



HHS Public Access

Author manuscript

IEEE/ACM Trans Comput Biol Bioinform. Author manuscript; available in PMC 2018 February 10.

Published in final edited form as:

IEEE/ACM Trans Comput Biol Bioinform. 2018 ; 15(1): 337–342. doi:10.1109/TCBB.2016.2604812.

Species tree inference from gene splits by Unrooted STAR methods

Elizabeth S. Allman,

Department of Mathematics and Statistics, University of Alaska Fairbanks, Fairbanks, AK 99775

James H. Degnan, and

Department of Mathematics and Statistics, The University of New Mexico, Albuquerque, NM 87131

John A. Rhodes

Department of Mathematics and Statistics, University of Alaska Fairbanks, Fairbanks, AK 99775

Abstract

The NJ_{st} method was proposed by Liu and Yu to infer a species tree topology from unrooted topological gene trees. While its statistical consistency under the multispecies coalescent model was established only for a 4-taxon tree, simulations demonstrated its good performance on gene trees inferred from sequences for many taxa. Here we prove the statistical consistency of the method for an arbitrarily large species tree. Our approach connects NJ_{st} to a generalization of the STAR method of Liu, Pearl and Edwards, and a previous theoretical analysis of it. We further show NJ_{st} utilizes only the distribution of splits in the gene trees, and not their individual topologies. Finally, we discuss how multiple samples per taxon per gene should be handled for statistical consistency.

Index Terms

coalescent model; STAR algorithm; NJ_{st} ; species tree

1 INTRODUCTION

With the growing feasibility of building large multilocus data sets of genetic sequences, questions of how to best infer ancestral relationships between taxa have increasingly been viewed in the light of the multispecies coalescent model. This model describes the formation of *gene trees* (or genealogies) relating orthologous loci within a *species tree* composed of populations. It thus brings into phylogenetics an important model of population genetics, in order to capture the phenomenon of incomplete lineage sorting and allow incongruence across gene trees to be used to more accurately infer species trees.

In principle it is straightforward to combine standard models of sequence evolution with the multispecies coalescent for inference of species trees under either a maximum likelihood or Bayesian framework. In practice, though, this is both computationally intensive and requires some additional assumptions — most importantly, a means of relating the coalescent and mutation time scales — that may or may not be reasonable. Such assumptions are not always

highlighted in data analysis, even though they may include 1) a molecular clock operating for each gene tree, 2) constant population sizes over each branch of the species tree, and 3) a common mutation rate across gene trees, or variants of these. It is also not clear to what extent existing software implementations have been applied to simulated data violating such assumptions, in order to understand their robustness. Finally, even accepting these assumptions, analyzing a many-gene many-taxon data set can be computationally infeasible using a standard approach.

Some inference approaches simplify the problem by first inferring individual gene trees by established phylogenetic methods, and then using these to infer a species tree. Of course this introduces a new source of error, as the inferred gene trees may differ from the unknown true ones, with implications we will discuss later. Nonetheless, by breaking the inference problem into pieces in this way, one can gain significant computational advantages.

From the inferred gene trees, one might use metric information, or only topologies, with or without a root. If one views the gene tree topologies as more robustly inferable than metric edge lengths, then two methods, STAR [1] and NJ_{st} [2], are especially attractive. By not using any metric information from the gene trees, they elegantly circumvent issues of how one should relate the coalescent and mutational time scales. They both encode gene tree topologies through special distance matrices, in what one might call a remetrization step, with STAR requiring rooted trees, and NJ_{st} unrooted ones. The average of these matrices is then used as input to a standard metric tree-building algorithm to recover the species tree topology. (Though the tree-building process may return edge lengths as part of the species tree estimate, whether and how they might be used to recover the true lengths on the species tree is not currently known.) All computations are both simple and fast, and accuracy on large datasets is competitive with the best current methods, as shown by the recent implementation and extension of NJ_{st} in the software ASTRID [3].

In [1] and [2] arguments were given establishing the statistical consistency of STAR and NJ_{st} for certain 4-taxon species trees only, and not for larger trees. (Consistency here refers to applying the method to true gene trees, and does not account for any gene tree inference error.) In [4] a rigorous proof of consistency was given for STAR and variants of it on arbitrarily large trees, along with a theoretical exploration of how the algorithm actually only required the distribution of clades on the gene trees. This recasts STAR as a clade consensus method attuned specifically to the multispecies coalescent model.

In this work we obtain similar theoretical results for NJ_{st} . We first prove its statistical consistency under the multispecies coalescent model on arbitrary trees in Theorem 4.1. Our proof is built on relating NJ_{st} to a generalized STAR method as introduced in [4], and deducing our results from those on STAR. (An unpublished work by Kreidl [5] offered the first full proof of the consistency of NJ_{st} , arguing directly from the behavior of the coalescent model.) In Theorem 5.1 we show the method uses only information in the distribution of splits on gene trees, and not the more detailed information of the gene trees themselves. Thus we view it as a split consensus method designed specifically for inference of species trees under the multispecies coalescent model. In Section 6 we then discuss how one might apply the method to data that involves multiple samples from each taxon. We

show the approach suggested by [2] for such data can be problematic even in a simple case, but then give an alternative which is statistically consistent under certain sampling schemes.

