these values differs from the actual root state with probability at most q . Therefore, recursive application of the majority function— ℓ levels at a time—results in an estimator whose error is at most q for any number of levels.

Properties of recursive majority. We proceed with a formal definition of recursive majority. Let $\widehat{\text{Maj}} : \{-1, 1\}^d \rightarrow \{-1, 1\}$ be the function defined as follows

$$\widehat{\text{Maj}}(x_1, \dots, x_d) = \text{sign} \left(\sum_{i=1}^d x_i + 0.5\omega \right),$$

where ω is ± 1 with probability $1/2$ independently of the x_i 's. In other words, $\widehat{\text{Maj}}$ outputs the majority value of its input arguments, unless there is a tie in which case it outputs ± 1 with probability $1/2$ each. For consistency, we denote all statistical estimators with a hat. Note in particular that our notation differs from [22].

Next we define the “noise-reduction” property of majority. The function η below is meant to measure the noise level at the leaves.

Definition 2.1 (Correlation at the root) Let $T = (V, E)$ be a tree rooted at ρ with leaf set ∂T . For functions $\tilde{\theta} : E \rightarrow [0, 1]$ and $\tilde{\eta} : \partial T \rightarrow [0, 1]$, let $\text{CFN}(\tilde{\theta}, \tilde{\eta})$ be the CFN model on T where $\theta(e) = \tilde{\theta}(e)$ for all e which are not adjacent to ∂T , and $\theta(e) = \tilde{\theta}(e)\tilde{\eta}(v)$ for all $e = (u, v)$, with $v \in \partial T$. Let

$$\text{MajCorr}(\tilde{\theta}, \tilde{\eta}) = \mathbf{E} \left[+\widehat{\text{Maj}}(\sigma_{\partial T}) \mid \sigma_\rho = +1 \right] = \mathbf{E} \left[-\widehat{\text{Maj}}(\sigma_{\partial T}) \mid \sigma_\rho = -1 \right],$$

where σ is one sample drawn from $\text{CFN}(\tilde{\theta}, \tilde{\eta})$.

Proposition 2.2 (Noise reduction of majority [22]) *Let b and θ_{\min} be such that $b\theta_{\min}^2 > h^2 > 1$. Then there exist $\ell = \ell(b, \theta_{\min})$, $\alpha = \alpha(b, \theta_{\min}) > h^\ell > 1$ and $\beta = \beta(b, \theta_{\min}) > 0$, such that any $\text{CFN}(\theta, \eta)$ model on the ℓ -level b -ary tree satisfying $\min_{e \in E} \theta(e) \geq \theta_{\min}$ and $\min_{v \in \partial T} \eta(v) \geq \eta_{\min} > 0$ also satisfies:*

$$\text{MajCorr}(\theta, \eta) \geq \min\{\alpha\eta_{\min}, \beta\}. \tag{2}$$

In particular if $\eta_{\min} \geq \beta$ then:

$$\text{MajCorr}(\theta, \eta) \geq \beta. \tag{3}$$

General trees. Recursive application of Proposition 2.2 allows the reconstruction of the root value on any *balanced* binary tree with correlation at least β . However, below we consider *general* trees. In particular, when estimating the sequence at an internal node u of the phylogenetic tree, we wish to apply Proposition 2.2 to a subtree rooted at u and this subtree need not be balanced. This can be addressed by “completing” the subtree into a balanced tree (with number of levels a multiple of ℓ) and taking all added edges to have length $d(e) = 0$, that is, $\theta(e) = 1$. Fix ℓ and β so as to satisfy Proposition 2.2 with $b = 2$ and $\theta_{\min} = \theta$. (Recall that for the proof of Theorem 1.2 we assume $\theta(e) \geq \theta > \theta^*$, for all edges e .) Consider the following recursive function of $x = (x_1, x_2, \dots)$:

$$\begin{aligned} \widehat{\text{Maj}}_\ell^0(x_1) &= x_1, \\ \widehat{\text{Maj}}_\ell^j(x_1, \dots, x_{2^{j\ell}}) &= \widehat{\text{Maj}} \left(\widehat{\text{Maj}}_\ell^{j-1}(x_1, \dots, x_{2^{(j-1)\ell}}), \dots, \right. \\ &\quad \left. \widehat{\text{Maj}}_\ell^{j-1}(x_{2^{j\ell}-2^{(j-1)\ell}+1}, \dots, x_{2^{j\ell}}) \right), \end{aligned}$$

for all $j \geq 1$. Now, if $\sigma_{\partial T}$ is a character at the leaves of T , let us define the function $\widehat{\text{Anc}}_{\rho, T}(\sigma_{\partial T})$ that estimates the ancestral state at the root ρ of T using recursive majority as follows:

1. Let \tilde{T} be the tree T (minimally) completed with edges of θ -value 1 so that \tilde{T} is a complete binary tree with a number of levels a multiple of ℓ , say $J\ell$;
2. Assign to the leaves of \tilde{T} the value of their ancestor leaf in T under $\sigma_{\partial T}$;
3. Let $\tilde{\sigma}$ be the leaf states of \tilde{T} arranged in pre-order;
4. Compute

$$\widehat{\text{Anc}}_{\rho, T}(\sigma_{\partial T}) := \widehat{\text{Maj}}_\ell^J(\tilde{\sigma}).$$

From Proposition 2.2, we get:

Proposition 2.3 (Recursive Majority) *Let $T = (V, E)$ be a tree rooted at ρ with leaf set ∂T . Let $\sigma_{\partial T}$ be one sample drawn from $\text{CFN}(\theta, \eta)$ with $\theta(e) \geq \theta > \theta^*$ for all edges e and $\eta(v) = 1$ for all $v \in \partial T$. Then, we can choose ℓ and $\beta > 0$ so that*

$$\mathbb{P} \left[\widehat{\text{Anc}}_{\rho, T}(\sigma_{\partial T}) = \sigma_\rho \mid \sigma_\rho = +1 \right] = \mathbb{P} \left[\widehat{\text{Anc}}_{\rho, T}(\sigma_{\partial T}) = \sigma_\rho \mid \sigma_\rho = -1 \right] \geq \frac{1 + \beta}{2}. \tag{4}$$

In the remainder of the paper, we will use Proposition 2.3 with $\theta(e) = e^{-2d(e)}$. Also, if $\sigma_{\partial T} = (\sigma_{\partial T}^i)_{i=1}^k$ is a collection of k characters at the leaves of T , we extend the

function $\widehat{\text{Anc}}_{\rho, T}$ to collections of characters in the natural way:

$$\widehat{\text{Anc}}_{\rho, T}(\sigma_{\partial T}) := \left(\widehat{\text{Anc}}_{\rho, T}(\sigma_{\partial T}^i) \right)_{i=1}^k. \tag{5}$$

3 Distance estimation

Throughout this section, we fix a tree T on n leaves. Recall the definition of our path metric d from Sect. 1.1. We assume that $\sigma_{\partial T} = (\sigma_{\partial T}^t)_{t=1}^k$ are k i.i.d. samples (or characters) at the leaves of T generated by the CFN model with parameters $d(e) \leq g$ for all $e \in \mathcal{E}(T)$ (a lower bound on d is not required in this section, but it will be in the next one). We think of $\sigma_{\partial T} = ((\sigma_u^t)_{u \in \partial T})_{t=1}^k$ as n sequences of length k and we sometimes refer to $\sigma_u = (\sigma_u^t)_{t=1}^k$ as the “sequence at u ”.

Distance between leaves. As explained in Sect. 1.1, a basic concept in phylogenetic reconstruction is the notion of a metric on the leaves of the tree. The distance between two leaves is a measure of the correlation between their respective sequences. We begin with the definition of a natural distance estimator.

Definition 3.1 (Correlation between sequences) Let u, v be leaves of T . The following quantity

$$\widehat{\text{Dist}}(\sigma_u, \sigma_v) = -\frac{1}{2} \ln \left[\left(\frac{1}{k} \sum_{t=1}^k \sigma_u^t \sigma_v^t \right)_+ \right], \tag{6}$$

is an estimate of $d(u, v)$. In cases where the sum inside the log is non-positive, we define $\widehat{\text{Dist}}(\sigma_u, \sigma_v) = +\infty$.

The next proposition provides a guarantee on the performance of $\widehat{\text{Dist}}$. The proof follows from standard concentration inequalities. See, e.g., [9]. An important point to note is that, if we use *short* sequences of length $O(\log n)$, the guarantee applies only to short distances, that is, distances of order $O(1)$.

Proposition 3.2 (Accuracy of $\widehat{\text{Dist}}$) For all $\varepsilon > 0, M > 0$, there exists $c = c(\varepsilon, M; \gamma) > 0$, such that if the following conditions hold:

- [Short distance] $d(u, v) < M$,
- [Sequence length] $k = c' \log n$, for some $c' > c$,

then

$$\left| d(u, v) - \widehat{\text{Dist}}(\sigma_u, \sigma_v) \right| < \varepsilon,$$

with probability at least $1 - O(n^{-\gamma})$.

Proof By Azuma’s inequality we get

$$\begin{aligned}
 \mathbb{P} \left[\widehat{\text{Dist}}(\sigma_u, \sigma_v) > d(u, v) + \varepsilon \right] &= \mathbb{P} \left[\frac{1}{k} \sum_{t=1}^k \sigma_u^t \sigma_v^t < e^{-2d(u,v)-2\varepsilon} \right] \\
 &= \mathbb{P} \left[\frac{1}{k} \sum_{t=1}^k \sigma_u^t \sigma_v^t < e^{-2d(u,v)} - (1 - e^{-2\varepsilon})e^{-2d(u,v)} \right] \\
 &\leq \mathbb{P} \left[\frac{1}{k} \sum_{t=1}^k \sigma_u^t \sigma_v^t < \mathbb{E}[\sigma_u^1 \sigma_v^1] - (1 - e^{-2\varepsilon})e^{-2M} \right] \\
 &\leq \exp \left(-\frac{((1 - e^{-2\varepsilon})e^{-2M})^2}{2k(2/k)^2} \right) \\
 &= n^{-c'K},
 \end{aligned}$$

with K depending on M and ε . Above, we used that

$$d(u, v) = -\frac{1}{2} \ln \mathbb{E}[\sigma_u^1 \sigma_v^1].$$

A similar inequality holds for the other direction. □

Distance between internal nodes. Our reconstruction algorithm actually requires us to compute distances between *internal* nodes. Note that, in general, we do not know the true sequences at internal nodes of the tree. Therefore, we need to estimate distances by applying (6) to *reconstructed* sequences. An obvious issue with this idea is that reconstructed sequences are subject to a systematic bias.

Definition 3.3 (Bias) Suppose v is the root of a subtree T_v of T . Let $\widehat{\sigma}_v = \widehat{\text{Anc}}_{v, T_v}(\sigma_{\partial T_v})$. Then, the quantity

$$\lambda(T_v, v) = -\frac{1}{2} \ln(\mathbb{E}[\sigma_v^1 \widehat{\sigma}_v^1]_+),$$

is called the *reconstruction bias* at v on T_v . (Note the similarity to (6).) We denote

$$B(\tilde{g}) = \sup \lambda(T_v, v),$$

where the supremum is taken over all rooted binary trees with edge lengths at most \tilde{g} . Note that for all \tilde{g} , $0 < \tilde{g} < g^*$, we have by Proposition 2.3

$$0 < B(\tilde{g}) < +\infty.$$

We also denote

$$\beta(\tilde{g}) = e^{-2B(\tilde{g})}.$$




Fig. 1 Basic disjoint setup for estimator $\widehat{\text{Int}}$ and routines ISSPLIT in Fig. 9 and ISSHORT in Fig. 6. The subtrees $T_{v_1}, T_{w_1}, T_{v_2}, T_{w_2}$ are edge-disjoint (see Definition 3.9 below)

Since $B(g) > 0$, we cannot estimate exactly the internal sequences. Hence Proposition 3.2 cannot be used directly to estimate distances between internal nodes. We deal with this issue as follows. Consider the configuration in Fig. 1. More precisely:

Definition 3.4 (Basic Disjoint Setup (Preliminary Version)) Root T at an arbitrary vertex. Note that, by reversibility, the CFN model on a tree T can be rerooted arbitrarily without changing the distribution at the leaves. Denote by T_x the subtree of T rooted at x . We consider two internal nodes u_1, u_2 that are not descendants of each other in T . Denote by v_φ, w_φ the children of u_φ in T_{u_φ} , $\varphi = 1, 2$. We call this configuration the *Basic Disjoint Setup (Preliminary Version)*.


We need a few basic combinatorial definitions.

Definition 3.5 (Restricted Subtree) Let $V' \subseteq V$ be a subset of the vertices of T . The *subtree of T restricted to V'* is the tree T' obtained by 1) keeping only nodes and edges on paths between vertices in V' and 2) by then contracting all paths composed of vertices of degree 2, except the nodes in V' . We sometimes use the notation $T|_{V'} := T'$. See Fig. 2 for an example.

