



HHS Public Access

Author manuscript

IEEE/ACM Trans Comput Biol Bioinform. Author manuscript; available in PMC 2021 December 08.

Published in final edited form as:

IEEE/ACM Trans Comput Biol Bioinform. 2020 ; 17(6): 2107–2118. doi:10.1109/TCBB.2019.2917204.

Topological metrizations of trees, and new quartet methods of tree inference

John A. Rhodes

Department of Mathematics and Statistics, University of Alaska, Fairbanks

Abstract

Topological phylogenetic trees can be assigned edge weights in several natural ways, highlighting different aspects of the tree. Here the rooted triple and quartet metrizations are introduced, and applied to formulate novel methods of inferring large trees from rooted triple and quartet data. These methods lead to new statistically consistent procedures for inference of a species tree from gene trees under the multispecies coalescent model.

Keywords

Phylogeny; Genomics; Evolution

1 INTRODUCTION

The inference of a species trees, which shows the evolutionary relationships between a collection of taxa, from gene trees, which depict the joining of ancestral lineages for genes sampled from individuals of those taxa, is made difficult by the fact that gene tree topologies often differ from each other. One biological explanation for such gene tree incongruity is *incomplete lineage sorting* (ILS). In ILS, gene lineages may not share a common ancestor in the most recent ancestral population in which it is possible for them to do so. This allows the lineages to merge with those from more distantly-related species before they do so with closer ones. The formation of gene trees within species trees taking into account ILS is described by the *multispecies coalescent model* (MSC). Though many inference methods have been proposed to recover a species tree from a collection of gene trees under the MSC, both computational requirements and performance in simulation vary enough that no single approach has yet become clearly preferred.

In [21] and [20], Liu and coworkers proposed particularly interesting and fast methods for this inference problem, using as data collections of unrooted and rooted gene trees, respectively. These methods, called STAR and NJ_{st} , proceed by first discarding any metric information on the gene trees, and then remetrizing them in a way that reflects only their topological structure. For instance, in the second of these works, the metrization of an unrooted gene tree is to simply make all edges have length 1. A table of intertaxon distances is then constructed for each gene tree, the mean distance table across the gene trees is

computed, and this mean table is used to infer a species tree by a standard distance approach such as Neighbor Joining. In other words, average consensus [8], [18] is applied to the remetrized gene trees. Although far from intuitively clear, the statistical consistency of these methods has been established [2], [3]. Moreover, the strong performance of the methods on simulated data has been shown both in the original works, and by the implementation of NJ_{st} in the software ASTRID [34].

Although not discussed in [21] or [20], but as explained below, the two remetrizations of gene trees these works use can be understood by viewing them as related to the notions of clades and splits on trees. That is, the intertaxon distances on a remetrized gene tree provides a numerical summary of the tree built on these specific combinatorial notions. Since clades and splits are only two of the combinatorial tools useful for describing topological trees, natural questions are 1) what other combinatorial notions lead to metrizations that usefully capture topological information about trees?, and 2) how can these metrizations be used for tree inference? The goal of this note is to investigate these questions for the notions of displayed rooted triples (induced rooted 3-leaf trees) and quartets (induced unrooted 4-leaf trees).

The rooted triple and quartet metrizations developed here have an important feature: Intertaxon distances can be computed from a list of rooted triples or quartets displayed on a tree — without knowing the tree itself. This allows them to be used in new supertree methods, which take a collection of rooted triples or quartets, compute a pairwise distance table from them, and then construct a large tree which fits this distance table. If the input rooted triples or quartets are the full set of those displayed on a large tree, this recovers the tree. Even if some of the triples or quartets are erroneous or missing, approximations to the intertaxon distances on the tree are still obtained, and a tree can be quickly inferred by any of the well-known distance-based methods for tree building or selection that are robust to error. We call these methods Quartet Distance Supertree (QDS) and Rooted Triple Distance Supertree (RTDS).

This approach extends to give new statistically consistent methods of species tree inference from samples of topological gene trees drawn from the MSC model, adding to those few already known [2], [3], [13]. The key additional ingredient for this inference is the fact that under the MSC the most frequent rooted triple or quartet topology across a collection of independent genes reflects that of the species tree. From a collection of gene trees, one can tabulate frequencies of the displayed quartets on four taxa, choosing the most frequent as the inferred species quartet on those taxa. Then QDS can be used for statistically consistent inference of the full species tree. We refer to this method, and its analog using rooted triples, as Quartet Distance Consensus (QDC) and Rooted Triple Distance Consensus (RTDC).

While QDC for species tree inference should certainly not be expected to have the speed of NJ_{st} , due to its need to consider quartets individually, it offers another advantage. Specifically, it remains statistically consistent even when some gene trees have some taxa missing. Although the quartet-based scheme implemented in the software ASTRAL-III should also be less vulnerable to problems with missing taxa on gene trees, QDC may in some circumstances offer advantages over it as well. In particular, a theoretical complexity

analysis indicates that while QDC's running time has higher exponent on the number of taxa than does ASTRAL-III's, it has lower exponent on the number of gene trees. Thus for a moderate number of taxa but a large number of gene trees, an efficient implementation may offer superior runtimes.

Although the focus of this paper is on the development of the quartet and rooted triple metrizations, using simulated datasets previously used to evaluate ASTRID and ASTRAL-III we provide some evidence on the performance of QDC in comparison to NJ_{st} and ASTRAL-III. Determining whether this behavior is typical, however, requires further work. An efficient software implementation and large-scale simulation studies of performance are necessary for a more complete assessment. Regardless of the ultimate performance of these algorithms, though, the distances on which they are built may find other uses in phylogenetic theory and practice.

Basic definitions are given in Section 2. The clade and split metrizations are presented in detail in Section 3. The main theoretical contributions of this note are the two new metrizations associated to rooted triples and quartets developed in Sections 4 and 5. The applicability of this theory to supertree inference from rooted triple and quartets and to species tree inference from gene trees is then developed in Sections 6 and 7, where simulation results are also given. Section 8 concludes with a few general comments.

2 NOTATION AND TERMINOLOGY

Throughout this work X denotes a finite set of N taxa. Upper case letters A, B, \dots denote subsets of X , and lower case letters a, b, \dots elements of X .

A *split of X* is a bipartition $X = A \sqcup B$ of the taxa into non-empty subsets, and is denoted $A|B = B|A$. A *clade of X* is a non-empty subset $A \subseteq X$ of the taxa. A resolved *rooted triple of X* is a subset of three elements of X , partitioned into a pair a, b and a singleton c , and denoted $ab|c = ba|c$. To allow for multifurcations on trees, we will also have need for an unresolved rooted triple abc . A resolved *quartet of X* is a subset of four elements of X , partitioned into two pairs a, b and c, d , and denoted $ab|cd = ba|cd = \dots = cd|ab$. An unresolved quartet is $abcd$.

Suppose the taxa X bijectively label the leaves of a rooted tree T , or of an unrooted tree T , with the root of degree at least 2 and all other internal nodes of degree at least 3. Then T and T are said to be *phylogenetic trees on X* . All edges on a rooted tree T are directed away from the root, so, for instance, the root is ancestral to all leaves. Edges on an unrooted tree T are undirected. A phylogenetic tree is *binary* if the minimal degree conditions on the nodes are met, and is otherwise said to be *polytomous*.