Finally, we suggest a rechristening of NJ_{st} as USTAR/NJ, for “Unrooted STAR with Neighbor Joining.” This emphasizes both the close relationship of the two methods, and that one might perform the method with tree selection approaches other than Neighbor Joining. Any statistically consistent method for selecting a metric tree from possibly non-ultrametric distance tables could be used in its place. For instance, USTAR/BIONJ [6] uses a different purely algorithmic tree building method, while USTAR/FastME [7] performs a heuristic search to optimize the balanced minimum evolution criterion to select a tree. Indeed, the ASTRID software already allows one to apply such methods and [3] compares their performance.

We reiterate that the consistency of USTAR methods established in this paper assumes that gene trees are accurately known, rather than estimated with some error. Indeed, this is typical of current theoretical results on two-stage inference methods, where gene trees are first inferred by standard phylogenetic methods, and these are then used as ‘data’ to infer a species tree. While investigating consistency in the absence of gene tree error as we do here is an important step, understanding the combined two-stage procedure is a desirable goal. However we currently have little formalized understanding of gene tree inference error, much less its effect on a procedure such as USTAR.

One attempt at dealing with gene tree error in species tree inference [8] showed that the GLASS method is consistent for estimating the species tree assuming that all estimated coalescence times in gene trees are within $m/2$ of the true times, where m is the shortest branch in the species tree. However, [9] argued this condition is actually quite stringent, requiring unrealistically accurate gene trees for typical applications, and the condition is less likely to be met as the number of loci increases. Another approach [10] proves consistency of a rooted triple method for inferring species trees from estimated gene trees, assuming the gene tree error has a certain feature: For each gene and triple of taxa, the most likely inferred rooted triple gene tree matches the true rooted triple gene tree in topology. A similar assumption is used to establish consistency of quartet methods. We believe that developing a more elaborate descriptive model of gene tree inference error will provide a promising route to greater theoretical understanding of other two-stage procedures.

Of course empiricists work with data of finite length sequences from a finite number of loci, where any consistency result can only offer hope of good performance. With both understanding of gene tree error and results on convergence as the number of loci increases lacking, simulation studies are the main source of insight here, and have shown USTAR to be competitive with other methods [2], [3], [11]. Nonetheless, as pointed out in [10], there are parameter regimes in which an inconsistent method, such as concatenation with maximum likelihood [12], [13] appears to have better performance on finite data than some consistent methods. Practical performance can depend on many factors, such as alignment lengths of the genes, the number of genes, the length of branches in the species tree, and the extent to which model assumptions are violated. Much work remains to understand these many aspects of species tree inference.

2 NOTATION AND TERMINOLOGY

Let \mathcal{X} be a finite set of n taxa, which we denote by lower case letters a, b, c, \dots . For any specific gene, we denote a single sample from each taxon by the corresponding upper case letter A, B, C, \dots , with \mathcal{X}_g the full set of such gene samples. If $\mathcal{A} \subseteq \mathcal{X}$ is a subset of taxa, the corresponding subset of genes is $\mathcal{A}_g \subseteq \mathcal{X}_g$. For example $\{a, b, c\}_g = \{A, B, C\}$.

By a species tree $\sigma = (\psi, \lambda)$ on \mathcal{X} we mean a rooted topological tree with leaves bijectively labelled by \mathcal{X} , together with an assignment of edge weights λ to its internal edges. These edge weights are specified in coalescent units, so that the multispecies coalescent model on σ leads to a probability distribution on gene trees with leaves labelled by \mathcal{X}_g . (For a more precise definition of the multispecies coalescent as we use it, we direct the reader to [14].) The gene trees here are metric rooted trees, though this probability distribution, by marginalization, also leads to ones on metric or topological, rooted or unrooted, gene trees. We denote rooted topological gene trees by T^r and unrooted topological gene trees by T . The probability of an unrooted topological gene tree T under the multispecies coalescent on σ is denoted $\mathbb{P}_\sigma(T)$.

A metric tree is called *binary* if the underlying topological tree is binary and all internal edge lengths are positive.

A *split* of a set of taxa \mathcal{X} is a bipartition $\mathcal{A}|\mathcal{B}$ of \mathcal{X} in which neither \mathcal{A} nor \mathcal{B} is empty. Note $\mathcal{A}|\mathcal{B}$ is the same split as $\mathcal{B}|\mathcal{A}$. If $\sigma = (\psi, \lambda)$ is a species tree on \mathcal{X} then a *split on σ* is a split of \mathcal{X} formed by deleting a single edge of ψ and grouping taxa according to the connected components of the resulting graph. We similarly define splits of \mathcal{X}_g , and splits of \mathcal{X}_g on a specific gene tree.

3 USTAR METHODS

Given an unrooted topological gene tree T on \mathcal{X}_g , we may metrize it by giving all edges length 1. The distance $D_T(A, B)$ between any two gene samples A, B on T is then the number of edges in the path connecting them, *i.e.*, the graph-theoretic distance. Fixing an ordering of the taxa, it is convenient to think of D_T as an $n \times n$ matrix. In essence, we have simply encoded the topology of T by the numerical matrix D_T .

