

CONSISTENCY AND CONVERGENCE RATE OF PHYLOGENETIC INFERENCE VIA REGULARIZATION²

BY VU DINH^{*,1}, LAM SI TUNG HO^{†,1}, MARC A. SUCHARD[†] AND
FREDERICK A. MATSEN IV^{*,3}

Fred Hutchinson Cancer Research Center and University of California, Los Angeles[†]*

It is common in phylogenetics to have some, perhaps partial, information about the overall evolutionary tree of a group of organisms and wish to find an evolutionary tree of a specific gene for those organisms. There may not be enough information in the gene sequences alone to accurately reconstruct the correct “gene tree.” Although the gene tree may deviate from the “species tree” due to a variety of genetic processes, in the absence of evidence to the contrary it is parsimonious to assume that they agree. A common statistical approach in these situations is to develop a likelihood penalty to incorporate such additional information. Recent studies using simulation and empirical data suggest that a likelihood penalty quantifying concordance with a species tree can significantly improve the accuracy of gene tree reconstruction compared to using sequence data alone. However, the consistency of such an approach has not yet been established, nor have convergence rates been bounded. Because phylogenetics is a nonstandard inference problem, the standard theory does not apply. In this paper, we propose a penalized maximum likelihood estimator for gene tree reconstruction, where the penalty is the square of the Billera–Holmes–Vogtmann geodesic distance from the gene tree to the species tree. We prove that this method is consistent, and derive its convergence rate for estimating the discrete gene tree structure and continuous edge lengths (representing the amount of evolution that has occurred on that branch) simultaneously. We find that the regularized estimator is “adaptive fast converging,” meaning that it can reconstruct all edges of length greater than any given threshold from gene sequences of polynomial length. Our method does not require the species tree to be known exactly; in fact, our asymptotic theory holds for any such guide tree.

Received June 2016; revised February 2017.

¹These authors contributed equally to this work. Vu Dinh is now affiliated with the Department of Mathematical Sciences, University of Delaware, Newark, Delaware, USA. Lam Si Tung Ho is now affiliated with the Department of Mathematics and Statistics, Dalhousie University, Halifax, Nova Scotia, Canada.

²Supported in part by National Institutes of Health Grants R01 AI107034 and U54 GM111274 and National Science Foundation Grants DMS-1223057, DMS-1264153 and CISE 1564137.

³Supported in part by a Faculty Scholar grant from the Howard Hughes Medical Institute and the Simons Foundation.

MSC2010 subject classifications. Primary 05C05, 62F12; secondary 92B10, 92D15.

Key words and phrases. Phylogenetics, tree reconstruction, gene tree, species tree, maximum likelihood estimator, regularization.

1. Introduction. Molecular phylogenetics is the reconstruction of evolutionary history from molecular sequences, typically DNA sampled from the present day. One common means of performing molecular phylogenetics is likelihood-based “gene tree analysis,” in which the molecular sequence of a gene is assumed to evolve according to a continuous time Markov chain (CTMC) along the branches of an unknown phylogenetic tree. The likelihood function corresponding to this CTMC is then used as the sole criterion to choose the “best” tree.

However, sometimes there is insufficient signal in the gene sequences to deliver a confident estimate. This may happen because of strong genetic conservation, insufficient evolutionary time, short sequences or other processes obscuring the tree signal. Insufficient signal, in turn, results in an error-prone high variance estimator in which small modifications of the input data can result in rather different inferred trees.

The introduction of an additional penalty parameter in the optimality function through regularization is a common statistical response to such ill-posed problems. This penalty parameter can introduce extra information into the problem, resulting in a lower variance estimator which in many cases can be shown to be an improvement over the raw estimator. Regularization-type strategies have made a recent appearance in phylogenetics, via the fact that gene trees evolve within an organism-level species tree. The gene tree history may differ from a species tree due to processes such as gene duplications, losses and horizontal gene transfer; however, despite these processes there is significant signal bringing the various gene trees together to their shared species tree [Boussau and Daubin (2010)].

Phylogenetic regularization has been implemented using two approaches thus far. The model-based approach to phylogenetic regularization works to model the entire process of gene and species tree development, resulting in a comprehensive joint likelihood function for an ensemble of gene trees, either conditioning on a species tree or including it as part of the joint likelihood function [Åkerborg et al. (2009), Boussau et al. (2013), Heled and Drummond (2010), Liu and Pearl (2007), Rasmussen and Kellis (2011), Mahmudi et al. (2013), Szöllősi et al. (2015)]. Consistency of such approaches is a matter of continuing research [Roch and Warnow (2015)]. Another approach is to use a penalty function describing the level of divergence of a given gene tree from a prespecified species tree [Bansal et al. (2015), David and Alm (2011), Scornavacca, Jacox and Szöllősi (2015), Wu et al. (2013)], an approach more typical of the regularization approaches common in statistics and machine learning. A relatively simple penalization approach can give performance competitive with inference under a full probabilistic model at a substantially smaller computational cost [Wu et al. (2013)].

In another direction, probabilists have built a powerful theory showing consistency of, and describing sequence length requirements for, accurate phylogenetic inference. For individual gene tree inference without additional information, this has recently reached a high-water mark with the demonstration that accurate maximum likelihood (ML) inference can be done with sequences of length polynomial

in the number of taxa [Roch and Sly (2017)], which is asymptotically tight given previous results [Mossel (2003)].

However, we are not aware of any work giving proofs of consistency for penalty-based methods of gene tree inference, or more generally giving sequence length requirements for gene tree inference in the presence of a species tree. In addition, the work that has been done showing polynomial sequence length requirements for ML inference has reduced the inference problem to a purely combinatorial one by discretizing branch lengths rather than treating branch lengths as continuous parameters.

In this paper, we propose a regularization framework for gene tree reconstruction and prove consistency of our estimation method as sequence length increases to infinity. Furthermore, we provide the corresponding sequence length requirements for accurate estimation. The convergence theory is simultaneously for continuous branch lengths and tree topology. Our penalty is in terms of distance from a “guide tree,” often taken as, and in this paper called, the species tree, although we make no assumptions about how the gene tree was generated from the guide tree. In fact, the guide tree can be arbitrarily chosen, although of course better gene tree reconstructions can be expected from species trees closer to the correct gene tree.

Our penalty is in terms of geodesic distance in the Billera–Holmes–Vogtmann (BHV) space, which represents the discrete and continuous aspects of phylogenetic trees as an ensemble of orthants glued together along their edges [Billera, Holmes and Vogtmann (2001)]. The discrete changes induced by the geodesic distance are nearest-neighbor-interchanges between subtrees separated by zero length branches, modeling horizontal transfer, lineage sorting or deep coalescence and gene duplication/extinction [Maddison (1997)]. The BHV distance can be computed in polynomial time [Owen (2011)], in contrast with other distances on tree spaces which are typically NP-hard to compute. BHV space has been previously used to extend classical statistical methods for phylogenetic inference. For example, Nye (2011) proposes an algorithm to identify principal paths in BHV space analogously to the standard principal component analysis.

Using this framework, we are able to prove that the regularized ML estimator is consistent, and derive its *global error* in various settings. Our results show a polynomial sequence length requirement to recover the true topology, which is the best possible sample complexity of any tree reconstruction methods. Moreover, unlike most of the previous algorithms that perform poorly on trees with indistinguishable edges, the regularized estimator can reconstruct all edges of length greater than any given threshold. We can also derive a confidence region for the estimator and use that to assess the support of tree splits from data.

The paper is organized as follows. Section 2 introduces our mathematical framework including the BHV space of phylogenetic trees and the regularized maximum likelihood estimator for tree reconstruction using geodesic BHV distance. We investigate properties of the geodesic distance in the BHV space and the likelihood

in Section 3 and Section 4, respectively. We prove that the regularized maximum likelihood estimator is consistent and obtain its convergence rate in Section 5. Section 6 derives an explicit bound of the convergence rate and provides the sample complexity of our method to recover the true topology under the Jukes–Cantor model. We apply our method to reconstruct the gene tree for eight yeast species in Section 7. Finally, Section 8 discusses our results in the context of related work.

2. Mathematical framework.

2.1. *Phylogenetic tree.* Throughout this paper, the term *phylogenetic tree* refers to a tree T with leaves labeled by a set of taxon (i.e., organism) names. Phylogeneticists often call these trees *unrooted* to highlight the undirected nature of the edges. Let $E(T)$ denote the edges of T and $V(T)$ denote the vertices. Any edge adjacent to a leaf is called a pendant edge, and any other edge is called an internal edge. Each edge e will be associated with a nonnegative number w_e called the branch length. We define the *diameter* of T as $\max_{u,v \in V(T)} \sum_{e \in \text{path}(u,v)} w_e$. A tree is said to be resolved if it is bifurcating and all branch lengths are positive.

The topological distance $d_T(u, v)$ between vertices u and v in a tree is the number of edges in the path between vertices u and v . For an edge e of T , let T_1 and T_2 be the two rooted subtrees of T obtained by deleting edge e from T , and for $i = 1, 2$, let $d_i(e)$ be the topological distance from the root of T_i to its nearest leaf in T_i . The *depth* of T is defined as $\max_e \max\{d_1(e), d_2(e)\}$, where e ranges over all internal edges in T . Given an unrooted phylogenetic tree T on a finite set X of taxa, any subset Y of X induces a phylogenetic tree on taxon set Y , denoted $T|_Y$, which, is the subtree of T that connects the taxa in Y only.

We will also follow the standard independent and identically distributed (IID) setting for likelihood-based phylogenetics with a finite number of sites [see Felsenstein (2004), for more details]. Let \mathcal{S} denote the set of possible molecular sequence character states and let $r = |\mathcal{S}|$; for convenience, we assume that the states have indices 1 to r . We assume that mutation events occur according to a continuous-time Markov chain (CTMC) on states \mathcal{S} . Specifically, the probability of ending in state y after time t given that the site started in state x is given by the xy th entry of $P(t)$, where $P(t)$ is the matrix valued function $P(t) = e^{Qt}$, and the matrix Q is the instantaneous rate matrix of the CTMC evolutionary model. The branch lengths represent the times t during which the mutation process operates. We assume that the rate matrix Q is reversible with respect to a stationary distribution π on the set of states \mathcal{S} .

To explore the set of all trees with a given number of taxa and to describe trees that are “near” to each other, we use the class of nearest neighbor interchange (NNI) moves [Robinson (1971)]. A NNI move is defined as a transformation that collapses an interior edge to zero and then expands the resulting degree 4 vertex into an edge and two degree 3 vertices in a new way (see Figure 1). Two trees τ_1 and τ_2 are said to be NNI-adjacent if there exists a single NNI move that transform τ_1 into τ_2 .

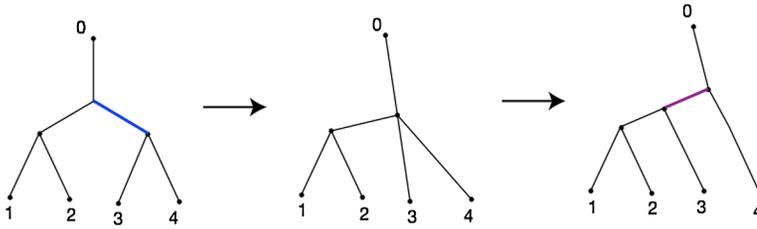


FIG. 1. A NNI move collapses an interior edge to zero and then expands the resulting degree 4 vertex into an edge and two degree 3 vertices in a new way.

2.2. *BHV space.* The BHV space employs a cubical complex as the geometric model of tree space \mathcal{T} on n taxa as follows:

1. \mathcal{T} consists of a collection of orthants, each isomorphic to $\mathbb{R}_{\geq 0}^{2n-3}$.
2. Each orthant itself corresponds uniquely to a tree topology, and the coordinates in each orthant parameterize the branch lengths for the corresponding tree.
3. The adjacent orthants of the complex with the same dimension correspond to NNI-adjacent trees.