A tree T *displays the clade A* if the most recent common ancestor (MRCA) on T of the taxa in A has as its descendants in X precisely the set A . Thus clades displayed on a rooted tree correspond to its nodes, and if the tree is binary, it displays exactly $2N - 1$ clades, including all singleton clades and the clade X . We say T *displays the rooted triple $ab|c$* if the MRCA of a and b is a proper descendent of the MRCA of a, b , and c . In the case of a rooted polytomous tree, we say the unresolved rooted triple abc is displayed if the MRCAs of the

three pairs a, b , a, c , and b, c coincide. If $N = 3$, a rooted tree on X thus displays $\binom{N}{3}$ rooted triples, one for each choice of three taxa.

Similarly, for an unrooted tree T on taxa X , we say T displays the split $A|B$ if the bipartition is obtained by removing an edge of T and partitioning X according to connected components of the resulting graph. If $N = 2$, a binary tree displays $2N - 3$ splits. We say T displays the quartet $ab|cd$ if on the induced 4-leaf tree relating a, b, c , and d the split $\{a, b\}|\{c, d\}$ is displayed. A tree displays the unresolved quartet $abcd$ if the induced tree relating them is a star tree. If $N = 4$ a tree thus displays $\binom{N}{4}$ quartets, one for each subset of four taxa.

Suppose positive weights are somehow assigned to the edges of T' or T , so the tree is now a *metric tree*. Any such edge weighting scheme W induces a metric d_W on X , using the sum of edge weights along paths between pairs of taxa. As is well known, however, a metric d on X need not arise from such a weighting. If $d = d_W$ for some W on T or T' , then we say d is a *tree metric* on T or T' with weighting W' .

3 SPLIT AND CLADE METRIZATIONS OF TREES

For completeness and perspective on what is to follow, we present two topological metrizations of trees that have been used in other works.

Given an unrooted topological tree T on X , we may assign weights $w(e) = 1$ to all edges e . The resulting tree metric on X is just the usual graph-theoretic distance along T . However, by the correspondence between displayed splits and edges on the tree, the distance between taxa x and y can also be described as the number of splits $A|B$, displayed on T that *separate* x, y , in the sense that $x \in A, y \in B$. For this reason, we denote the weighting scheme with all weights 1 by Sp , and say it gives the *split metrization* of T . This is essentially the metrization used in [20] for the NJ_{st} algorithm for species tree inference, which was renamed U-STAR/NJ in [3] since the same distance approach generalizes to U-STAR/M for any distance method M of tree construction or selection. The ASTRID software [34] is an implementation of these U-STAR methods.

For a rooted topological tree T' on X , assign numbers to the internal nodes of the tree as follows: To the root assign N , to its children that are internal nodes assign $N - 1$, to their children that are internal nodes assign $N - 2$, and so on, decreasing by 1 for each parent-to-child step. Assign 0 to all leaves. Then assign edge weights $w(e)$ as the positive difference of the numbers on the endpoints of e . Thus all internal edges are weighted 1, but terminal edges are weighted with possibly different numbers between 2 and N . All leaf-to-root distances are N , so the tree is ultrametric. We denote this weighting scheme by Cl , say it gives the *clade metrization* of T' , and denote the induced metric on X by d_{Cl} . The name is justified by the observation in [2] that for $x, y \in X$,

$$d_{Cl}(x, y) = 2(1 + N - |C_{x, y}|)$$

where $C_{x,y}$ is the set of clades displayed on T that contain both x and y . This follows directly by the correspondence between nodes on the path between the tree root and $\text{MRCA}(x, y)$ to displayed clades containing x and y . This clade metrization was introduced in [21] for the STAR algorithm for species tree inference.

Remark 1.

As shown in [2], there are generalizations of the clade metrization which can be used for consistent inference of species trees from gene trees following the plan of [21]. These generalizations, which allow for the weight of an edge to depend on the number of edges between it and the root, are not used in this paper.

Remark 2.

One might propose an alternative clade metrization, Cl' defined by assigning unit weights to all edges in a rooted tree. Then the intertaxon distance $d_{Cl'}(x, y)$ is the number of clades displayed on the tree that contain one of x, y but not both. However, $d_{Cl'}$ lacks the ultrametricity of d_{Cl} . It can also be obtained by restricting the split distance on the larger unrooted tree obtained by attaching to the root a single edge leading to an extra taxon.

4 ROOTED TRIPLE METRIZATION OF A ROOTED TREE

With T a rooted phylogenetic tree on X , we may assign edge weights to T as follows: First number each node of the tree, including leaves, with the number of taxa descended from it. Leaves are numbered 1, as they are considered their own descendants, and the root is numbered N , the total number of taxa. Then assign weights $w(e)$ as the positive difference of the numbers on the endpoints of e . That is, for any edge $e = (u, v)$ directed away from the root with u the parent of v , the edge weight is

$$w_{RT}(e) = |\text{descendants}(u)| - |\text{descendants}(v)|,$$

the decrease in number of descendants across e . We refer to this as the *rooted triple metrization*, a name justified by Theorem 1 below, and denote the weighting scheme RT . It results in an ultrametric tree, with the root at distance $N - 1$ from every leaf, and more generally every internal node u at distance $|\text{descendants}(u)| - 1$ from its leaf descendants.

Example 3.

With the rooted triple metrization, a rooted caterpillar tree

$$(\dots(((a_1, a_2), a_3), a_4), \dots, a_N)$$

will have all internal edges of weight 1. The pendant edges, listed from the cherry toward the root, will have weights 1, 1, 2, 3, 4, ..., $N - 1$.

Example 4.

With the rooted triple metrization, a balanced tree

$$(\dots((a_1, a_2), (a_3, a_4)), \dots, ((a_{N-3}, a_{N-2}), (a_{N-1}, a_N))\dots)$$

on $N = 2^k$ taxa will have pendant edges of weight 1, and as one moves toward the root, internal edges of weight 2, 4, 8, ..., 2^{k-1} .

Theorem 1.

Suppose a rooted phylogenetic tree T is given the rooted triple metrization. Then the resulting tree metric d_{RT} satisfies, for all $x, y \in X$, $x \neq y$,

$$d_{RT}(x, y) = 2|R_{x, y}| + 2,$$

where $R_{x, y}$ is the set of rooted triples displayed on T of the form $xz|y$, $yz|x$, or xyz .

More informally, for a binary tree T the distance $d_{RT}(x, y)$ is, up to a simple transformation, the number of rooted triples displayed on T in which x, y are *separated*. This remains true for trees with polytomies as long as unresolved triples are viewed as separating all taxa in them.