In [2], the internodal distance, *i.e.*, the number of nodes on the path in the tree between two taxa, is used to define a similar distance table. The graph-theoretic distance between taxa is always one more than the internodal distance, and it is straightforward to check that this difference between them has no essential impact on anything we do in this paper. We use the graph-theoretic distance here for its simple interpretation in terms of assigning edge lengths of 1, and its more direct connection to the notion of splits on the tree.

For a probability distribution μ on unrooted gene trees on \mathcal{X}_g , the expected value

$$D := \mathbb{E}_\mu(D_T) = \sum_T \mu(T) D_T$$

defines a dissimilarity function on \mathcal{X}_g . Identifying \mathcal{X} with \mathcal{X}_g , we call this the *USTAR dissimilarity on \mathcal{X}* with respect to μ . For an empirically-obtained collection of gene trees, this dissimilarity is just the mean of the matrices D_T for trees in the sample.

In this paper, we focus on the particular choice $\mu = \mathbb{P}_\sigma$ i.e., we use the probability of unrooted gene trees arising from the multispecies coalescent on a specific species tree σ , or an empirical distribution describing a sample from this theoretical one.

From the USTAR dissimilarity D obtained from a gene tree distribution, one can construct or choose a tree on \mathcal{X} , using any of a variety of well-known methods — e.g., UPGMA, Neighbor Joining, BIONJ, Balanced Minimum Evolution, etc. Discarding any edge lengths that might have been produced in the course of applying the tree selection method, yields a topological tree on \mathcal{X} . Thus we have a family of methods whose input is a theoretical or empirical distribution of unrooted topological gene trees, and output is a single unrooted topological tree on the taxa. In particular, USTAR/NJ is the method obtained when Neighbor Joining is used, and coincides with NJ_{st} . The output of such a method can be viewed as an estimate of the species tree.

USTAR methods can be helpfully viewed as related to generalized STAR methods developed in [4], building on [1]. STAR methods of inferring a species tree from *rooted* gene trees similarly involve metrizing the gene trees and averaging the resulting pairwise distance matrices over a gene tree distribution. However the metrization is done as follows: For n taxa, first choose a non-increasing sequence of node numbers $a_0 \ a_1 \ a_2 \ \dots \ a_{n-2} \ 0$, with at least one of these inequalities strict. Assign a_0 to the root, a_1 to its non-leaf children, a_2 to their non-leaf children, etc. Then interpret the assigned numbers as distances from the leaves in an ultrametric tree.

For the particular case of node numbers $n - 3/2, n - 2, n - 3, n - 4, \dots$, the generalized STAR metrization has the effect of giving length 1 to all internal edges of the rooted gene tree, except those incident to the root. However, if suppressing the root leads to a new internal edge in the unrooted version, the total length of that edge is 1. Thus after suppressing the root, all internal edges of the gene tree are given the same length as they would be by USTAR. However, lengths of pendant edges are different, as they are all 1 under USTAR and they vary to achieve ultrametricity under STAR.

Example 3.1

Consider the 5-taxon gene tree $T = ((A, B), ((C, D), E))$ shown in Figure 1. Viewing T as an unrooted tree, with taxa ordered alphabetically, USTAR leads to the distance matrix

$$D_T = \begin{pmatrix} 0 & 2 & 4 & 4 & 3 \\ 2 & 0 & 4 & 4 & 3 \\ 4 & 4 & 0 & 2 & 3 \\ 4 & 4 & 2 & 0 & 3 \\ 3 & 3 & 3 & 3 & 0 \end{pmatrix}.$$

Separating the contributions from internal and pendant edges, we can write

$$D_T = D_T^{ui} + D_T^{up}$$

$$= \begin{pmatrix} 0 & 0 & 2 & 2 & 1 \\ 0 & 0 & 2 & 2 & 1 \\ 2 & 2 & 0 & 0 & 1 \\ 2 & 2 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 2 & 2 & 2 & 2 \\ 2 & 0 & 2 & 2 & 2 \\ 2 & 2 & 0 & 2 & 2 \\ 2 & 2 & 2 & 0 & 2 \\ 2 & 2 & 2 & 2 & 0 \end{pmatrix}.$$

Here ‘*ui*’ and ‘*up*’ refer to the ‘unrooted internal’ and ‘unrooted pendant’ edge contributions.

Viewing T as a rooted tree, T^r , STAR with the node numbering given above, leads to the distance matrix

$$D_{T^r}^r = \begin{pmatrix} 0 & 6 & 7 & 7 & 7 \\ 6 & 0 & 7 & 7 & 7 \\ 7 & 7 & 0 & 4 & 6 \\ 7 & 7 & 4 & 0 & 6 \\ 7 & 7 & 6 & 6 & 0 \end{pmatrix}.$$

Again separating the contributions from internal and pendant edges of the unrooted tree, we have

$$D_{T^r}^r = D_T^{ui} + D_{T^r}^{rp}$$

$$= \begin{pmatrix} 0 & 0 & 2 & 2 & 1 \\ 0 & 0 & 2 & 2 & 1 \\ 2 & 2 & 0 & 0 & 1 \\ 2 & 2 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 6 & 5 & 5 & 6 \\ 6 & 0 & 5 & 5 & 6 \\ 5 & 5 & 0 & 4 & 5 \\ 5 & 5 & 4 & 0 & 5 \\ 6 & 6 & 5 & 5 & 0 \end{pmatrix},$$

where ‘*rp*’ refers to the ‘rooted pendant’ edge contributions. For a general tree, the rooted pendant edge contributions may include some that arise from an internal edge incident to the root that becomes part of a pendant edge when the root is suppressed (such as when there is a single outgroup on the tree).

