



Research article

Convergence of maximum likelihood supertree reconstruction

Vu Dinh¹ and Lam Si Tung Ho^{2,*}

¹ Department of Mathematical Sciences, University of Delaware, Newark, Delaware, USA

² Department of Mathematics and Statistics, Dalhousie University, Halifax, Nova Scotia, Canada

* **Correspondence:** Email: Lam.Ho@dal.ca.

Abstract: Supertree methods are tree reconstruction techniques that combine several smaller gene trees (possibly on different sets of species) to build a larger species tree. The question of interest is whether the reconstructed supertree converges to the true species tree as the number of gene trees increases (that is, the consistency of supertree methods). In this paper, we are particularly interested in the convergence rate of the maximum likelihood supertree. Previous studies on the maximum likelihood supertree approach often formulate the question of interest as a discrete problem and focus on reconstructing the correct topology of the species tree. Aiming to reconstruct both the topology and the branch lengths of the species tree, we propose an analytic approach for analyzing the convergence of the maximum likelihood supertree method. Specifically, we consider each tree as one point of a metric space and prove that the distance between the maximum likelihood supertree and the species tree converges to zero at a polynomial rate under some mild conditions. We further verify these conditions for the popular exponential error model of gene trees.

Keywords: supertree; maximum likelihood estimator; species tree reconstruction; convergence rate; exponential model

Mathematics Subject Classification: Primary 05C05, 62F12; secondary 92B10, 92D15

1. Introduction

High-throughput sequencing is making large collections of sequences available to researchers at a low cost. These genomic data represent a broad spectrum of life and motivates studies of the problem of reconstructing large phylogenetic trees using statistical methods. Those data, however, also come from different sources, cover different genomic regions on which the evolutionary processes happen very differently, and may not be collected on the same set of species. Thus, combining trees on different, overlapping sets of species into a “supertree” has become a popular approach for reconstructing large species trees. Over the years, several methods of supertree reconstruction have been developed [8],

and analyses of supertrees continue to play more and more important roles in the search for answers of many fundamental evolutionary questions.

While supertree methods have been of great interest in phylogenetics, little is known about their theoretical properties, especially in non-asymptotic settings. In cases when the individual trees are gene trees and the set of taxa are the same across the input trees, several statistical consistent methods have been derived [5, 7, 11, 14, 16–18, 20, 21, 29]. However, most proofs of statistical consistency have been analyzed under the condition that each of the gene trees can be estimated accurately, and do not provide a guarantee of robustness to gene tree estimation errors [25]. This is a critical concern because when the species tree is constructed using sequences taken from large genomic regions, those regions have a high chance of involving some recombination, which violates one of the main assumptions of the multi-species coalescent model. On the other hand, limiting analyses to short regions increase gene tree estimation error (GTEE) and various summary methods had impaired accuracy when the error was high [10, 19]. It is later proved in [24] that when the sequence length of each locus is bounded and the gene tree cannot be estimated reliably, most summary methods that estimate the species tree by combining gene trees are not statistically consistent.

Because some coalescent-based summary methods sometimes produce less accurate estimates than concatenation [3, 19, 22], seemingly as a result of GTEE, the question of whether provable guarantees can be established in the presence of GTEE (for both species trees and supertree) naturally arises. In the context of species tree estimation, [25] establish statistical consistency of the Rooted Triplet Consensus method and the Maximum Pseudo-likelihood for Estimating Species Trees method [18] under GTEE and provide bounds on the sampling complexity for these methods to construct the correct species tree with high probability.

The maximum likelihood (ML) supertree method is proposed by [28] based on a probability model that permits “errors” in gene tree topologies and allows the species tree to be estimated even if there is topological conflict amongst gene trees. [28] shows that ML estimate of the species tree is topologically consistent under fairly general conditions and also shows that the method of Matrix Representation with Parsimony [2] may be inconsistent under these same conditions. However, while the ML estimate is topologically consistent, no results on the convergence rate of the estimator are obtained. In this paper, we propose a new analytic approach to study the convergence of the ML supertree method. Based on embedding trees into a metric space, we establish the conditions for which the convergence rate of the ML supertree can be obtained. We verify these conditions for the popular exponential error model of gene trees, thus obtain a polynomial convergence rate for the ML supertree to the true species tree under this model.

2. Mathematical framework

In this paper, the term *phylogenetic tree* refers to a tree \mathbb{T} with leaves labeled by a set of species. Each branch of \mathbb{T} is associated with a non-negative branch length. A tree is said to be resolved if it is bifurcating and all branch lengths are positive. Given an unrooted phylogenetic tree \mathbb{T} on a finite set \mathcal{L} of species, any subset \mathcal{L}' of \mathcal{L} induces a phylogenetic tree on \mathcal{L}' , denoted $\mathbb{T}|_{\mathcal{L}'}$, which is the subtree of \mathbb{T} that connects the species in \mathcal{L}' only.

We regard the tree space as a metric space (\mathcal{T}, d) where \mathcal{T} is the set of all phylogenetic trees with branch lengths bounded from above by a positive constant g_0 and d is a continuous metric.

For simplicity of presentation, we will assume that d is the BHV distance on the set of trees with n species [4], but note that the analysis of the paper can be extended to any continuous and locally-Euclidean distance, including the branch-score distance [15] and the AGPS distance [1].

To describe trees that are “near” to each other, the BHV distance use the class of nearest neighbor interchange (NNI) moves [23]. An NNI move is defined as a transformation that collapses an interior branch to zero and then expands the resulting degree 4 vertex into a branch in a different way. The BHV space models the set of trees \mathcal{T} on n species as a cubical complex consisting of a collection of orthants, each isomorphic to $\mathbb{R}_{\geq 0}^{2n-3}$. Each orthant of \mathcal{T} corresponds uniquely to a tree topology, and the coordinates in each orthant parameterize the branch lengths for the corresponding tree. The adjacent orthants of the complex with the same dimension correspond to NNI-adjacent trees.