Proof. Let $v = \text{MRCA}(x, y)$, and k be the number of leaf descendants of v (i.e., k is the size of the smallest displayed clade containing x, y). Since $d_{RT}(x, y)$ is the sum of edge weights on the path between x and y , we find that

$$\begin{aligned} d_{RT}(x, y) &= d_{RT}(x, v) + d_{RT}(v, y) \\ &= (k-1) + (k-1) = 2k-2 \end{aligned}$$

The number of rooted triples of the forms $xz|y$, $yz|x$, or xyz is the number of taxa z descended from v , excluding x and y . Thus

$$|R_{x, y}| = k - 2.$$

Eliminating k from these two equations yields the claim. \square

Remark 5.

Combining Theorem 1 with the fact that one can determine a rooted ultrametric tree from intertaxon distances on it gives an alternative proof of the well-known result that the collection of rooted triples displayed on a tree determines the rooted tree topology.

5 QUARTET METRIZATION OF A UNROOTED TREE

Let T be an unrooted binary tree on X . Each internal edge of T determines a partition of X into 4 non-empty blocks, X_1, X_2, X_3, X_4 where the split associated to the edge is $X_1 \cup X_2 | X_3 \cup X_4$, and the splits associated to the 4 adjacent edges all have an X_i as one split set. We refer to this partition as the *quartet partition* associated to an internal edge, and denote it by $X_1, X_2 | X_3, X_4$. Assign an internal edge e with quartet partition $X_1, X_2 | X_3, X_4$ the weight

$$w_Q(e) = |X_1||X_2| + |X_3||X_4|.$$

For a pendant edge to leaf x , the non-leaf endpoint determines a tripartition of the taxa, $\{x\}$, X_1 , X_2 . Assign to such an edge e the weight

$$w_Q(e) = |X_1||X_2|.$$

This weighting scheme can be extended to polytomous trees as follows: Suppose an edge $e = \{u, v\}$ is incident to m_u other edges e_i at u , and m_v other edges \tilde{e}_j at v . Note that for a leaf v we allow $m_v = 0$. Now let $X_1, X_2, \dots, X_{m_u}, \tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_{m_v}$ be the partition of X with X_i (respectively \tilde{X}_j) the set of taxa connected to e by a path through e_i (respectively \tilde{e}_j). Then assign to e the edge weight

$$w_Q(e) = \sum_{1 \leq i < i' \leq m_u} |X_i||X_{i'}| + \sum_{1 \leq j < j' \leq m_v} |\tilde{X}_j||\tilde{X}_{j'}|.$$

Interpreting an empty sum as 0, this agrees with the definition above for binary trees.

We refer to this as the *quartet metrization*, due to Theorem 2 below, and denote the weighting scheme Q . Trees with the quartet metrization are usually not ultrametric, as examples show.

Example 6.

An unrooted caterpillar tree

$$(\dots((a_1, a_2), a_3), a_4), \dots, a_N)$$

will have internal edges inducing quartet partitions $X_1, X_2|X_3, X_4$ with sets of size

$$|X_1| = k - 1, |X_2| = 1, |X_3| = 1, |X_4| = N - k - 1,$$

for $k = 2, 3, \dots, N - 2$. Under the quartet metrization, the internal edge weights will thus all be $N - 2$. The pendant edges to taxa a_1, a_2, a_{N-1}, a_N will also have weights $N - 2$. Pendant edges to taxa a_k , for $k = 3, \dots, N - 2$, will have weights $(k - 1)(N - k)$. A 16-taxon illustration is shown in Figure 1.

Example 7.

An unrooted balanced tree

$$(\dots((a_1, a_2), (a_3, a_4)), \dots, ((a_{N-3}, a_{N-2}), (a_{N-1}, a_N)) \dots)$$

on $N = 2^k$ taxa will have pendant edges of weight $N - 2$. With $\ell \in \{1, 2, \dots, k - 1\}$ denoting the minimal number of edges needed to connect a given internal edge to a leaf, the central internal edge, for which $\ell = k - 1$, has weight $2^{k-2} \cdot 2^{k-2} + 2^{k-2} \cdot 2^{k-2} = 2^{2k-3}$, while other internal edges, with $1 \leq \ell \leq k - 2$, are of weight

$$(2^{\ell-1})(2^{\ell-1}) + (2^{\ell})(2^k - 2^{\ell+1}) = 2^{k+\ell-7} \cdot 2^{2\ell-2}.$$

A 16-taxon illustration is shown in Figure 1.

Theorem 2.

Suppose an unrooted phylogenetic tree T has been given the quartet metrization. Then the resulting tree metric d_Q satisfies for all $x, y \in X$, $x \neq y$,

$$d_Q(x, y) = 2|Q_{x,y}| + 2N - 4 \quad (1)$$

where $Q_{x,y}$ is the set of quartets displayed on T of the form $xz|yw$ or $xy|zw$.

More informally, for a binary T the distance $d_Q(x, y)$ is, up to a simple transformation, the number of quartets displayed on T in which x, y are *separated*. This remains true for trees with polytomies as long as unresolved quartets are viewed as separating all their taxa.

Although Theorem 2 can be deduced from Theorem 1 by summing its formula over all placements of the root on pendant edges of T , a more direct argument is given here.

Proof. Fix taxa $x \neq y$, and let P denote the path in T between them. Any node $v \in P$ determines a partition of the taxa $X = A_v \sqcup B_v \sqcup C_v^1 \dots \sqcup C_v^{k-2}$ as follows: If v has degree k , deleting v and its incident edges partitions X into k non-empty subsets according to the connected components of the resulting graph. Let A_v denote the partition set containing x , B_v the one containing y , and C_v^i the $k - 2$ remaining ones. Thus $C_v = \cup C_v^i$ contains all those taxa z for which a path from z to x or y joins P at v .

Now any quartet $xu|yz$ or $xy|uz$ that is displayed on the tree T determines a node v on P at which the path from u to y joins P . Suppose v has degree $k = k(v)$. Then the number of quartets of these forms that are displayed on T and determine the same node v in this way is

$$\begin{aligned} & \sum_{1 \leq i \leq k} |C_v^i| \left(\sum_{i < j \leq k} |C_v^j| + |B_v| - 1 \right) \\ &= \sum_{1 \leq i \leq k} |C_v^i| \left(\sum_{i < j \leq k} |C_v^j| + |B_v| \right) - |C_v|. \end{aligned}$$

Thus

$$\begin{aligned}
& |Q_{x,y}| \\
&= \sum_{\substack{v \text{ on } P \\ v \neq x,y}} \left(\sum_{1 \leq i \leq k(v)} |C_v^i| \left(\sum_{i < j \leq k(v)} |C_v^j| + |B_v| \right) - |C_v| \right) \\
&= \sum_{\substack{v \text{ on } P \\ v \neq x,y}} \left(\sum_{1 \leq i \leq k(v)} |C_v^i| \left(\sum_{i < j \leq k(v)} |C_v^j| + |B_v| \right) \right) - N + 2.
\end{aligned} \tag{2}$$

Interchanging the roles of x and y we also find

$$\begin{aligned}
& |Q_{x,y}| \\
&= \sum_{\substack{w \text{ on } P \\ w \neq x,y}} \left(\sum_{1 \leq i \leq k(w)} |C_w^i| \left(\sum_{i < j \leq k(w)} |C_w^j| + |A_w| \right) \right) - N + 2.
\end{aligned} \tag{3}$$

Adding equations (2) and (3) and expressing the sums over v, w as a single sum over edges $e = (w, v)$ in P shows

$$2|Q_{x,y}| = \sum_{e \in P} w_Q(e) - 2N + 4 = d_Q(x,y) - 2N + 4,$$

which yields the claim. \square

Remark 8.