Note that the same contributions appear from the internal edges of the unrooted tree in both the USTAR and STAR distance matrices. Our analysis of USTAR in the proof of Theorem 4.2 below will be based in the fact that, for the particular STAR numbering scheme where branches incident to the root have length 1/2 and all other internal branches have length 1, the distance matrices differ only in contributions from pendant edges.

4 STATISTICAL CONSISTENCY

Our goal in this section is to prove the following:

Theorem 4.1

Let M denote any method of obtaining an unrooted topological tree from a dissimilarity function satisfying

1. M applied to a tree metric returns the unique tree fitting it, and
2. M is continuous at tree metrics arising from binary trees.

Let $\sigma = (\psi, \lambda)$ be a binary species tree on \mathcal{X} . Then USTAR/ M is a statistically consistent method of inference of the unrooted topology of ψ from gene trees under the multispecies coalescent model on σ .

Informally, the continuity required of the method M in condition (2) means that when M is applied to a sufficiently small perturbation of a binary tree metric, it returns the correct tree topology, and edge lengths close to those underlying the tree metric. As NJ is known [15] to satisfy conditions (1) and (2), we see that in particular USTAR/NJ is consistent. Since UPGMA does not, in general, satisfy condition (1) for non-ultrametric trees, the theorem does not apply to USTAR/UPGMA.

Theorem 4.1 is a consequence of the following.

Theorem 4.2

The USTAR dissimilarity on \mathcal{X} with respect to the probability distribution on unrooted topological gene trees arising from multispecies coalescent model on $\sigma = (\psi, \lambda)$, $D = \mathbb{E}_{\sigma}(D_{\mathcal{T}})$, exactly fits the unrooted species tree topology of ψ .

Proof—Let T denote a rooted gene tree topology. Consider the generalized STAR number scheme for rooted gene trees with node numbering sequence $n-3/2, n-2, n-3, n-4, \dots$. As discussed previously, when the root is suppressed on the STAR remetrized rooted gene tree T , all internal edges on the resulting unrooted tree have length 1. Using this node numbering scheme, let D_{Tr}^r denote the STAR distance matrix for a remetrized rooted tree T , and $\mathbb{E}_{\sigma}(D_{Tr}^r)$ its expected value under the distribution on rooted topological gene trees arising from the coalescent.

We now relate the STAR dissimilarities $D^r = \mathbb{E}_{\sigma}(D_{Tr}^r)$ to those of USTAR, $D = \mathbb{E}_{\sigma}(D_{\mathcal{T}})$. Since both the rooted and unrooted schemes give each internal edge length 1 in the unrooted gene tree topology we can write

$$D_T = D_T^{ui} + D_T^{up}, \quad D_T^{ui} = D_{Tr}^r - D_{Tr}^{rp} \quad (1)$$

where D_T^{ui} contains the contributions to distances from internal edges of the unrooted tree, D_T^{up} contains contributions from pendant edges of the unrooted scheme, and D_{Tr}^{rp} contains contributions from (unrooted tree) pendant edges in the rooted scheme. Equations (1) thus imply

$$D_T = D_{T^r}^r + D_T^{up} - D_{T^r}^{rp}. \quad (2)$$

Now the matrix D_T^{up} is independent of T and has a simple structure; all diagonal entries are 0, and all off-diagonal entries are $1+1 = 2$. The matrix $D_{T^r}^{rp}$, however, does depend on T^r . While it also has 0 in every diagonal entry, the off-diagonal entry in row x , column y is $w_x + w_y$, where w_x, w_y are the lengths assigned to the pendant edges to taxa x, y after the root is suppressed on the remetrized ultrametric T^r .

Passing to expected values, we have from equation (2) that

$$D = D^r + \mathbb{E}_\sigma(D_T^{up}) - \mathbb{E}_\sigma(D_{T^r}^{rp}). \quad (3)$$

By Theorem 3.2 of [4], D^r exactly fits the topology of the rooted species tree (ultrametrically), and hence for each choice of 4 taxa, with some permutation of their labels the 4-point condition

$$\begin{aligned} D^r(a, c) + D^r(b, d) &= D^r(a, d) + D^r(b, c) \\ &\geq D^r(a, b) + D^r(c, d) \end{aligned} \quad (4)$$

holds. Now in the case that a, b, c, d are all distinct, this implies

$$\begin{aligned} D(a, c) + D(b, d) &= D(a, d) + D(b, c) \\ &\geq D(a, b) + D(c, d), \end{aligned} \quad (5)$$

since by equation (3), we have only added $4 - \mathbb{E}(w_a + w_b + w_c + w_d)$ to the three sums in (4) to obtain (5).