The BHV space is equipped with a natural metric distance: the shortest path lying in the BHV space between the points. If two points lie in the same orthant, this distance is the usual Euclidean distance. If two points are in different orthants, they can be joined by a sequence of straight segments, with each segment lying in a single orthant. We can then measure the length of the path by adding up the lengths of the segments. The distance between the two trees \mathbb{T}_1 and \mathbb{T}_2 on the BHV space is defined as the minimum of the lengths of such segmented paths joining the two points.

Throughout the paper, we assume that the evolution of the species has not involved reticulate processes, and there exists an underlying “true species tree” in \mathcal{T} , denoted by \mathbb{T}^* . Furthermore, \mathbb{T}^* is a resolved tree with leaf set \mathcal{L}^* . In a supertree reconstruction problem, we observe a sequence of gene trees $(\mathbb{T}_1, \mathbb{T}_2, \dots, \mathbb{T}_k)$, where \mathbb{T}_i has leaf set \mathcal{L}_i , and wish to combine these trees into a phylogenetic tree $\hat{\mathbb{T}}_k$ on the union of the leaf sets

$$\hat{\mathcal{L}}_k = \cup_{i=1}^k \mathcal{L}_i \subset \mathcal{L}^*.$$

In this paper, we consider a tree-generating probability model $(\mathcal{P}_{\mathbb{T}})_{\mathbb{T} \in \mathcal{T}}$ such that for each set of species $\mathcal{L} \subset \mathcal{L}^*$, $\mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\cdot)$ is a distribution of trees forming by species in \mathcal{L} . We assume that the observed gene trees $\{\mathbb{T}_i\}_{i=1}^k$ are independently distributed according to $\{\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}(\cdot)\}_{i=1}^k$ respectively. In other words, the joint probability of the observed gene trees is

$$\mathcal{P}_{\mathbb{T}^*}(\mathbb{T}_1, \mathbb{T}_2, \dots, \mathbb{T}_k) = \prod_{i=1}^k \mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}(\mathbb{T}_i).$$

Remark 2.1. We note that the tree-generating probability model $\mathcal{P}_{\mathbb{T}^*}$ may not necessarily be nested. That is, given two nested sets of leaves $\mathcal{L}_1 \subset \mathcal{L}_2 \subset \mathcal{L}^*$, although the probability $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_2}$ (which is used to generate trees with leaf set \mathcal{L}_2) also induces a natural distribution on the set of trees with leaf set \mathcal{L}_1 , this probability $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_2}|_{\mathcal{L}_1}$ may not be the same as $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_1}$.

Given the observed gene trees $(\mathbb{T}_1, \mathbb{T}_2, \dots, \mathbb{T}_k)$, the ML supertree is defined as

$$\hat{\mathbb{T}}_k = \operatorname{argmax}_{\mathbb{T} \in \mathcal{T}} \ell_k(\mathbb{T})$$

where $\ell_k(\mathbb{T})$ denotes the log-likelihood function

$$\ell_k(\mathbb{T}) = \sum_{i=1}^k \log \mathcal{P}_{\mathbb{T}}^{\mathcal{L}_i}(\mathbb{T}_i).$$

To enable theoretical analyses of the ML supertree method, we make the following assumptions.

Assumption 2.1 (Weak covering property). *The sequence of subsets of \mathcal{L}^* satisfies the weak covering property: there exists $c_1 > 0, \gamma > 1/2$, and $K > 0$ such that for each subset \mathcal{L} of taxa from \mathcal{L}^* of size 4,*

$$\frac{1}{k^\gamma} |\{i \leq k : \mathcal{L} \subset \mathcal{L}_i\}| \geq c_1 \quad \forall k \geq K.$$

Here, $|\mathcal{A}|$ denotes the number of elements in the set \mathcal{A} .

We note that the weak covering property ensures that as the number of gene trees increases, all quartets (subtrees with 4 leaves) of \mathbb{T}^* are visited with enough frequency to enable a reliable estimate for each of the quartets. Assumption 2.1 is a direct generalization of the *covering property* introduced in [28], which can be obtained from Assumption 2.1 by setting $\gamma = 1$.

Assumption 2.2 (Model identifiability). *For all $\mathcal{L} \subset \mathcal{L}^*$, the distribution $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}(\cdot)$ is identifiable. That is, if $d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}}) > 0$, then $KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) > 0$.*

Assumption 2.2 guarantees that it is at least possible to reconstruct the restriction $\mathbb{T}^*|_{\mathcal{L}}$ of \mathbb{T}^* to the subset \mathcal{L} with a complete knowledge of $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}$. This assumption is similar to, but distinct from, the condition of *basic centrality* introduced by [28], which requires that for all subsets of $\mathcal{L} \subset \mathcal{L}^*$ of size 4,

$$\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}[\mathbb{T}^*|_{\mathcal{L}}] \geq (1 + \eta)\mathcal{P}_{\mathbb{T}}^{\mathcal{L}}[\mathbb{T}']$$

for all trees \mathbb{T}' on leaf set \mathcal{L} that are different from $\mathbb{T}^*|_{\mathcal{L}}$, and where $\eta > 0$:

- On one hand, the basic centrality condition implies that if $d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}}) > 0$, the modes of the distributions $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}$ and $\mathcal{P}_{\mathbb{T}}^{\mathcal{L}}$ are different. From this, we can deduce that $KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) > 0$. Thus, our identifiability assumption is somewhat weaker than this condition. Assumption 2.2 also does not impose any assumption on the family of distribution themselves and thus can be applied to a wider class of probabilistic models.
- On the other hand, the basic centrality condition only concerns leaf sets of size 4, while Assumption 2.2 impose restrictions on leaf sets of all sizes. However, we note that conditions on subset of leaves (such as the basic centrality condition) only work under the implicit assumption that there are some connection between $\mathcal{P}_{\mathbb{T}}^{\mathcal{L}_2}|_{\mathcal{L}_1}$ and $\mathcal{P}_{\mathbb{T}}^{\mathcal{L}_1}$ for $\mathcal{L}_1 \subset \mathcal{L}_2$. Since our framework does not assume any nested structure in the probability model, Assumption 2.2 is more appropriate.