Proposition 9 of [5] shows that $|Q_{x,y}| + 1$ for $x \neq y$ yields a tree metric on T , which is equivalent to the right hand side of equation (1) defining a tree metric on T . However, edge weights associated to the tree metric are not investigated in that paper. Moreover, applications of the result to tree inference, such as those discussed in Sections 6 and 7 of this work, seem not to have been pursued in intervening years.

Remark 9.

Combining Theorem 2 with the fact that one can determine an unrooted metric tree from its intertaxon distances gives an alternative proof of the well-known result that the collection of quartets displayed on a tree determines the unrooted tree topology.

Remark 10.

As a heuristic, the ASTRAL-II software [23] introduced a similarity on taxa that counts quartets *not* separating two taxa on a gene tree. By Theorem 2 this is essentially equivalent to the quartet metrization for that gene tree. While ASTRAL-II's goal is species tree inference, it uses this similarity quite differently from the quartet metrization's use in the statistically consistent approach to inference presented in Section 7 below.

6 QUARTET DISTANCE SUPERTREE

Theorems 1 and 2 lead to new supertree methods for finding a tree from certain collections of rooted triples or quartets. We present this fully for quartets, indicating the small modifications for rooted triples in a remark.

Suppose we are given a collection \mathcal{Q} of unweighted quartets on a set of taxa X . We take the viewpoint that most of the given quartets show the correct phylogenetic relationship between the taxa, though some are in error. Ideally \mathcal{Q} contains exactly one quartet for each subset of 4 taxa.

We choose a distance-based method M of tree building or selection that when applied to a tree metric on an unrooted tree T returns T , even if T is not ultrametric. We further require that its output topology is robust to small errors in the input distances at a tree metric. Possible choices for M include NJ and BioNJ (but not UPGMA) for tree building [17], [30], and Balanced Minimum Evolution for tree selection. In practice, a heuristic implementation known to perform well, such as FastME [14], may be chosen.

Algorithm 1.

(QDS/M) Quartet Distance Supertree with distance method M

Input: A collection \mathcal{Q} of quartets on taxa in X

1. For each pair $x, y \in X$ of taxa, $x \neq y$, count the number $q(x, y)$ of quartets in \mathcal{Q} separating x, y , and define the distance $\hat{d}_{\mathcal{Q}}(x, y) = 2q(x, y) + 2N - 4$.
2. Use the distance method M to build or select an unrooted tree from $\hat{d}_{\mathcal{Q}}$.

Remark 11.

The name "Quartet Distance Supertree" has been chosen to emphasize its key uses of 1) displayed quartets on the input trees and 2) a distance method for constructing the supertree. Unfortunately, the term "quartet distance" is often used to refer to a distance between two trees (based on the two trees' displayed quartets) and not the intertaxon distance (based on a collection of quartets). Since supertree methods based on finding a median tree under such an intertree distance have been explored (e.g., [29] for a rooted-triple example), there is some potential confusion with the name chosen here. Nonetheless, the name has not been used before, and provides an accurate brief description.

Remark 12.

If \mathcal{Q} contains either no quartets for some sets of 4 taxa, or multiple quartets on them, one might view this as additional error, and modify the algorithm slightly. For instance, one might use all quartets on a given set of 4 taxa by weighting them by their relative frequency. Omitted 4-taxon subsets might be left out of counting when determining intertaxon distances, or treated as the 3 possible quartets on those taxa, each weighted by 1/3, in counting. However, these are simply heuristic adjustments. Developing any theoretical

justification for them would require some model of the way in which the quartets were produced or omitted.

Remark 13.

For Rooted Triple Distance Supertree with M, one instead counts the number $r(x, y)$ of rooted triples in a set \mathcal{R} that separate x, y , and defines $\hat{d}_{RT} = 2r(x, y) + 2$. The method M can now be chosen to assume ultrametricity (e.g., UPGMA), since \hat{d}_{RT} approximates the ultrametric tree metric d_{RT} . If so, then a rooted tree will be returned.

Although we refer to the method of Algorithm 1 as Quartet Distance Supertree (QDS), and the variant for Rooted Triples as Rooted Triple Distance Supertree (RTDS), for a complete specification it is necessary to also indicate the distance method M used for tree construction or selection. If none of the quartets in \mathcal{Q} are erroneous or omitted, QDS recovers the correct tree. However, how much error, and of what form, can occur in \mathcal{Q} with the desired tree still accurately recovered may depend on the particular distance method M used. Since theoretical guarantees on toleration of error by distance methods tend to be much weaker than results seen in simulation studies, performance of QDS/M needs to be judged through simulation.

Assuming QDS is applied to a list of $\binom{N}{4}$ quartets, one for each 4-taxon subset of a N taxa, the running time to produce the quartet distance matrix will be $\mathcal{O}(N^4)$, since considering each quartet in turn, one can increment counts for the 4 pairs of taxa that quartet separates. If Neighbor Joining, with time $\mathcal{O}(N^3)$, is then used to build a tree, the total complexity remains $\mathcal{O}(N^4)$, which is the best one can achieve for any method that had the same input. While this may be too slow for some large applications, some scheme by which subsets of the quartets are sampled randomly to estimate quartet distances might still give a reasonable approximation to the distance.

The following simulations, performed in R using the ape package [26], give a first indication of the performance of QDS. Using the two extreme topologies of caterpillar and balanced trees on 16 taxa, the set of all displayed quartets was formed. Error was then introduced into the quartets in one of two ways. In the first scenario, for choices of probability $0 < p \leq .5$ of quartet error, each quartet was modified with probability p to one of the two resolved alternatives on the same taxa (with equal probability). In the second scenario, for choices of probability $0 < p \leq .9$, the quartet was removed from the set. For each of these modified quartet sets, QDS/NJ was used to construct a tree. In the second scenario, omitted quartets were simply left out of the counting that determines intertaxon distances. The Robinson-Foulds distance was then computed between the inferred QDS/NJ tree and the original tree. This was repeated 100 times, with results summarized in the plots of Figure 2. Similar results (not shown) were obtained using the FastME heuristic for balanced minimum evolution in place of NJ.

These results show that even with about a quarter of the quartets incorrect, on average the correct tree was recovered to within an RF distance of 2 (i.e, all but 1 of the 13 nontrivial

splits were recovered correctly) for the caterpillar tree. And even with about half of the caterpillar's quartets omitted, results were similarly accurate. The balanced tree topology was even more robustly recovered than the caterpillar tree, allowing quite large amounts of quartet error under both scenarios.

Of course one should interpret these results cautiously, as empirical quartet error may not have the simple form of the simulation. In empirical settings it is unlikely that all quartets would be equally likely to be incorrect or omitted, or that in the case of an incorrect quartet that both alternatives would be equally likely. Nonetheless, these simulations strongly indicate the need for more realistic simulation studies to investigate performance.