A simple visualization of a part of the BHV space of all trees with 5 taxa is provided in Figure 2. The BHV space is not a manifold, but it is equipped with

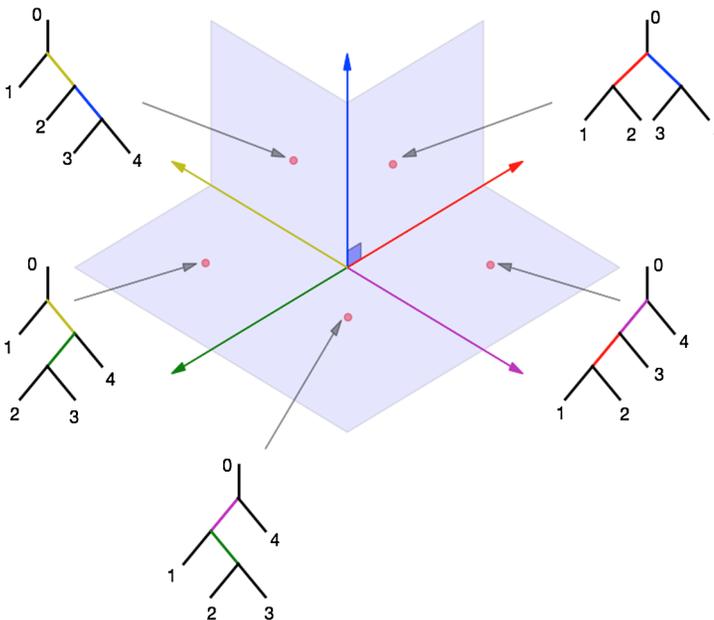


FIG. 2. The BHV space is a cubical complex, where each orthant corresponds uniquely to a tree topology and the coordinates in each orthant parameterize its branch lengths. Adjacent orthants represent trees that can be transformed to each other via a single NNI move.

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 points is defined as the minimum of the lengths of such segmented paths joining the two points.

A segmented path giving the smallest distance between two points is called a *geodesic*. The geodesic connecting any two points in BHV space is unique and can be computed in polynomial time [Owen (2011)]. The space is clearly locally compact.

In this paper, beside the BHV geodesic distance itself, we will also consider the branch-score (BS) distance. This distance was proposed by Kuhner and Felsenstein (1994) and shown by Amenta et al. (2007) to be an approximation to the BHV distance. We say that a bipartition or *split* of $V(T)$ (into two disjoint subset A, B) is in the tree T if there is some edge $g \in E(T)$ such that all elements of A lie on one side of g and all elements of B lie on the other side. The branch-score distance between two trees T_1 and T_2 is defined as $d_{BS}(T_1, T_2) = \sqrt{\sum_e (w_1(e) - w_2(e))^2}$ where $w_i(e)$ is the length of the corresponding edge if the split e is in the tree T_i , and $w_i(e) = 0$ otherwise [Kuhner and Felsenstein (1994)]. The branch-score distance is equivalent to the BHV distance $d_{BS}(T_1, T_2) \leq d_{BHV}(T_1, T_2) \leq \sqrt{2}d_{BS}(T_1, T_2)$, $\forall T_1, T_2$ and can be computed in $O(n)$ time [Amenta et al. (2007)].

2.3. *Phylogenetic likelihood and the forward operator.* For a fixed topology τ , vector of branch lengths q and observed sequences $\mathbf{Y}_k = (Y_1, Y_2, \dots, Y_k)$ in $\mathcal{S}^{n \times k}$ of length k over n taxa, the likelihood of observing \mathbf{Y}_k given τ and q has the form

$$L_k(\tau, q) = \prod_{s=1}^k \sum_{a^{(s)}} \left(\prod_{(u,v) \in E(\tau, q)} P_{a_u^{(s)} a_v^{(s)}}^{uv} (q_{uv}) \right) \pi(a_{u_0}^{(s)}),$$

where $a^{(s)}$ ranges over all extensions of Y_s to the internal nodes of the tree, $a_u^{(s)}$ denotes the assigned state of node u by $a^{(s)}$, $E(\tau, q)$ denotes the set of tree edges, P^{uv} is the probability transition matrix on edge (u, v) and u_0 is an arbitrary internal node.

We will also denote $\ell_k = \log(L_k)$ and refer to it as the log-likelihood function given the observed sequence. It is known that ℓ_k is continuous on \mathcal{T} and is smooth (up to the boundary) on each orthant of \mathcal{T} . However, there is no general notion of differentiability of ℓ_k on the whole tree space.

Each tree (τ, q) generates a joint distribution on the single-site patterns at the leaf nodes, hereafter denoted by $P_{\tau, q}$. Define the *forward operator*

$$F : \mathcal{T} \rightarrow \mathbb{X}$$

$$(\tau, q) \mapsto P_{\tau, q},$$

where \mathbb{X} denotes the space of all possible single-site distributions on the leaves. Throughout the paper, we will use the notation $\text{KL}(P_1, P_2)$ for the Kullback–Leibler divergence of $P_2 \in \mathbb{X}$ from $P_1 \in \mathbb{X}$. We note that the likelihood function can be rewritten in term of the forward operator as

$$L_k(\tau, q) = \prod_{s=1}^k P_{\tau, q}(Y_s).$$

In this paper, we mostly focus on the case when the sequence length k increases while the number of taxa n is fixed, with the exception of Section 6 where we quantify the distance between our estimate and the true tree in both k and n . We will make the following assumptions.

ASSUMPTION 2.1 (Identifiability). $P_{\tau, q} = P_{\tau', q'} \Leftrightarrow \tau = \tau'$ and $q = q'$.

ASSUMPTION 2.2. We assume that the data \mathbf{Y}_k are generated from a true tree (τ^*, q^*) , a resolved tree with bounded branch lengths: let $e, f > 0$ be the lower bound of branch lengths of pendant edges and inner edges, respectively, and $g > 0$ be the upper bound of all branch lengths. Note that e, f and g may depend on the number of taxa. We further assume that e and g are known.

From now on, we denote $\mathcal{T}_{e, g}$ be the set of all trees for which edge lengths are bounded from above by g and pendant edge lengths are bounded from below by e .

2.4. *Phylogenetic inference via regularization.* Given a species tree (τ°, q°) and regularization parameter $\alpha_k > 0$, the regularized estimator $(\hat{\tau}_k, \hat{q}_k)$ of the tree is defined as the minimizer of the BHV-penalized phylogenetic likelihood function:

$$(2.1) \quad (\hat{\tau}_k, \hat{q}_k) := \operatorname{argmin}_{(\tau, q) \in \mathcal{T}} -\frac{1}{k} \ell_k(\tau, q) + \alpha_k R(\tau, q).$$

Here, $R(\tau, q) = d((\tau, q), (\tau^\circ, q^\circ))^2$ where $d(s, s')$ denotes the geodesic distance between the trees s and s' on the BHV space [Billera, Holmes and Vogtmann (2001)]. This formulation is analogous to the ridge estimator in the standard linear regression setting [Hoerl (1962)]. The existence of a minimizer as in (2.1) is guaranteed by the following lemma.

LEMMA 2.1 (Existence of the estimator). *For fixed $\beta > 0$ and observed sequence \mathbf{Y}_k , there is a (τ, q) minimizing $Z_{\beta, \mathbf{Y}_k}(\tau, q) = -\ell_k(\tau, q) + \beta R(\tau, q)$.*

PROOF. Let (τ_n, q_n) be a sequence such that $Z_{\beta, \mathbf{Y}_k}(\tau_n, q_n)$ converges to $\inf_{(\tau, q)} Z_{\beta, \mathbf{Y}_k}(\tau, q)$. Note that $d((\tau_n, q_n), (\tau^\circ, q^\circ))$ is bounded from above by Assumption 2.2, so the sequence lies inside a compact set. We deduce that a subsequence (τ_m, q_m) converges to some $(\tau^*, q^*) \in \mathcal{T}$. Since the likelihood and the penalty R are continuous, we deduce that (τ^*, q^*) is a minimizer of Z_{β, \mathbf{Y}_k} . \square

3. Properties of the regularization penalty. In this section, we investigate the analytical properties of the penalty in our regularization problem. We recall the definition of a strongly convex distance on a geodesic space.

DEFINITION 3.1 (Strongly convex distance). A distance d on a space \mathcal{X} is strongly convex if for any geodesic $\gamma : [0, 1] \rightarrow \mathcal{X}$, $x \in \mathcal{X}$ and $t \in [0, 1]$, we have $d(x, \gamma(t))^2 \leq (1 - t)d(x, \gamma(0))^2 + td(x, \gamma(1))^2 - t(1 - t)d(\gamma(0), \gamma(1))^2$.

We have the following lemma about the convexity of the aforementioned distances.

LEMMA 3.1. *The BHV geodesic distance is strongly convex on the whole tree space.*

PROOF. It is known that the BHV space with its distance is a CAT(0) space [Billera, Holmes and Vogtmann (2001)]. Standard results about distance on CAT(0) spaces [e.g., see Bačák (2013, 2014)] imply strong convexity of the geodesic distance. \square

Recall that in our formulation, $R(\tau, q) = d((\tau, q), (\tau^\circ, q^\circ))^2$. Letting $x^\circ := (\tau^\circ, q^\circ)$ denote our species tree with branch lengths, we note that $R(x)$ is locally Lipschitz by two applications of the triangle inequality:

$$(3.1) \quad \begin{aligned} |R(y) - R(x)| &= |d(y, x^\circ) - d(x, x^\circ)|(d(y, x^\circ) + d(x, x^\circ)) \\ &\leq d(y, x)(d(y, x) + 2d(x, x^\circ)). \end{aligned}$$

We will use the notation $\gamma_{x,y}(t)$, $t \in (0, 1)$ to denote the linear parameterization of the geodesic going from x to y in \mathcal{T} . Thus

$$d(x, \gamma_{x,y}(t)) = td(x, y), \quad \text{and} \quad d(y, \gamma_{x,y}(t)) = (1 - t)d(x, y).$$

The following lemma describes a directional derivative ω of R along geodesics, as well as a function D that will be useful later on for analyzing convergence properties of the regularization term.

LEMMA 3.2. *Define*

$$\omega(x, y) := \lim_{t \rightarrow 0^+} \frac{R(\gamma_{x,y}(t)) - R(x)}{t} \quad \forall x, y \in \mathcal{T}, \quad \text{and}$$

$$D(x, y) := R(y) - R(x) - \omega(x, y).$$

Then:

1. $\omega(x, y)$ is well defined.
2. For any $x \in \mathcal{T}$, $|\omega(x, y)| \leq 2d(x, x^\circ)d(x, y) \quad \forall y \in \mathcal{T}$.
3. For all $x, y \in \mathcal{T}$, $D(x, y) \geq d(x, y)^2$.

PROOF. 1. For any $x, y \in \mathcal{T}$, we recall that the geodesic between x and y is composed of straight segments. Consider the function

$$g(t) = \frac{R(\gamma_{x,y}(t)) - R(x)}{t}, \quad t \in (0, 1).$$

Let $\lambda \in (0, 1)$. By convexity of R , we have

$$R(\gamma_{x,y}(\lambda t)) = R(\gamma_{x,\gamma_{x,y}(t)}(\lambda)) \leq (1 - \lambda)R(x) + \lambda R(\gamma_{x,y}(t)).$$

This implies that

$$g(\lambda t) = \frac{R(\gamma_{x,y}(\lambda t)) - R(x)}{\lambda t} \leq \frac{R(\gamma_{x,y}(t)) - R(x)}{t} = g(t).$$

Hence, $g(t)$ decreases as $t \downarrow 0$. Since R is locally Lipschitz, we also have that $g(t)$ is bounded. Thus the limit of $g(t)$ as $t \downarrow 0$ exists and $\omega(x, y)$ is well defined.

2. By (3.1), we have

$$\begin{aligned} |\omega(x, y)| &\leq \lim_{t \rightarrow 0^+} \frac{|R(\gamma_{x,y}(t)) - R(x)|}{t} \\ &\leq \lim_{t \rightarrow 0^+} \frac{(td(x, y) + 2d(x, x^\circ))td(x, y)}{t} = 2d(x, x^\circ)d(x, y). \end{aligned}$$

3. Lemma 3.1 also implies

$$\frac{R(\gamma_{x,y}(t)) - R(x)}{t} \leq R(y) - R(x) - (1 - t)d(x, y)^2.$$

By letting t go to 0, we obtain $R(y) - R(x) - \omega(x, y) \geq d(x, y)^2$. \square

4. Properties of the likelihood. Next, we investigate properties of an “expected per-site log likelihood”

$$\phi(\tau, q) := \mathbb{E}_{\psi \sim P(\tau^*, q^*)}[\log P_{\tau, q}(\psi)]$$

and construct probabilistic bounds on deviation of the empirical average per-site log likelihood from this expected quantity; these bounds are uniform over $\mathcal{T}_{e, g}$. Recalling Assumption 2.2, we have the following lemma.