Finally, we impose the following regularity conditions on the tree-generating probability model:

Assumption 2.3. (Regularity)

- (a) For all $\mathbb{T} \in \mathcal{T}$, $\mathcal{L} \subset \mathcal{L}^*$, and any tree \mathbb{T}' with leaf set \mathcal{L} ,
- $\mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\mathbb{T}') > 0$,
 - $\log \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\mathbb{T}')$ is a locally-Lipchitz function with respect to \mathbb{T} , and the Lipchitz constant does not depend on \mathbb{T}' .
- (b) There exist $c_2, c_3 > 0$, $m \geq 2$ such that for any leaf set \mathcal{L} and tree \mathbb{T} , if $d(\mathbb{T}^*, \mathbb{T}) \leq c_2$, then

$$KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) \geq c_3 d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}})^m.$$

Remark 2.2. *If for each tree \mathbb{T}' with leaf set \mathcal{L} , the probability density function $\mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\mathbb{T}')$ is an analytic function with respect to \mathbb{T} in a neighborhood of \mathbb{T}^* , then Assumption 2.3(b) holds.*

Proof. Note that the function $\mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} \left[\log \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\mathbb{T}') \right]$ (with respect to \mathbb{T}) is analytic in a neighborhood \mathcal{U} of \mathbb{T}^* . Define

$$\mathcal{A} = \{ \mathbb{T} \in \mathcal{U} : \mathbb{E}_{\mathbb{T}' \sim \mathcal{P}_{\mathbb{T}^*}} \left[\log \mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}(\mathbb{T}') \right] = \mathbb{E}_{\mathbb{T}' \sim \mathcal{P}_{\mathbb{T}^*}} \left[\log \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\mathbb{T}') \right] \}.$$

For all $\mathbb{T} \in \mathcal{A}$, we have

$$\text{KL}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) = \mathbb{E}_{\mathbb{T}' \sim \mathcal{P}_{\mathbb{T}^*}} \left[\log \mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}(\mathbb{T}') \right] - \mathbb{E}_{\mathbb{T}' \sim \mathcal{P}_{\mathbb{T}^*}} \left[\log \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\mathbb{T}') \right] = 0.$$

By Assumption 2.2, we conclude that $\mathbb{T}|_{\mathcal{L}} = \mathbb{T}^*|_{\mathcal{L}}$ for all $\mathbb{T} \in \mathcal{A}$. Applying Łojasiewicz inequality [13, Theorem 1], we deduce that there exists $c_3 > 0$ and $m \geq 2$ such that for any \mathbb{T} in the neighborhood,

$$\mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} \left[\log \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}(\mathbb{T}') \right] - \mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} \left[\log \mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}(\mathbb{T}') \right] \geq c_3 d(\mathbb{T}, \mathcal{A})^m \geq c_3 d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}})^m.$$

This implies the result. \square

Throughout this paper, we will assume that Assumptions 2.1, 2.2, and 2.3 hold. The main point of Assumption 2.1 is to guarantee that every subset of 4 species is sufficiently represented by the data. It is worth noticing that 4 is the smallest number of species that carries the information of the topology of an unrooted tree. Therefore, if a set of 4 species is only found in few trees, we will not be able to reconstruct their phylogenetic relationship. While Assumption 2.2 is technical, its goal is to guarantee that there is a unique “true” supertree. Finally, Assumption 2.3 ensures that the model has enough information to distinguish between the “true” supertree and any alternative with substantial accuracy.

In the next section, we will establish the following convergence rate of the ML supertree.

Theorem 1. *Under Assumptions 2.1, 2.2, and 2.3, for any $\delta > 0$, there exist constants $m > 0$, $C_\delta > 0$ and $K_\delta > 0$ such that for all $k \geq K_\delta$, we have*

$$d(\hat{\mathbb{T}}_k, \mathbb{T}^*) \leq C_\delta \left(\frac{\log k}{k^{(\gamma-1/2)}} \right)^{1/m}$$

with probability at least $1 - \delta$.

Theorem 1 establishes that under fairly mild regularity conditions, the ML supertree is consistent and has a polynomial convergence rate. Notably, the result holds for all identifiable family of locally-analytic distributions (Remark 2.2). The degree of this (polynomial) convergence rate depends on the sampling scheme of the leaves (characterized by the covering coefficient γ) and on the geometry of the probabilistic model (characterized by the constant m in Assumption 2.3(b)). We further note that Assumption 2.3(b) is only required for establishing the convergence rate of the ML estimator and the absence of this condition does not affect the proof of consistency.