If at most 3 of the taxa in the 4-point condition are distinct this last argument is not valid. However, if, say, $c = d$, the 4-point condition we need to establish degenerates to

$$D(a, c) + D(b, c) \geq D(a, b).$$

That this holds follows from the fact the corresponding inequality holds for every tree metric, and in particular for each USTAR remetrization D_{T^r} , and hence for the expected value as well.

Thus the four point condition holds for D for every set of 4 taxa, and it yields the same unrooted quartet topology as does D^r . Thus by standard results in [16] D exactly fits the same unrooted tree topology as D^r , which is that of the species tree.

Proof of Theorem 4.1—As the size of a sample of gene trees from the multispecies coalescent model on σ increases, the empirical distribution of unrooted gene tree topologies approaches the exact one with probability 1, and thus the empirical USTAR dissimilarity approaches the theoretical one, D . Since Theorem 4.2 and condition (2) ensures the method M returns the correct tree when applied to D , condition (1) then implies with probability 1 USTAR/ M returns the correct unrooted species tree topology as the sample size (i.e., number of loci) increases to infinity.

5 USTAR AND SPLITS

Here we establish a relationship between the USTAR expected distance matrix and split probabilities, analogous to that given in [4] for STAR expected distances and clade probabilities.

As a consequence of this relationship, it is natural to view USTAR methods as a type of *split consensus method*. Specifically, USTAR methods use only information on probabilities of splits on gene trees, and not the finer information of the gene tree topologies themselves.

The fact that USTAR uses only split frequencies, yet can produce statistically consistent inference for the coalescent model is notable, as other split methods lack this feature. For instance greedy consensus [17] accepts splits in order of decreasing frequency, if they are compatible with previously accepted splits. Greedy consensus on clades has been proven inconsistent [18], though STAR can be viewed as a consistent clade consensus method [4]. The arguments in [18] can be modified to give a similar result for greedy consensus on splits, with signs of inconsistent behavior also observed in simulations [11]. For consistency, a consensus method must be attuned to the model of tree variation, with USTAR and STAR being appropriate for the coalescent.

Given any two leaves A, B of a gene tree T , let $S_T^{A,B}$ denote the set of splits of T in which A and B are separated (i.e., in different bipartition sets). Elements of $S_T^{A,B}$ correspond to the edges of T lying on the path from the A to B , so

$$D_T(A, B) = |S_T^{A,B}|. \quad (6)$$

This means on an individual gene tree the distances used in USTAR are simply counts of ‘separating’ splits, with gene samples being judged further apart when there are more splits on T which separate them. Thus graph-theoretic distance might also be called ‘split separation distance.’

Now for any distribution μ of gene trees, if $\mathcal{A}|\mathcal{B}$ is a split of \mathcal{X} , and $\mathbb{P}_\mu(\mathcal{A}|\mathcal{B})$ denotes the probability of the event that an observed gene tree displays split $\mathcal{A}|\mathcal{B}_g$, then

$$\mathbb{P}_\mu(\mathcal{A}|\mathcal{B}) = \sum_{T \text{ displaying } \mathcal{A}_g|\mathcal{B}_g} \mathbb{P}_\mu(T).$$

Theorem 5.1

For any distribution μ of gene trees, the collection of split probabilities $\{\mathbb{P}_\mu(\mathcal{A}|\mathcal{B})\}$ determines $\mathbb{E}_\mu(D_T)$.

Proof—Define indicator functions

$$I_{\mathcal{A}|\mathcal{B}}(T) = \begin{cases} 1 & \text{if } T \text{ displays } \mathcal{A}_g|\mathcal{B}_g, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$JI_{A,B}(\mathcal{A}|\mathcal{B}) = \begin{cases} 1 & \text{if } A, B \text{ separated in } \mathcal{A}_g|\mathcal{B}_g, \\ 0 & \text{otherwise.} \end{cases}$$

Then using equation (6),

$$\begin{aligned} \mathbb{E}_\mu(D_T(A, B)) &= \sum_T \mathbb{P}_\mu(T) D_T(A, B) \\ &= \sum_T \mathbb{P}_\mu(T) |S_T^{A,B}| \\ &= \sum_T \mathbb{P}_\mu(T) \left(\sum_{\substack{\text{splits} \\ \mathcal{A}|\mathcal{B}}} I_{\mathcal{A}|\mathcal{B}}(T) J_{\mathcal{A},\mathcal{B}}(\mathcal{A}|\mathcal{B}) \right) \\ &= \sum_{\substack{\text{splits} \\ \mathcal{A}|\mathcal{B}}} \left(\sum_T \mathbb{P}_\mu(T) I_{\mathcal{A}|\mathcal{B}}(T) \right) J_{A,B}(\mathcal{A}|\mathcal{B}) \\ &= \sum_{\substack{\text{splits} \\ \mathcal{A}|\mathcal{B}}} \mathbb{P}_\mu(\mathcal{A}|\mathcal{B}) J_{A,B}(\mathcal{A}|\mathcal{B}), \end{aligned} \tag{7}$$

so the USTAR dissimilarity is computable from split probabilities.