Definition 3.6 (Quartet) A *quartet* is a set of four leaves. More generally, for any set $\mathcal{Q} = \{v_1, w_1, v_2, w_2\}$ of four nodes of T , we think of \mathcal{Q} as a quartet on $T|_{\mathcal{Q}}$. We say that a quartet $\mathcal{Q} = \{v_1, w_1, v_2, w_2\}$ is nondegenerate if none of the nodes in \mathcal{Q} are on a path between two other nodes in \mathcal{Q} . There are three possible (leaf-labeled) topologies on a nondegenerate quartet, called *quartet splits*, one for each partition of the four leaves into two pairs. In Fig. 1, the correct quartet split on \mathcal{Q} is $\{\{v_1, w_1\}\{v_2, w_2\}\}$ which we denote by $v_1 w_1 | v_2 w_2$.

To compute the distance between u_1 and u_2 , we think of the path between u_1 and u_2 as the internal edge of the quartet $\mathcal{Q} = \{v_1, w_1, v_2, w_2\}$. The reconstructed sequences at $x \in \mathcal{Q}$ also suffer from a systematic bias. However, we prove in Proposition 3.8 that the bias does not affect the computation of the length of the *internal* edge of the

Fig. 2 Restricting the top tree to its white nodes



quartet. Indeed, as depicted in Fig. 3, we treat the systematic errors introduced by the reconstruction as “extra edges” attached to the nodes in $\mathcal{Q} = \{v_1, w_1, v_2, w_2\}$, with corresponding “endpoints” $\widehat{\mathcal{Q}} = \{\hat{v}_1, \hat{w}_1, \hat{v}_2, \hat{w}_2\}$.

Our estimator $\widehat{\text{Int}}$ is then obtained from a classical distance-based computation—known in phylogenetics as the *Four-Point Method* [4]—applied to the “extra nodes” $\widehat{\mathcal{Q}}$ (see (7) below). For this idea to work, it is crucial that the systematic error in the reconstructed sequences at $x \in \mathcal{Q}$ be *independent* given the true sequences at $x \in \mathcal{Q}$. This is the case when u_1 and u_2 are not descendant of each other because of the Markov property.

We now define our distance estimator $\widehat{\text{Int}}$. We let $\widehat{\sigma}_x = \widehat{\text{Anc}}_{x, T_x}(\sigma_{\partial T_x})$ be the reconstructed sequence at x .

Definition 3.7 (*Distance Estimator $\widehat{\text{Int}}$*) Consider the basic disjoint setup (preliminary version). Then, we let

$$\widehat{\text{Int}}(\widehat{\sigma}_{v_1}, \widehat{\sigma}_{w_1}; \widehat{\sigma}_{v_2}, \widehat{\sigma}_{w_2}) = \frac{1}{2} \left(\widehat{\text{Dist}}(\widehat{\sigma}_{v_1}, \widehat{\sigma}_{v_2}) + \widehat{\text{Dist}}(\widehat{\sigma}_{w_1}, \widehat{\sigma}_{w_2}) - \widehat{\text{Dist}}(\widehat{\sigma}_{v_1}, \widehat{\sigma}_{w_1}) - \widehat{\text{Dist}}(\widehat{\sigma}_{v_2}, \widehat{\sigma}_{w_2}) \right). \tag{7}$$




Fig. 3 Basic setup for estimator $\widehat{\text{Int}}$ where reconstruction bias is represented by *dashed edges*

If one of the $\widehat{\text{Dist}}$ quantities is $+\infty$, the function $\widehat{\text{Int}}$ is set to $+\infty$.

The next proposition provides a guarantee on the performance of $\widehat{\text{Int}}$. As in Proposition 3.2, the guarantee applies only to $O(1)$ distances, if we use sequences of length $O(\log n)$.

Proposition 3.8 (Accuracy of $\widehat{\text{Int}}$) *Consider the basic disjoint setup (preliminary version). Let $\mathcal{Q} = \{v_1, w_1, v_2, w_2\}$. For all $\varepsilon > 0, M > 0$, there exists $c = c(\varepsilon, M; \gamma, g) > 0$ such that, if the following hold:*

- [Edge lengths] *For all $e \in \mathcal{E}(T)$, $d(e) \leq g < g^*$;*
- [Short distances] *We have $d(x, y) < M, \forall x, y \in \mathcal{Q}$;*
- [Sequence length] *The sequences used have length $k = c' \log n$, for some $c' > c$;*

then, with probability at least $1 - O(n^{-\gamma})$,

$$|d(u_1, u_2) - \widehat{\text{Int}}(\widehat{\sigma}_{v_1}, \widehat{\sigma}_{w_1}; \widehat{\sigma}_{v_2}, \widehat{\sigma}_{w_2})| < \varepsilon.$$

Proof First note that, for all $t \in \{1, \dots, k\}$, conditioned on $\{\sigma_x^t\}_{x \in \mathcal{Q}}$, the values $\{\widehat{\sigma}_x^t\}_{x \in \mathcal{Q}}$ are independent by the Markov property. Therefore, as we observed before, the systematic errors introduced by the reconstruction process can be treated as “extra edges”, with endpoints $\widehat{\mathcal{Q}} = \{\widehat{v}_1, \widehat{w}_1, \widehat{v}_2, \widehat{w}_2\}$ (see Fig. 3). Moreover, from (4), it follows that, $\forall x \in \mathcal{Q}, \forall t \in \{1, \dots, k\}$,

$$\mathbb{P}[\widehat{\sigma}_x^t = \sigma_x^t | \sigma_x^t] \geq \frac{1 + \beta(g)}{2}, \tag{8}$$

where $\beta(g) > 0$ is specified by Definition 3.3. From (8), the lengths of the “extra edges”—equal to the reconstruction biases—are at most $B(g)$.

Appealing to Proposition 3.2 with constants $\frac{\varepsilon}{2}$ and $M + 2B(g)$ and the union bound, we argue that with probability at least $1 - O(n^{-\gamma})$ for all $x, y \in \mathcal{Q}$

$$|\widehat{\text{Dist}}(\widehat{\sigma}_x, \widehat{\sigma}_y) - d(\widehat{x}, \widehat{y})| < \frac{\varepsilon}{2},$$

Algorithm DISTANCEESTIMATE
Input: Two nodes u_1, u_2 ; a rooted forest \mathcal{F} ; accuracy radius \widehat{R}_{acc} ;
Output: Distance estimate ν ;

- [Children] Let w_φ, v_φ be the children of u_φ in \mathcal{F} for $\varphi = 1, 2$ (if u_φ is a leaf, set $w_\varphi = v_\varphi = u_\varphi$);
- [Sequence Reconstruction] For $x \in \{v_1, w_1, v_2, w_2\}$, set $\widehat{\sigma}_x = \widehat{\text{Anc}}_{x, T_x^{\mathcal{F}}}(\sigma_{\partial T_x^{\mathcal{F}}})$;
- [Accuracy Cutoff] If there is $\{x, y\} \subseteq \{v_1, w_1, v_2, w_2\}$ such that

$$\widehat{\text{Dist}}(\widehat{\sigma}_x, \widehat{\sigma}_y) > \widehat{R}_{\text{acc}},$$
 return $+\infty$;
- [Distance Computation] Otherwise return

$$\nu = \widehat{\text{Int}}(\widehat{\sigma}_{v_1}, \widehat{\sigma}_{w_1}; \widehat{\sigma}_{v_2}, \widehat{\sigma}_{w_2}).$$

Fig. 4 Routine DISTANCEESTIMATE

where we extend d to the “extra nodes” \hat{x}, \hat{y} by setting

$$d(x, \hat{x}) = -\frac{1}{2} \ln(\mathbb{E}[\sigma_x \widehat{\sigma}_x]_+),$$

and likewise for \hat{y} (see Fig. 3). Observing that $d(u_1, u_2) = \frac{1}{2}(d(\hat{v}_1, \hat{v}_2) + d(\hat{w}_1, \hat{w}_2) - d(\hat{v}_1, \hat{w}_1) - d(\hat{v}_2, \hat{w}_2))$ and using the above we get

$$|d(u_1, u_2) - \widehat{\text{Int}}(\widehat{\sigma}_{v_1}, \widehat{\sigma}_{w_1}; \widehat{\sigma}_{v_2}, \widehat{\sigma}_{w_2})| < \varepsilon.$$

□

Distances between restricted subtrees. In fact, we need to apply Proposition 3.8 to *restricted subtrees* of the true tree. Indeed, in the reconstruction algorithm, we maintain a “partially reconstructed subforest” of the true tree, that is, a collection of restricted subtrees of T . For this more general setup, we use the routine DISTANCEESTIMATE detailed in Fig. 4.

To generalize Proposition 3.8, we need a few definitions. First, the notion of edge disjointness is borrowed from [23].


Definition 3.9 (Edge Disjointness) Denote by $\text{path}_T(x, y)$ the path (sequence of edges) connecting x to y in T . We say that two restricted subtrees T_1, T_2 of T are *edge disjoint* if

$$\text{path}_T(x_1, y_1) \cap \text{path}_T(x_2, y_2) = \emptyset,$$

for all $x_1, y_1 \in \mathcal{L}(T_1)$ and $x_2, y_2 \in \mathcal{L}(T_2)$. We say that T_1, T_2 are *edge sharing* if they are not edge disjoint. See Fig. 5 for an example. (If T_1 and T_2 are directed, we take this definition to refer to their underlying undirected version.)

Definition 3.10 (Legal Subforest) We say that a tree is a *rooted full binary tree* if all its internal nodes have degree 3 except for the root which has degree 2. A restricted subtree T_1 of T is a *legal subtree* of T if:

Fig. 5 The subtrees $T|_{\{u_1, u_2, u_3, u_8\}}$ and $T|_{\{u_4, u_5, u_6, u_7\}}$ are edge-disjoint. The subtrees $T|_{\{u_1, u_5, u_6, u_8\}}$ and $T|_{\{u_2, u_3, u_4, u_7\}}$ are edge-sharing



1. It is a rooted full binary tree, that is, if it has a unique node of degree 2—its root;
2. And if all its leaves are leaves of T .

Moreover, we say that a forest

$$\mathcal{F} = \{T_1, \dots, T_\alpha\},$$

is *legal subforest* of T if the T_φ 's are *edge-disjoint* legal subtrees of T . We denote by $\rho(\mathcal{F})$ the set of roots of \mathcal{F} .

Definition 3.11 (Dangling Subtrees) We say that two edge-disjoint legal subtrees T_1, T_2 of T are *dangling* if there is a choice of root ρ^* for T *not in* T_1 or T_2 that is consistent with the rooting of both T_1 and T_2 , that is, pointing the edges of T away from ρ^* is such that the edges of T_1 and T_2 are directed away from their respective roots.

See Fig. 7 below for an example where two legal, edge-disjoint subtrees are *not* dangling. We generalize the basic configuration of Fig. 1 as follows:

Definition 3.12 (Basic Disjoint Setup (Dangling)) Let $T_1 = T_{u_1}$ and $T_2 = T_{u_2}$ be two legal subtrees of T rooted at u_1 and u_2 respectively. Assume further that T_1 and T_2 are *edge-disjoint* and *dangling*. Denote by v_φ, w_φ the children of u_φ in $T_\varphi, \varphi = 1, 2$. If u_φ is a leaf, we let instead $v_\varphi = w_\varphi = u_\varphi$. We call this configuration the *Basic Disjoint Setup (Dangling)*.

We make another important remark about the more general setup considered here. Note that even though the true tree satisfies the “short edge length” condition (that is, $d(e) \leq g$) by assumption, partially reconstructed subtrees *may not*—because their edges are actually *paths* in the true tree. Therefore, an important step of the reconstruction algorithm is to make sure that all restricted edges of the partially reconstructed subforest are short enough for recursive majority to be accurate. (This explains why we need to use two constants $g < g'$ below g^* .) Refer to routine LOCALCHERRY in Fig. 13 in Sect. 5.

We then obtain the following generalization of Proposition 3.8. Note that the routine DISTANCEESTIMATE in Fig. 4 uses an “accuracy cutoff” \widehat{R}_{acc} . As we show in the proof below, this ensures that distances that are too long are rejected, in which case the value $+\infty$ is returned.