LEMMA 4.1 (Limit likelihood). (τ^*, q^*) is the unique maximizer of ϕ , and $\ell_k(\tau, q | \mathbf{Y}_k) / k \rightarrow \phi(\tau, q), \forall (\tau, q) \in \mathcal{T}$.

PROOF. Note that the sequences \mathbf{Y}_k are IID samples of $P(\tau^*, q^*)$. By the strong law of large numbers,

$$\frac{1}{k} \ell_k(\tau, q | \mathbf{Y}_k) = \frac{1}{k} \sum_{i=1}^k \log P_{\tau, q}(Y_i) \longrightarrow \phi(\tau, q),$$

for all $(\tau, q) \in \mathcal{T}$. Moreover,

$$\begin{aligned} \phi(\tau, q) - \phi(\tau^*, q^*) &= \mathbb{E}_{\psi \sim P(\tau^*, q^*)} [\log P_{\tau, q}(\psi) - \log P_{\tau^*, q^*}(\psi)] \\ &= -\text{KL}(F(\tau^*, q^*), F(\tau, q)) \leq 0 \end{aligned}$$

with equality if and only if $P_{\tau, q} = P_{\tau^*, q^*}$. Hence, Assumption 2.1 guarantees that (τ^*, q^*) is the unique maximizer of ϕ . \square

Next, we employ the ‘‘covering number’’ technique from machine learning [see, e.g., Cucker and Smale (2002), Cuong, Ho and Dinh (2013)] to obtain uniform bounds for the distance between the likelihood and its expected value.

LEMMA 4.2. *Let γ be the element of largest magnitude in the rate matrix Q . There exists a constant $C > 0$ such that for any $k \geq 3$, $\delta > 0$, we have*

$$\begin{aligned} &\left| \frac{1}{k} \ell_k(\tau, q | \mathbf{Y}_k) - \phi(\tau, q) \right| \\ &\leq C \left(\frac{\log k}{k} \right)^{1/2} \left(n^2 \log \frac{1}{\delta} + n^3 \log(n\gamma) + n^4 \log \frac{4}{\lambda_e} \right)^{1/2} \log \frac{1}{\lambda_e} \end{aligned}$$

for all $(\tau, q) \in \mathcal{T}_{e, g}$ with probability greater than $1 - \delta$. Here, λ_e is a constant depending on e .

PROOF. For every assignment ψ of states to the tree taxa, we have

$$P_{\tau, q}(\psi) = \sum_{a(\psi)} \left(\prod_{(u, v) \in E(\tau, q)} P_{a_u a_v}^{uv}(q_{uv}) \right) \pi(a_{u_0}),$$

where a ranges over all extensions of ψ to the internal nodes of the tree. Note that

$$\left| \frac{\partial P(\psi)}{\partial q_i}(q) \right| \leq \sum_{a(\psi)} \left(\prod_{(u, v) \in E(\tau, q) \setminus i} P_{a_u a_v}^{uv}(q_{uv}) \left| \sum_s Q_{a_u s} P_{s a_v}(q_i) \right| \right) \pi(a_{u_0}) \leq \gamma 4^n,$$

where the bound is obtained from the fact that all probability terms are at most 1, and that the total number of assignments a to the $n - 1$ inner nodes is 4^{n-1} .

On the other hand, since the lengths of all pendant edges are bounded from below by e (Assumption 2.2), $P_{\tau, q}(\psi)$ is bounded from below by the probability of the assignment when we set the value at the internal nodes to be a constant i_0 , thus

$$P_{\tau, q}(\psi) \geq \left(\min_{t \geq 0} P_{i_0 i_0}(t) \right)^{n-3} \left(\min_i \min_{t \geq e} P_{i_0 i}(t) \right)^n \pi(i_0).$$

We define

$$\lambda_e = \left(\min_{t \geq 0} P_{i_0 i_0}(t) \right)^{(n-3)/n} \left(\min_i \min_{t \geq e} P_{i_0 i}(t) \right) \pi(i_0)^{1/n}.$$

By the mean value theorem, we have

$$|\log P_{\tau,q}(\psi) - \log P_{\tau,q'}(\psi)| \leq c_1 \|q - q'\| \quad \forall \tau, q, q', \psi,$$

where $c_1 := \gamma 4^n / \lambda_e^n$, and $\|\cdot\|$ is the Euclidean distance in \mathbb{R}^{2n-3} . This implies

$$(4.1) \quad \left| \frac{1}{k} \ell_k(\tau, q | \mathbf{Y}_k) - \frac{1}{k} \ell_k(\tau, q' | \mathbf{Y}_k) \right| \leq c_1 \|q - q'\|$$

and

$$(4.2) \quad |\phi(\tau, q) - \phi(\tau, q')| \leq c_1 \|q - q'\|$$

for all τ, q, q' .

For each $(\tau, q) \in \mathcal{T}$, define the events

$$A(\tau, q, k, \varepsilon) = \left\{ \left| \frac{1}{k} \ell_k(\tau, q | \mathbf{Y}_k) - \phi(\tau, q) \right| > \varepsilon \right\}$$

and

$$B(\tau, q, k, \varepsilon)$$

$$= \left\{ \exists q' \text{ such that } \|q' - q\| \leq \frac{\varepsilon}{2c_1} \text{ and } \left| \frac{1}{k} \ell_k(\tau, q' | \mathbf{Y}_k) - \phi(\tau, q') \right| > 2\varepsilon \right\}$$

then we have $B(\tau, q, k, \varepsilon) \subset A(\tau, q, k, \varepsilon)$ by the triangle inequality, (4.1), and (4.2). We note that $0 \geq \log P_{\tau,q}(\psi) \geq n \log \lambda_e, \forall \tau, q, \psi$. By Hoeffding’s inequality [Hoeffding (1963)], we have

$$(4.3) \quad \mathbb{P}[A(\tau, q, k, \varepsilon)] \leq 2 \exp\left(\frac{-2\varepsilon^2 k}{n^2(\log \lambda_e)^2}\right).$$

Note that the total number of balls of radius $\varepsilon/2c_1$ required to cover $[0, g]^{2n-3}$ is bounded above by

$$(4.4) \quad \left(\frac{2gc_1}{\varepsilon}\right)^{2n-3} (2n-3)!! = \left(\frac{2gc_1}{\varepsilon}\right)^{2n-3} e^{O(n \log n)}.$$

We deduce that for some $c_2 > 0$,

$$\begin{aligned} & \mathbb{P}\left[\exists(\tau, q) \in \mathcal{T}_{e,g} : \left| \frac{1}{k} \ell_k(\tau, q | \mathbf{Y}_k) - \phi(\tau, q) \right| > 2\varepsilon\right] \\ & \leq 2 \exp\left(\frac{-2\varepsilon^2 k}{n^2(\log \lambda_e)^2}\right) \left(\frac{2gc_1}{\varepsilon}\right)^{2n-3} e^{c_2 n \log n}. \end{aligned}$$

To make the left-hand side less than δ , we need

$$\begin{aligned} & \log(2) + (2n-3) \log(2gc_1) + c_2 n \log n + \log \frac{1}{\delta} \\ & \leq \frac{2\varepsilon^2 k}{n^2(\log \lambda_e)^2} - (2n-3) \log \frac{1}{\varepsilon}. \end{aligned}$$

We will choose ε such that

$$\frac{\varepsilon^2 k}{n^2 (\log \lambda_e)^2} \geq \log(2) + (2n - 3) \log(2g c_1) + c_2 n \log n + \log \frac{1}{\delta}$$

and $\varepsilon^2 k / [n^2 (\log \lambda_e)^2] \geq (2n - 3) \log(1/\varepsilon)$, which is valid if we choose

$$\varepsilon = \sqrt{2 + c_2} \sqrt{\frac{\log k}{k}} \left(n^2 \log \frac{1}{\delta} + n^3 \log(n g \gamma) + n^4 \log \frac{4}{\lambda_e} \right)^{1/2} \log \frac{1}{\lambda_e}. \quad \square$$

LEMMA 4.3. For any $0 < \delta < 1$, denote

$$C_{n,\delta,\gamma,e,g} = C \left(n^2 \log \frac{\pi^2}{6\delta} + n^3 \log(n g \gamma) + n^4 \log \frac{4}{\lambda_e} \right)^{1/2} \log \frac{1}{\lambda_e},$$

where C is the constant in Lemma 4.2. Then

$$\left| \frac{1}{k} \ell_k(\tau, q | \mathbf{Y}_k) - \phi(\tau, q) \right| \leq \frac{C_{n,\delta,\gamma,e,g} \log k}{\sqrt{k}} \quad \forall (\tau, q) \in \mathcal{T}_{e,g}, k \geq 3$$

with probability greater than $1 - \delta$.

PROOF. Denote

$$A_k = \left\{ \left| \frac{1}{k} \ell_k(\tau, q | \mathbf{Y}_k) - \phi(\tau, q) \right| \leq C_{n,\delta/k^2,\gamma,e,g} \sqrt{\frac{\log k}{k}}, \forall (\tau, q) \in \mathcal{T}_{e,g} \right\}.$$

By Lemma 4.2, $\mathbb{P}(A_k^c) \leq 6\delta / (\pi^2 k^2)$ where A_k^c is the complement set of A_k . Therefore,

$$\begin{aligned} \mathbb{P}\left(\bigcap_{k=3}^{\infty} A_k\right) &= 1 - \mathbb{P}\left(\bigcup_{k=3}^{\infty} A_k^c\right) \geq 1 - \sum_{k=3}^{\infty} \mathbb{P}(A_k^c) \\ (4.5) \qquad \qquad \qquad &\geq 1 - \sum_{k=3}^{\infty} \frac{6\delta}{\pi^2 k^2} = 1 - \delta. \end{aligned} \quad \square$$

5. Asymptotic theory of the regularized estimator.

5.1. Consistency.

THEOREM 5.1 (Consistency). Assume that the sequences \mathbf{Y}_k are generated from a tree $(\tau^*, q^*) \in \mathcal{T}_e$ and that

$$\alpha_k \rightarrow 0 \quad \text{and} \quad \alpha_k = \Omega\left(\frac{\log k}{\sqrt{k}}\right).$$

Then $(\hat{\tau}_k, \hat{q}_k)$ converges to (τ^*, q^*) almost surely.

PROOF. From the definition of $(\hat{\tau}_k, \hat{q}_k)$, we have

$$-\frac{1}{k}\ell(\hat{\tau}_k, \hat{q}_k | \mathbf{Y}_k) + \alpha_k R(\hat{\tau}_k, \hat{q}_k) \leq -\frac{1}{k}\ell(\tau^*, q^* | \mathbf{Y}_k) + \alpha_k R(\tau^*, q^*),$$

which implies that

$$\begin{aligned} \alpha_k R(\hat{\tau}_k, \hat{q}_k) &\leq -\frac{1}{k}\ell(\tau^*, q^* | \mathbf{Y}_k) + \phi(\tau^*, q^*) - \phi(\tau^*, q^*) \\ &\quad + \phi(\hat{\tau}_k, \hat{q}_k) - \phi(\hat{\tau}_k, \hat{q}_k) + \frac{1}{k}\ell(\hat{\tau}_k, \hat{q}_k | \mathbf{Y}_k) + \alpha_k R(\tau^*, q^*). \end{aligned}$$

By Lemma 4.3, with probability greater than $1 - \delta$,

$$\begin{aligned} \alpha_k R(\hat{\tau}_k, \hat{q}_k) &\leq -\phi(\tau^*, q^*) + \phi(\hat{\tau}_k, \hat{q}_k) + \frac{2C_{n,\delta,\gamma,e,g} \log k}{\sqrt{k}} + \alpha_k R(\tau^*, q^*) \\ (5.1) \quad &\leq \frac{2C_{n,\delta,\gamma,e,g} \log k}{\sqrt{k}} + \alpha_k R(\tau^*, q^*), \quad \forall k \in \mathbb{N} \end{aligned}$$

since (τ^*, q^*) is the maximizer of ϕ (Lemma 4.1).