Corollary 1. *Under Assumptions 2.1, 2.2, and 2.3(a), the ML supertree is consistent.*

Remark 2.3. *Theorem 1 is still valid if the universal Assumption 2.3(b) is replaced by the following condition:*

There exist $c_2, c_3 > 0$, $m \geq 2$ such that if $d(\mathbb{T}^*, \mathbb{T}) \leq c_2$, then

$$\text{KL}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}_i}) \geq c_3 d(\mathbb{T}^*|_{\mathcal{L}_i}, \mathbb{T}|_{\mathcal{L}_i})^m \quad \forall i \in \mathbb{N}.$$

Remark 2.4. *We note that the bound of the convergence rate in Theorem 1 is not sharp.*

3. Convergence of maximum likelihood supertree

To enable the analysis of convergence, we define

$$R_k(\mathbb{T}) = \frac{1}{k} \ell_k(\mathbb{T}^*) - \frac{1}{k} \ell_k(\mathbb{T})$$

and

$$\mathcal{R}_k(\mathbb{T}) := E_{\mathcal{P}_{\mathbb{T}^*}}[R_k(\mathbb{T})] = \frac{1}{k} \left(E_{\mathcal{P}_{\mathbb{T}^*}}[\ell_k(\mathbb{T}^*)] - E_{\mathcal{P}_{\mathbb{T}^*}}[\ell_k(\mathbb{T})] \right) = \frac{1}{k} \text{KL}(\mathcal{P}_{\mathbb{T}^*}, \mathcal{P}_{\mathbb{T}}).$$

We refer to R_k and $E_{\mathcal{P}_{\mathbb{T}^*}}[R_k(\mathbb{T})]$ as the *empirical risk function* and the *expected risk function*, respectively.

An intuitive argument for the consistency of the ML estimator can be described as follows. As the number of gene trees increases, we have

$$|R_k(\mathbb{T}) - \mathcal{R}_k(\mathbb{T})| \rightarrow 0$$

with sufficiently high probability. Therefore, the ML estimator will converge to the optimal value of the risk function, which is attained at the true species tree \mathbb{T}^* . This simple argument is formalized by a lower bound of the expected risk function (Section 3.1) and a uniform concentration bound on the deviation of the empirical risk function and its expectation (Section 3.2).

3.1. Lower bound of the expected risk

Lemma 3.1. *There exist a neighborhood \mathcal{U} of \mathbb{T}^* and $c_4 > 0$ such that*

$$\mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}}[\ell_k(\mathbb{T}^*) - \ell_k(\mathbb{T})] \geq c_4 k^\gamma d(\mathbb{T}^*, \mathbb{T})^m \quad \forall \mathbb{T} \in \mathcal{U}.$$

Proof. We note that

$$\mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}}[\ell_k(\mathbb{T}^*) - \ell_k(\mathbb{T})] = \text{KL}(\mathcal{P}_{\mathbb{T}^*}, \mathcal{P}_{\mathbb{T}}).$$

Consider an arbitrary leaf set $\mathcal{L} \subset \mathcal{L}^*$ of size 4 and define $\mathcal{I} = \{i \leq k : \mathcal{L} \subset \mathcal{L}_i\}$. By the weak covering property, we have $|\mathcal{I}| \geq c_1 k^\gamma$. Thus,

$$\text{KL}(\mathcal{P}_{\mathbb{T}^*}, \mathcal{P}_{\mathbb{T}}) \geq \sum_{i \in \mathcal{I}} \text{KL}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}_i}) \geq \sum_{i \in \mathcal{I}} c_3 d(\mathbb{T}^*|_{\mathcal{L}_i}, \mathbb{T}|_{\mathcal{L}_i})^m \geq c_1 c_3 k^\gamma (2n-3)^{m/2} d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}})^m$$

for all \mathbb{T} in some neighborhood U around \mathbb{T}^* . Here, the second inequality comes from Assumption 2.3(b).

On the other hand, [9] (Lemma 6.2 (i)) proved that for some leaf set $\mathcal{L} \subset \mathcal{L}^*$ of size 4, we have

$$d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}}) \geq \frac{1}{(2n-3)} d(\mathbb{T}^*, \mathbb{T}).$$

We deduce that

$$\text{KL}(\mathcal{P}_{\mathbb{T}^*}, \mathcal{P}_{\mathbb{T}}) \geq \frac{1}{(2n-3)^{m/2}} c_1 c_3 k^\gamma d(\mathbb{T}^*, \mathbb{T})^m$$

□

Lemma 3.2. For any leaf set \mathcal{L} and any $x > 0$, there exists $C_x > 0$ such that

$$KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) \geq C_x$$

when $d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}}) \geq x$.

Proof. Assume that there exists a sequence of tree $\{\mathbb{T}_i\}$ such that $d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}_i|_{\mathcal{L}}) \geq x$ and $KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}_i}^{\mathcal{L}}) \rightarrow 0$. Since the tree space is compact, we can extract a sub-sequence $\{\mathbb{T}_{i_j}\}$ of $\{\mathbb{T}_i\}$ that converges to some tree \mathbb{T}_0 . We deduce that $KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}_0}^{\mathcal{L}}) = 0$ and $\mathbb{T}^*|_{\mathcal{L}} \neq \mathbb{T}_0|_{\mathcal{L}}$. This contradicts Assumption 2.2. \square

Lemma 3.3. Let \mathcal{V} be a neighborhood of \mathbb{T}^* . There exist $c_{\mathcal{V}} > 0$ such that for any $\mathbb{T} \notin \mathcal{V}$, we have

$$KL(\mathcal{P}_{\mathbb{T}^*}, \mathcal{P}_{\mathbb{T}}) \geq c_{\mathcal{V}} k^{\gamma}.$$

Proof. By [9] (Lemma 6.2 (i)), there exists $\mathcal{L} \subset \mathcal{L}^*$ of size 4 such that

$$d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}}) \geq \frac{1}{(2n-3)} d(\mathbb{T}^*, \mathbb{T}) \geq C_{\mathcal{V}} > 0$$

since $\mathbb{T} \notin \mathcal{V}$. Define $\mathcal{I} = \{i \leq k : \mathcal{L} \subset \mathcal{L}_i\}$, we note that for all $i \in \mathcal{I}$

$$d(\mathbb{T}^*|_{\mathcal{L}_i}, \mathbb{T}|_{\mathcal{L}_i}) \geq (2n-3)^{1/2} d(\mathbb{T}^*|_{\mathcal{L}}, \mathbb{T}|_{\mathcal{L}}) \geq C'_{\mathcal{V}},$$

which implies $KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}_i}) \geq C''_{\mathcal{V}}$ for some $C''_{\mathcal{V}} > 0$ by Lemma 3.2.