Of course the distribution μ we have in mind here is either the one arising from the multispecies coalescent model, or an empirical one from a sample of gene trees from that model.

From Theorems 5.1 and 4.2 we immediately obtain the following:

Corollary 5.2

The unrooted species tree topology is identifiable from split probabilities under the multispecies coalescent.

It is known [14] that the *rooted* species tree topology is identifiable from the distribution of *unrooted* gene tree topologies. It is also known that the rooted species tree topology is identifiable from clade probabilities. Thus a natural question is whether the split probabilities, the unrooted analogues of clade probabilities, can further identify the root on the species tree. Though our investigation here does not seem to shed light on this, we plan to address it in another work.

6 USTAR WITH MULTIPLE SAMPLES PER TAXON

When NJ_{st} was introduced in [2], a suggestion was given for how one might deal with gene trees relating multiple lineages sampled from each taxon. For a collection \mathcal{T} of gene trees, if $T \in \mathcal{T}$ relates $m_a(T)$ lineages sampled from taxon a and $m_b(T)$ lineages from taxon b , then intertaxon distances were defined (up to an additive constant) as an average

$$D(a, b) = \frac{\sum_{T \in \mathcal{T}} \sum_{\substack{1 \leq i \leq m_a(T) \\ 1 \leq j \leq m_b(T)}} D_T(A_i, B_j)}{\sum_{T \in \mathcal{T}} m_a(T) m_b(T)}, \quad (8)$$

where $D_T(A_i, B_j)$ denotes the graph theoretic distance on tree T between the i th sample of gene A and the j th of gene B . Unfortunately, this approach is not statistically consistent. In fact, as the size of the sample of gene trees is increased, the probability of inferring the correct species tree can approach 0. After demonstrating this, we propose a different method of handling multiple samples per taxon, one that is statistically consistent.

To investigate the behavior of formula (8), consider the species tree

$$((a, b), (c, d)),$$

with all branch lengths long enough that incomplete lineage sorting between different taxa is vanishingly rare. Sample lineages for a large number of genes as follows: For 50% of the genes, sample 3 lineages in each of taxa a, b and 1 lineage in each of taxa c, d . In the other 50% of genes, sample 1 lineage in taxa a, b and 3 lineages in taxa c, d . For sufficiently long branch lengths on the species tree, the coalescent model gives that the sampled genes trees will be approximately equally of topologies

$$(((A, A), A), ((B, B), B)), (C, D))$$

and

$$((A, B), (((C, C), C), ((D, D), D))),$$

with an arbitrarily small fraction of gene trees of other topologies. For the first of these gene tree topologies, after unrooting and assigning all edges the length 1, the different interlineage USTAR distances are

$$\begin{aligned} D_T(A_i, B_j) &= 6, 6, 6, 6, 5, 5, 5, 5, 4; \\ D_T(x, y) &= 5, 5, 4, \text{ for } x=A_i, B_j, y=C, D; \\ D_T(C, D) &= 2. \end{aligned}$$

For the second tree, the same distances arise, but with the roles of A, B interchanged with C, D . Then formula (8) gives intertaxon distances arbitrarily close to

$$\begin{aligned} D(a, b) &= D(c, d) = \frac{.5(6+6+6+6+5+5+5+5+4) + .5(2)}{.5(9) + .5(1)} = 5 \\ D(x, y) &= \frac{.5(5+5+4) + .5(5+5+4)}{.5(3) + .5(3)} = \frac{14}{3} \text{ for } x=a, b, y=c, d. \end{aligned}$$

These intertaxon distances do not fit any unrooted topological tree, as they do not satisfy the four-point condition [16]. In fact, selection of a tree topology by applying (part of) the four-point condition requires computing

$$\begin{aligned} D(a, b) + D(c, d) &= 5 + 5 = 10, \\ D(a, c) + D(b, d) &= \frac{14}{3} + \frac{14}{3} = \frac{28}{3}, \\ D(a, d) + D(b, c) &= \frac{14}{3} + \frac{14}{3} = \frac{28}{3} \end{aligned}$$

and choosing the smallest to determine the cherries of the tree. Here the smallest is a tie, yielding the two incorrect topologies, $((a, c), (b, d))$ and $((a, d), (b, c))$. Neighbor Joining, which is built upon this selection criterion, would choose either of the incorrect topologies with equal probability, and then go on to compute positive lengths for the edges, obtaining either of the unrooted metric trees $((a:2.333, c:2.333):0.167, b:2.333, d:2.333)$ or $((a:2.333, d:2.333):0.167, b:2.333, c:2.333)$.

Finite length edges on the species tree will only produce intertaxon distances arbitrarily close to those in the calculations above, with probability approaching 1 as the number of gene trees increases. However, continuity of the Neighbor Joining algorithm at these distances implies that the output of Neighbor Joining will be the wrong topology with probability approaching 1.

A different approach to averaging than the one used in formula (8) can however lead to statistically consistent inference of the species tree.

First, suppose multiple samples are drawn from taxa in exactly the same number for each gene. That is, there are integers $m_x \geq 1$ so that each gene tree has m_x leaves X_1, X_2, \dots, X_{m_x} for each $x \in \mathcal{X}$, for a total of $\sum_{x \in \mathcal{X}} m_x$ leaves. We will refer to a specific choice of the numbers $(m_x)_{x \in \mathcal{X}}$ as a *multisample scheme*.