Proposition 3.13 (Accuracy of DISTANCEESTIMATE) *Consider the basic disjoint setup (dangling). Let $\mathcal{Q} = \{v_1, w_1, v_2, w_2\}$. For all $\varepsilon > 0$, $M > 0$, and $\widehat{R}_{\text{acc}} > M + 2B(g')$, there exists $c = c(\varepsilon, M, \widehat{R}_{\text{acc}}; \gamma, g') > 0$ such that, if the following hold:*

- [Edge length] *It holds that $d(e) \leq g' < g^*$, $\forall e \in \mathcal{E}(T_x)$, $x \in \mathcal{Q}$;*
 - [Sequence length] *The sequences used have length $k = c' \log n$, for some $c' > c$;*
- then, with probability at least $1 - O(n^{-\gamma})$, the following holds: letting ν be the output of DISTANCEESTIMATE in Fig. 4, we have that if one of the following holds*

1. [Short distances] *We have $d(x, y) < M$, $\forall x, y \in \mathcal{Q}$;*
2. [Finite estimate] *$\nu < +\infty$;*

then

$$|d(u_1, u_2) - \nu| < \varepsilon.$$

Proof The first half of the proposition follows immediately from Propositions 3.2 and 3.8. Refer to Fig. 3.

The second part follows from a “double window” argument as in [9, Theorem 9]. Let $0 < \tau < \frac{1}{2}$ such that

$$\widehat{R}_{\text{acc}} = M + 2B(g') - \frac{1}{2} \ln(2\tau).$$

By assumption, $\nu < +\infty$ and therefore for all pairs $\{x, y\} \subseteq \{v_1, w_1, v_2, w_2\}$ we have

$$\widehat{\text{Dist}}(\widehat{\sigma}_x, \widehat{\sigma}_y) \leq M + 2B(g') - \frac{1}{2} \ln(2\tau). \tag{9}$$

It follows directly from Azuma’s inequality that a pair $\{x, y\}$ such that

$$d(\widehat{x}, \widehat{y}) \geq M + 2B(g') - \frac{1}{2} \ln \tau,$$

satisfies (9) with probability at most

$$\begin{aligned} & \mathbb{P} \left[\widehat{\text{Dist}}(\widehat{\sigma}_x, \widehat{\sigma}_y) \leq M + 2B(g') - \frac{1}{2} \ln 2\tau \right] \\ &= \mathbb{P} \left[\frac{1}{k} \sum_{t=1}^k \widehat{\sigma}_x^t \widehat{\sigma}_y^t \geq 2\tau e^{-2[M+2B(g')]} \right] \\ &\leq \mathbb{P} \left[\frac{1}{k} \sum_{t=1}^k \widehat{\sigma}_x^t \widehat{\sigma}_y^t \geq e^{-2d(\widehat{x}, \widehat{y})} + \tau e^{-2[M+2B(g')]} \right] \\ &\leq \exp \left(-\frac{\left(\tau e^{-2[M+2B(g')]} \right)^2}{2k(2/k)^2} \right) \\ &= n^{-c'K}, \end{aligned}$$

Algorithm ISSHORT
Input: Two pairs of nodes $(v_1, w_1), (v_2, w_2)$; rooted forest \mathcal{F} ; accuracy radius \widehat{R}_{acc} ; tolerance ε ;
Output: Boolean value and length estimate;

- [Sequence Reconstruction] For $x \in \{v_1, w_1, v_2, w_2\}$, set $\widehat{\sigma}_x = \widehat{\text{Anc}}_{x, T_{\mathcal{F}}}(\sigma_{\partial T_{\mathcal{F}}})$;
- [Accuracy Cutoff] If there is $\{x, y\} \subseteq \{v_1, w_1, v_2, w_2\}$ such that

$$\widehat{\text{Dist}}(\widehat{\sigma}_x, \widehat{\sigma}_y) > \widehat{R}_{\text{acc}},$$
 return $+\infty$;
- [Internal Edge Length] Set

$$\nu = \widehat{\text{Int}}(\widehat{\sigma}_{v_1}, \widehat{\sigma}_{w_1}; \widehat{\sigma}_{v_2}, \widehat{\sigma}_{w_2});$$
- [Test] If $\nu < g + \varepsilon$ return (TRUE, ν), o.w. return (FALSE, 0);

Fig. 6 Routine ISSHORT

for a constant $K > 0$ depending on M, g', τ . The result follows by applying the first part of the proposition with M replaced by $M + 2B(g') - \frac{1}{2} \ln \tau$. (Note that $d(x, y) \leq d(\widehat{x}, \widehat{y})$, for all $x, y \in \mathcal{Q}$.) □

We also need a simpler variant of DISTANCEESTIMATE (Fig. 4) whose purpose is to test whether the internal path of a quartet is longer than g . We record this variant in Fig. 6.


Detecting long distances when T_1 and T_2 are not dangling. Roughly speaking, our reconstruction algorithm works by progressively merging subtrees that are close in the true tree. (See Sect. 5 for further details.) Hence, the algorithm needs to tell whether or not two subtrees are sufficiently close to be considered for this merging operation. However, as we explain in the next sections, we cannot guarantee that the basic disjoint setup (dangling) in Proposition 3.13 applies to all situations encountered during the execution of the algorithm. Instead we use a special “multiple test” to detect long distances. This test is performed by the routine DISTORTEDMETRIC detailed in Fig. 8.

The routine has a further important property. During the course of the algorithm, since we only have partial information about the structure of the tree, we may not always know whether or not two subtrees are dangling—and therefore whether or not DISTANCEESTIMATE in Fig. 4 returns accurate distance estimates. The “multiple test” in DISTORTEDMETRIC is such that, if the routine returns a *finite* estimate, that estimate is accurate. We proceed with an explanation of these properties.

We further generalize the basic configuration of Fig. 1 as follows:

Definition 3.14 (Basic Disjoint Setup (General)) Let $T_1 = T_{x_1}$ and $T_2 = T_{x_2}$ be two restricted subtrees of T rooted at x_1 and x_2 respectively. Assume further that T_1 and T_2 are *edge-disjoint*, but not necessarily *dangling*. Denote by y_φ, z_φ the children of x_φ in $T_\varphi, \varphi = 1, 2$. Let w_φ be the node in T where the path between T_1 and T_2 meets $T_\varphi, \varphi = 1, 2$. Note that w_φ may not be in T_φ since T_φ is *restricted*, $\varphi = 1, 2$. If $w_\varphi \neq x_\varphi$, assume without loss of generality that w_φ is in the subtree of T rooted at z_φ , or on the edge $(x_\varphi, z_\varphi), \varphi = 1, 2$. We call this configuration the *basic disjoint setup (general)*. See Fig. 7. Let $d(T_1, T_2)$ be the length of the path between w_1 and w_2 (in the path metric d).

Fig. 7 Basic Disjoint Setup (General). The rooted subtrees T_1, T_2 are edge-disjoint but are not assumed to be dangling. The white nodes may not be in the restricted subtrees T_1, T_2 . The case $w_1 = x_1$ and/or $w_2 = x_2$ is possible. Note that if we root the tree at any node along the dashed path, the subtrees rooted at y_1 and y_2 are edge-disjoint and dangling (unlike T_1 and T_2)



The key point to note is that when computing the distance between y_1 and y_2 rather than the distance between x_1 and x_2 , then the assumptions of Proposition 3.13 are satisfied (as the claim holds in this case by rooting the tree at any node along the path connecting w_1 and w_2). Hence, if T_1 and T_2 are far apart, the distance between y_1 and y_2 is correctly estimated as being large. On the other hand, if T_1 and T_2 are dangling and close, then distances between all pairs in $\{y_1, z_1\} \times \{y_2, z_2\}$ are accurately estimated as being short.

Proposition 3.15 (Accuracy of DISTORTEDMETRIC) *Consider the basic disjoint setup (general) with $\mathcal{F} = \{T_1, T_2\}$ and $\mathcal{Q} = \{y_1, z_1, y_2, z_2\}$. For all $\varepsilon > 0, M > 0$, and $\hat{R}_{acc} > M + 2B(g') + 4g'$ there exists $c = c(\varepsilon, M, \hat{R}_{acc}; \gamma, g') > 0$ such that, if the following hold:*

- [Edge length] *It holds that $d(e) \leq g' < g^*, \forall e \in \mathcal{E}(T_x), x \in \mathcal{Q}$; also, $d(x_\varphi, y_\varphi), d(x_\varphi, z_\varphi) \leq g', \varphi = 1, 2$;*
- [Weight estimates] *We are given weight estimates*

$$h : \mathcal{E}(\mathcal{F}) \rightarrow \mathbb{R}_+,$$

such that $|h(e) - d(e)| < \varepsilon/16, \forall e \in \mathcal{E}(\mathcal{F})$;

- [Sequence length] *The sequences used have length $k = c' \log n$, for some $c' > c$;*
- then, with probability at least $1 - O(n^{-\gamma})$, the following holds: letting v be the output of DISTORTEDMETRIC in Fig. 8, we have that if one of the following hold*

1. [Dangling case] T_1 and T_2 are dangling and $d(T_1, T_2) < M$;
2. [Finite estimate] $v < +\infty$;

then

$$|v - d(x_1, x_2)| < \varepsilon.$$

Proof Choose c as in Proposition 3.13 with parameters $\varepsilon/16$ and $M + 4g'$. As in Fig. 8, let

Algorithm DISTORTEDMETRIC
Input: Two nodes x_1, x_2 ; a rooted forest \mathcal{F} ; edge lengths $\{h(e)\}_{e \in \mathcal{F}}$; accuracy radius \widehat{R}_{acc} ; tolerance ε ;
Output: Distance ν ;

- [Children] Let y_φ, z_φ be the children of x_φ in \mathcal{F} for $\varphi = 1, 2$ (if x_φ is a leaf, set $z_\varphi = y_\varphi = x_\varphi$);
- [Distance Computations] For all pairs $(r_1, r_2) \in \{y_1, z_1\} \times \{y_2, z_2\}$, compute

$$\mathcal{D}'(r_1, r_2) := \text{DISTANCEESTIMATE}(r_1, r_2; \mathcal{F}; \widehat{R}_{acc}) - h(x_1, r_1) - h(x_2, r_2);$$
- [Multiple Test] If

$$\max \left\{ \left| \mathcal{D}'(r_1^{(1)}, r_2^{(1)}) - \mathcal{D}'(r_1^{(2)}, r_2^{(2)}) \right| : (r_1^{(\varphi)}, r_2^{(\varphi)}) \in \{y_1, z_1\} \times \{y_2, z_2\}, \varphi = 1, 2 \right\} < \varepsilon/2,$$
 return $\nu := \mathcal{D}'(z_1, z_2)$, otherwise return $\nu := +\infty$ (return $\nu := +\infty$ if any of the distances above is $+\infty$).

Fig. 8 Routine DISTORTEDMETRIC

$$\mathcal{D}'(r_1, r_2) := \text{DISTANCEESTIMATE}(r_1, r_2; \mathcal{F}; \widehat{R}_{acc}) - h(x_1, r_1) - h(x_2, r_2),$$

for all pairs $(r_1, r_2) \in \{y_1, z_1\} \times \{y_2, z_2\}$ (where DISTANCEESTIMATE is defined in Fig. 4). In *Case 1*, the result follows directly from Proposition 3.13 and the remarks above Proposition 3.15. In particular, note that, with probability $1 - O(n^{-\nu})$,

$$\left| \mathcal{D}'(r_1, r_2) - d(x_1, x_2) \right| < \frac{3\varepsilon}{16},$$

for all $(r_1, r_2) \in \{y_1, z_1\} \times \{y_2, z_2\}$ and therefore

$$\begin{aligned} & \max \left\{ \left| \mathcal{D}'(r_1^{(1)}, r_2^{(1)}) - \mathcal{D}'(r_1^{(2)}, r_2^{(2)}) \right| : (r_1^{(\varphi)}, r_2^{(\varphi)}) \in \{y_1, z_1\} \times \{y_2, z_2\}, \varphi = 1, 2 \right\} \\ & < 2 \left(\frac{3\varepsilon}{16} \right) < \varepsilon/2, \end{aligned}$$

and DISTORTEDMETRIC returns a finite value (namely $\nu := \mathcal{D}'(z_1, z_2)$; see Fig. 8) that is accurate within $3\varepsilon/16 < \varepsilon$.

In *Case 2*, the condition $\nu < +\infty$ implies in particular that all four distance estimates are equal up to $\varepsilon/2$. From the remark above the statement of the proposition, at least one such distance, say $\mathcal{D}'(y_1, y_2)$ w.l.o.g., is computed within the basic disjoint setup (dangling) and we can therefore apply Proposition 3.13 again. In particular, we have

$$\left| \mathcal{D}'(y_1, y_2) - d(x_1, x_2) \right| < \frac{3\varepsilon}{16},$$

and therefore

$$|\nu - d(x_1, x_2)| < \frac{\varepsilon}{2} + \frac{3\varepsilon}{16} < \varepsilon,$$

(where again $\nu := \mathcal{D}'(z_1, z_2)$). □

Remark Note that DISTORTEDMETRIC outputs $+\infty$ in two very distinct situations. In particular, if DISTORTEDMETRIC outputs $+\infty$, then either the subtrees are too far to obtain an accurate estimate or T_1 and T_2 are not dangling (or both). This convention will turn out to be convenient in the statement of the full algorithm.