By the weak covering property, we have $|\mathcal{I}| \geq c_1 k^{\gamma}$. Thus,

$$KL(\mathcal{P}_{\mathbb{T}^*}, \mathcal{P}_{\mathbb{T}}) \geq \sum_{i \in \mathcal{I}} KL(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}_i}) \geq c_1 k^{\gamma} C''_{\mathcal{V}}$$

which completes the proof. \square

3.2. Uniform concentration bound

Lemma 3.4 (Concentration bound). For any $\delta > 0$, $k \geq 3$, there exists $c_5(\delta)$ such that

$$|R_k(\mathbb{T}) - \mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T})]| \leq \frac{c_5 \log k}{\sqrt{k}} \quad \forall \mathbb{T},$$

with probability at least $1 - \delta$.

Proof. Since the functions $\log \mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}(\mathbb{T}_i)$ is locally Lipschitz with respect to \mathbb{T} (Assumption 2.3(a)) and number of species is finite, there exists $C_1 > 0$ such that

$$|\log \mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}_i}(\mathbb{T}_i) - \log \mathcal{P}_{\mathbb{T}'}^{\mathcal{L}_i}(\mathbb{T}_i)| \leq C_1 d(\mathbb{T}, \mathbb{T}') \quad \forall i, \mathbb{T}_i, \mathbb{T}, \mathbb{T}'.$$

Since the tree space is compact, $C_1 d(\mathbb{T}, \mathbb{T}')$ is bounded by a constant C_2 . Using Hoeffding's inequality [12], we obtain

$$\mathbb{P} \left[|R_k(\mathbb{T}) - \mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T})]| \geq \frac{x}{\sqrt{k}} \right] \leq 2 \exp \left(-\frac{2x^2}{kC_2^2} \right).$$

On the other hand, we have $|R_k(\mathbb{T}) - R_k(\mathbb{T}')| \leq C_1 d(\mathbb{T}, \mathbb{T}')$ and $|\mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T})] - \mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T}')]| \leq C_1 d(\mathbb{T}, \mathbb{T}')$. Thus, if we define the events

$$\mathcal{A}(x, k, \mathbb{T}) = \left\{ \left| R_k(\mathbb{T}) - \mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T})] \right| \geq \frac{x}{2\sqrt{k}} \right\}$$

and

$$\mathcal{B}(x, k, \mathbb{T}) = \left\{ \exists \mathbb{T}' : d(\mathbb{T}, \mathbb{T}') \leq \frac{x}{4C_1\sqrt{k}} \text{ and } \left| R_k(\mathbb{T}') - \mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T}')] \right| \geq \frac{x}{\sqrt{k}} \right\}$$

then $\mathcal{B}(x, k, \mathbb{T}) \subset \mathcal{A}(x, k, \mathbb{T})$. Note that the total number of balls of radius $x/(4C_1\sqrt{k})$ required to cover the tree space is bounded above by

$$C_n \left(\frac{4g_0 C_1 \sqrt{k}}{x} \right)^{2n-3} (2n-3)!!$$

where g_0 is the upper bound of branch lengths and C_n is a constant depending on the number of species. To obtain the desired inequality, we will choose x such that

$$C_n \left(\frac{4g_0 C_1 \sqrt{k}}{x} \right)^{2n-3} (2n-3)!! \times 2 \exp\left(-\frac{2x^2}{kC_2^2}\right) \leq \delta$$

which can be done with $x = C(\delta, n, C_2, g_0) \log k$. □

3.3. Proof of Theorem 1

First, we establish that the ML estimator is consistent.

By Lemma 3.4, we have

$$\mathcal{R}(\hat{\mathbb{T}}_k) \leq R_k(\hat{\mathbb{T}}_k) + c_5 \frac{\log(k)}{\sqrt{k}} \leq c_5 \frac{\log(k)}{\sqrt{k}}$$

with probability at least $1 - \delta$, since $\hat{\mathbb{T}}_k$ is the maximizer of the empirical risk function.

On the other hand, let \mathcal{V} be a neighborhood of \mathbb{T}^* , by Lemma 3.3, we have

$$\mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T}')] \geq c_{\mathcal{V}} \frac{1}{k^{1-\gamma}} \quad \forall \mathbb{T}' \notin \mathcal{V}.$$

Since $\gamma > 1/2$, there exists $K_{\delta, \mathcal{V}}$ such that for $k \geq K_{\delta, \mathcal{V}}$, we have

$$c_5 \frac{\log(k)}{\sqrt{k}} \leq c_{\mathcal{V}} \frac{1}{k^{1-\gamma}}$$

We deduce that for $k \geq K_{\delta, \mathcal{V}}$, $\hat{\mathbb{T}}_k \in \mathcal{V}$ with probability at least $1 - \delta$. This proves that the ML estimator is consistent.