Therefore, again with probability greater than $1 - \delta$,

$$(5.2) \quad 0 \leq R(\hat{\tau}_k, \hat{q}_k) \leq \frac{2C_{n,\delta,\gamma,e,g} \log k}{\alpha_k \sqrt{k}} + R(\tau^*, q^*), \quad \forall k \in \mathbb{N}.$$

We also deduce from the first inequality of equation (5.1) that

$$\lim_{k \rightarrow \infty} |\phi(\hat{\tau}_k, \hat{q}_k) - \phi(\tau^*, q^*)| = 0.$$

By the assumption on α_k , the right-hand side of (5.2) is bounded above. Thus the sequence $\{(\hat{\tau}_k, \hat{q}_k)\}$ lies inside a compact set, and there is a subsequence $\{(\tau_{k_m}, q_{k_m})\}$ which converges to some $(\tau', q') \in \mathcal{T}$. The continuity of the likelihood function implies that $\phi(\tau', q') = \phi(\tau^*, q^*)$. By Lemma 4.1, we deduce that $(\tau', q') = (\tau^*, q^*)$. We can repeat this argument for every subsequence of $\{(\hat{\tau}_k, \hat{q}_k)\}$ and deduce that

$$(5.3) \quad \lim_{k \rightarrow \infty} (\hat{\tau}_k, \hat{q}_k) = (\tau^*, q^*).$$

Because (5.3) holds with probability larger than $1 - \delta$ and δ can be arbitrarily small, we deduce that (5.3) holds almost surely. \square

Equation (5.1) gives us the following corollary.

COROLLARY 5.1. *There exists $C_{n,\delta,\gamma,e,g} > 0$ such that with probability larger than $1 - \delta$,*

$$\phi(\tau^*, q^*) - \phi(\hat{\tau}_k, \hat{q}_k) + \alpha_k R(\hat{\tau}_k, \hat{q}_k) \leq \frac{C_{n,\delta,\gamma,e,g} \log k}{\sqrt{k}} + \alpha_k R(\tau^*, q^*), \quad \forall k \in \mathbb{N}.$$

5.2. *Convergence rate.* While regularized estimators are consistent in general, it is well known that without any a priori information, their convergence can be arbitrarily slow [Hofmann and Yamamoto (2010), Kazimierski (2010)]. Uniform error bounds necessarily require further regularity assumptions on the (asymptotic) minimizing solution $x^* = (\tau^*, q^*)$, that is, the data source [Engl, Hanke and Neubauer (1996)]. Such conditions are referred to as source conditions and have played a central role in analyses of convergence of regularized estimators in various settings. In recent publications, starting with Hofmann et al. (2007), variational inequalities have become an increasingly popular way to formulate source conditions, especially in the case of nonsmooth operators [Hohage and Weidling (2015)]. In this section, we will follow the approach of Hofmann et al. (2007) to derive such a variational inequality, and from that, to obtain a global error estimate for our nonsmooth regularization problem.

Denote $x = (\tau, q)$ and $x^* = (\tau^*, q^*)$. For all evolutionary models, $\phi(x)$ is an analytic function on each orthant of $\mathcal{T}_{e,g}$. By the Łojasiewicz inequality [Ji, Kollár and Shiffman (1992), Theorem 1], under any identifiable evolutionary model there exists an integer $m \geq 2$, a neighborhood U around x^* , and some constant $C_n > 0$ such that

$$(5.4) \quad \phi(x^*) - \phi(x) \geq C_n d(x^*, x)^m \quad \forall x \in U.$$

We will also bound $d(x^*, x)$ in an intermediate set, defined by

$$(5.5) \quad U' = \{x : d(x^*, x) \leq 4d(x^*, x^\circ)\} \setminus U,$$

using $\eta_{x^*} = \inf_{x \in U'} (\phi(x^*) - \phi(x)) = \phi(x^*) - \sup_{x \in U'} \phi(x) > 0$. We will use these to bound the directional derivative ω of R defined in Lemma 3.2.

LEMMA 5.1 (Variational source condition). *There exists $\beta > 0$ such that*

$$-\omega(x^*, x) \leq \frac{1}{2} D(x^*, x) + \beta(\phi(x^*) - \phi(x))^{1/m} \quad \forall x \in \mathcal{T}.$$

PROOF. The proof comes from controlling $d(x^*, x)$ in three regimes: in U using (5.4), when $d(x^*, x)$ is large with Lemma 3.2, and in the intermediate regime U' . If U' is nonempty, we have

$$(5.6) \quad \phi(x^*) - \phi(x) \geq \frac{\eta_{x^*}}{[4d(x^*, x^\circ)]^m} d(x^*, x)^m, \quad \forall x \in U'.$$

Moreover, when $d(x, x^*) > 4d(x^*, x^\circ)$, we have

$$(5.7) \quad d(x, x^*) < \frac{1}{4d(x^*, x^\circ)} d(x^*, x)^2.$$

Combining equations (5.4), (5.6) and (5.7) and using Lemma 3.2, we deduce that $-\omega(x^*, x) \leq 2d(x^*, x^\circ)d(x^*, x) \leq D(x^*, x)/2 + \beta(\phi(x^*) - \phi(x))^{1/m}$ where $\beta = 2d(x^*, x^\circ) \max\{\frac{1}{C_n^{1/m}}, \frac{4d(x^*, x^\circ)}{\eta_{x^*}^{1/m}}\}$. \square

This variational inequality enables us to derive the following error estimates.

THEOREM 5.2 (Global error of regularized estimator). *Let $x_k = (\hat{\tau}_k, \hat{q}_k)$ be the minimizer of (2.1) on $\mathcal{T}_{e,g}$. Denote*

$$C_{n,\delta,\gamma,e,g,m,x^*,x^\circ} = \sqrt{2} \left(C_{n,\delta,\gamma,e,g} + \frac{(m-1)\beta^{m/(m-1)}}{m^{m/(m-1)}} \right)^{1/2},$$

where $\beta = 2d(x^*, x^\circ) \max\{\frac{1}{C_n^{1/m}}, \frac{4d(x^*, x^\circ)}{\eta_{x^*}^{1/m}}\}$ and $C_{n,\delta,\gamma,e,g}$ is the constant in Lemma 4.3. For any $\delta > 0$, we have

$$d(x^*, x_k) \leq C_{n,\delta,\gamma,e,g,m,x^*,x^\circ} \left(\frac{\log k}{\alpha_k \sqrt{k}} + \alpha_k^{1/(m-1)} \right)^{1/2}$$

with probability greater than $1 - \delta$.

PROOF. From Corollary 5.1, we note that with probability greater than $1 - \delta$, we have

$$\phi(x^*) - \phi(x_k) + \alpha_k R(x_k) \leq \frac{C_{n,\delta,\gamma,e,g} \log k}{\sqrt{k}} + \alpha_k R(x^*), \quad \forall k \in \mathbb{N}$$

which implies

$$(5.8) \quad \begin{aligned} & \phi(x^*) - \phi(x_k) + \alpha_k D(x^*, x_k) \\ & \leq \frac{C_{n,\delta,\gamma,e,g} \log k}{\sqrt{k}} + \alpha_k (R(x^*) - R(x_k) + D(x^*, x_k)), \quad \forall k \in \mathbb{N}. \end{aligned}$$

By Lemma 5.1, we have

$$(5.9) \quad \begin{aligned} & R(x^*) - R(x_k) + D(x^*, x_k) \\ & = -\omega(x^*, x) \leq \frac{1}{2} D(x^*, x_k) + \beta (\phi(x^*) - \phi(x_k))^{1/m}. \end{aligned}$$

Combining (5.8) and (5.9), we obtain

$$\frac{\alpha_k}{2} D(x^*, x_k) \leq \frac{C_{n,\delta,\gamma,e,g} \log k}{\sqrt{k}} + \alpha_k \beta (\phi(x^*) - \phi(x_k))^{1/m} - (\phi(x^*) - \phi(x_k)).$$

Denote $\bar{m} = m/(m-1)$. Applying Young's inequality, we get

$$\frac{1}{m} ((\phi(x^*) - \phi(x_k))^{1/m})^m + \frac{1}{\bar{m}} \left(\frac{\alpha_k \beta}{m} \right)^{\bar{m}} \geq \frac{\alpha_k \beta}{m} (\phi(x^*) - \phi(x_k))^{1/m}.$$

We deduce that

$$D(x^*, x_k) \leq 2 \left(\frac{C_{n,\delta,\gamma,e,g} \log k}{\alpha_k \sqrt{k}} + \frac{(m-1)}{\alpha_k} \left(\frac{\alpha_k \beta}{m} \right)^{\bar{m}} \right)$$

and the proof is complete by another application of Lemma 3.2. \square

By choosing $\alpha_k = k^{(1-m)/(2m)}$ and applying Theorem 5.2, we have the following corollary.

COROLLARY 5.2. For any $\delta > 0$, we have $d(x^*, x_k) = \mathcal{O}((\log k)/(\sqrt[4m]{k}))$ with probability greater than $1 - \delta$.

Next, we describe a sufficient condition for m achieving its minimum of 2, which we then show holds everywhere for some simple phylogenetic models. Define the expected Fisher information matrix $I_\tau(q)$ by

$$I_\tau(q)_{ij} = -\mathbb{E}_{\psi \sim P(\tau^*, q^*)} \left[\frac{\partial^2}{\partial q_i \partial q_j} \log P_{\tau, q}(\psi) \right].$$

Since ϕ is analytic around (τ^*, q^*) and $\nabla \phi(\tau^*, q^*) = 0$, we deduce that around (τ^*, q^*) ,

$$\phi(\tau^*, q) - \phi(\tau^*, q^*) = -(q - q^*)^t I_{\tau^*}(q^*)(q - q^*) + o(\|q^* - q\|^2).$$

This gives a sufficient condition for $m = 2$.

LEMMA 5.2. Assume that the Fisher information matrix $I_{\tau^*}(q)$ is positive-definite at q^* . Then there exists a neighborhood U around (τ^*, q^*) and some constant $C_n > 0$ such that $\phi(\tau^*, q^*) - \phi(\tau^*, q) \geq C_n \|q^* - q\|^2 \forall (\tau^*, q) \in U$.

The condition that the Fisher information matrix is positive-definite at q^* is not restrictive. Indeed, the matrix $I = I_{\tau^*}(q^*)$ can be rewritten in the form

$$I_{ij} = \mathbb{E} \left[\left(\frac{\partial}{\partial q_i} \log P_{\tau^*, q^*}(\psi) \right) \left(\frac{\partial}{\partial q_j} \log P_{\tau^*, q^*}(\psi) \right) \right].$$

For every nonzero vector $u \in \mathcal{R}^n$, we have

$$\begin{aligned} \sum_{ij} u_i I_{ij} u_j &= \sum_{ij} u_i \mathbb{E} \left[\left(\frac{\partial}{\partial q_i} \log P_{\tau^*, q^*}(\psi) \right) \left(\frac{\partial}{\partial q_j} \log P_{\tau^*, q^*}(\psi) \right) \right] u_j \\ (5.10) \quad &= \mathbb{E} \left[\left(\sum_i u_i \frac{\partial}{\partial q_i} \log P_{\tau^*, q^*}(\psi) \right) \left(\sum_j u_j \frac{\partial}{\partial q_j} \log P_{\tau^*, q^*}(\psi) \right) \right] \\ &= \mathbb{E} \left[\left(\sum_i u_i \frac{\partial}{\partial q_i} \log P_{\tau^*, q^*}(\psi) \right)^2 \right] \geq 0 \end{aligned}$$

with equality when

$$(5.11) \quad \sum_i u_i \frac{\partial}{\partial q_i} \log P_{\tau^*, q^*}(\psi) = 0 \quad \forall \psi.$$

If we consider the mapping

$$(5.12) \quad \begin{aligned} G: \quad \Delta &\rightarrow \mathbb{R}^{4^n} \\ [G(\tau, q)]_\psi &= \log P_{\tau, q}(\psi). \end{aligned}$$

on some compact neighborhood $\Delta \subset \mathcal{T}$ of (τ^*, q^*) that is contained in the same orthant as (τ^*, q^*) , then (5.11) implies that each column of the Jacobian of the map G at q^* is orthogonal to u . Thus, in this case the Jacobian $J_G(q^*)$ is rank-deficient. Hence, in order to verify the criteria of Lemma 5.2, we just need to prove that $J_G(q^*)$ has full rank. This condition can be verified for various classes of model including the r -state symmetric models [Semple and Steel (2003)] and the popular Felsenstein 1984 (F84) model first implemented in the DNAML software program [Felsenstein (1984)].