4 Quartet tests

Before giving a complete description of the reconstruction algorithm, we introduce some important combinatorial tools. As we discussed previously, we make use of a key concept from phylogenetics—the notion of a quartet. In the previous section, we showed how to estimate accurately distances between internal nodes of the tree. In this section, we explain how to use such estimates to perform topological tests on quartets. Those tests form the basic building blocks of the combinatorial component of the algorithm. As before, we fix a tree T on n leaves. Moreover, we assume that our path metric d on T satisfies $d(e) \geq f$ for all $e \in T$. Note that an upper bound on d is not explicitly required in this section since we assume that appropriate distance estimates are handed to us.

Splits. The routine ISSPLIT in Fig. 9 performs a classic test to decide the correct split of a quartet. It is based on the so-called *Four-Point Method* [4].

Proposition 4.1 guarantees the correctness of ISSPLIT. Its proof is omitted. We consider once again the basic disjoint setup (dangling) of Sect. 3.

Proposition 4.1 (Correctness of ISSPLIT) *Consider the basic disjoint setup (dangling). Let*

$$\mathcal{Q} = \{v_1, w_1, v_2, w_2\},$$

and let \mathcal{D} be the distance matrix on the four nodes of \mathcal{Q} passed to ISSPLIT in Fig. 9. Assume that $d(e) \geq f$ for all edges on the subtree restricted to \mathcal{Q} . If

$$|d(x, y) - \mathcal{D}(x, y)| < \frac{f}{4}, \quad \forall x, y \in \mathcal{Q},$$

then the call $\text{ISSPLIT}((v_1, w_1), (v_2, w_2); \mathcal{D})$ returns TRUE, whereas the calls $\text{ISSPLIT}((v_1, w_2), (v_2, w_1); \mathcal{D})$ and $\text{ISSPLIT}((v_1, v_2), (w_1, w_2); \mathcal{D})$ return FALSE.

Algorithm ISSPLIT

Input: Two pairs of nodes $(v_1, w_1), (v_2, w_2)$; a distance matrix \mathcal{D} on these four nodes;

Output: TRUE or FALSE;

- [Internal Edge Length] Set

$$\nu = \frac{1}{2} (\mathcal{D}(w_1, w_2) + \mathcal{D}(v_1, v_2) - \mathcal{D}(w_1, v_1) - \mathcal{D}(w_2, v_2));$$

(set $\nu = +\infty$ if any of the distances is $+\infty$)

- [Test] If $\nu < f/2$ return FALSE, o.w. return TRUE.

Fig. 9 Routine ISSPLIT

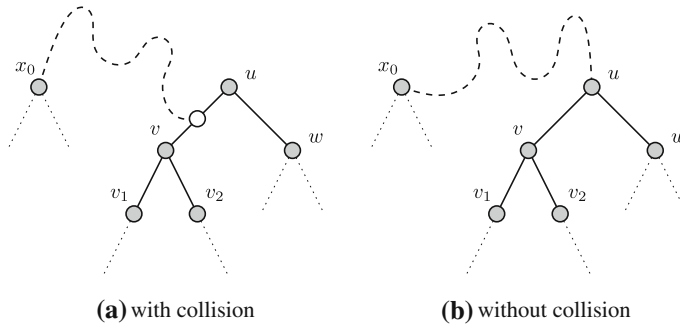


Fig. 10 Basic setup of the ISCOLLISION routine in Fig. 11. The subtrees T_{x_0} and T_u are edge-disjoint. The configuration on the *left* shows a collision, whereas the one on the *right* does not. Note that in the collision case, the internal path of the quartet $\{x_0, w, v_1, v_2\}$ is shorter than the edge (u, v) . This is the basic observation behind the routine ISCOLLISION. We call x_0 the “reference point”

Collisions. As we discussed before, for routines DISTANCEESTIMATE (Fig. 4) and ISSPLIT (Fig. 9) to work, we need a configuration as in Fig. 1 where two edge-disjoint subtrees are connected by a path that lies “above” them. However, for reasons that will be explained in later sections (see also the discussion in Sect. 3), in the course of the reconstruction algorithm we have to deal with configurations where two subtrees are not dangling, as in Fig. 7. We use the following definitions. Recall the basic disjoint setup (general) in Fig. 7.

Definition 4.2 (Collisions) Suppose that T_1 and T_2 are legal subtrees of T and suppose they are *not* dangling. We say that T_1 *collides into* T_2 at edge $e_2 = (u_2, v_2)$ (u_2 is the parent of v_2 in T_2), if the path

$$\text{path}_T(\rho(T_1), \rho(T_2))$$

has non-empty intersection with edge e_2 (i.e., with $\text{path}_T(u_2, v_2)$) but with no other edge in the subtree of T_2 rooted at v_2 . See Fig. 7. We sometimes say that the trees T_1, T_2 *collide*. We say that *the collision is within distance M* if $d(T_1, T_2) \leq M$. (See Definition 3.14.)

An important step of the reconstruction algorithm is to detect collisions, at least when they are within a short distance. Routine ISCOLLISION defined in Fig. 11 and analyzed in Proposition 4.4 below performs this task. We consider the following configuration, which we call the “basic collision setup.”

Definition 4.3 (Basic Collision Setup) We have two legal subtrees T_{x_0} and T_u rooted at x_0 and u in T . We let v, w be the children of u . We assume that we are in either of the configurations depicted in Fig. 10, that is:

- a. Either the path between T_{x_0} and T_u attaches in the middle of the edge (u, v) —in other words, T_{x_0} collides into T_u at edge (u, v) ;
- b. Or it goes through u —in other words, T_{x_0} and T_u are dangling.

We call x_0 the “reference point”.

Algorithm ISCOLLISION

Input: Four nodes x_0, v, w, u ; an edge length h ; a rooted forest and a distance matrix $(\mathcal{F}, \mathcal{D})$;

Output: TRUE or FALSE;

- [Children] Let v_1, v_2 be the children of v in \mathcal{F} (or $v_1 = v_2 = v$ if v is a leaf);
- [Internal Edge Length] Set

$$\nu = \frac{1}{2} (\mathcal{D}(v_1, x_0) + \mathcal{D}(v_2, w) - \mathcal{D}(v_1, v_2) - \mathcal{D}(x_0, w));$$

(set $\nu = +\infty$ if any of the distances is $+\infty$)

- [Test] If $(h - \nu) > f/2$ return TRUE, else return FALSE;

Fig. 11 Routine ISCOLLISION

The purpose of ISCOLLISION is to distinguish between the two configurations above.

Proposition 4.4 (Correctness of ISCOLLISION) *Consider the Basic Collision Setup. In particular, assume that one of the two configurations in Fig. 10 holds. Let $\mathcal{Q} = \{x_0, v_1, v_2, w\}$ and let \mathcal{D} and h be the parameters passed to ISCOLLISION in Fig. 11. Assume that all edges in the tree satisfy $d(e) \geq f$. If*

$$|d(x, y) - \mathcal{D}(x, y)| < \frac{f}{8}, \quad \forall x, y \in \mathcal{Q},$$

and

$$|h - d(u, v)| < \frac{f}{4},$$

then the routine ISCOLLISION returns TRUE if and only if T_{x_0} collides into T_u at edge (u, v) .

Proof Whether or not there is a collision, the quantity ν computed in the routine is the length of the internal path of the quartet split $x_0 w | v_1 v_2$. If there is not a collision, then this path is actually equal to the path corresponding to the edge (u, v) . Otherwise, the length of the path is shorter than $d(u, v)$ by at least f by assumption. (See Fig. 10.) The proof follows. □

5 Reconstruction algorithm

We proceed with a formal description of the reconstruction algorithm. A detailed example can be found in Sect. 5.2. The reader may want to take a look at the example before reading the details of the algorithm.

5.1 Description of the algorithm

Cherry picking. Recall that in a 3-regular tree a *cherry* is a pair of leaves at graph distance 2. Roughly speaking, our reconstruction algorithm proceeds from a simple idea: it builds the tree one layer of cherries at a time. To see how this would work, imagine

that we had access to a “cherry oracle,” that is, a function $C(u, v, T)$ that returns the parent of the pair of leaves $\{u, v\}$ if the latter forms a cherry in the tree T (and say 0 otherwise). Then, we could perform the following “cherry picking” algorithm:

- Currently undiscovered tree: $T' := T$;
- Repeat until T' is empty,
 - For all $(u, v) \in \mathcal{L}(T') \times \mathcal{L}(T')$, if $w := C(u, v, T') \neq 0$, set $\text{Parent}(u) := \text{Parent}(v) := w$;
 - Remove from T' all cherries discovered at this step.

Unfortunately, the cherry oracle cannot be simulated from short sequences at the leaves. Indeed, as we discussed in Sect. 3, short sequences provide only “local” metric information on the structure of the tree. See the example in Sect. 5.2 for an illustration of the problems that may arise. Nevertheless, the above scheme can be roughly followed by making a number of modifications which we now describe.

The high-level idea of the algorithm, which we call **BLINDFOLDED CHERRY PICKING (BCP)**, is to apply the cherry picking scheme above with two important differences:

- [Reconstructed sequences] Leaf sequences provide only local metric information “around the leaves.” To gain information about higher, internal nodes of the tree, we reconstruct sequences at the internal nodes of our partially reconstructed subforest and compute local metric information “around these nodes.” By repeating this process, we gain information about higher and higher parts of the tree.
- [Fake cherries] Moreover, because of the local nature of our information, some of the cherries we pick may in fact turn out *not* to be cherries—that is, they correspond to a path made of more than two edges in the true tree. (See Sect. 5.2 for an example.) As it turns out, this only becomes apparent once a larger fraction of the tree is reconstructed, at which point a subroutine detects the “fake” cherries and removes them.

The full algorithm is detailed in Figs. 12, 13, 14, and 15. We now explain its main components. The parameters k , \widehat{R}_{acc} , and ε will be set in Sect. 6.

Adding a cherry. The algorithm BCP maintains a partially reconstructed subforest of the true tree, or more precisely, a legal subforest \mathcal{F} of T . The main operation we employ to “grow” our partially reconstructed subforest is the merging of two subtrees of \mathcal{F} at their roots—an operation we call “adding a cherry,” in reference to the cherry picking algorithm above. Suppose the current forest \mathcal{F} contains two edge-disjoint legal subtrees T_x and T_y . We merge them by creating a new node z and adding the edges (z, x) and (z, y) as in Fig. 16. We sometimes denote the pair of edges $\{(z, x), (z, y)\}$ by (x, z, y) .

We call this operation *adding cherry* (x, z, y) to \mathcal{F} .

Identifying “local” cherries. As we explained above, we cannot hope to identify with certainty the cherries of the unexplored part of the tree from *short* sequences at the *leaves*. Instead, we settle for detecting what we refer to as “local” cherries, roughly, cherries in a “local” neighbourhood around the roots of the current reconstructed subforest. More precisely, a “local” cherry is a pair of roots of the current subforest that passes a series of tests as detailed below.

Algorithm BLINDFOLDED CHERRY PICKING

Input: Samples at the leaves $\{\sigma_u\}_{u \in [n]}$;

Output: Estimated topology;

• **0) Initialization:**

- [Iteration Counter] $i := 0$;
- [Rooted Subforest] $\mathcal{F}_0 := [n]$;
- [Local Metric] For all $u, v \in \mathcal{F}_0$, set $\mathcal{D}_0(u, v) = \text{DISTANCEESTIMATE}(u, v; \mathcal{F}_0; \widehat{R}_{acc})$;

• **1) Local Cherry Identification:**

- Iteration: i ;
- Set $\mathcal{F}_{i+1} := \mathcal{F}_i$;
- For all $(v_1, w_1) \in \binom{\rho(\mathcal{F}_i)}{2}$,
 - * [Main Step] Compute $(\text{IsCherry}, l_v, l_w) := \text{LOCALCHERRY}((v_1, w_1); (\mathcal{F}_i, \mathcal{D}_i))$;
 - * If $\text{IsCherry} = \text{TRUE}$,
 - [Update Forest] Create new node u_1 and add cherry (v_1, u_1, w_1) to \mathcal{F}_{i+1} ;
 - [Edge Lengths] Set $h(u_1, v_1) := l_v$ and $h(u_1, w_1) := l_w$;

• **2) Collision Removal:**

- [Update Metric] For all $x_1, x_2 \in \rho(\mathcal{F}_{i+1})$, for all $u_\varphi \in T_{x_\varphi}^{\mathcal{F}_{i+1}}$, $\varphi = 1, 2$, set

$$\mathcal{D}_{i+1}(u_1, u_2) = \text{DISTORTEDMETRIC}(u_1, u_2; \mathcal{F}_{i+1}; \{h(e)\}_{e \in \mathcal{F}_{i+1}}; \widehat{R}_{acc}; \varepsilon)$$
- Set $\mathcal{F} := \mathcal{F}_{i+1}$ and $\mathcal{D} := \mathcal{D}_{i+1}$;
- For all $(u_0, u_1) \in \rho(\mathcal{F}) \times \rho(\mathcal{F})$,
 - * Set $\text{HasCollision} := \text{FALSE}$;
 - * If u_1 is not a leaf,
 - [Main Step] Compute $(\text{HasCollision}, z) := \text{DETECTCOLLISION}((u_0, u_1); (\mathcal{F}, \mathcal{D}))$;
 - * If $\text{HasCollision} = \text{TRUE}$,
 - [Update Forest] Compute $\mathcal{F}_{i+1} := \text{REMOVECOLLISION}(z; \mathcal{F}_{i+1})$;
- [Second Pass] Set $\mathcal{F} := \mathcal{F}_{i+1}$ and repeat the previous step;

• **3) Termination:**

- If $|\rho(\mathcal{F}_{i+1})| \leq 3$,
 - * Join nodes in $\rho(\mathcal{F}_{i+1})$ (star if 3, single edge if 2);
 - * Return (tree) \mathcal{F}_{i+1} ;
- Else, set $i := i + 1$, and go to Step 1.