Remark 14.

A potential drawback of QDS for general tree inference from quartets is that its theoretical basis assumes one has a quartet in \mathcal{Q} for every subset of 4 taxa, and no weights can be supplied expressing relative confidence in those quartets. This differs from the maximum quartet consistency framework in which one seeks to maximize an objective function expressing the total weights of quartets displayed on the tree. While the above simulations suggest uniformly missing quartets may be of less concern, confidence weighting seems to be desirable, at least with quartets inferred by Maximum Likelihood, as discussed in [28]. However, in some applications, and especially for species tree inference from gene trees as described in the next section, these aspects of QDS may not be a great disadvantage.

Remark 15.

It is possible that new distance methods could be developed that are more finely tuned to QDS than those existing now. Since the distance \hat{d}_Q approximates distances on an unknown tree T endowed with the quartet metrization, a tree building or selection method that takes that specific metrization into account may improve performance. Current distance methods are general, making no assumption about a tree's edge lengths as related to its topology.

7 SPECIES TREE INFERENCE BY QUARTET DISTANCE CONSENSUS

We next show how QDS and RTDS can be applied to the problem of inferring a species tree from a collection of gene trees. This provides new consensus methods that are statistically consistent under the multispecies coalescent model, beyond those surveyed in [13]. For simplicity, we focus on the application of QDS.

For inference from multilocus sequence data, this can be used in a two-step procedure in which gene trees are first inferred from gene sequences, and then these inferred gene trees are treated as data for inference of a species tree. As is common for such two-stage schemes, the second stage of this method is provably statistically consistent, in the sense that if the gene trees were sampled without error under the multispecies coalescent model, then as the number of gene trees increases the probability of inferring the correct species tree approaches 1. In practice, however, there may be some inference error in the gene trees, as well as violations of the coalescent model, such as horizontal gene transfer.

For the algorithm, we assume we already have in hand a collection of binary trees on X , but allow some missing taxa on each tree. However, for good performance it is desirable that each 4-taxon subset appears on many of the gene trees.

Algorithm 2.

(QDC/M) Quartet Distance Consensus with distance method M

(Input: A collection of binary trees on subsets of taxa X)

1. For each subset of four taxa $x, y, z, w \in X$, determine the *dominant* (i.e., most frequent) quartet $xy|zw$, $xz|yw$, or $xw|yz$ displayed on the input trees. In the case of a tie, choose from the most frequent uniformly at random.
2. With \mathcal{Q} the set of dominant quartets, apply QDS with M.

Straightforward modifications lead to a formulation of Rooted Triple Distance Consensus (RTDC).

While QDC is similar to the clade-distance based STAR of [21], and split-distance based NJ_{st} (a.k.a. U-STAR/NJ) of [20], both of those average distances across gene trees, while QDC/M instead chooses the dominant quartets across gene trees to define a distance. Note that Rooted Triple Consensus of [15] similarly chooses the dominant rooted triple for rooted species tree inference, and inference of a population tree by the BUCKy software [19] proceeds through choosing dominant quartets, though neither utilizes a distance.

Next we establish statistical consistency of QDC for species tree inference under the multispecies coalescent model. This model has parameters specified by a rooted metric species tree σ^r , as described, for instance, in [1], and gives a probability distribution on binary metric gene trees. After marginalization over branch lengths and root location, one obtains a distribution on unrooted binary topological gene trees T . The structure of the model is such that one may view the generation of gene trees on subsets of taxa Y as either generating gene trees under the multispecies coalescent on the induced species tree σ_Y^r on that subset, or generating gene trees T on the full species tree σ^r and then passing to the induced gene trees T_Y on the subset of taxa. We take the second approach, as it is more convenient for our argument.

By a *taxon deletion model* for X we mean a random variable taking as its values subsets Y of X . Given any tree T on X , we apply the deletion model to T by passing to the induced tree T_Y on Y . In this formulation, the deletion model is independent of the tree it is applied to. We call a deletion model *quartet informative* if for each 4-taxon subset F of X the event FC_Y has positive probability.

The statistical consistency of QDC/M is then established by the following.

Theorem 3.

Let σ^r denote a rooted binary metric species tree, with positive branch lengths, and fix any quartet-informative taxon deletion model. Consider a sample S_n of n gene trees obtained by

first independently drawing gene trees on X from the multispecies coalescent model on σ^r and then applying the deletion model independently to each tree. Let M denote any tree building or selection algorithm that given pairwise distances fitting a (not necessarily ultrametric) tree will return that tree. Then with the unrooted topological species tree and $\hat{\sigma}_n$ the unrooted topological tree inferred by QDC/M from the sample S_n ,

$$\lim_{n \rightarrow \infty} \mathbb{P}(\hat{\sigma}_n = \sigma) = 1.$$

Proof. Consider a subset $\{a, b, c, d\}$ of 4 taxa in X . Let $S'_n \subseteq S_n$ be those trees in the sample on which the 4 appear. Because the taxon deletion model is quartet informative, with probability 1 we have $|S'_n| \rightarrow \infty$ as $n \rightarrow \infty$. Moreover, if $ab|cd$ is displayed on σ^r , the quartet on the four displayed on any tree in S'_n is a trinomial random variable [1] with parameters satisfying

$$p_{ab|cd} > p_{ac|bd} = p_{ad|bc}.$$

(While the precise values of these probabilities depend on branch lengths in σ^r , only the inequality and equality shown are needed for our argument.) Then with

$$c(n) = (c_{ab|cd}(n), c_{ac|bd}(n), c_{ad|bc}(n))$$

denoting the vector of counts of the quartets displayed in S'_n we have that as $n \rightarrow \infty$, $c(n)/|S'_n|$ converges in probability to $(p_{ab|cd}, p_{ac|bd}, p_{ad|bc})$. This implies

$$\mathbb{P}(c_{ab|cd} > c_{ac|bd}, c_{ad|bc}) \rightarrow 1.$$

That is, with probability approaching 1 the dominant quartet displayed on the gene trees matches that displayed on the species tree.

Since there are a finite number of subsets of 4 taxa, this implies that as $n \rightarrow \infty$ the probability approaches 1 that for *all* sets of 4 taxa the dominant gene tree quartet is displayed on the species tree σ . But if all the dominant quartets are those displayed on the species tree σ , then the algorithm computes the quartet distance on σ , so it returns $\hat{\sigma}_n = \sigma$. \square

The ability to deal with missing taxa is potentially an advantage of species tree inference by QDC over the U-STAR approach. Although simulations [34] have shown good performance of U-STAR with taxa missing from gene trees uniformly at random, it is unclear how relevant that pattern of missing-ness is to empirical data. The consistency of U-STAR under a uniform deletion model is investigated in [25], but the proof given there is flawed (with a correction in preparation [24]). However, in the case of non-uniform patterns of missing taxa statistical consistency seems unlikely. Indeed, it is relatively easy to construct small examples with non-uniform missing taxa where the U-STAR distance does not exactly fit

any tree, though a formal proof of inconsistency would have to establish that whatever distance method of tree construction is used cannot overcome this.