For a single fixed multisample scheme, the results of previous sections apply if we replace the species tree by one where m_x edges are attached to the leaf formerly labeled x with the new leaves labeled x_1, x_2, \dots, x_{m_x} . (This is called the *extended species tree* in [19].) While this tree is not binary, one can consider binary perturbations of it, and use continuity to conclude that the expected USTAR dissimilarity on $\sum_x m_x$ taxa will exactly fit this tree. If one then defines $D(a, b)$ as the expectation of $D_T(A_1, B_1)$ for each $a, b \in \mathcal{X}$, or as the expectation of the average of $D_T(A_i, B_j)$ over $1 \leq i \leq m_a$ and $1 \leq j \leq m_b$, the expected dissimilarity on \mathcal{X} is the same, as the A_i lineages for various i are exchangeable under the coalescent model. Since this expected dissimilarity must exactly fit the unrooted topology relating only the X_1 for $x \in \mathcal{X}$, it thus fits the unrooted topology relating the taxa in \mathcal{X} . Thus either retaining only one sample per taxon, or averaging over the lineages sampled from each taxon will lead to consistent inference. Since data sets have only a finite number of gene trees, by averaging the empirical $D_T(A_i, B_j)$ one would hope to improve one's estimate of the expected value, so we choose to do so. Moreover, one could obtain the same dissimilarity by averaging over samples for each gene tree T individually, creating a USTAR dissimilarity matrix for \mathcal{X} from one tree at a time, and then averaging over these.

Now suppose we specify a finite number of multisample schemes, as well as probabilities of using each one for any gene. Given a data set of gene trees obtained from such an approach, as described in the last paragraph one could apply a USTAR method averaged over multiple lineage samples to each subcollection of trees with the same multisample scheme. But since the dissimilarity for each such subcollection in expectation approaches one fitting the species tree as the number of gene trees increases, then any weighted average of them over the multisample schemes does as well. This is a consequence of the dissimilarity arising from each subcollection satisfying the same four-point condition equality and inequalities, so a convex linear combination of them does also. Thus with multisample schemes $(m_{x,s})_{x \in \mathcal{X}}$ for $1 \leq s \leq S$, and any non-negative weighting constants α_s , if we define an empirical dissimilarity as

$$\hat{D}(a, b) = \sum_{1 \leq s \leq S} \alpha_s \sum_{\substack{T \text{ displaying} \\ (m_{x,s})}} \frac{1}{m_{a,s} m_{b,s}} \sum_{\substack{1 \leq i \leq m_{a,s} \\ 1 \leq j \leq m_{b,s}}} D_T(A_i, B_j) \quad (9)$$

then we will have consistent inference provided the number of gene trees for each scheme in the sum all go to infinity. Choosing $\alpha_s = 1/|\mathcal{T}|$ where \mathcal{T} is the collection of gene trees yields our suggested formula

$$\hat{D}(a, b) = \frac{1}{|\mathcal{T}|} \sum_{T \in \mathcal{T}} \frac{1}{m_a(T)m_b(T)} \sum_{\substack{1 \leq i \leq m_a(T) \\ 1 \leq j \leq m_b(T)}} D_T(A_i, B_j). \quad (10)$$

Note that the formula (9) cannot be specialized to give formula (8). Taking $\alpha_s = m_{a,s}m_{b,s}/\sum_s m_{a,s}m_{b,s}$ does make them agree for the single comparison of a and b , but will not for other pairs of taxa (unless $m_{x,s}$ is independent of x).

The essential difference between the formulas (10) and (8) is how the product $m_a(T)m_b(T)$ appears in them. In formula (8) all $D_T(A_i, B_j)$ are treated on an equal basis, whether they come from the same locus and are therefore correlated, or from different loci and thus independent trials of the coalescent process. Formula (10) can be viewed as first constructing an intertaxon distance matrix for each locus by averaging pairwise distances over choices of alleles, and then averaging these over loci, to create a final intertaxon distance matrix.

We emphasize that using the consistency of formulas (9) and (10) to justify their use in applying USTAR to finite data sets hinges on an assumption that every multisample scheme that appears in a data set appears many times. Particularly for data sets assembled from several earlier studies, there may be little commonality in the sampling scheme from one gene to the next. Simulations are needed to explore whether our formulas behave well under such circumstances.

Simulations in [3] testing the performance of USTAR methods did not explore multisample schemes at all. However, in that work a new variant of a USTAR method that allows for gene trees missing some taxa was studied — in the notation above the $m_x(T)$ could be 1 or 0. Although such USTAR methods were reported to perform well on simulated data under these circumstances, theoretical justification for the particular approach taken has yet to be developed. Moreover, one should be cautious that if the test simulations involve random deletion of taxa from gene trees, they may not be relevant to empirical data sets in which taxa are missing in more patterned ways.