Fig. 12 Algorithm BLINDFOLDED CHERRY PICKING

To determine whether two roots v_1, w_1 of the current forest form a “local” cherry, our routine LOCALCHERRY² in Fig. 13 performs three tests:

1. [Short distance] The nodes v_1, w_1 are at a short distance (roughly $2g$);
2. [Local cherry] For all pairs of roots v_2, w_2 at short distance (roughly $5g$), the quartet

$$\mathcal{Q} = \{v_1, w_1, v_2, w_2\},$$

admits the split $v_1 w_1 | v_2 w_2$;

² In [7], the routine was called CHERRYID.

Algorithm LOCALCHERRY*Input:* Two nodes (v_1, w_1) ; current forest and distance matrix $(\mathcal{F}, \mathcal{D})$;*Output:* Boolean value and length estimates;

- Set $\text{IsCherry} := \text{TRUE}$ and $l_v = l_w = 0$;
- [Short Distance]
 - If $\mathcal{D}(v_1, w_1) > 2g + \varepsilon$, then $\text{IsCherry} := \text{FALSE}$;
- [Local Cherry]
 - Set $\mathcal{N} = \left\{ (v_2, w_2) \in \binom{\rho(\mathcal{F})}{2} : \overline{\mathcal{D}}(\{v_1, w_1, v_2, w_2\}) \leq 5g + \varepsilon \right\}$;
 - If \mathcal{N} is empty, then $\text{IsCherry} := \text{FALSE}$; Else, for all $(v_2, w_2) \in \mathcal{N}$,
 - * If $\text{ISPLIT}((v_1, w_1), (v_2, w_2); \mathcal{D}) = \text{FALSE}$ then set $\text{IsCherry} := \text{FALSE}$ and *break*;
- [Edge Lengths]
 - If $\text{IsCherry} = \text{TRUE}$,
 - * Let x_1, x_2 be the children of v_1 in \mathcal{F} (or let $x_1 = x_2 = v_1$ if v_1 is a leaf);
 - * Let z_0 be the closest node to v_1 in $\rho(\mathcal{F}) - \{v_1, w_1\}$ under \mathcal{D} ;
 - * Set $(b_v, l_v) := \text{ISHORT}((x_1, x_2), (w_1, z_0); \mathcal{F}; \widehat{R}_{\text{acc}}; \varepsilon/16)$;
 - * Repeat previous steps switching the roles of v_1 and w_1 ;
 - * Set $\text{IsCherry} := b_v \wedge b_w$;
- Return $(\text{IsCherry}, l_v, l_w)$;

Fig. 13 Routine LOCALCHERRY**Algorithm DETECTCOLLISION***Input:* Two roots u_0, u_1 ; directed forest and distance matrix $(\mathcal{F}, \mathcal{D})$;*Output:* Boolean value and node;

- Set $\text{HasCollision} := \text{FALSE}$ and $z := 0$;
- Let x_0, y_0 be the children of u_0 in \mathcal{F} ;
- Scan through all nodes v in T_{u_1} (except u_1) in reverse BFS manner,
 - Let $w := \text{Sister}_{\mathcal{F}}(v)$ and $u := \text{Parent}_{\mathcal{F}}(v)$;
 - [Collision Test] Compute

$$b_x := \text{ISCOLLISION}(x_0, v, w, u; h(u, v); (\mathcal{F}, \mathcal{D})),$$
 and

$$b_y := \text{ISCOLLISION}(y_0, v, w, u; h(u, v); (\mathcal{F}, \mathcal{D}));$$
 - If $\text{HasCollision} := b_x \wedge b_y = \text{TRUE}$ then set $z := v$ and *break*;
- Return $(\text{HasCollision}, z)$;

Fig. 14 Routine DETECTCOLLISION

3. [Edge lengths] The edges connecting v_1, w_1 to their (hypothetical) parent are short (roughly g).

The routine LOCALCHERRY has three key properties, proved in Sect. 6:

1. [Edge disjointness] It preserves the *edge-disjointness* of the current forest;
2. [Restricted subforest] It builds a forest that is always a *legal restriction* of the true tree;
3. [Short edges] It guarantees that all edges of the restricted forest are *short* (smaller than g').

Algorithm REMOVECOLLISION


Input: Node v ; rooted forest \mathcal{F} ;

Output: Rooted forest;

- If v is not in \mathcal{F} or v is a root in \mathcal{F} , return \mathcal{F} ;
- Let z_0 be the root of the subtree of \mathcal{F} in which v lies;
- Set $x := v$;
- While $x \neq z_0$,
 - Set $x := \text{Parent}_{\mathcal{F}}(x)$;
 - Remove node x and its adjacent edges below it;
- Return the updated forest \mathcal{F} .

Fig. 15 Routine REMOVECOLLISION

Fig. 16 Adding a cherry (x, z, y) to $\mathcal{F} \supseteq \{T_x, T_y\}$



These properties are crucial for the proper operation of the algorithm (in particular for the correctness of routines such as DISTORTEDMETRIC (Fig. 8) and ISSHORT (Fig. 6) as seen in Sect. 3).

Finally, another important property of LOCALCHERRY is that it is guaranteed to detect *true* cherries—at least those that are “locally witnessed.” We now define this notion more precisely. For a distance matrix \mathcal{D} and a set of nodes \mathcal{N} , we denote

$$\overline{\mathcal{D}}(\mathcal{N}) = \max \{ \mathcal{D}(x, y) : \{x, y\} \subseteq \mathcal{N} \}.$$

Definition 5.1 (Witnessed Cherry) Let \mathcal{F} be a forest with path metric \mathcal{D} . We say that a pair of leaves $\{u, v\}$ is an M -witnessed cherry in \mathcal{F} if $\{u, v\}$ is a cherry in \mathcal{F} and there are at least two other leaves u', v' s.t.

$$\overline{\mathcal{D}}(\mathcal{Q}) \leq M,$$

where $\mathcal{Q} = \{u, v, u', v'\}$ (the leaves u', v' will act as “witnesses” of the cherry $\{u, v\}$).

Detecting collisions. The merging of subtrees through “local” cherries that are *not* actual cherries eventually results in *collisions* between subtrees of the current forest




Fig. 17 The REMOVECOLLISION in Fig. 15 routine removes all *white nodes* and *dashed edges*. The *dotted curvy line* indicates the location of a collision to be removed

such as in Fig. 7. Such configurations are undesirable since they do not allow to complete the reconstruction by simple merging of the subtrees at their roots. Therefore, we seek to detect these collisions using the routine³ in Fig. 14.

After adding a new layer of “local” cherries, DETECTCOLLISION checks whether new collisions can be found. For this, we use routine ISCOLLISION in Fig. 11 from Sect. 4. We actually perform *two* such ISCOLLISION tests for each candidate edge in the target subtree T_{u_1} (see Fig. 14). This is done precisely for the same reasons that we did a “multiple test” in routine DISTORTEDMETRIC (Fig. 8) in Sect. 3: it guarantees that at least one test is performed under the basic disjoint setup (dangling), which is needed for ISCOLLISION to be correct. See Sect. 6 for details. Also, to make sure that we are always in either of the configurations in Fig. 10, we scan the nodes of the target subtree T_{u_1} in “reverse breath-first search (BFS) order”:

1. Order the vertices in subtree T_{u_1} according to breath-first search v_1, \dots, v_s ;
2. Scan the list in reverse order v_s, \dots, v_1 .

This way, we never find ourselves “above” a collision in T_{u_1} (because otherwise we would have encountered the collision first).

Removing collisions. Once a collision is identified we remove it using the routine in Fig. 15.⁴ As seen in Fig. 17, the routine essentially removes the path from the collision all the way to the root of the corresponding subtree.

For a rooted forest \mathcal{F} , we use the notation $\text{Parent}_{\mathcal{F}}(x)$ to indicate the unique parent of node x .

³ In [7], the routine was called FAKECHERRY.

⁴ In [7], the routine was called BUBBLE.

We note that we actually go through the “Collision Identification/Removal” step *twice* to make sure that we detect collisions occurring on edges adjacent to the roots. See the proof of Lemma 6.5.

We prove in Proposition 7.4 below that, at each iteration of BCP, at least one cherry is added that will not be removed at a later stage. Hence, the algorithm makes progress at every iteration and eventually recovers the full tree.

Implementation. We argue that the running time of the algorithm is $O(n^5k)$ and that with the use of appropriate data structures it can be reduced further to $O(n^3k)$. Let us start with the naive analysis. The distance estimations in the initialization step of the algorithm take overall time $O(n^2k)$, since there are $O(n^2)$ pairs of leaves and each distance estimation between leaves takes linear time in the sequence length k . Now, in every iteration of the algorithm:

- The Local Cherry Identification step takes overall time $O(n^4k)$, since we consider $O(n^2)$ pairs of roots in the for-loop of this step, and each call of LOCALCHERRY requires $O(n^2 + nk)$ time— $O(n^2)$ time for the [Local Cherry] step and $O(nk)$ time for the [Edge Lengths] step, in which the ISSHORT call involves $O(1)$ ancestral sequence reconstructions on trees of size $O(n)$.
- The Collision Removal step requires $O(n^3k)$ time in each iteration. Indeed, it performs $O(n^2)$ distance computations and each of these takes $O(nk)$ time, since it requires $O(1)$ ancestral sequence reconstructions on trees of size $O(n)$. It also performs $O(n^2)$ calls to DETECTCOLLISION and REMOVECOLLISION, each of which is a linear time operation on trees of size $O(n)$.

Since each iteration requires $O(n^4k)$ time and there are $O(n)$ iterations (see proof in Sect. 7), the overall running time is $O(n^5k)$. The above analysis is wasteful in (1) not reusing the already reconstructed ancestral sequences and in (2) performing various tests on pairs of nodes that are far apart in the tree. With the use of appropriate data structures, we could keep track of the “local” neighborhood of each node and restrict many computations to these neighborhoods. This should bring down the running time to $O(n^3k)$ with a constant that would depend explicitly on f, g . The details are left to the reader.

5.2 Example

We give a detailed example of the execution of the algorithm. Consider the tree depicted in Fig. 18a.

It is made of a large complete binary tree (on the right) with a small 3-level complete binary tree attached to its root (on the left). All edges have length g , except (v, x) , (x, v') and the edges adjacent to the root of the small tree which have length $g/2$. In the figure, the subtree currently discovered by BCP is made of solid arrows and full circles. The remaining (undiscovered) tree is in dotted lines and empty circles. Assume that the length of the sequences at the leaves allows us to estimate accurately distances up to $5g$ (the actual constants used by the algorithm can be found later).