To analyze the running time of QDC, suppose QDC is applied to a list of n gene trees, all on a set of N taxa. As stated in [34], one can compute the matrix of pairwise split distances for a gene tree in time $\mathcal{O}(N^2)$. From this, for any 4 distinct taxa one can use the 4-point condition to determine which quartet is displayed, in constant time. Determining all displayed quartets on a single gene tree can thus be done in time $\mathcal{O}(N^4)$, and counting all quartets on all n gene trees in time $\mathcal{O}(N^4n)$. Once this is done, the dominant quartet for each set of 4 taxa separates 4 pairs of taxa, and so contributes to 4 pairwise quartet distances. Considering each quartet, then, we obtain the pairwise distance matrix in additional time $\mathcal{O}(N^4)$. Using, say, NJ for tree construction requires time $\mathcal{O}(N^3)$, so the total time complexity of QDC/NJ is $\mathcal{O}(N^4n)$.

This theoretical time complexity of QDC compares poorly with $\mathcal{O}(N^2n + N^3)$ for NJ_{st} stated in [34]. The comparison to the quartet-based ASTRAL-III, however, is more interesting: In [35], it is stated that ASTRAL-III has time complexity $\mathcal{O}((Nm)^{2.726})$ for input of binary gene trees. Thus ASTRAL-III may have better runtimes when $N \gg n$ (more taxa than gene trees), but QDC may be faster for $n \gg N$ (more gene trees than taxa). Since QDC has currently only been programmed in R, in a form unlikely to optimize runtime, a better implementation of QDC is needed for a fair practical speed comparison to ASTRAL-III.

An R implementation of QDC/NJ, using the ape package [26], on a data set of 1000 gene trees on 30 taxa took approximately 30 minutes to run on a desktop Macintosh with a 3.2GHz processor. (Of this time, over 28 minutes was spent simply tallying the displayed quartets on all the gene trees.) Although this compares poorly to approximately 22 seconds for ASTRAL-III and approximately 3 seconds for USTAR/NJ implemented in R, it still places it well within feasibility for data analysis. Indeed, the computational time to infer a large number of gene trees to serve as input will dwarf QDC's runtime. Moreover, recoding the algorithm is likely to give substantial speed improvement.

Simulations.

For a first look at the possible performance of QDC for species tree inference, it was applied to Avian simulated data sets of [7], which were also analyzed in [34]. These data sets are simulated on a fixed species tree of 48 taxa, drawn from a study of avian species. Samples of 1000 gene trees were simulated under the multispecies coalescent model on the species tree (scaling factor 1), and on rescalings of it by .5 (more incomplete lineage sorting) and 2 (less ILS). 20 replicate data sets were produced for each scaling factor. In addition to these samples of gene trees from the coalescent, sequences of length 500bp were simulated on each gene tree, and an estimated species tree inferred from them, which introduces inference error. More details on branch lengths, population sizes, and mutation rates can be found in the original publication.

For our simulation study, in order to reduce computational time, we reduced the number of taxa to 30, by deleting 18 taxa to obtain the species tree shown in Figure 3. By restricting sampled gene trees on 48 taxa to the 30 chosen ones, we obtain a valid sample of gene trees from the coalescent on the restricted species tree. However, by restricting the estimated gene

trees in the same way, we may not have obtained the same estimated gene tree topologies that would have been obtained from the 30 simulated sequences. Nonetheless, for a first look at performance, we expect the difference to be minor.

Data sets with missing taxa from some gene trees were also derived from these. Two subsets of the taxa, chosen as shown in Figure 3 were designated. For deletion probabilities $p = 0, 0.05, 0.1$, gene trees were chosen to have a group deleted with probability $2p$, with the particular group deleted equally likely. Thus the expected proportions of gene trees missing no taxa was $1 - 2p$, and missing each group was p , with no trees missing both. This pattern of missing taxa was chosen to roughly mimic what might occur in empirical data sets. In particular, if the source of missing taxa on some gene trees is the biological process of gene loss, it is likely to affect closely related taxa, and if it is due to uncollected data, it might be more likely in outgroups that were not the primary focus of data collection. While our specific model of missing data is quite crude, it is perhaps more relevant than uniform-at-random deletion of individual taxa, such as was used in the simulations of [34].

As with restricting to the 30 taxa from 48, further deletion of taxa from estimated gene trees may mean they do not agree topologically with estimated gene trees from the smaller set of sequences, but differences are likely to be minor.

On these simulated data sets, we compare the performance of 5 methods based on topological features of gene trees: QDC with both balanced FastME and NJ for tree construction, U-STAR with balanced FastME and NJ, and ASTRAL-III. QDC was implemented in R (code available on request), as was U-STAR, while the more complex ASTRAL-III software was used directly.

Results are shown in Figure 4 using a gene tree sample under the MSC model, and in Figure 5 using estimated gene trees. Note that the vertical scales on all plots differ between the two figures, as species tree inference is more reliable using the sampled gene trees.

In both figures one sees that QDC performs as well or better with NJ than with FastME in all conditions. When there are no missing taxa (top row of each figure), QDC/NJ, the U-STAR methods, and ASTRAL-III perform quite similarly. Given the extra computational time QDC and ASTRAL-III require, however, these simulations show that when no taxa are missing there is no reason to prefer QDC/NJ or ASTRAL-III to U-STAR.

When gene trees have missing taxa, however, the conclusion is quite different. At either level of missing taxa investigated (second and third rows of figures), the U-STAR methods are the poorest performing, as is in line with our earlier comments. Indeed, the plots suggest a lack of statistical consistency of U-STAR under these conditions. ASTRAL-III's and QDC's performance, however, do not show any pronounced changes across these missing taxon simulations, so they are clearly to be preferred to U-STAR. Presumably, the robustness of both QDC and ASTRAL-III to missing taxa arises from their common approach of basing inference on quartets, so that if some 4-taxon sets are not on all trees, one still gets a good estimate of their relationship from the remaining trees. Between ASTRAL-III and QDC/NJ there is little difference in performance, with no clear pattern as to which was more accurate.

While the simulations done here are by no means sufficient to draw final conclusions on the performance of QDC relative to other methods, they do indicate its potential.

Remark 16.

The algorithm presented and used in the simulations above implicitly assumes the species tree is binary, so that for every choice of four taxa there will be a single most probable quartet. For non-binary species trees, one may instead have all three quartets equiprobable (but not have a two-way tie for most probable). Using Theorem 2 one could modify the algorithm to allow for non-binary species trees, making some choice of cutoff for judging “near equality” in quartet frequencies.

Remark 17.