Next, we will derive the convergence rate of the ML supertree. From Lemma 3.1, there exist a neighborhood \mathcal{U} of \mathbb{T}^* and $c_4, m > 0$ such that

$$\mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [R_k(\mathbb{T})] = \frac{1}{k} \mathbb{E}_{\mathcal{P}_{\mathbb{T}^*}} [\ell_k(\mathbb{T}^*) - \ell_k(\mathbb{T})] \geq c_4 \frac{1}{k^{1-\gamma}} d(\mathbb{T}^*, \mathbb{T})^m \quad \forall \mathbb{T} \in \mathcal{U}.$$

For $k \geq K_{\delta, \mathcal{U}}$, we have $\hat{\mathbb{T}}_k \in \mathcal{U}$ with probability at least $1 - \delta$. By Lemma 3.4, with probability $1 - 2\delta$, we obtain

$$\begin{aligned} c_4 \frac{1}{k^{1-\gamma}} d(\mathbb{T}^*, \hat{\mathbb{T}}_k)^m &\leq \mathcal{R}_k(\hat{\mathbb{T}}_k) - R_k(\hat{\mathbb{T}}_k) \\ &\leq c_5 \frac{\log(k)}{\sqrt{k}}. \end{aligned}$$

Here,

$$R_k(\hat{\mathbb{T}}_k) = \frac{1}{k} \ell_k(\mathbb{T}^*) - \frac{1}{k} \ell(\hat{\mathbb{T}}_k) \leq 0$$

because $\hat{\mathbb{T}}_k$ is the maximizer of R_k . This completes the proof.

4. Applications: convergence under the exponential model

In this section, we consider the *exponential model*, a simple model of gene tree estimation errors in which the probability of observing a given tree decreases exponentially with its distance from the species tree. For a detailed discussion on its validity and applicability, we refer the readers to [28] and the references therein.

Suppose d is some metric on the set of trees, in the exponential model, the probability of reconstructing any tree \mathbb{T}' with a leaf set \mathcal{L} , when \mathbb{T}^* is the generating tree is proportional to an exponentially decaying function of the distance from \mathbb{T}' to $\mathbb{T}^*|_{\mathcal{L}}$:

$$\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}(\mathbb{T}') = \alpha_{\mathbb{T}^*, \mathcal{L}} \exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^*|_{\mathcal{L}})).$$

Here $\beta_{\mathcal{L}}$ is a constant that depends only on the set of leaves \mathcal{L} , while $\alpha_{\mathbb{T}^*, \mathcal{L}}$ is the normalizing constant to ensure that $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}$ is a density function.

In this section, we verify the identifiability and regularity conditions for the exponential model when d is any continuous tree distance such that if \mathbb{T} and \mathbb{T}' have the same topology, then $d(\mathbb{T}, \mathbb{T}')$ is the Euclidean distance.

Theorem 2. *Under the exponential model with Assumptions 2.1, for any $\delta > 0$, $C_{\delta} > 0$ and $K_{\delta} > 0$ such that for all $k \geq K_{\delta}$,*

$$d(\hat{\mathbb{T}}_k, \mathbb{T}^*) \leq C_{\delta} \left(\frac{\log k}{k^{(\gamma-1/2)}} \right)^{1/m}$$

with probability at least $1 - \delta$, where

$$m = 4 \sup_i |\mathcal{L}_i| - 4.$$

Proof. First, we note that for all \mathbb{T}' with leaf set \mathcal{L} , $\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}(\mathbb{T}')$ is a positive and locally Lipschitz function. To obtain an upper bound on the convergence rate for the ML supertree, we need to verify Assumption 2.2 and Assumption 2.3(b)/Remark 2.3.

Identifiability. Since it is sufficient to prove either $\text{KL}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) > 0$ or $\text{KL}(\mathcal{P}_{\mathbb{T}}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}) > 0$, we can assume that $\alpha_{\mathbb{T}^*, \mathcal{L}} \geq \alpha_{\mathbb{T}, \mathcal{L}}$.

We denote

$$\mathcal{A}_c = \{\mathbb{T}' : \beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^*|_{\mathcal{L}}) \leq c\}$$

and pick c small enough such that $\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}|_{\mathcal{L}}) \geq 2c$. We note that in A_c ,

$$\exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^*|_{\mathcal{L}})) \geq \exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}|_{\mathcal{L}})).$$

Since $\alpha_{\mathbb{T}^*, \mathcal{L}} \geq \alpha_{\mathbb{T}, \mathcal{L}}$, we have

$$\begin{aligned} \alpha_{\mathbb{T}^*, \mathcal{L}} \exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^*|_{\mathcal{L}})) - \alpha_{\mathbb{T}, \mathcal{L}} \exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}|_{\mathcal{L}})) \\ \geq \alpha_{\mathbb{T}, \mathcal{L}} [\exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^*|_{\mathcal{L}})) - \exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}|_{\mathcal{L}}))]. \end{aligned}$$

By Pinsker's inequality,

$$\begin{aligned} \text{KL}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) &\geq 2d_{TV}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}})^2 \\ &\geq 2\alpha_{\mathbb{T}, \mathcal{L}}^2 \left[\int_{\mathcal{A}_c} \exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^*|_{\mathcal{L}})) - \exp(-\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}|_{\mathcal{L}})) \right]^2 \\ &\geq 2\alpha_{\mathbb{T}, \mathcal{L}}^2 \mu(\mathcal{A}_c)^2 (e^c - e^{2c})^2 > 0, \end{aligned}$$

which establishes the identifiability of the exponential model. Here, μ is the Lebesgue measure.