LEMMA 5.3 (Information-regularity). *The Fisher information matrix $I_\tau(q)$ is everywhere positive-definite for r -state symmetric models (e.g., Jukes–Cantor), and F84.*

PROOF. Wang (2004) proves that $I_\tau(q)$ is positive definite for F84. For r -state symmetric models, we note that G [as defined in (5.12)] is continuous and injective. This shows that G^{-1} is well defined and continuous. If G^{-1} is smooth, then the Jacobian of G at q^* has full rank. This can be shown by the following commutative diagram:

$$\begin{array}{ccc} \mathcal{A} & \overset{\text{Hadamard}}{\leftrightarrow} & \mathcal{B} \\ \downarrow r_1 & & \downarrow r_2 \\ \mathcal{T} & \xrightarrow{G} & \text{Im}[G(\cdot)] \end{array}$$

where r_1 maps (τ, q) to its edge length spectrum $Q(\tau, q) \in \mathcal{A}$, r_2 maps $G(\tau, q)$ to its corresponding sequence spectrum $\Phi \in \mathcal{B}$, and the Hadamard-exponential conjugation [Felsenstein (2004), Semple and Steel (2003)] provides a smooth conversion between the two spectra. (A detailed description of the Hadamard-exponential conjugation will be provided in the next section.) We deduce that G^{-1} is smooth and the result holds. \square

6. Special case: Jukes–Cantor model. In this section, we aim to quantify the two constants C_n and η_{x^*} defined before Lemma 5.1 to derive an explicit bound on the rate of convergence of the regularized estimator under the Jukes–Cantor model using Theorem 5.2. From that we will obtain a sequence length bound to recover the true tree topology. While we provide details for the Jukes–Cantor model, our results can be generalized to r -state symmetric models. We will take an information-theoretic perspective, working in terms of KL divergences rather than the equivalent formulation using differences in the expected per-site log likelihood function ϕ .

Without loss of generality, we assume that the substitution rate $\mu = 1$. We recall that by Lemma 5.3, equation (5.4) holds with $m = 2$.

6.1. *Convergence rate of the regularized estimator under the Jukes–Cantor model.* Recall that the constants C_n and η_{x^*} in Theorem 5.2 are defined in such a way that $\text{KL}(P_{x^*}, P_x) = \phi(x^*) - \phi(x) \geq C_n d(x^*, x)^m, \forall x \in U$, and

$$\eta_{x^*} = \inf_{x \in U'} \text{KL}(P_{x^*}, P_x) = \inf_{x \in U'} (\phi(x^*) - \phi(x)) = \phi(x^*) - \sup_{x \in U'} \phi(x),$$

for some neighborhood U around x^* and U' as defined in (5.5).

In order to quantify these constants, we need to derive lower bounds of KL divergence between the site expected site pattern frequencies of trees. The underlying ideas behind the following somewhat technical proofs can be sketched as follows. We define $U = \{x : d(x, x^*) < f\}$ where f denotes the length of the shortest edge of x^* , and consider two different scenarios:

(i) If $x \in U$, then x has the same topology as x^* . We then prove (in Lemma 6.2) that we can find four taxa such that the restriction of x^* and x to these four taxa induces two quartets s^* and s of the same topology such that

$$\frac{1}{2n - 3} d(x^*, x) \leq d(s^*, s) \leq d(x^*, x), \quad \text{and} \quad \text{KL}(P_{x^*}, P_x) \geq \text{KL}(P_{s^*}, P_s)$$

by which a bound for the 4-taxon case gives a bound for the general case. This case provides a lower bound for C_n .

(ii) If $x \in \mathcal{T}_{e,g} \setminus U$, we can construct two quartets s^* and s such that

$$\text{diam}(s^*) \leq 4g \log_2(n), \quad \text{and} \quad d(s, s^*) \geq f/(2n - 3)$$

which also reduces our bound to the 4-taxon case. The resulting bound on η_{x^*} need not depend on the distance $d(x^*, x)$. Details are provided in Lemma 6.4.

In both cases, a lower bound of $\text{KL}(P_{s^*}, P_s)$ by $d(s, s^*)$ is needed. This bound is partially derived in Lemma 6.1 under the assumption that the diameters of s^* and s are bounded from above. While this assumption works fine for case (i), an upper bound for $\text{diam}(s)$ cannot be derived for case (ii), which prompts us to consider two different subcases for the value of $\text{diam}(s)$ in the proof of Lemma 6.4.

Before proceeding to provide the detailed proofs, we note that a lower bound of $\text{KL}(P_{s^*}, P_s)$ by $d(s, s^*)$ cannot be obtained without the assumption that the diameters of the quartets are bounded. This can be seen by considering the phylogenetic “orange” [Kim (2000), Moulton and Steel (2004)]: the space of all leaf-node distributions $P_{\tau,q}$ for $(\tau, q) \in \mathcal{T}$. As sets of branch lengths become large, the corresponding point on the orange moves toward the top point of this compact set, which is that induced by independent samples from the stationary distribution on states. For that reason, without the bound on the diameter of the trees, one can easily construct two trees x_1 and x_2 such that P_{x_1} and P_{x_2} are arbitrarily close and $d(x_1, x_2)$ is arbitrarily large, rendering the bound impossible (Figure 3).

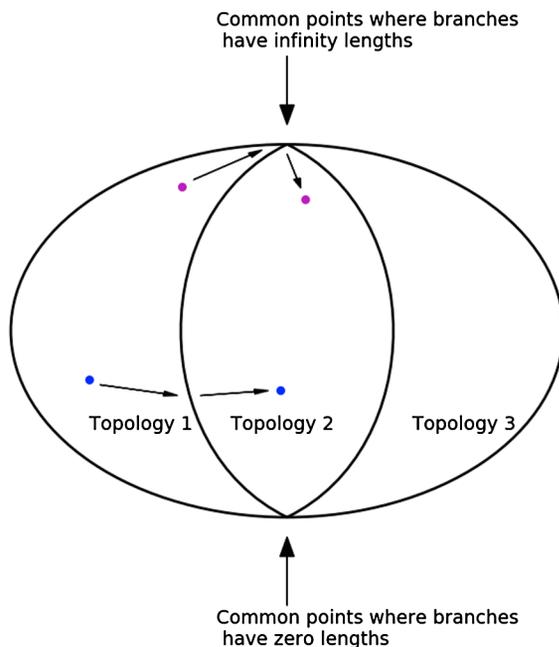


FIG. 3. Visualization of the phylogenetic orange: without the bound on the diameter of the trees, one can easily construct two trees x_1 and x_2 such that P_{x_1} and P_{x_2} are arbitrarily close and $d_{\text{BHV}}(x_1, x_2)$ is arbitrarily large.

6.1.1. *Lemmas on quartets.* However, when the diameter of quartets is bounded above, we can use Hadamard conjugation [Hendy and Penny (1993), Steel, Hendy and Penny (1998), Semple and Steel (2003), Felsenstein (2004)] to bound KL divergence below in terms of the BHV distance as follows.

LEMMA 6.1. Let \mathcal{K}_b be the BHV space of all quartets with diameter bounded from above by b . We have

$$\text{KL}(P_{\tau_1, q_1}, P_{\tau_2, q_2}) \geq 256e^{-8b} d((\tau_1, q_1), (\tau_2, q_2))^2$$

for all $(\tau_1, q_1), (\tau_2, q_2) \in \mathcal{K}_b$.

PROOF. First, we recall that the *branch length spectrum* \mathbf{a} of a quartet τ is defined as

$$\mathbf{a}_A = \begin{cases} q_e & \text{if } e \in E(\tau) \text{ induces the split } A | (\{1, 2, 3, 4\} \setminus A), \\ - \sum_{e \in E(\tau)} q_e & \text{if } A = \emptyset, \\ 0 & \text{otherwise,} \end{cases}$$

for any subset $A \subset \{1, 2, 3\}$. We will use r_1 to denote the invertible map from a set of branch lengths to the branch length spectrum. The *sequence spectrum*

is derived from the vector of single-site patterns at the leaf nodes by summing the probabilities for site patterns that induce identical splits. The mapping from an expected site pattern frequency (the image of G) to its sequence spectrum is an invertible transformation we will call r_2 [see Section 2.2.2, Steel, Hendy and Penny (1998)].

Hadamard exponential conjugation induces the commutative diagram mentioned previously:

$$\begin{array}{ccc} \mathcal{A} & \overset{\text{Hadamard}}{\longleftrightarrow} & \mathcal{B} \\ \updownarrow r_1 & & \updownarrow r_2 \\ \mathcal{K}_b & \xrightarrow{G} & \text{Im}[G(\cdot)] \end{array}$$

where r_1 maps (τ, q) to its edge length spectrum $Q(\tau, q) \in \mathcal{A} = \text{Im}[r_1(\cdot)]$, r_2 maps $G(\tau, q)$ to its corresponding sequence spectrum $\Phi \in \mathcal{B} = \text{Im}[r_2(\cdot)]$ and the Hadamard-exponential conjugation provides a smooth conversion between the two spectra by the formula

$$(6.1) \quad \mathbf{a} = \mathbf{H}^{-1} \log(\mathbf{H}\mathbf{b}), \quad \text{and} \quad \mathbf{b} = \mathbf{H}^{-1} \exp(\mathbf{H}\mathbf{a}),$$

where $\mathbf{a} \in \mathcal{A}$, $\mathbf{b} \in \mathcal{B}$ are vectors of dimension 8, \log and \exp are applied componentwise and \mathbf{H} is a Hadamard matrix of rank 8 defined as follows:

$$\mathbf{H}_1 = [1], \quad \mathbf{H}_{i+1} = \begin{bmatrix} \mathbf{H}_i & \mathbf{H}_i \\ \mathbf{H}_i & -\mathbf{H}_i \end{bmatrix}, \quad \text{and} \quad \mathbf{H} = \mathbf{H}_4.$$

Hence, if $(\tau_1, q_1), (\tau_2, q_2) \in \mathcal{K}_b$, then the entries of $\mathbf{H}\mathbf{a}$ satisfy

$$[\mathbf{H}\mathbf{a}]_i = \sum_j H_{ij} a_j \geq \sum_j -|a_j| = -2 \sum_{e \in E(\tau)} q_e \geq -4 \text{diam}(\tau, q) \geq -4b \quad \forall i.$$

The first inequality holds because $|H_{ij}| = 1$ for all i and j .