We further note that averaging the quartet distances, or rooted triple distances, across gene trees as is done in STAR and U-STAR would not lead to consistency under the coalescent model. In fact, an example is already given by [2] in which an inconsequential variant of the rooted triple metrization averaged across gene trees is shown to exactly fit an incorrect species tree for infinite sample size. A similar example for the quartet metrization is as follows: Consider the rooted caterpillar species tree $(((((a, b):x, c):y, d):z, e):w, f)$ with $x, z, w = \infty, y = 0$. Then under the multispecies coalescent model the gene trees $(((((a, b), c), d), e), f)$, $(((((a, b), d), c), e), f)$, and $((((a, b), (c, d)), e), f)$ each have probability $1/3$, and all others have probability 0. Unrooting these gene trees and applying the quartet metrization, with alphabetical ordering of taxa we obtain the three distance matrices

$$\begin{pmatrix} 0 & 8 & 14 & 18 & 20 & 20 \\ 0 & 14 & 18 & 20 & 20 & 20 \\ 0 & 16 & 18 & 18 & 20 & 20 \\ 0 & 14 & 14 & 18 & 20 & 20 \\ 0 & 8 & 18 & 18 & 20 & 20 \\ 0 & 0 & 18 & 18 & 20 & 20 \end{pmatrix}, \begin{pmatrix} 0 & 8 & 18 & 14 & 20 & 20 \\ 0 & 18 & 14 & 20 & 20 & 20 \\ 0 & 16 & 14 & 14 & 20 & 20 \\ 0 & 18 & 18 & 14 & 20 & 20 \\ 0 & 8 & 18 & 18 & 20 & 20 \\ 0 & 0 & 18 & 18 & 20 & 20 \end{pmatrix},$$

$$\begin{pmatrix} 0 & 8 & 18 & 18 & 18 & 18 \\ 0 & 18 & 18 & 18 & 18 & 18 \\ 0 & 8 & 18 & 18 & 18 & 18 \\ 0 & 18 & 18 & 18 & 18 & 18 \\ 0 & 8 & 18 & 18 & 18 & 18 \\ 0 & 0 & 18 & 18 & 18 & 18 \end{pmatrix}.$$

Weighting the matrices by $1/3$ and summing yields

$$\begin{pmatrix} 0 & 8 & 50/3 & 50/3 & 58/3 & 58/3 \\ 0 & 50/3 & 50/3 & 58/3 & 58/3 & 58/3 \\ 0 & 40/3 & 50/3 & 50/3 & 58/3 & 58/3 \\ 0 & 50/3 & 50/3 & 58/3 & 58/3 & 58/3 \\ 0 & 8 & 58/3 & 58/3 & 58/3 & 58/3 \\ 0 & 0 & 58/3 & 58/3 & 58/3 & 58/3 \end{pmatrix},$$

which exactly agrees with distances on the unrooted tree

$$((a:4, b:4):17/3, (c:20/3, d:20/3):1/3, (e:4, f:4):17/3).$$

This tree does not have the same unrooted topology as the species tree.

To construct an example with a binary species tree with no zero or infinite edge lengths, we perturb the above one slightly. If the distances on the original species tree are chosen so x , z , w are very large and y is very small but positive, the average of the gene tree quartet metrizations will change only slightly. Thus it cannot fit the topology of the species tree.

8 DISCUSSION

To place QDS in the context of other quartet supertree methods, note that the most common framework in current quartet methods of inferring trees — maximum quartet consistency — is to minimize an objective function measuring conflict between the given quartets and the tree. Alternative quartet-based tree construction approaches are given in [9], [22]. (Note that although our metrizations can be viewed as based in an instance of the “isolation weighting” of [9], in that work there is no concept of a true tree that one seeks to infer.)

While the optimization problem for maximum quartet consistency should be addressed by a search over all possible trees, in practice heuristic searches are usually necessary. The number of taxa or the search space may be limited in order to achieve acceptable performance and runtimes [4], [6], [23], [27], [31], [32], [33], [35]. As reasonable as this broad framework is, however, it is important to remember that the objective functions used are not ones deduced from theory. In fact, no such theory is even possible without an explicit model of error in the quartets, and it does not appear any attempt has been made to justify current approaches in such a way. Instead, simulations which incorporate inference error in the quartets are used for evaluation and comparison of methods.

A rather different notion of fitting a tree to quartets underlies QDS/M, whether the distance method M is a tree building algorithm or optimization of a distance-based objective function. By constructing a distance from the quartets, the selection of a “best” tree to fit the quartets is transferred to selecting one that fits the distance. Unfortunately no current theory can guide us as to whether this is better or worse than previous approaches. Extensive simulation studies are needed to judge the practical effectiveness of the new methods proposed here. Moreover, since the quartet error involved in different applications may have different features, simulation studies must reflect this and be targeted at specific applications. For instance, the effectiveness of QDC for species tree inference from full N -taxon gene trees inferred by Maximum Likelihood (ML) may be different from that of inference by QDS of a single gene tree from quartet trees inferred by ML.

Quartet methods have played a role in recent progress in phylogenetics in using algebraic methods for tree inference from sequence data, in work by [10], [11], [12] and [16], and the ideas presented here may be useful for those applications. Using these methods one can infer a quartet tree very quickly under very general models. However, technical issues complicate inference of larger trees directly. If QDS works well with the quartet trees these methods

produce, then the significant advantages they offer, in speed and the generality of the underlying substitution model, may be broadened to include quick inference of larger trees as well.

For the specific problem of inference of species trees from gene trees, rooted triple and quartet approaches have been taken before by [15], [19] and [35], with this last work presenting the highly-developed ASTRAL-III software. While the simulations of [34] suggested ASTRAL-III's accuracy is only comparable to U-STAR, the modified analysis presented here using the same simulated data suggests that its performance is significantly more robust to missing taxa on gene trees than is U-STAR. Indeed, this feature of quartet approaches should, we believe, be more appreciated as a justification for their development. Nonetheless, in the limited simulations presented here, QDC/NJ appears to have similar performance to ASTRAL-III whether or not taxa are missing. Moreover, complexity analysis suggests that when the number of genes far exceeds the number of taxa, an efficient implementation of QDC might achieve shorter runtimes than ASTRAL-III.

Inference of species trees under the multispecies coalescent model is made possible by our growing ability to assemble large data sets, comprised of many genetic loci, each with its own particular genealogical history. While Bayesian methods are conceptually attractive and have been implemented to address the simultaneous inference of gene trees and species trees, with current methodology there is little hope of them giving acceptable runtimes for data sets with many taxa and loci. The QDC method proposed here takes an alternative approach, through summarizing inferred gene trees by their displayed quartets. Compared to methods able to handle similar sized datasets, it may especially offer some gain in accuracy in the face of missing taxa. It is based in the new QDS method of tree inference from quartets, which itself is worthy of investigation as an alternative to methods based on the standard optimization formulation of maximum quartet consistency. While further testing of performance of these algorithms is needed, both in simulation and on empirical datasets, they offer hope for improving phylogenetic inference, and thus for helping address the many biological questions for which that is a key ingredient.

Acknowledgements

This work was 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. Thanks to Elizabeth Allman and Mike Steel, and anonymous reviewers, for comments and discussion on an earlier draft.

Biography



John A. Rhodes received his Ph.D. degree in Mathematics from the Massachusetts Institute of Technology in 1986. Since 2003, his research has focus on mathematical phylogenetics, especially using algebraic perspectives to analyze statistical models. In 2005, he moved to

the University of Alaska Fairbanks, where he has been department chair of Mathematics and Statistics, and is affiliated with the Institute of Arctic Biology.