Regularity. Consider a fixed tree $\mathbb{T} \in \mathcal{T}$ and leaf set \mathcal{L} , we first assume that $\alpha_{\mathbb{T}^*, \mathcal{L}} \geq \alpha_{\mathbb{T}, \mathcal{L}}$. If we define

$$\mathcal{A} = \{\mathbb{T}' : \beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^*|_{\mathcal{L}}) \leq \frac{1}{3} d(\mathbb{T}|_{\mathcal{L}}, \mathbb{T}^*|_{\mathcal{L}})\},$$

then

$$\beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}|_{\mathcal{L}}) \geq \frac{2}{3} d(\mathbb{T}|_{\mathcal{L}}, \mathbb{T}^*|_{\mathcal{L}}) \quad \forall \mathbb{T}' \in \mathcal{A}.$$

Using the same argument as above, we have

$$d_{TV}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) \geq \alpha_{\mathbb{T}, \mathcal{L}} \mu(\mathcal{A}) \left[e^{-\frac{1}{3} d(\mathbb{T}|_{\mathcal{L}}, \mathbb{T}^*|_{\mathcal{L}})} - e^{-\frac{2}{3} d(\mathbb{T}|_{\mathcal{L}}, \mathbb{T}^*|_{\mathcal{L}})} \right].$$

It can be verified that there exists $C > 0$ such that for all $x \in [0, C]$, we have

$$e^{-x} - e^{-2x} \geq \frac{x}{2}. \quad (4.1)$$

Moreover since d is Euclidean inside each orthant, if $d(\mathbb{T}, \mathbb{T}^*)$ is small enough, we have

$$\mu(\mathcal{A}) = C_{2|\mathcal{L}|-3} \left(\frac{d(\mathbb{T}|_{\mathcal{L}}, \mathbb{T}^*|_{\mathcal{L}})}{3\beta_{\mathcal{L}}} \right)^{2|\mathcal{L}|-3}. \quad (4.2)$$

Thus, let \mathcal{W} be a neighborhood in the same topology of \mathbb{T}^* such that

$$\alpha_{\mathbb{T}, \mathcal{L}} \geq \frac{1}{2} \alpha_{\mathbb{T}^*, \mathcal{L}} \quad \forall \mathbb{T} \in \mathcal{W},$$

and both Eqs (4.1) and (4.2) are satisfied, we have

$$\text{KL}(\mathcal{P}_{\mathbb{T}^*}^{\mathcal{L}}, \mathcal{P}_{\mathbb{T}}^{\mathcal{L}}) \geq C d(\mathbb{T}, \mathbb{T}^*|_{\mathcal{L}})^{4|\mathcal{L}|-4} \quad \forall \mathbb{T} \in \mathcal{W},$$

for some constant $C > 0$ independent of \mathbb{T} .

For the case when $\alpha_{\mathbb{T}^*, \mathcal{L}} < \alpha_{\mathbb{T}, \mathcal{L}}$ we define

$$\mathcal{A}' = \{\mathbb{T}' : \beta_{\mathcal{L}} d(\mathbb{T}', \mathbb{T}^* |_{\mathcal{L}}) \leq \frac{1}{3} d(\mathbb{T}' |_{\mathcal{L}}, \mathbb{T} |_{\mathcal{L}})\},$$

and the argument proceeds similarly. This validates Assumption 2.3(b) for exponential model with $m = 4n - 4$. However, if we use the regularity condition in Remark 2.3, we can obtain the result with

$$m = 4 \sup_i |\mathcal{L}_i| - 4.$$

□

5. Application to yeast species-tree reconstruction

We apply our method to reconstruct the yeast species-tree from 106 genes of eight yeast species obtained from the R package `phangorn` [27]. Specifically, we reconstruct the gene-tree for each gene using UPGMA (unweighted pair group method with arithmetic mean) method implemented in `phangorn`. To replicate the setting of this paper, we randomly remove two species from each reconstructed gene-tree. In other words, the inputs are 106 reconstructed gene-trees, each has 6 species (see Figure 1 for examples).




Figure 1. Examples of the reconstructed gene-trees that we use as the inputs.

We apply our method to estimate the full species-tree using the exponential model with BHV distance. The R function `dist.multiPhylo` in the package `distory` [6] is used to compute the BHV

distance. The reconstructed species-tree (see Figure 2) is consistent with the well-known result of [26]. This confirms the applicability of our method to reliably estimate supertrees.




Figure 2. The reconstructed species-tree obtained by applying the ML supertree method to 106 reconstructed gene-trees.

6. Discussion and conclusions

In this paper, we propose a novel analytic approach to analyze the convergence of the ML supertree method. Instead of focusing on reconstructing the correct discrete topology of the species tree as in previous studies [25, 28, e.g.], we employ a continuous model of the tree space and analyze the ML estimator on this metric space, aiming at recovering both the topology and the branch lengths of the species tree. This framework enables us to use tools from statistical learning and information theory to establish the convergence rate of the ML estimator and at the same time, to weaken the conditions to obtain consistency and convergence of the estimator. Our *weak covering property* is an extension of the classical covering property [28] and provides a considerable relaxation on the sampling schemes for supertree estimation. Our identifiability condition is also more intuitive and generalizable than the well-known basic centrality condition and does not impose constraints on the shape of the probabilistic model of gene tree estimation errors. Our information-theoretical approach to analyze statistical estimator on tree spaces is of independent interest and can be extended to other problems in phylogenetics.

There are several avenues for future directions for this work. The first direction is extending our results to other practical models of phylogenetic errors, including the multiple-coalescent model (along with a detailed model of the effects of short sequence length on the accuracy in estimating the individual gene trees). Second, while our result provides a polynomial bound on the convergence rate, the power

of the convergence (characterized by the geometric constant m) is not sharp. A sharper bound of the convergence rate would be of great interest to the field (from both theoretical and applied perspective) and would require further understanding of the tree-generating probabilistic model.