Note that

$$(6.2) \quad \begin{aligned} \|\mathbf{H}\mathbf{x}\|_1 &\leq 8\|\mathbf{x}\|_1, \quad \text{and} \quad \|\mathbf{H}^{-1}\mathbf{x}\|_1 = \left\| \frac{1}{8}\mathbf{H}\mathbf{x} \right\|_1 \leq \|\mathbf{x}\|_1, \quad \forall \mathbf{x} \in \mathbb{R}^8 \\ [\mathbf{H}\mathbf{b}]_i &= \exp([\mathbf{H}\mathbf{a}]_i) \geq e^{-4b} \quad \forall i, \end{aligned}$$

where $\|\cdot\|_1$ is the standard L_1 norm. From Amenta et al. (2007), we have

$$d((\tau_1, q_1), (\tau_2, q_2)) \leq \sqrt{2}d_{\text{BS}}((\tau_1, q_1), (\tau_2, q_2)) \leq \sqrt{2}\|\mathbf{a}_1 - \mathbf{a}_2\|_1,$$

where $\|\mathbf{a}_1 - \mathbf{a}_2\|_1$ denotes the L_1 distance between \mathbf{a}_1 and \mathbf{a}_2 . Thus

$$(6.3) \quad \begin{aligned} d((\tau_1, q_1), (\tau_2, q_2)) &= \sqrt{2}\|\mathbf{a}_1 - \mathbf{a}_2\|_1 = \sqrt{2}\|\mathbf{H}^{-1}(\log(\mathbf{H}\mathbf{b}_1)) - \mathbf{H}^{-1}(\log(\mathbf{H}\mathbf{b}_2))\|_1 \\ &\leq \sqrt{2}\|\log(\mathbf{H}\mathbf{b}_1) - \log(\mathbf{H}\mathbf{b}_2)\|_1 \leq \sqrt{2}e^{4b}\|\mathbf{H}\mathbf{b}_1 - \mathbf{H}\mathbf{b}_2\|_1 \\ &\leq 8\sqrt{2}e^{4b}\|\mathbf{b}_1 - \mathbf{b}_2\|_1 = 16\sqrt{2}e^{4b}d_{\text{TV}}(P_{\tau_1, q_1}, P_{\tau_2, q_2}) \\ &\leq 16e^{4b}\sqrt{\text{KL}(P_{\tau_1, q_1}, P_{\tau_2, q_2})}, \end{aligned}$$

where d_{TV} is the total variation distance and the last line is by Pinsker’s inequality. \square

LEMMA 6.2. *Let $U = \{x : d(x, x^*) < f\}$ where f denotes the length of the shortest edge of x^* :*

(i) *If $x \in U$, then there exist four taxa such that the restriction of x^* and x to these four taxa induces two quartets s^* and s of the same topology such that*

$$\text{diam}(s^*) \leq 4g \log_2(n) \quad \text{and} \quad \frac{1}{2n-3}d(x^*, x) \leq d(s^*, s) \leq d(x^*, x).$$

(ii) *If $x \in \mathcal{T}_{e,g} \setminus U$, then there exist four taxa such that the restriction of x^* and x to these four taxa induces two quartets s^* and s such that*

$$\text{diam}(s^*) \leq 4g \log_2(n) \quad \text{and} \quad f/(2n-3) \leq d(s, s^*).$$

Moreover, in both cases, we have $\text{KL}(P_{x^*}, P_x) \geq \text{KL}(P_{s^*}, P_s)$.

PROOF. (i) It is obvious that x has the same topology as x^* if $x \in U$. For every internal edge i , we define a quartet s_i^* by the following procedure: delete i and the edges touching it to obtain 4 rooted subtrees, then select one leaf that is closest to the root within each subtree to form the quartet s_i^* . Following the notation of Erdős et al. (1999), we refer to such a quartet as a *representative quartet* of the tree x . Let s_i be the quartet of tree x formed of the same four leaves. We call s_i the *associated quartet* of s_i^* .

Since x and x^* have the same topology, s_i^* and s_i contain the same set of edges. Moreover, every edge of x^* belongs to at least one s_i^* . Hence,

$$d(x^*, x) \leq \sum_i d(s_i^*, s_i).$$

Therefore, there exists a quartet s^* and its associated quartet s such that

$$\frac{1}{2n-3}d(x^*, x) \leq d(s^*, s).$$

By the construction, the pendant edge lengths of the representative quartet s^* are bounded from above by $g(\text{depth}(x^*) + 1)$ [Erdős et al. (1999)]. Moreover, we can bound $\text{depth}(x^*)$ above by $\log_2(n)$ [Erdős et al. (1999), Csurös (2002), Mossel (2004)]. Hence,

$$\text{diam}(s^*) \leq 2(\text{depth}(x^*) + 1)g + g \leq g(2 \log_2(n) + 3) \leq 4g \log_2(n).$$

It is worth noting that while our notation of depth of a tree is the same as those defined in Erdős et al. (1999), Csurös (2002), Mossel (2004), our notation of diameter incorporates branch lengths. Thus g appears in the bound on $\text{diam}(s^*)$.

(ii) For every $x \in \mathcal{T}(e, g) \setminus U$, let s^* and s be constructed as per the previous case if x has the same topology as x^* . Otherwise, using Lemma 2 in [Erdős et al.

(1999)], we deduce that there exists a set of four leaves M such that the topologies of $s^* = x^*|_M$ and $s = x|_M$ are different and that s^* is a representative quartet of the tree x^* . Using the same arguments as in part (i), we have $\text{diam}(s^*) \leq 4g \log_2(n)$.

In both scenarios, we can choose s^* and s such that

$$\text{diam}(s^*) \leq 4g \log_2(n), \quad \text{and} \quad f/(2n - 3) \leq d(s, s^*).$$

Finally, from the ‘‘information processing inequality’’ [see, e.g., Theorem 9 of van Erven and Harremoës (2014)], we have $\text{KL}(P_{x^*}, P_x) \geq \text{KL}(P_{s^*}, P_s)$. \square

6.1.2. *Convergence rate.* Next, we will bound $d(x^*, x_k)$ in terms of sequence length by obtaining values for the constants C_n and η_{x^*} . The following lemma states that we can take $C_n = 64n^{-32g/\log 2 - 8\sqrt{3}f - 2}$.

LEMMA 6.3. *Let $U = \{x : d(x, x^*) < f\}$ where f denotes the length of the shortest edge of x^* ; we have $\text{KL}(P_{x^*}, P_x) \geq 64n^{-32g/\log 2 - 8\sqrt{3}f - 2}d(x^*, x)^2$, for all $x \in U$.*

PROOF. Since $d(s^*, s) \leq d(x^*, x) \leq f$, we have $\text{diam}(s) \leq \text{diam}(s^*) + \sqrt{3}f$. Lemma 6.2 gives quartets s and s^* with diameter $\text{diam}(s^*) \leq 4g \log_2(n)$ and

$$\text{diam}(s) \leq \text{diam}(s^*) + \sqrt{3}f \leq 4g \log_2(n) + \sqrt{3}f =: b.$$

This bound can be rearranged to take the form $e^{-8b} \geq n^{-32g/\log 2 - 8\sqrt{3}f}$, which when combined with Lemmas 6.1 and 6.2 gives

$$\text{KL}(P_{x^*}, P_x) \geq \text{KL}(P_{s^*}, P_s) \geq 64n^{-32g/\log 2 - 8\sqrt{3}f - 2}d(x^*, x)^2. \quad \square$$

Next, we obtain a value for η_{x^*} , which bounds $\text{KL}(P_{x^*}, P_x)$ below in U' .

LEMMA 6.4. *Let $U = \{x : d(x, x^*) < f\}$ where f denotes the length of the shortest edge of x^* , we have $\text{KL}(P_{x^*}, P_x) \geq C_f n^{-64g/\log 2 - 2}$, $\forall n \geq 4$ and $x \in \mathcal{T}(e, g) \setminus U$, where $C_f = \min\{64f^2, 72[1 - 4^{-16f/(3\log 2)}]^2\}$.*

PROOF. If $\text{diam}(s) \leq \text{diam}(s^*) + 4g \log_2(n)$, then both $\text{diam}(s)$ and $\text{diam}(s^*)$ are bounded from above by $8g \log_2(n)$. Thus

$$(6.4) \quad \text{KL}(P_{x^*}, P_x) \geq \text{KL}(P_{s^*}, P_s) \geq 64n^{-64g/\log 2 - 2}f^2$$

by Lemma 6.1 and Lemma 6.2. If not, there exist two taxa S_1 and S_2 such that $\iota := \iota_x(S_1, S_2) > \text{diam}(s^*) + 4g \log_2(n)$, and $\iota^* := \iota_{x^*}(S_1, S_2) \leq \text{diam}(s^*)$ where $\iota_x(S_1, S_2)$ is the sum of the branch lengths between S_1 and S_2 on tree x . We have

$$\text{KL}(P_{x^*}, P_x) \geq \text{KL}(P_{x^*|_{\{S_1, S_2\}}}, P_{x|_{\{S_1, S_2\}}}) \geq 2d_{\text{TV}}(P_{x^*|_{\{S_1, S_2\}}}, P_{x|_{\{S_1, S_2\}}})^2.$$

Under the Jukes–Cantor model,

$$P_{x|_{\{S_1, S_2\}}}(uv) = \begin{cases} \frac{1}{4} + \frac{3}{4}e^{-4t/3} & \text{if } u = v \\ \frac{1}{4} - \frac{1}{4}e^{-4t/3} & \text{if } u \neq v. \end{cases}$$

Therefore,

$$\begin{aligned} & d_{\text{TV}}(P_{x^*|_{\{S_1, S_2\}}}, P_{x|_{\{S_1, S_2\}}}) \\ &= 6(e^{-4t^*/3} - e^{-4t/3}) \\ (6.5) \quad & \geq 6n^{-16g/(3 \log 2)} [1 - n^{-16g/(3 \log 2)}] \\ & \geq 6n^{-16g/(3 \log 2)} [1 - 4^{-16f/(3 \log 2)}]. \end{aligned}$$

The lemma follows directly from (6.4) and (6.5). \square

By Lemma 6.3, Lemma 6.4 and Theorem 5.2, we get the following theorem.

THEOREM 6.1. *Let $x_k = (\hat{\tau}_k, \hat{q}_k)$ be the minimizer of (2.1) on $\mathcal{T}_{e,g}$. For any $\delta > 0$,*

$$d(x^*, x_k) \leq \sqrt{2} \left(C_{n,\delta,\gamma,e,g} + \frac{\beta^2}{4} \right)^{1/2} \left(\frac{\log k}{\alpha_k \sqrt{k}} + \alpha_k \right)^{1/2}$$

with probability greater than $1 - \delta$, where $C_{n,\delta,\gamma,e,g}$ is the constant in Lemma 4.3, C_f is the constant in Lemma 6.4 and

$$\beta = \max \left\{ 32d(x^*, x^\circ) n^{16g/\log 2 + 4\sqrt{3}f + 1}, \frac{8}{\sqrt{C_f}} d(x^*, x^\circ)^2 n^{32g/\log 2 + 1} \right\}.$$

If we have a lower bound f for all inner edges, all constants in Theorem 6.1 can be evaluated. Hence, we can construct a conservative confidence region for the regularized ML estimator. This enables us to assess the support of tree splits from data. We also note that the term $d(x^*, x^\circ)^2$ appears in the constants of Theorem 6.1, illustrating the fact that informative prior knowledge (characterized by a good regularizing tree) makes inference better. On the other hand, since the NNI-diameter of the tree space is of order $\mathcal{O}(n \log n)$ [Roch and Sly (2017)], $d(x^*, x^\circ)$ is upper bound by $(\mathcal{O}(gn \log n))$. This implies that even a horrid regularizing tree would help stabilizing an estimator globally without greatly affecting the convergence rate.

6.2. Sample complexity for recovering the true tree topology. We can use the explicit convergence bound to derive a sequence length requirement for recovering

the true topology. From Theorem 6.1, if the length of the observed sequence k satisfies

$$\sqrt{2} \left(C_{n,\delta,\gamma,e,g} + \frac{\beta^2}{4} \right)^{1/2} \left(\frac{\log k}{\alpha_k \sqrt{k}} + \alpha_k \right)^{1/2} \leq f,$$

then the topology of reconstructed tree using the regularized estimator is correct with probability greater than $1 - \delta$. By choosing $\alpha_k = k^{-1/4}$, we have the following.

THEOREM 6.2. *For any $\delta > 0$ and $0 < \nu < 1/4$, if the length of the observed sequence k satisfies*

$$k \sim \left(\left[\frac{1}{f^2} \left(\log \frac{1}{\lambda_e} \right)^2 \left(\log \frac{4g\gamma}{\lambda_e \delta} \right) n^4 + \frac{\beta(n, g, f)^2}{4f^2} \right]^{\frac{4}{1-4\nu}} \right),$$

where $\beta(n, g, f)$ is defined as in Theorem 6.1, then the topology of the reconstructed tree using the regularized estimator is correct with probability greater than $1 - \delta$.