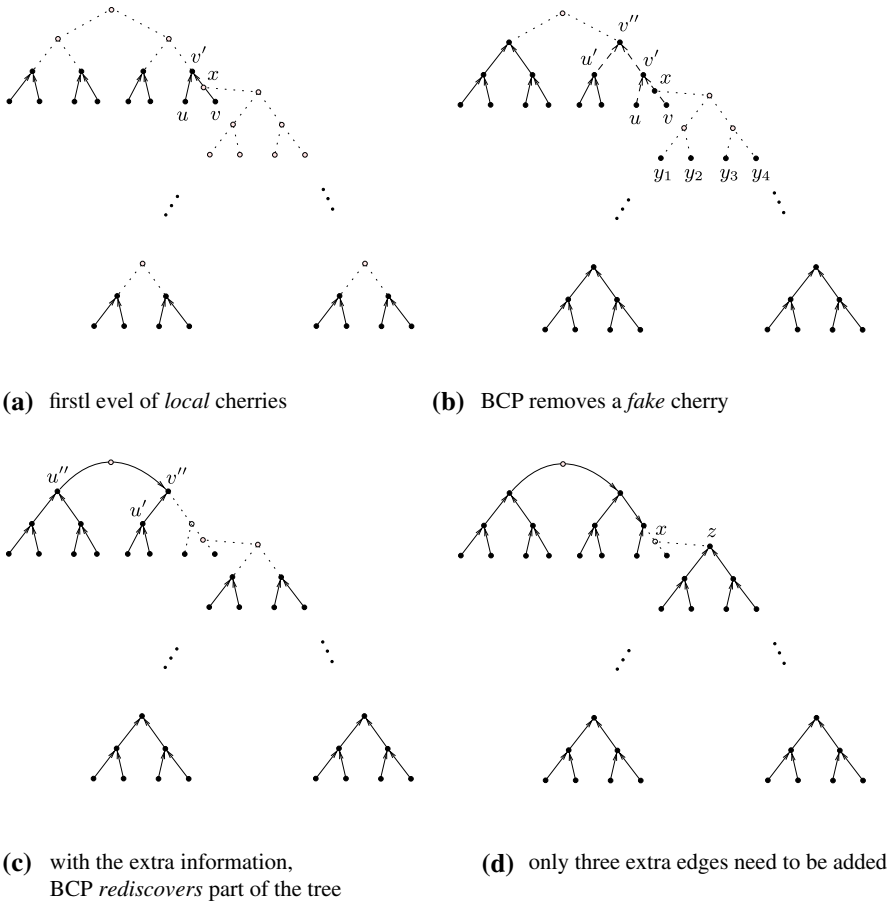


Fig. 18 Illustration of BCP's unraveling

First level. We first join into cherries all pairs of leaves that “look” like g -cherries in this “local” metric. We are guaranteed to find all true g -cherries. However, consider pairs of leaves such as u, v for which there is no “local evidence” that it does *not* form a cherry. Even though u, v is not a cherry, it is joined into a cherry by BCP. Figure 18a depicts the current forest after the first iteration of BCP. Before proceeding further, we apply our estimator $\widehat{\text{Anc}}$ from Sect. 2 to obtain reconstructed sequences at the roots of the current forest and recompute the “local” metric.

Removing a fake cherry. We subsequently proceed to join “local” g -cherries one layer at a time, reconstructing internal sequences as we do so. After many iterations, we find ourselves in the situation of Fig. 18b where most of the large complete tree has been reconstructed (assume for now that edges (u', v'') , (v', v'') , (u, v') , (v, v') represented in dashed lines are present). Now, the new information coming from sequences at y_1, \dots, y_4 provides evidence that (u, v', v) is *not* a cherry and that there is in fact a

node x on edge (v, v') . For example, the quartet $\{y_1, y_2, u, v\}$ suggests that u, v forms a cherry with a $3g/2$ -edge, which cannot hold in a true cherry. At this point, we remove the “fake” cherry (u, v', v) as well as all cherries built upon it, here only (u', v'', v') . Note that we have removed parts of the tree that were in fact reconstructed correctly (e.g., the path between u and u').

Rediscovering removed parts. Subsequently, BCP continues to join “local” cherries and “rediscovers” the parts of the tree that were removed. For instance, in Fig. 18c, the edge (u', v'') is reconstructed again but this time it forms a cherry with (u'', v'') rather than (v', v'') .

Final step. Eventually, the full tree is correctly reconstructed except maybe for a few (at most 3) remaining edges. Those can be added separately. For example in Fig. 18d only the three edges around x remain to be uncovered. Note that the reconstructed tree has a root which is different from that of the original tree.

6 Analysis I: induction step

In this section and the next, we establish that BCP reconstructs the phylogeny correctly. In this section, we establish a number of combinatorial properties of the current forest \mathcal{F}_i grown by BCP. Then, in the next section, we prove that the “correctly reconstructed subforest” of \mathcal{F}_i increases in size at every iteration.

Parameters. Let $\delta > 0$ be the error tolerance. Each application of Proposition 3.15 has an error of $O(n^{-\gamma})$. We will do $O(n^3)$ distance estimations so that by the union bound we require $O(n^{3-\gamma}) \leq \delta$. Let $\varepsilon < \frac{1}{8} \min\{f, g' - g\}$ and $\widehat{R}_{\text{col}} > 6g$. Set $M > \widehat{R}_{\text{col}} + 4g'$; take $\widehat{R}_{\text{acc}} > M + 2B(g') + 4g'$; and choose k to be equal to the maximum sequence length requirement for Proposition 3.15 with parameters ε, M and \widehat{R}_{acc} , for Proposition 3.2 with parameters ε and M , and for Proposition 3.13 with parameters $\varepsilon/16, M$ and \widehat{R}_{acc} .

Induction step. The following proposition establishes a number of properties of the forest grown by BCP. We assume here that the conclusion of Proposition 3.15 holds at every iteration of the algorithm, until the tree is fully recovered. In the next section, we will prove that the latter is indeed true with high probability.

Proposition 6.1 (Properties of \mathcal{F}_i) *Denote by $\mathcal{F}_i = \{T_u : u \in \rho(\mathcal{F}_i)\}$ the current forest at the beginning of BCP’s i -th iteration. Also, assume that the conclusion of Proposition 3.15 holds for each call of the routine DISTORTEDMETRIC throughout the execution of the algorithm. Then, $\forall i \geq 0$ (until the algorithm stops),*

1. [Legal subforest] \mathcal{F}_i is a legal subforest of T ;
2. [Edge lengths] $\forall u \in \rho(\mathcal{F}_i)$, T_u has edge lengths at most g' ;
3. [Weight estimation] The estimated lengths $\{h(e)\}_{e \in \mathcal{E}(\mathcal{F}_i)}$ of the edges in \mathcal{F}_i are within $\varepsilon/16$ of their right values;
4. [Collisions] There is no collision within distance \widehat{R}_{col} . (See Definition 4.2.)

Proof of Proposition 6.1 The proof is by induction on i .

$i = 0$: The set $\rho(\mathcal{F}_0)$ consists of the leaves of T . The claims are therefore trivially true.

$i > 0$: Assume the claims are true at the beginning of the j -th iteration for all $j \leq i$. By doing a step-by-step analysis of the i -th iteration, we show that the claims are still true at the beginning of the $(i + 1)$ -st iteration.

We first analyze the routine LOCALCHERRY (Fig. 13). For a legal subforest \mathcal{F} of T , we denote the “remaining” forest by

$$\tilde{\mathcal{F}} = T - \mathcal{F}.$$

More precisely, if $\mathcal{F} = \{T_1, \dots, T_\alpha\}$ then $\tilde{\mathcal{F}}$ is the forest obtained from T as follows:

1. Remove all edges of T in the union of the trees T_1, \dots, T_α . In particular, for those edges of a T_i representing a path in T , we remove all corresponding edges of T .
2. The nodes of $\tilde{\mathcal{F}}$ are all the endpoints of the remaining edges of T . All other nodes of T are discarded.

Note that the set $\tilde{\mathcal{F}}$ is in fact a subforest of T .

Lemma 6.2 (Local Cherry Identification) *Let \mathcal{F}_i be the current forest at the beginning of the i -th iteration. Then we have the following:*

- If $\{v_1, w_1\}$ is a $5g$ -witnessed cherry in $\tilde{\mathcal{F}}_i$, then it passes all tests in LOCALCHERRY;
- If $\{v_1, w_1\}$ passes all the tests in LOCALCHERRY, then

$$d(u_1, v_1) \leq g + 2\varepsilon,$$

and

$$d(u_1, w_1) \leq g + 2\varepsilon,$$

where u_1 is the parent of $\{v_1, w_1\}$ (as defined by LOCALCHERRY and the [Update Forest] step in the main loop). Moreover the length estimates satisfy

$$|d(u_1, v_1) - l_v| < \frac{\varepsilon}{16},$$

and

$$|d(u_1, w_1) - l_w| < \frac{\varepsilon}{16}.$$

- If two pairs $\{v_1, w_1\}$ and $\{v_2, w_2\}$ both pass all tests in LOCALCHERRY, then it must be that the paths $\text{path}_T(v_1, w_1)$ and $\text{path}_T(v_2, w_2)$ are non-intersecting.

Proof Suppose first that $\{v_1, w_1\}$ is a $5g$ -witnessed cherry in $\tilde{\mathcal{F}}_i$ with witness $\{v_2, w_2\}$. Since there is no collision within $\widehat{R}_{\text{col}} > 6g$ by assumption, we are in the dangling case. Moreover, all edge weights below $\{v_1, w_1, v_2, w_2\}$ have been estimated within

$\varepsilon/16$ and all corresponding edge weights (possibly corresponding to paths) are $\leq g'$. Therefore, by Propositions 3.15 and 4.1, DISTORTEDMETRIC (Fig. 8) is accurate within ε , ISSPLIT (Fig. 9) returns the correct answer and ISSHORT (Fig. 6) is accurate within $\varepsilon/16$. Hence, $\{v_1, w_1\}$ passes the three tests in LOCALCHERRY.

Conversely, suppose $\{v_1, w_1\}$ passes all tests in LOCALCHERRY. By our assumptions, when DISTORTEDMETRIC returns a finite value, it is accurate within ε . Let z_0 be as in Fig. 13 (that is, z_0 is the closest node to v_1 in $\rho(\mathcal{F}) - \{v_1, w_1\}$ under the distorted metric). The fact that $\{v_1, w_1\}$ previously passed the [Local Cherry] test implies in particular that the diameter of $\{v_1, w_1, z_0\}$ must be less than distance $5g + 2\varepsilon$. In particular, there is no collision between the subtrees rooted at v_1, w_1 , and z_0 since $\widehat{R}_{\text{col}} > 6g$. Let u_1 be the intersection of $\{v_1, w_1, z_0\}$. Then, ISSHORT returns an estimate within $\varepsilon/16 < \varepsilon$ which in turn implies that $d(u_1, v_1)$ and $d(u_1, w_1)$ are $\leq g + 2\varepsilon$.

For the third part, assume by contradiction that the paths $\text{path}_T(v_1, w_1)$ and $\text{path}_T(v_2, w_2)$ intersect. Then by the triangle inequality, the diameter of $\{v_1, w_1, v_2, w_2\}$ is at most $5g$. In particular, when LOCALCHERRY is applied to $\{v_1, w_1\}$, the pair $\{v_2, w_2\}$ is considered in the [Local Cherry] test and, since there is no collision within $\widehat{R}_{\text{col}} > 6g$, ISSPLIT correctly returns FALSE. That is a contradiction. \square

We can now prove Claims 1, 2, and 3 of Proposition 6.1.

Lemma 6.3 (Claims 1, 2, and 3) *Let \mathcal{F}_{i+1} be the current forest at the beginning of the $(i + 1)$ -st iteration. The Claims 1, 2, and 3 of the induction hypothesis hold for \mathcal{F}_{i+1} .*

Proof Since REMOVECOLLISION (Fig. 15) only removes edges from the current forest, it is enough to prove that after the completion of the Local Cherry Identification step the resulting forest satisfies Claims 1, 2, and 3.

Claim 1. By the induction hypothesis, \mathcal{F}_i is legal. We only need to check that \mathcal{F}_{i+1} is edge-disjoint. Suppose on the contrary that the forest is not edge-disjoint. Also, suppose that, along the execution of the Local Cherry Identification step, the forest stops being edge-disjoint when cherry (v_1, u_1, w_1) is added to \mathcal{F}_i . Then one of the following must be true:

1. There is an “old” root $z \in \rho(\mathcal{F}_i)$ such that $\text{path}_T(v_1, w_1)$ is edge-sharing with $T_z^{\mathcal{F}_i}$. But then there is a collision in \mathcal{F}_i within distance $2g + 2\varepsilon$ which contradicts the induction hypothesis (Claim 4).
2. There is a “new” root $z \in \rho(\mathcal{F}_{i+1}) \setminus \rho(\mathcal{F}_i)$ with corresponding cherry (x, z, y) such that $\text{path}_T(v_1, w_1)$ is edge-sharing with $\text{path}_T(x, y)$. In that case $v_1 w_1 | x y$ is not the correct split and

$$\bar{d}(\{v_1, w_1, x, y\}) \leq 5g,$$

by the triangle inequality and Proposition 3.15. But then, by Lemma 6.2, LOCALCHERRY (Fig. 13) rejects $\{v_1, w_1\}$ when performing the [Local Cherry] test with witness $\{x, y\}$ —a contradiction.

Claim 2 and 3. Follows from the induction hypothesis and Lemma 6.2. \square

It remains to prove Claim 4 of Proposition 6.1. This follows from the following analysis of DETECTCOLLISION (Fig. 14). Note that, since Claim 4 holds for all iterations $j \leq i$, it must be the case that any new collision between two trees involves

at least one of the new edges of these trees added in the *Local Cherry Identification* step. We call an edge *deep* if it is not adjacent to a root in the current forest. Otherwise we call the edge a *top* edge. We show that the first pass of the *Collision Removal* step removes all collisions into deep edges. At the beginning of the second pass, all collisions (if any) must be into top edges. We show that the second pass cleans those up. We first prove that there are no false positives in the Collision Removal step.