Acknowledgments

LSTH was supported by startup funds from Dalhousie University, the Canada Research Chairs program, the NSERC Discovery Grant RGPIN-2018-05447, and the NSERC Discovery Launch Supplement DGEGR-2018-00181. VD was supported by a startup fund from University of Delaware and National Science Foundation grant DMS-1951474.

Conflict of interest

The author declares no conflict of interest in this paper.

References

1. N. Amenta, M. Godwin, N. Postarnakevich, K. S. John, Approximating geodesic tree distance, *Inform. Process. Lett.*, **103** (2007), 61–65.
2. B. R. Baum, Combining trees as a way of combining data sets for phylogenetic inference, and the desirability of combining gene trees, *Taxon*, **41** (1992), 3–10.
3. M. S. Bayzid, T. Warnow, Naive binning improves phylogenomic analyses, *Bioinformatics*, **29** (2013), 2277–2284.
4. L. J. Billera, S. P. Holmes, K. Vogtmann, Geometry of the space of phylogenetic trees, *Adv. Appl. Math.*, **27** (2001), 733–767.
5. D. Bryant, R. Bouckaert, J. Felsenstein, N. A. Rosenberg, A. RoyChoudhury, Inferring species trees directly from biallelic genetic markers: bypassing gene trees in a full coalescent analysis, *Mol. Biol. Evol.*, **29** (2012), 1917–1932.
6. J. Chakerian, S. Holmes, DISTORY: Distance between phylogenetic histories. *R package version*, **1** (2013).
7. J. Chifman, L. Kubatko, Quartet inference from SNP data under the coalescent model, *Bioinformatics*, **30** (2014), 3317–3324.
8. J. A. Cotton, M. Wilkinson, Majority-rule supertrees, *Syst. Biol.*, **56** (2007), 445–452.
9. V. Dinh, L. S. T. Ho, M. A. Suchard, F. A. Matsen IV, Consistency and convergence rate of phylogenetic inference via regularization, *Ann. Stat.*, **46** (2018), 1481.
10. J. Gatesy, M. S. Springer, Phylogenetic analysis at deep timescales: unreliable gene trees, bypassed hidden support, and the coalescence/concatalescence conundrum, *Mol. Phylogenet. Evol.*, **80** (2014), 231–266.
11. J. Heled, A. J. Drummond, Bayesian inference of species trees from multilocus data, *Mol. Biol. Evol.*, **27** (2009), 570–580.

12. W. Hoeffding, Probability inequalities for sums of bounded random variables, *J. Am. Stat. Assoc.*, **58** (1963), 13–30.
13. S. Ji, J. Kollár, B. Shiffman, A global Łojasiewicz inequality for algebraic varieties, *T. Am. Math. Soc.*, **329** (1992), 813–818.
14. L. S. Kubatko, B. C. Carstens, L. L. Knowles, STEM: species tree estimation using maximum likelihood for gene trees under coalescence, *Bioinformatics*, **25** (2009), 971–973.
15. M. K. Kuhner, J. Felsenstein, A simulation comparison of phylogeny algorithms under equal and unequal evolutionary rates, *Mol. Biol. Evol.*, **11** (1994), 459–468.
16. B. R. Larget, S. K. Kotha, C. N. Dewey, C. Ané, BUCKy: gene tree/species tree reconciliation with bayesian concordance analysis, *Bioinformatics*, **26** (2010), 2910–2911.
17. L. Liu, L. Yu, Estimating species trees from unrooted gene trees, *Syst. Biol.*, **60** (2011), 661–667.
18. L. Liu, L. Yu, S. V. Edwards, A maximum pseudo-likelihood approach for estimating species trees under the coalescent model, *BMC Evol. Biol.*, **10** (2010), 302.
19. S. Mirarab, M. S. Bayzid, B. Boussau, T. Warnow, Statistical binning enables an accurate coalescent-based estimation of the avian tree, *Science*, **346** (2014), 1250463.
20. S. Mirarab, R. Reaz, M. S. Bayzid, T. Zimmermann, M. S. Swenson, T. Warnow, ASTRAL: genome-scale coalescent-based species tree estimation, *Bioinformatics*, **30** (2014), i541–i548.
21. E. Mossel, S. Roch, Incomplete lineage sorting: consistent phylogeny estimation from multiple loci, *IEEE ACM T. Comput. Bi.*, **7** (2008), 166–171.
22. S. Patel, R. T. Kimball, E. L. Braun, Error in phylogenetic estimation for bushes in the tree of life, *Journal of Phylogenetics & Evolutionary Biology*, (2013).
23. D. F. Robinson, Comparison of labeled trees with valency three, *J. Comb. Theory B*, **11** (1971), 105–119.
24. S. Roch, M. Nute, T. Warnow, Long-branch attraction in species tree estimation: inconsistency of partitioned likelihood and topology-based summary methods, *Syst. Biol.*, **68** (2019), 281–297.
25. S. Roch, T. Warnow, On the robustness to gene tree estimation error (or lack thereof) of coalescent-based species tree methods, *Syst. Biol.*, **64** (2015), 663–676.
26. A. Rokas, B. L. Williams, N. King, S. B. Carroll, Genome-scale approaches to resolving incongruence in molecular phylogenies, *Nature*, **425** (2003), 798–804.
27. K. P. Schliep, phangorn: phylogenetic analysis in r, *Bioinformatics*, **27** (2011), 592–593.
28. M. Steel, A. Rodrigo, Maximum likelihood supertrees, *Syst. Biol.*, **57** (2008), 243–250.
29. P. Vachaspati, T. Warnow, ASTRID: accurate species trees from internode distances, *BMC genomics*, **16** (2015), 1–13.