PROOF. The topology of reconstructed tree x_k using the regularized estimator is correct with probability greater than $1 - \delta$ if

$$(6.6) \quad (2/f^2)[C_{n,\delta,\gamma,e,g} + (\beta^2/4)] \leq k^{1/4}/(\log k + 1).$$

Moreover, for all $0 < \nu < 1/4$, there exists C_ν such that $k^{1/4}/(\log k + 1) \geq C_\nu k^{\frac{1}{4}-\nu}$. Thus (6.6) holds if

$$k \geq \left[\frac{2}{C_\nu f^2} \left(C_{n,\delta,\gamma,e,g} + \frac{\beta^2}{4} \right) \right]^{\frac{4}{1-4\nu}}.$$

This completes the proof. \square

When the model parameters (including e and g) and the level of confidence δ are fixed, the sample complexity to recover the tree topology is thus

$$k = \mathcal{O} \left(\frac{n^{\mathcal{O}(g)}}{f^{16/(1-4\nu)}} \right)$$

for any $0 < \nu < 1/4$. Since our regularized-ML algorithm can reconstruct the topology from input sequences of length polynomial in the number of terminal taxa n , it falls in to the class of so-called *fast converging* tree reconstruction algorithms [Huson, Nettles and Warnow (1999), Warnow, Moret and John (2001)].

Gronau, Moran and Snir (2012) point out that the main drawback of most fast converging algorithms is that they perform poorly on trees with indistinguishable/very short edges whereas a single short edge may affect the ability of the algorithm to support any other long edges of the generating tree. Such algorithms

share an “all or nothing” nature: when failing to reconstruct an edge, some of these algorithms do not return any tree, while others return a tree that may contain many faulty edges [Huson, Nettles and Warnow (1999), Mossel (2007), Gronau, Moran and Snir (2012)]. Gronau, Moran and Snir (2012) introduce the concept of the *adaptive fast converging algorithm* to refer to those that can reconstruct all edges of length greater than any given threshold from sequence of polynomial length and construct the first algorithm with that property. In our case, from Theorem 6.1, we deduce that if the length k of the observed sequences satisfies

$$\sqrt{2} \left(C_{n,\delta,\gamma,e,g} + \frac{\beta^2}{4} \right)^{1/2} \left(\frac{\log k}{\alpha_k \sqrt{k}} + \alpha_k \right)^{1/2} \leq t_0,$$

then any edge of length greater than t_0 can be reconstructed by the regularized estimator with high probability.

This leads to the following theorem, which shows that the regularized estimator is also adaptive fast-converging.

THEOREM 6.3. *For any $\delta > 0$, $t_0 > 0$ and $0 < \nu < 1/4$, if the length of the observed sequence k satisfies*

$$k \sim \left(\left[\frac{1}{t_0^2} \left(\log \frac{1}{\lambda_e} \right)^2 \left(\log \frac{4g\gamma}{\lambda_e \delta} \right) n^4 + \frac{\beta^2(n, g, f)}{t_0^2} \right]^{\frac{4}{1-4\nu}} \right)$$

then the regularized estimator correctly reconstruct all edges of length greater than t_0 in the true tree (τ^, q^*) with probability greater than $1 - \delta$.*

In other words, if the model parameters (including e and g) and the level of confidence δ are fixed, then the sample complexity to recover all edges of length greater than t_0 is

$$k = \mathcal{O} \left(\frac{n^{\mathcal{O}(g)}}{(t_0 f)^{8/(1-4\nu)}} \right)$$

for any $0 < \nu < 1/4$.

7. Application to yeast gene-tree reconstruction. In this section, we illustrate the applicability of our method in practical contexts by using the newly constructed estimator to analyze 106 widely distributed orthologous genes from the genomes of eight yeast species. This data set was originally studied in Rokas et al. (2003) and later incorporated to the R package phangorn [Schliep (2011)] as a standard data set for phylogenetic analyses.

The previous results from Rokas et al. (2003) indicate that data sets obtained from a single gene or by concatenating a small number of genes have a significant probability of supporting conflicting topologies. By contrast, analyses of the entire

data set of concatenated genes (by both maximum likelihood and maximum parsimony) yield a single, fully resolved species tree with high support. Comparable results were obtained with a concatenation of a minimum of (randomly chosen) 20 genes; substantially more genes than commonly used. One possible explanation for the success of using randomly chosen concatenated sequences to recover the species tree and the failure of the use of single-gene to do so is that there might not be sufficient information in single genes for maximum likelihood (or maximum parsimony) methods to reliably recover the correct tree.

The analyses from previous sections suggest a regularized maximum likelihood estimator of the form

$$(\hat{\tau}_k, \hat{q}_k) := \operatorname{argmin}_{(\tau, q) \in \mathcal{T}} -\frac{1}{k} \ell_k(\tau, q) + C \frac{1}{k^{1/4}} R(\tau, q),$$

where $R(\tau, q) = d((\tau, q), (\tau^\circ, q^\circ))^2$ and $d(s, s')$ denotes the BHV distance between the trees s and s' on the tree space.

To obtain the species tree (τ°, q°) , we concatenate all genes in the data set and infer the species tree by maximum likelihood using package `phangorn` [Schliep (2011)]. As explained in Rokas et al. (2003), this tree receives strong and stable support from many standard reconstruction methods. We will use the sequence data for the YKL120W gene as an example, noting that the reconstructed maximum likelihood tree is different from the species tree for that gene. We note that analyses of other genes in the data set can also be obtained in a similar manner, but the incongruence between the ML tree and the species tree is of greater interest. We compute the BHV distance using the R function `dist.multiPhylo` in the `distory` package [Chakerian and Holmes (2013)] and modify the likelihood computation subroutine of the `phangorn` package to incorporate the penalty term into the optimization procedure.

While our theoretical results help provide insights about the convergence and asymptotic properties of the regularized estimator, they offer little about how the estimator performs when the tuning parameter is chosen in a data-dependent and stochastic way, or what is a proper procedure to tune such parameter. There are several possible techniques for choosing the penalty parameter; however, the recommended (and default) technique for selecting C in regularization problems is to choose C such that the estimator minimizes the cross-validation estimate of the risk [Homrighausen and McDonald (2013), Zhang (1993)]. In our example, we will follow this standard procedure to select the constant C for the penalty.

The procedure for choosing C can be summarized as follows. We first randomly divide the sequence data into three parts, consider each part as validation data once and use the remaining two parts to reconstruct the tree with various values of $C = 0, 0.25, 0.5, 0.75, 1, 1.25, 1.5$. Here, $C = 0$ corresponds to maximum likelihood estimator. This process is repeated 1000 times and the average log-likelihood per site for each tree on the validation data is used to choose the constant C . The result

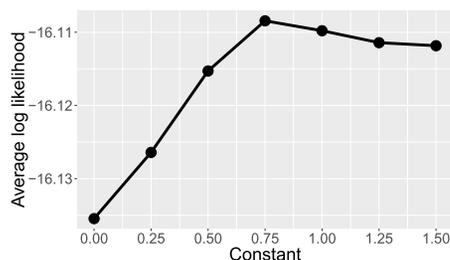


FIG. 4. Average log likelihood from 3-fold cross validation.

is summarized in Figure 4. We observe that large values of C (e.g., 1, 1.25, 1.5) are preferred over small values (0, 0.25, 0.5), which indicates that the regularized method has better performance compared to the standard ML approach. However, over-penalizing the distance to the species tree does not lead to better estimates.

Using the value of $C = 0.75$ that maximizes the average log-likelihood, we reconstruct the regularized maximum likelihood tree, which is visualized in Figure 5(a). This gene tree has the same topology as the species tree, while the gene tree reconstructed by maximum likelihood method has a different topology [Figure 5(b)]. This result illustrates that while there might not be enough information from the (short) sequence data to reliably estimate the correct tree, additional knowledge from the species tree may help produce a better estimate (in terms of likelihood value on the validation data). We note that this is consistent with the findings in Rokas et al. (2003), whereas adding more data from a different source improves the estimator. In our case, the additional information from other genes is compressed in the guiding species tree and has the same effect.

8. Discussion. In this paper, we propose a tree reconstruction method via regularization using the BHV geodesic distance and analyze its asymptotic properties. We prove that the regularized maximum likelihood estimator is consistent and derive global error estimates of the estimator in various settings. Using these estimates, we are able to provide an upper bound on the sample complexity of the estimator to recover the true topology.

Steel and Székely (2002, 2009) provide a bound on the sample complexity of the maximum likelihood estimator, which requires that sequence length k increase exponentially in the number of species n in order to reconstruct the true topology. With the same requirement, Atteson (1997) proves that the neighbor-joining method can recover the true topology for the Neyman 2-state model on a binary trees for which the transition probability on all edges is bounded from below and above. With the same assumptions in Atteson (1997), Erdős et al. (1999) prove that the dyadic closure method only needs k to grow polynomially with respect to the number of species n in order to reconstruct the true topology. The work

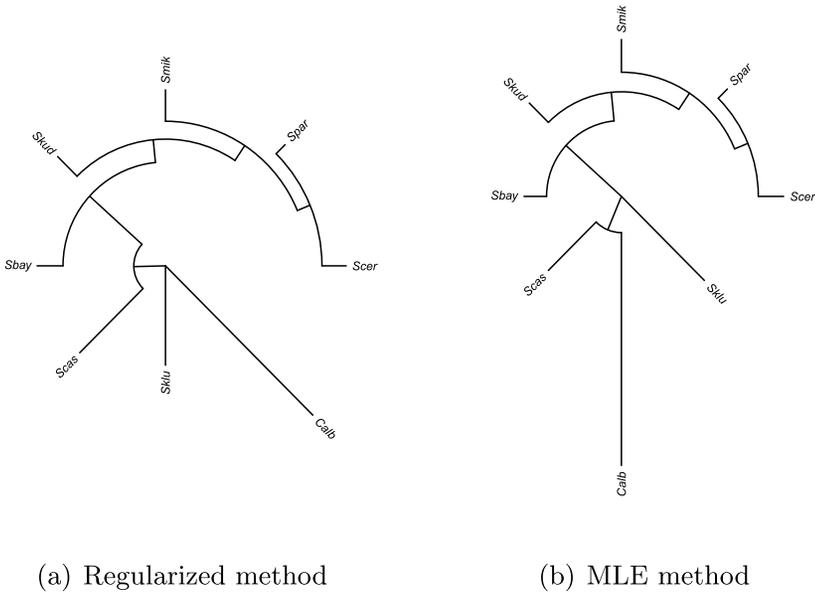


FIG. 5. Reconstructed gene tree for the YKL120W gene.

of Daskalakis, Mossel and Roch (2011), Gronau, Moran and Snir (2012), Huson, Nettles and Warnow (1999), Mossel (2007) extend the result in Erdős et al. (1999) to construct various fast converging algorithms for tree reconstruction. Recently, Roch and Sly (2017) prove that for r -state symmetric models, the sample complexity for MLE is $\mathcal{O}(\log n)$ if all the branch lengths are less than $\log \sqrt{2}$. Otherwise, the sample complexity is $\text{poly}(n)$. Mossel (2003) gives a lower bound of order $\mathcal{O}(n^{\mu g / \log 2 - 1 - o(1)})$ for the sample complexity of any tree reconstruction method on binary trees when $g > \log \sqrt{2}$.

Theorem 6.2 proves that the regularized maximum likelihood estimator recovers the true topology if $k = \text{poly}(n)$, which is the best possible sample complexity of other methods. However, unlike most of the previous algorithms that perform poorly on trees with indistinguishable/very short edges, the regularized estimator can reconstruct all edges of length greater than any given threshold from sequences of polynomial length. Moreover, when information about the model configuration is available, we can derive a conservative confidence region for the estimator and use that to assess the support of tree splits from data. We also do not assume a lower bound for internal edges as previous work has done.

The uniform bounds on the convergence of the empirical likelihood, provided by Lemma 4.2 and Lemma 4.3, are of independent interest. These results allow us to consider the space of phylogenies as a continuous object instead of discretizing the branch lengths as has been done in previous analyses of sample complexity for tree reconstruction [e.g., Roch and Sly (2017)]. In a similar manner, we only as-

sume that the lower bound for pendant edges is known and do not impose this condition on the inner edges. This minor relaxation allows smooth transition between tree topologies and enables the uses of analytical methods to analyze phylogenetic algorithms.