Lemma 6.4 (Collision Removal: No False Positive) *Let \mathcal{F}_{i+1} be the current forest at the beginning of the first or second pass of the Collision Removal step of the i -th iteration, and let $u_0, u_1 \in \rho(\mathcal{F}_{i+1,1})$ be the roots of the trees $T_0 = T_{u_0}^{\mathcal{F}_{i+1,1}}$ and $T_1 = T_{u_1}^{\mathcal{F}_{i+1,1}}$. Let v be a node in T_1 , and suppose that T_0 does not collide into T_1 below v or on the edge immediately above it. Then no collision is detected in the corresponding step of DETECTCOLLISION.*

Proof Let x_0, y_0 be the children of u_0 . It suffices to show that either b_x is FALSE or b_y is FALSE (see Fig. 14). Without loss of generality we can assume that the path connecting u_0 to u_1 does not pass through x_0 . In particular, we are in the case b) of Fig. 10. If any of the distances passed to ISCOLLISION is $+\infty$, ISCOLLISION returns FALSE. Otherwise, by our assumption on the output of DISTORTEDMETRIC, the assumptions of Proposition 4.4 are satisfied. Hence, ISCOLLISION returns FALSE in that case as well. □

Lemma 6.5 (Collision Removal) *The first and second passes of the Collision Removal step satisfy:*

1. *Let $\mathcal{F}_{i+1,1}$ be the current forest at the beginning of the first pass of the Collision Removal step of the i th iteration, and let $u_0, u_1 \in \rho(\mathcal{F}_{i+1,1})$. Suppose $T_0 = T_{u_0}^{\mathcal{F}_{i+1,1}}$ collides into $T_1 = T_{u_1}^{\mathcal{F}_{i+1,1}}$ within distance \widehat{R}_{col} on a deep edge $e = (u, v)$ of T_1 . Then DETECTCOLLISION in Fig. 14 correctly detects the collision.*
2. *Let $\mathcal{F}_{i+1,2}$ be the current forest at the beginning of the second pass of the Collision Removal step of the i -th iteration, and let $u_0, u_1 \in \rho(\mathcal{F}_{i+1,2})$. Suppose $T_0 = T_{u_0}^{\mathcal{F}_{i+1,2}}$ collides into $T_1 = T_{u_1}^{\mathcal{F}_{i+1,2}}$ within distance \widehat{R}_{col} . Then DETECTCOLLISION correctly detects the collision.*

Proof 1. Denote by x_φ, y_φ the children of $u_\varphi, \varphi = 0, 1$. Consider the Basic Collision Setup of Sect. 4. By the remark above the statement of the lemma, the path coming from e enters T_0 through a top edge, or at u_0 . (See Fig. 19)


Let

$$A_{0 \rightarrow 1} = \{z \in \mathcal{V}(T_1) : e \text{ is not in the subtree of } T_1 \text{ rooted at } z\},$$

and

$$B_{0 \rightarrow 1} = \mathcal{V}(T_1) \setminus A_{0 \rightarrow 1}.$$

Fig. 19 Configuration in the proof of Lemma 6.5



Denote by z the current-node variable used by DETECTCOLLISION as it scans the tree T_1 in a reverse BFS manner. Observe that, for all $z \in A_{0 \rightarrow 1} - \{v\}$, the Basic Collision Setup of Proposition 4.4 holds for both x_0 and y_0 . Hence, by Lemma 6.4, ISCOLLISION in Fig. 11 returns the correct answer. Furthermore, in the case of v , ISCOLLISION returns TRUE for both x_0 and y_0 since the collision is within \widehat{R}_{col} and

$$\bar{d}(\{x_0, y_0, v_1, v_2, w\}) \leq \widehat{R}_{\text{col}} + 4g' \leq M,$$

by Claim 2. For w , ISCOLLISION returns FALSE because in that case $h(u, w) - v \leq -f + 3\epsilon \leq f/2$, where v is the estimated length of the internal edge of $\{x_0, w_1, w_2, v\}$ or $\{y_0, w_1, w_2, v\}$ with w_1, w_2 the children of w . Finally, by scanning the nodes in reverse BFS order, we guarantee that v and w are encountered before any node in $B_{0 \rightarrow 1}$. Hence, DETECTCOLLISION identifies correctly the collision on edge e .

2. Observe that, after a call to REMOVECOLLISION, the set of edges in the remaining forest is a subset of what it used to be. In particular, the subset of these edges involved in a collision decreases in size. Moreover, from the first part of the lemma, at the end of the first pass there is no collision remaining into deep edges. Given this, the argument above implies that any remaining collision within distance \widehat{R}_{col} will be found and removed in the second pass. □

7 Analysis II: tying it all together

In the previous section, we showed that provided Proposition 3.15 holds at every iteration the forest built is “well-behaved.” Below, we finish the proof of our main theorem by showing that Proposition 3.15 indeed holds until termination and that the algorithm eventually converges. We also discuss the issues involved in extending $O(\log n)$ -reconstruction beyond the Δ -Branch Model in Sect. 7.3.

7.1 Quantifying progress

Our main tool to assess the progress of the algorithm is the notion of a fixed subforest—in words, a subforest of the current forest that will not be modified again under the normal operation of the algorithm.

Definition 7.1 (Fixed Subforest) Let \mathcal{F} be a legal subforest of T . Let $u \in \mathcal{V}(\mathcal{F})$. We say that u is fixed if $T_u^{\mathcal{F}}$ is fully reconstructed, or in other words, $T_u^{\mathcal{F}}$ can be obtained from T by removing (at most) one edge adjacent to u and keeping the subtree containing u . Note that descendants of a fixed node are also fixed. We denote by \mathcal{F}^* the rooted subforest of \mathcal{F} made of all fixed nodes of \mathcal{F} . We say that \mathcal{F}^* is the maximal fixed subforest of \mathcal{F} .

Let $\mathcal{F}_i = \{T_1, \dots, T_\alpha\}$ be the current forest at iteration i with remaining forest $\tilde{\mathcal{F}}_i = \{T'_1, \dots, T'_\beta\}$. Assuming the conclusion of Proposition 6.1 holds, each leaf v in $\tilde{\mathcal{F}}_i$ satisfies exactly one of the following:

- *Fixed root*: v is a root of a fully reconstructed tree $T_a \in \mathcal{F}_i$ (that is, T_a is in \mathcal{F}_i^*);
- *Colliding root*: v is a root of a tree $T_a \in \mathcal{F}_i$ that contains a collision (that is, T_a is not in \mathcal{F}_i^*);
- *Collision node*: v belongs to a path connecting two vertices in $T_a \in \mathcal{F}_i$ but is not the root of T_a (that is, it lies in the “middle” of an edge of T_a).

Note in particular that the fixed roots of $\tilde{\mathcal{F}}_i$ are roots in \mathcal{F}_i^* (although not all roots of the maximal fixed subforest \mathcal{F}_i^* are fixed roots in $\tilde{\mathcal{F}}_i$ as they may lie below a colliding root). We also need the notion of a fixed bundle—in words, a cherry (along with witnesses) that will be picked by the algorithm at the next iteration and remain until termination.

Definition 7.2 (Fixed Bundle) A bundle in $\tilde{\mathcal{F}}_i$ is a group of four leaves in $\tilde{\mathcal{F}}_i$ such that:

- Any two leaves in the bundle are at topological distance at most 5 in T ;
- It includes at least one cherry of $\tilde{\mathcal{F}}_i$.

A fixed bundle is a bundle in $\tilde{\mathcal{F}}_i$ whose leaves are fixed roots.

The following lemma is the key to our convergence argument. It ensures that a fixed bundle always exists and hence that progress is made at every iteration.

Proposition 7.3 (Existence of a Fixed Bundle) Assume $\tilde{\mathcal{F}}_i$ has at least 4 leaves and satisfies the conclusion of Proposition 6.1. Then, $\tilde{\mathcal{F}}_i$ contains at least one fixed bundle.

Proof To avoid confusion between the forests \mathcal{F}_i and $\tilde{\mathcal{F}}_i$, we will refer to the forest $\tilde{\mathcal{F}}_i$ as the anti-forest, to its trees as anti-trees and to its leaves as anti-leaves. We first make a few observations:

1. The anti-forest $\tilde{\mathcal{F}}_i$ is binary, that is, all its nodes have degrees in $\{1, 3\}$. Indeed, note first that since T is binary, one cannot obtain nodes of degree higher than 3 by removing edges from T . Assume by contradiction that there is a node, u , of degree 2 in $\tilde{\mathcal{F}}_i$. Let w_1, w_2, w_3 be the neighbors of u in T and assume that

- $w_3 \notin \tilde{\mathcal{F}}_i$. By construction of the anti-forest $\tilde{\mathcal{F}}_i$ the edge (u, w_3) is in the forest \mathcal{F}_i (possibly included in an edge of \mathcal{F}_i corresponding to a path of T). Moreover, u is a node of degree 1 in \mathcal{F}_i . But this is a contradiction because the only nodes of degree 1 in a legal subforest of T are the leaves of T and u cannot be a leaf of T .
2. A binary tree (or anti-tree) T_0 with 4 or more leaves (or anti-leaves) contains at least one bundle. Indeed, let L_0 be the leaves of T_0 . We call $R_0 = L_0$ the *level-0 leaves* of T_0 . Now remove all cherries from T_0 and let R_1 be the roots of the removed cherries—the *level-1 leaves*. Denote by T_1 the tree so obtained and note that its leaves L_1 contain R_1 as well as some remaining level-0 leaves. Consider the cherries of T_1 (there is at least one). If any such cherry is made of two level-1 leaves, then the corresponding (descendant) level-0 leaves of T_0 form a bundle in T_0 and we are done. (In that case, the diameter of the bundle is 4.) Suppose there is no such cherry. Note that there is no cherry in T_1 formed from two level-0 leaves as those were removed in constructing T_1 . Hence, all remaining cherries of T_1 must contain at least one level-1 leaf. Now, merge all such cherries to obtain T_2 . Denote by R_2 the roots of the removed cherries—the *level-2 leaves*. Once again, by the argument above, any cherry of T_2 contains at least one level-2 leaf. Any such cherry (there is at least one) provides a bundle by considering its descendant level-0 leaves in T_0 . (If the second leaf in the cherry is level-0, the diameter of the bundle is 4. If it is level-1 or level-2, the diameter is 5.) This proves the claim.
 3. By Claim 4 in Proposition 6.1, there is no collision within $\hat{R}_{\text{col}} > 6g$ (see Definition 4.2). In particular, collision nodes are at distance at least $\hat{R}_{\text{col}} > 6g$ from any other anti-leaf (collision nodes, fixed roots, colliding roots) in $\tilde{\mathcal{F}}_i$. Therefore, if an anti-tree in $\tilde{\mathcal{F}}_i$ contains a collision node, then it has ≥ 4 anti-leaves and, from the previous observation, it contains at least one bundle. Moreover, this bundle cannot contain a collision node since in a bundle all anti-leaves are at distance at most $5g$ and collision nodes are at distance at least $\hat{R}_{\text{col}} > 6g$ from all other anti-leaves in $\tilde{\mathcal{F}}_i$.
 4. From the previous observations, we get the following: if a tree in $\tilde{\mathcal{F}}_i$ contains a collision, then either it has a fixed bundle or it has at least one colliding root.

We now proceed with the proof. Assume first that there is no collision node in $\tilde{\mathcal{F}}_i$. Then, there cannot be any colliding root either because by definition a colliding root is the root of a tree containing a collision. In particular, $\tilde{\mathcal{F}}_i$ is composed of a single anti-tree whose anti-leaves are all fixed roots. Then, since by assumption $\tilde{\mathcal{F}}_i$ has at least 4 anti-leaves, there is a fixed bundle by Observation 2 above.

Assume instead that there is a collision node. Let \tilde{T}^0 be an anti-tree in $\tilde{\mathcal{F}}_i$ with such a collision node, say c^0 . Then by Observation 4, either 1) \tilde{T}^0 has a fixed bundle in which case we are done, or 2) one of \tilde{T}^0 's anti-leaves is a colliding root, say r^0 . In the latter case, let T^0 be the tree in \mathcal{F}_i whose root is r^0 . The tree T^0 contains at least one collision node (included in an edge corresponding to a path of T). This collision node, say c^1 , is also contained in an anti-tree in $\tilde{\mathcal{F}}_i$, say \tilde{T}^1 . Repeat the argument above on \tilde{T}^1 , and so on.