Acknowledgments

This work was begun while EA and JR were Short-term Visitors and JD was a Sabbatical Fellow at the National Institute for Mathematical and Biological Synthesis, an institute sponsored by the National Science Foundation, the U.S. Department of Homeland Security, and the U.S. Department of Agriculture through NSF Award #EF-0832858, with additional support from the University of Tennessee, Knoxville. It was further supported by the National Institutes of Health grant R01 GM117590, awarded under the Joint DMS/NIGMS Initiative to Support Research at the Interface of the Biological and Mathematical Sciences.

References

1. Liu L, Yu L, Pearl DK, Edwards SV. Estimating species phylogenies using coalescence times among sequences. *Syst Biol.* 2009; 58:468–477. [PubMed: 20525601]
2. Liu L, Yu L. Estimating species trees from unrooted gene trees. *Syst Biol.* 2011; 60:661–667. [PubMed: 21447481]

3. Vachaspati P, Warnow T. ASTRID: Accurate Species TREes from Internode Distances. *BMC Genomics*. 2015; 16(Suppl 10):S3.
4. Allman ES, Degnan JH, Rhodes JA. Species tree inference by the STAR method, and generalizations. *J Comput Biol*. 2013; 20(1):50–61. [PubMed: 23294273]
5. Kreidl M. Note on expected internode distances for gene trees in species trees. 2011 arXiv: 1108.5154.
6. Gascuel O. BIONJ: an improved version of the NJ algorithm based on a simple model of sequence data. *Mol Biol Evol*. 1997; 14(7):685–695. [PubMed: 9254330]
7. Lefort V, Desper R, Gascuel O. FastME 2.0: A comprehensive, accurate, and fast distance-based phylogeny inference program. *Mol Biol Evol*. 2015; 32(10):2798–2800. [PubMed: 26130081]
8. Mossel E, Roch S. Incomplete lineage sorting: consistent phylogeny estimation from multiple loci. *IEEE/ACM Trans Comput Biol Bioinf*. 2010; 7(1):166–171.
9. DeGiorgio M, Degnan JH. Robustness to divergence time underestimation when inferring species trees from estimated gene trees. *Syst Biol*. 2014; 63:66–82. [PubMed: 23988674]
10. Roch S, Warnow T. On the robustness to gene tree estimation error (or lack thereof) of coalescent-based species tree methods. *Syst Biol*. 2015:syv016.
11. Mirarab S, Bayzid MS, Warnow T. Evaluating summary methods for multilocus species tree estimation in the presence of incomplete lineage sorting. *Syst Biol*.
12. Kubatko LS, Degnan JH. Inconsistency of phylogenetic estimates from concatenated data under coalescence. *Systematic Biology*. 2007; 56(1):17–24. [PubMed: 17366134]
13. Roch S, Steel M. Likelihood-based tree reconstruction on a concatenation of aligned sequence data sets can be statistically inconsistent. *Theoretical population biology*. 2015; 100:56–62.
14. Allman ES, Degnan JH, Rhodes JA. Identifying the rooted species tree from the distribution of unrooted gene trees under the coalescent. *J Math Biol*. 2011; 62(6):833–862. [PubMed: 20652704]
15. Atteson K. The performance of neighbor-joining methods of phylogeny reconstruction. *Algorithmica*. 1999; 25(2):251–278.
16. Semple, C., Steel, M. *Phylogenetics*, ser. Oxford Lecture Series in Mathematics and its Applications. Vol. 24. Oxford: Oxford University Press; 2003.
17. Bryant D. A classification of consensus methods for phylogenetics. DIMACS series in discrete mathematics and theoretical computer science. 2003; 61:163–184.
18. Degnan JH, DeGiorgio M, Bryant D, Rosenberg NA. Properties of consensus methods for inferring species trees from gene trees. *Syst Biol*. 2009; 58:35–54. [PubMed: 20525567]
19. Allman ES, Degnan JH, Rhodes JA. Determining species tree topologies from clade probabilities under the coalescent. *J Theor Biol*. 2011; 289:96–106. [PubMed: 21867714]

Biographies



Elizabeth S. Allman earned her Ph.D. in Mathematics from the University of California at Los Angeles in 1995, with a thesis in Brauer Groups. Co-authoring an undergraduate textbook in Mathematical Biology while teaching at the University of Southern Maine led to a new interest in phylogenetics, which is now her primary focus. Currently at the University of Alaska Fairbanks, her work generally uses algebraic perspectives to analyze complex statistical models.



James H. Degnan received the PhD degree in statistics in 2005 from the University of New Mexico. He did postdoctoral research at Harvard and the University of Michigan. After spending five years at the University of Canterbury in Christchurch, New Zealand, he moved to the Department of Mathematics and Statistics, University of New Mexico, where he primarily works in statistical phylogenetics.



John A. Rhodes received his Ph.D. degree in Mathematics from the Massachusetts Institute of Technology in 1986, in Number Theory. A curriculum development project with E.S. Allman led to an undergraduate textbook on Mathematical Models in Biology, and a switch in research focus to the mathematics of phylogenetics. In 2005, he moved to the University of Alaska Fairbanks, where he recently ended a term as department chair of Mathematics and Statistics.




Fig. 1.

A gene tree $((A, B), ((C, D), E))$ (left) and its metrizations for the generalized STAR method discussed in the text (center), and USTAR (right).