There are several venues for improvement. First, in our current setting, we still assume that the upper bound on the branch lengths are known a priori. This condition can be relaxed by extending Lemma 4.3 to the case of unbounded edges and by highlighting the fact that the regularization term already imposes some implicit constraints on the branch lengths. Second, we currently use the BHV geodesic distance as the penalty for regularization. While the geodesic distance is the most natural distance on the BHV space and is strongly convex, this distance cannot be computed in linear time. We would like to extend the result to other (computationally cheaper) distances on tree space such as those proposed in Amenta et al. (2007), or to other nonstrongly convex penalties. Relating to the species tree/gene tree problem, it is also worthwhile to consider penalty metrics based on explicit models of discordance-generating events, such as the duplication-loss cost investigated in Wu et al. (2013).

REFERENCES

- ÅKERBERG, Ö., SENBLAD, B., ARVESTAD, L. and LAGERGREN, J. (2009). Simultaneous Bayesian gene tree reconstruction and reconciliation analysis. *Proc. Natl. Acad. Sci. USA* **106** 5714–5719.
- AMENTA, N., GODWIN, M., POSTARNAKEVICH, N. and JOHN, K. S. (2007). Approximating geodesic tree distance. *Inform. Process. Lett.* **103** 61–65. [MR2322069](#)
- ATTESON, K. (1997). The performance of neighbor-joining algorithms of phylogeny reconstruction. In *Computing and Combinatorics (Shanghai, 1997)*. *Lecture Notes in Computer Science* **1276** 101–110. Springer, Berlin. [MR1616308](#)
- BAČÁK, M. (2013). The proximal point algorithm in metric spaces. *Israel J. Math.* **194** 689–701. [MR3047087](#)
- BAČÁK, M. (2014). Computing medians and means in Hadamard spaces. *SIAM J. Optim.* **24** 1542–1566. [MR3264572](#)
- BANSAL, M. S., WU, Y.-C., ALM, E. J. and KELLIS, M. (2015). Improved gene tree error correction in the presence of horizontal gene transfer. *Bioinformatics* **31** 1211–1218.
- BILLERA, L. J., HOLMES, S. P. and VOGTMANN, K. (2001). Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.* **27** 733–767. [MR1867931](#)
- BOUSSAU, B. and DAUBIN, V. (2010). Genomes as documents of evolutionary history. *Trends Ecol. Evol.* **25** 224–232.
- BOUSSAU, B., SZÖLLŐSI, G. J., DURET, L., GOUY, M., TANNIER, E. and DAUBIN, V. (2013). Genome-scale coestimation of species and gene trees. *Genome Res.* **23** 323–330.
- CHAKERIAN, J. and HOLMES, S. (2013). *distory: Distance Between Phylogenetic Histories*. R package version 1.4.2.
- CSURÓS, M. (2002). Fast recovery of evolutionary trees with thousands of nodes. *J. Comput. Biol.* **9** 277–297.
- CUCKER, F. and SMALE, S. (2002). On the mathematical foundations of learning. *Bull. Amer. Math. Soc. (N.S.)* **39** 1–49. [MR1864085](#)

- CUONG, N. V., HO, L. S. T. and DINH, V. (2013). Generalization and robustness of batched weighted average algorithm with V-geometrically ergodic Markov data. In *Algorithmic Learning Theory. Lecture Notes in Computer Science* **8139** 264–278. Springer, Heidelberg. [MR3133071](#)
- DASKALAKIS, C., MOSSEL, E. and ROCH, S. (2011). Phylogenies without branch bounds: Contracting the short, pruning the deep. *SIAM J. Discrete Math.* **25** 872–893. [MR2817535](#)
- DAVID, L. A. and ALM, E. J. (2011). Rapid evolutionary innovation during an Archaean genetic expansion. *Nature* **469** 93–96.
- ENGL, H. W., HANKE, M. and NEUBAUER, A. (1996). *Regularization of Inverse Problems. Mathematics and Its Applications* **375**. Kluwer Academic, Dordrecht. [MR1408680](#)
- ERDŐS, P. L., STEEL, M. A., SZÉKELY, L. A. and WARNOW, T. J. (1999). A few logs suffice to build (almost) all trees. I. *Random Structures Algorithms* **14** 153–184. [MR1667319](#)
- FELSENSTEIN, J. (1984). DNAML in PHYLIP 2. University of Washington, Seattle. 6.
- FELSENSTEIN, J. (2004). *Inferring Phylogenies*. Sinauer Associates, Sunderland.
- GRONAU, I., MORAN, S. and SNIR, S. (2012). Fast and reliable reconstruction of phylogenetic trees with indistinguishable edges. *Random Structures Algorithms* **40** 350–384. [MR2900143](#)
- HELED, J. and DRUMMOND, A. J. (2010). Bayesian inference of species trees from multilocus data. *Mol. Biol. Evol.* **27** 570–580.
- HENDY, M. D. and PENNY, D. (1993). Spectral analysis of phylogenetic data. *J. Classification* **10** 5–24.
- HOEFFDING, W. (1963). Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* **58** 13–30. [MR0144363](#)
- HOERL, A. E. (1962). Application of ridge analysis to regression problems. *Chem. Eng. Prog.* **58** 54–59.
- HOFMANN, B. and YAMAMOTO, M. (2010). On the interplay of source conditions and variational inequalities for nonlinear ill-posed problems. *Appl. Anal.* **89** 1705–1727. [MR2683677](#)
- HOFMANN, B., KALTENBACHER, B., PÖSCHL, C. and SCHERZER, O. (2007). A convergence rates result for Tikhonov regularization in Banach spaces with non-smooth operators. *Inverse Probl.* **23** 987–1010. [MR2329928](#)
- HOHAGE, T. and WEIDLING, F. (2015). Verification of a variational source condition for acoustic inverse medium scattering problems. *Inverse Probl.* **31** 075006, 14. [MR3366755](#)
- HOMRIGHAUSEN, D. and McDONALD, D. (2013). The lasso, persistence, and cross-validation. In *Proceedings of the 30th International Conference on Machine Learning* 1031–1039.
- HUSON, D. H., NETTLES, S. M. and WARNOW, T. J. (1999). Disk-covering, a fast-converging method for phylogenetic tree reconstruction. *J. Comput. Biol.* **6** 369–386.
- Ji, S., KOLLÁR, J. and SHIFFMAN, B. (1992). A global Łojasiewicz inequality for algebraic varieties. *Trans. Amer. Math. Soc.* **329** 813–818. [MR1046016](#)
- KAZIMIERSKI, K. S. (2010). *Aspects of Regularization in Banach Spaces*. Logos Verlag, Berlin.
- KIM, J. (2000). Slicing hyperdimensional oranges: The geometry of phylogenetic estimation. *Mol. Phylogenet. Evol.* **17** 58–75.
- KUHNER, M. K. and FELSENSTEIN, J. (1994). A simulation comparison of phylogeny algorithms under equal and unequal evolutionary rates. *Mol. Biol. Evol.* **11** 459–468.
- LIU, L. and PEARL, D. K. (2007). Species trees from gene trees: Reconstructing Bayesian posterior distributions of a species phylogeny using estimated gene tree distributions. *Syst. Biol.* **56** 504–514.
- MADDISON, W. P. (1997). Gene trees in species trees. *Syst. Biol.* **46** 523–536.
- MAHMUDI, O., SJÖSTRAND, J., SENNBLAD, B. and LAGERGREN, J. (2013). Genome-wide probabilistic reconciliation analysis across vertebrates. *BMC Bioinform.* **14** 1–11.
- MOSSEL, E. (2003). On the impossibility of reconstructing ancestral data and phylogenies. *J. Comput. Biol.* **10** 669–676.
- MOSSEL, E. (2004). Phase transitions in phylogeny. *Trans. Amer. Math. Soc.* **356** 2379–2404. [MR2048522](#)

- MOSSEL, E. (2007). Distorted metrics on trees and phylogenetic forests. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **4** 108–116.
- MOULTON, V. and STEEL, M. (2004). Peeling phylogenetic ‘oranges’. *Adv. in Appl. Math.* **33** 710–727. [MR2095862](#)
- NYE, T. M. W. (2011). Principal components analysis in the space of phylogenetic trees. *Ann. Statist.* **39** 2716–2739. [MR2906884](#)
- OWEN, M. and PROVAN, J. S. (2011). A fast algorithm for computing geodesic distances in tree space. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **8** 2–13.
- RASMUSSEN, M. D. and KELLIS, M. (2011). A Bayesian approach for fast and accurate gene tree reconstruction. *Mol. Biol. Evol.* **28** 273–290.
- ROBINSON, D. F. (1971). Comparison of labeled trees with valency three. *J. Combin. Theory Ser. B* **11** 105–119. [MR0294170](#)
- ROCH, S. and SLY, A. (2017). Phase transition in the sample complexity of likelihood-based phylogeny inference. *Probab. Theory Related Fields* **169** 3–62. [MR3704765](#)
- ROCH, S. and WARNOW, T. (2015). On the robustness to gene tree estimation error (or lack thereof) of coalescent-based species tree methods. *Syst. Biol.* **64** 663–676.
- ROKAS, A., WILLIAMS, B. L., KING, N. and CARROLL, S. B. (2003). Genome-scale approaches to resolving incongruence in molecular phylogenies. *Nature* **425** 798–804.
- SCHLIEP, K. (2011). phangorn: Phylogenetic analysis in R. *Bioinformatics* **27** 592–593.
- SCORNAVACCA, C., JACOX, E. and SZÖLLÖSI, G. J. (2015). Joint amalgamation of most parsimonious reconciled gene trees. *Bioinformatics* **31** 841–848.
- SEMPLE, C. and STEEL, M. (2003). *Phylogenetics. Oxford Lecture Series in Mathematics and Its Applications* **24**. Oxford Univ. Press, Oxford. [MR2060009](#)
- STEEL, M., HENDY, M. D. and PENNY, D. (1998). Reconstructing phylogenies from nucleotide pattern probabilities: A survey and some new results. *Discrete Appl. Math.* **88** 367–396. [MR1658533](#)
- STEEL, M. A. and SZÉKELY, L. A. (2002). Inverting random functions. II. Explicit bounds for discrete maximum likelihood estimation, with applications. *SIAM J. Discrete Math.* **15** 562–575. [MR1935839](#)
- STEEL, M. A. and SZÉKELY, L. A. (2009). Inverting random functions. III. Discrete MLE revisited. *Ann. Comb.* **13** 365–382. [MR2557045](#)
- SZÖLLÖSI, G. J., TANNIER, E., DAUBIN, V. and BOUSSAU, B. (2015). The inference of gene trees with species trees. *Syst. Biol.* **64** e42–e62.
- VAN ERVEN, T. and HARREMOËS, P. (2014). Rényi divergence and Kullback–Leibler divergence. *IEEE Trans. Inform. Theory* **60** 3797–3820. [MR3225930](#)
- WANG, Q. (2004). Maximum Likelihood Estimation of Phylogenetic Tree with Evolutionary Parameters Ph.D. thesis, The Ohio State University. Available at https://etd.ohiolink.edu/!etd.send_file?accession=osu1083177084&disposition=inline.
- WARNOW, T., MORET, B. M. E. and JOHN, K. S. (2001). Absolute convergence: True trees from short sequences. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (Washington, DC, 2001)* 186–195. SIAM, Philadelphia, PA. [MR1958406](#)
- WU, Y.-C., RASMUSSEN, M. D., BANSAL, M. S. and KELLIS, M. (2013). TreeFix: Statistically informed gene tree error correction using species trees. *Syst. Biol.* **62** 110–120.
- ZHANG, P. (1993). Model selection via multifold cross validation. *Ann. Statist.* **21** 299–313. [MR1212178](#)

V. DINH
 F. A. MATSEN IV
 PROGRAM IN COMPUTATIONAL BIOLOGY
 FRED HUTCHINSON CANCER RESEARCH CENTER
 SEATTLE, WASHINGTON 98122
 USA
 E-MAIL: vdinh@fredhutch.org
matsen@fredhutch.org

L. S. T. HO
 DEPARTMENT OF BIostatISTICS
 UNIVERSITY OF CALIFORNIA
 LOS ANGELES, CALIFORNIA 90995
 USA
 E-MAIL: lamho@ucla.edu

M. A. SUCHARD
DEPARTMENTS OF BIOMATHEMATICS
BIOSTATISTICS AND HUMAN GENETICS
UNIVERSITY OF CALIFORNIA
LOS ANGELES, CALIFORNIA 90995
USA
E-MAIL: msuchard@ucla.edu