We claim that this process generates a simple path P in T starting from the node c^0 above, passing through an alternating sequence of colliding roots and collision nodes $r^0, c^1, r^1, c^2, \dots$, and eventually reaching a fixed bundle. Indeed, since there is no

cycle in T , the only way for P not to be simple is for it to “reverse on itself.” But note that by definition, for all j we have $c^j \neq r^j$ and $r^j \neq c^{j+1}$. Moreover the simple paths $c^j \rightarrow r^j$ and $r^j \rightarrow c^{j+1}$ belong respectively to the anti-forest $\tilde{\mathcal{F}}_i$ and the forest \mathcal{F}_i (possibly with subpaths collapsed into edges). In particular, their edges (in T) cannot intersect. Hence, P is simple. Finally, since T is finite, this path cannot be infinite, and we eventually find a fixed bundle. \square

7.2 Proof of the main theorem

Consider now the Δ -Branch Model. The proof of convergence works by considering first the hypothetical case where all distance estimates computed by the algorithm are perfectly accurate, that is the case where we have “perfect” local information. We denote by $\{(\mathcal{F}_0)_i\}_{i \geq 0}$ the sequence of forests built under this assumption. Note that, up to arbitrary choices (tie breakings, orderings, etc.), this sequence is *deterministic*. We now show that it terminates in a polynomial number of steps with the correct tree.

Proposition 7.4 (progress under perfect local information) *Assume $(\mathcal{F}_0)_i = \{T_u : u \in \rho((\mathcal{F}_0)_i)\}$ is the current forest at the beginning of BCP’s i th iteration under perfect local information with corresponding maximal fixed subforest $(\mathcal{F}_0)_i^*$. Then for all $i \geq 0$ (before the termination step), $(\mathcal{F}_0)_i^* \subseteq (\mathcal{F}_0)_{i+1}^*$ and $|\mathcal{V}((\mathcal{F}_0)_{i+1}^*)| > |\mathcal{V}((\mathcal{F}_0)_i^*)|$.*

Proof of Proposition 7.4 We first argue that $(\mathcal{F}_0)_i^* \subseteq (\mathcal{F}_0)_{i+1}^*$. Note that the only routine that removes edges is REMOVECOLLISION in Fig. 15. Since REMOVECOLLISION only removes edges above identified collisions and $(\mathcal{F}_0)_i^*$ is fully reconstructed, it suffices to show that collisions identified by DETECTCOLLISION in Fig. 14 are actual collisions. This follows from Lemma 6.4.

To prove $|\mathcal{V}((\mathcal{F}_0)_{i+1}^*)| > |\mathcal{V}((\mathcal{F}_0)_i^*)|$, assume $(\mathcal{F}_0)_i = \{T_1, \dots, T_\alpha\}$ with remaining forest $(\tilde{\mathcal{F}}_0)_i = \{T'_1, \dots, T'_\beta\}$. From Proposition 7.3, it follows that $(\tilde{\mathcal{F}}_0)_i$ contains at least one fixed bundle. This immediately implies the second claim. Indeed, by Lemma 6.2 note that the cherry in the fixed bundle is found by LOCALCHERRY in Fig. 13 during the $(i + 1)$ st iteration and is not removed by the *Collision Removal* step from the argument above. \square

Proof of Theorem 1.2 To prove our main theorem under the Δ -Branch Model, we modify our reconstruction algorithm slightly by rounding the estimates in Proposition 3.15 to the closest multiple of Δ . Also, we choose a number of samples large enough so that the distance estimation error is smaller than $\Delta/3$. In that case, *we simply mimic the algorithm under perfect local information*. Note that by Proposition 7.4 there are at most $O(n)$ iterations until termination under perfect local information. By a union bound, it follows that Proposition 3.15 holds true for all pairs of subtrees in $\{(\mathcal{F}_0)_i\}_{i \geq 0}$ with high probability.

We can now conclude the proof. By Proposition 6.1, the current forest at each iteration is correctly reconstructed. By Proposition 7.4 after $O(n)$ iterations there remain at most three nodes in $\rho(\mathcal{F}_i)$ and at that point, from Proposition 6.1 Claim 4, we have that $\mathcal{F}_i \equiv \mathcal{F}_i^*$. Therefore the remaining task is to join the remaining roots and there

is only one possible topology. So when the BCP algorithm terminates, it outputs the tree T (as an undirected tree) with high probability.

The tightness of the value $g^* = \frac{\ln 2}{4}$ is justified by the polynomial lower bound [22] on the number of required characters if the mutation probability p on all edges of the tree satisfies $2(1 - 2p)^2 < 1$. \square

7.3 Beyond the Δ -branch model?

Extending Theorem 1.2 to *continuous* edge lengths appears far from trivial. The issue arises in the final union bound over all applications of Proposition 3.15 (more specifically, the ancestral state reconstruction step) which is valid except with inverse polynomial probability over the generated sequences (for a fixed subtree of T). Indeed, note that in general there are *super-polynomially* many restricted subtrees of the true tree T where all edges (paths in T) are shorter than g^* . Therefore, using only a simple union bound, we cannot hope to guarantee that ancestral state reconstruction is successful *simultaneously* on all possible partial reconstructed subtrees.

In the previous subsection, we avoided this problem by showing that under the Δ -BM assumption the algorithm follows a *deterministic* realization path of polynomial length. Moving beyond this proof would require proving that the ancestral state reconstruction can be performed on the *random* forests generated by the algorithm—but this is not straightforward as the partially reconstructed forest is generated by the same data that is used to perform the ancestral estimation. We conjecture that the correlation created by this “bootstrap” is mild enough to allow our algorithm to work in general, but we cannot provide a rigorous proof at this point.

We remark that Mossel’s earlier result [22] on the balanced case with continuous edge lengths is not affected by the issue above because, there, the reconstruction of the tree occurs one level at a time (there is no collision). Hence, ancestral state reconstruction is performed only on *fully* reconstructed subtrees—of which there are only polynomially many.

Note finally that, even under the discretization assumption made in this paper, achieving $\log(n)$ -reconstruction is nontrivial and does not follow from previous techniques. In particular, it can be shown that all previous rigorous reconstruction algorithms for general trees require polynomial sequence lengths even when all edge lengths are identical and below g^* . See [29] for a formal argument of this type.

8 Conclusion

The proof of Steel’s Conjecture [34] provides tight results for the phylogenetic reconstruction problem. However, many theoretical and practical questions remain:

- Can the discretization assumption be removed? We conjecture that the answer is yes.
- Can the results be extended to other mutation models? In subsequent work, Roch [30] showed that our results hold for so-called General Time-Reversible

(GTR) mutation models below the Kesten–Stigum bound. Can the results be extended to deal with “rates across sites” (see e.g. [11])?

- What is the optimal g -value for the Jukes–Cantor model? Is it identical to the critical value $g_{q=4}$ of the reconstruction problem for the so-called “Potts model” with $q = 4$? We note that it is a long standing open problem to find $g_{q=4}$. The best bounds known are given in [24,26]. See also [19,31].

Acknowledgments S.R. thanks Martin Nowak and the Program for Evolutionary Dynamics at Harvard University where part of this work was done. E.M. and S.R. thank Mike Steel for his enthusiastic encouragement for studying the connections between the reconstruction problem and phylogenetics. We thank Satish Rao and Allan Sly for interesting discussions. We also thank the reviewers for their helpful comments. C.D. and S.R. performed this work at UC Berkeley and Microsoft Research.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Borgs, C., Chayes, J.T., Mossel, E., Roch, S.: The Kesten–Stigum reconstruction bound is tight for roughly symmetric binary channels. In: FOCS, pp. 518–530 (2006)
2. Berger, N., Kenyon, C., Mossel, E., Peres, Y.: Glauber dynamics on trees and hyperbolic graphs. *Probab. Theory Rel.* **131**(3), 311–340 (2005). In: Kenyon, C., Mossel, E., Peres, Y. (eds.) Proceedings of 42nd IEEE Symposium on Foundations of Computer Science (FOCS), pp. 568–578 (2001)
3. Bleher, P.M., Ruiz, J., Zagrebnov, V.A.: On the purity of the limiting Gibbs state for the Ising model on the Bethe lattice. *J. Stat. Phys.* **79**(1–2), 473–482 (1995)
4. Buneman, P.: The recovery of trees from measures of dissimilarity. In: *Mathematics in the Archaeological and Historical Sciences*, pp. 187–395. Edinburgh University Press, Edinburgh (1971)
5. Cavender, J.A.: Taxonomy with confidence. *Math. Biosci.* **40**(3–4) (1978)
6. Chang, J.T.: Full reconstruction of Markov models on evolutionary trees: identifiability and consistency. *Math. Biosci.* **137**(1), 51–73 (1996)
7. Daskalakis, C., Mossel, E., Roch, S.: Optimal phylogenetic reconstruction. In: STOC’06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, pp. 159–168. ACM, New York (2006)
8. Evans, W.S., Kenyon, C., Peres, Y., Schulman, L.J.: Broadcasting on trees and the Ising model. *Ann. Appl. Probab.* **10**(2), 410–433 (2000)
9. Erdős, P.L., Steel, M.A., Székely, L.A., Warnow, T.A.: A few logs suffice to build (almost) all trees (part 1). *Random Struct. Algor.* **14**(2), 153–184 (1999)
10. Farris, J.S.: A probability model for inferring evolutionary trees. *Syst. Zool.* **22**(4), 250–256 (1973)
11. Felsenstein, J.: *Inferring Phylogenies*. Sinauer, New York (2004)
12. Georgii, H.O.: Gibbs measures and phase transitions, volume 9 of de Gruyter Studies in Mathematics. Walter de Gruyter & Co., Berlin (1988)
13. Higuchi, Y.: Remarks on the limiting Gibbs states on a $(d + 1)$ -tree. *Publ. Res. Inst. Math. Sci.* **13**(2), 335–348 (1977)
14. Ioffe, D.: On the extremality of the disordered state for the Ising model on the Bethe lattice. *Lett. Math. Phys.* **37**(2), 137–143 (1996)
15. Jukes, T. H., Cantor, C.: Mammalian protein metabolism. In: Munro, H. N. (ed.) *Evolution of Protein Molecules*, pp. 21–132. Academic, London (1969)
16. Janson, S., Mossel, E.: Robust reconstruction on trees is determined by the second eigenvalue. *Ann. Probab.* **32**, 2630–2649 (2004)
17. Kesten, H., Stigum, B.P.: Additional limit theorems for indecomposable multidimensional Galton–Watson processes. *Ann. Math. Stat.* **37**, 1463–1481 (1966)
18. Lyons, R.: The Ising model and percolation on trees and tree-like graphs. *Comm. Math. Phys.* **125**(2), 337–353 (1989)

19. Mézard, M., Montanari, A.: Reconstruction on trees and spin glass transition. *J. Stat. Phys.* **124**(6), 1317–1350 (2006)
20. Mossel, E.: Recursive reconstruction on periodic trees. *Random Struct. Algor.* **13**(1), 81–97 (1998)
21. Mossel, E.: Reconstruction on trees: beating the second eigenvalue. *Ann. Appl. Probab.* **11**(1), 285–300 (2001)
22. Mossel, E.: Phase transitions in phylogeny. *Trans. Am. Math. Soc.* **356**(6), 2379–2404 (2004)
23. Mossel, E.: Distorted metrics on trees and phylogenetic forests. *IEEE/ACM Trans. Comput. Bio. Bioinform.* **4**(1), 108–116 (2007)
24. Mossel, E., Peres, Y.: Information flow on trees. *Ann. Appl. Probab.* **13**(3), 817–844 (2003)
25. Mossel, E., Steel, M.: A phase transition for a random cluster model on phylogenetic trees. *Math. Biosci.* **187**(2), 189–203 (2004)
26. Martinelli, F., Sinclair, A., Weitz, D.: Fast mixing for independent sets, colorings, and other models on trees. In: *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms*, pp. 449–458, 2004
27. Martinelli, F., Sinclair, A., Weitz, D.: Glauber dynamics on trees: boundary conditions and mixing time. *Comm. Math. Phys.* **250**(2), 301–334 (2004)
28. Neyman, J.: Molecular studies of evolution: a source of novel statistical problems. In: Gupta, S. S., Yackel, J. (eds.) *Statistical Decision Theory and Related Topics*, pp. 1–27. Academic Press, New York (1971)
29. Roch, S.: Sequence-length requirement for distance-based phylogeny reconstruction: breaking the polynomial barrier. In: *FOCS*, pp. 729–738 (2008)
30. Roch, S.: Phase transition in distance-based phylogeny reconstruction. Preprint (2009)
31. Sly, A.: Reconstruction of symmetric Potts models. Preprint available at <http://arxiv.org/abs/0811.1208> (2008)
32. Spitzer, F.: Markov random fields on an infinite tree. *Ann. Probab.* **3**(3), 387–398 (1975)
33. Semple, C., Steel, M.: *Phylogenetics*. Mathematics and its Applications Series, vol. 22. Oxford University Press, NY (2003)
34. Steel, M.: My Favourite Conjecture. Preprint (2001)