REFERENCES

- [1]. Allman ES, Degnan JH, and Rhodes JA Identifying the rooted species tree from the distribution of unrooted gene trees under the coalescent. *J. Math. Biol.* 62(6):833–862, 2011. [PubMed: 20652704]
- [2]. Allman ES, Degnan JH, and Rhodes JA Species tree inference by the STAR method, and generalizations. *J. Comput. Biol.* 20(1):50–61, 2013. [PubMed: 23294273]
- [3]. Allman ES, Degnan JH, and Rhodes JA Species tree inference from gene splits by Unrooted STAR methods. *IEEE/ACM Trans. Comput. Biol. Bioinf.* 15:337–342, 2018.
- [4]. Anvi E, Cohen R, and Snir S Weighted quartets phylogenetics. *Syst. Biol.* 64(2):233–242, 2015. [PubMed: 25414175]
- [5]. Bandelt H-J and Dress A Reconstructing the shape of a tree from observed dissimilarity data. *Adv. Appl. Math.* 7:309–343, 1986.
- [6]. Baum BR Combining trees as a way of combining data sets for phylogenetic inference, and the desirability of combining gene trees. *Taxon*, 41(3–10), 1992.
- [7]. Bayzid MS, Mirarab S, and Warnow T Weighted statistical binning: enabling statistically consistent genome-scale phylogenetic analyses. *PLOS One*, 2014.
- [8]. Bryant D A classification of consensus methods for phylogenetics. *DIMACS series in discrete mathematics and theoretical computer science*, 61:163–184, 2003.
- [9]. Bryant D and Berry V A structured family of clustering and tree construction methods. *Adv. Appl. Math.* 27:705–732, 2001.
- [10]. Casanellas M and Fernández-Sánchez J Performance of a new in-variants method on homogeneous and non-homogeneous quartet trees. *Mol. Biol. Evol.* 24(1):288–293, 2007. [PubMed: 17053050]
- [11]. Chifman J and Kubatko L Quartet inference from SNP data under the coalescent model. *Bioinformatics*, 30(23):3317–3324, 2014. [PubMed: 25104814]
- [12]. Chifman J and Kubatko L Identifiability of the unrooted species tree topology under the coalescent model with time-reversible substitution processes, site-specific rate variation, and invariable sites. *J. Theor. Biol.* 374:35–47, 2015. [PubMed: 25791286]
- [13]. Degnan JH, DeGiorgio M, Bryant D, and Rosenberg NA Properties of consensus methods for inferring species trees from gene trees. *Syst. Biol.* 58(1):35–54, 2009. [PubMed: 20525567]
- [14]. Desper R and Gascuel O Fast and accurate phylogeny reconstruction algorithms based on the minimum-evolution principle. *J. Comput. Biol.* 9(5):687–705, 2002. [PubMed: 12487758]
- [15]. Ewing GB, Ebersberger I, Schmidt HA, and von Haeseler A Rooted triple consensus and anomalous gene trees. *BMC Evol. Biol.* 8:118, 2008. [PubMed: 18439266]
- [16]. Fernández-Sánchez J and Casanellas M Invariant versus classical quartet inference when evolution is heterogeneous across sites and lineages. *Syst. Biol.* 65(2):280–291, 2016. [PubMed: 26559009]
- [17]. Gascuel O BIONJ: an improved version of the NJ algorithm based on a simple model of sequence data. *Mol. Biol. Evol.* 14:685–695, 1997. [PubMed: 9254330]
- [18]. Lapointe JF and Cucumel G The average consensus procedure: combination of weighted trees containing identical or overlapping sets of taxa. *Syst. Biol.* 46(2):306–312, 1997.
- [19]. Larget BR, Kotha SK, Dewey CN, and Ané C BUCKy: Gene tree / species tree reconciliation with Bayesian concordance analysis. *Bioinformatics*, 26:2910–11, 2010. [PubMed: 20861028]
- [20]. Liu L and Yu L Estimating species trees from unrooted gene trees. *Syst. Biol.* 60:661–667, 2011. [PubMed: 21447481]
- [21]. Liu L, Yu L, Pearl DK, and Edwards SV Estimating species phylogenies using coalescence times among sequences. *Syst. Biol.* 58:468–477, 2009. [PubMed: 20525601]
- [22]. Ma B, Xin L, and Zhang K A new quartet approach for reconstructing phylogenetic trees: quartet joining method. *J. Comb. Optim.* 16(3):293–306, 2008.

- [23]. Mirarab S and Warnow T ASTRAL-II: coalescent-based species tree estimation with many hundreds of taxa and thousands of genes. *Bioinformatics*, 31:i44–i52, 2015. [PubMed: 26072508]
- [24]. Nute M personal communication. 2019.
- [25]. Nute M, Chou J, Molloy EK, and Warnow T The performance of coalescent-based species tree estimation methods under models of missing data. *BMC Genomics*, 19(Suppl 5):286, 2018. [PubMed: 29745854]
- [26]. Paradis E, Claude J, and Strimmer K APE: analyses of phylogenetics and evolution in R language. *Bioinformatics*, 20:289–290, 2004. [PubMed: 14734327]
- [27]. Ragan MA Phylogenetic inference based on matrix representation of trees. *Mol. Phylo. Evol.*, 1(53–58), 1992.
- [28]. Ranwez V and Gascuel O Quartet-based phylogenetic inference: Improvements and limits. *Mol. Biol. Evol.*, 18(6):1103–1116, 2001. [PubMed: 11371598]
- [29]. Ranwez Vincent, Criscuolo Alexis, and Emmanuel JP Douzery. Supertriplets: a triplet-based supertree approach to phylogenomics. *Bioinformatics*, 26(12):i115–i123, 2010. [PubMed: 20529895]
- [30]. Saitou N and Nei M The neighbor-joining method: A new method for reconstruction of phylogenetic trees. *Mol. Biol. Evol.*, 4:406–425, 1987. [PubMed: 3447015]
- [31]. Snir S and Rao S Quartets MaxCut: a divide and conquer quartets algorithm. *IEEE/ACM Trans. Comput. Biol. Bioinf.*, 7(4):704–718, 2010.
- [32]. Strimmer K and von Haeseler A Quartet puzzling: A quartet maximum-likelihood method for reconstructing tree topologies. *Mol. Biol. Evol.*, 13(7):964–969, 1996.
- [33]. Swenson MS, Suri R, Linder CR, and Warnow T An experimental study of quartets MaxCut and other supertree methods. *Alg. Mol. Biol.*, 6(1):7, 2011.
- [34]. Vachaspati P and Warnow T ASTRID: Accurate Species TRees from Internode Distances. *BMC Genomics*, 16(Suppl 10):S3, 2015.
- [35]. Zhang C, Rabiee M, Sayyari E, and Mirarab S ASTRAL-III: polynomial time species tree reconstruction from partially resolved gene trees. *BMC Bioinformatics*, 19 (Suppl 6)(153):15–30, 2018. [PubMed: 29343218]




Fig. 1. 16-taxon caterpillar and balanced trees, with edge lengths given by the quartet metrization

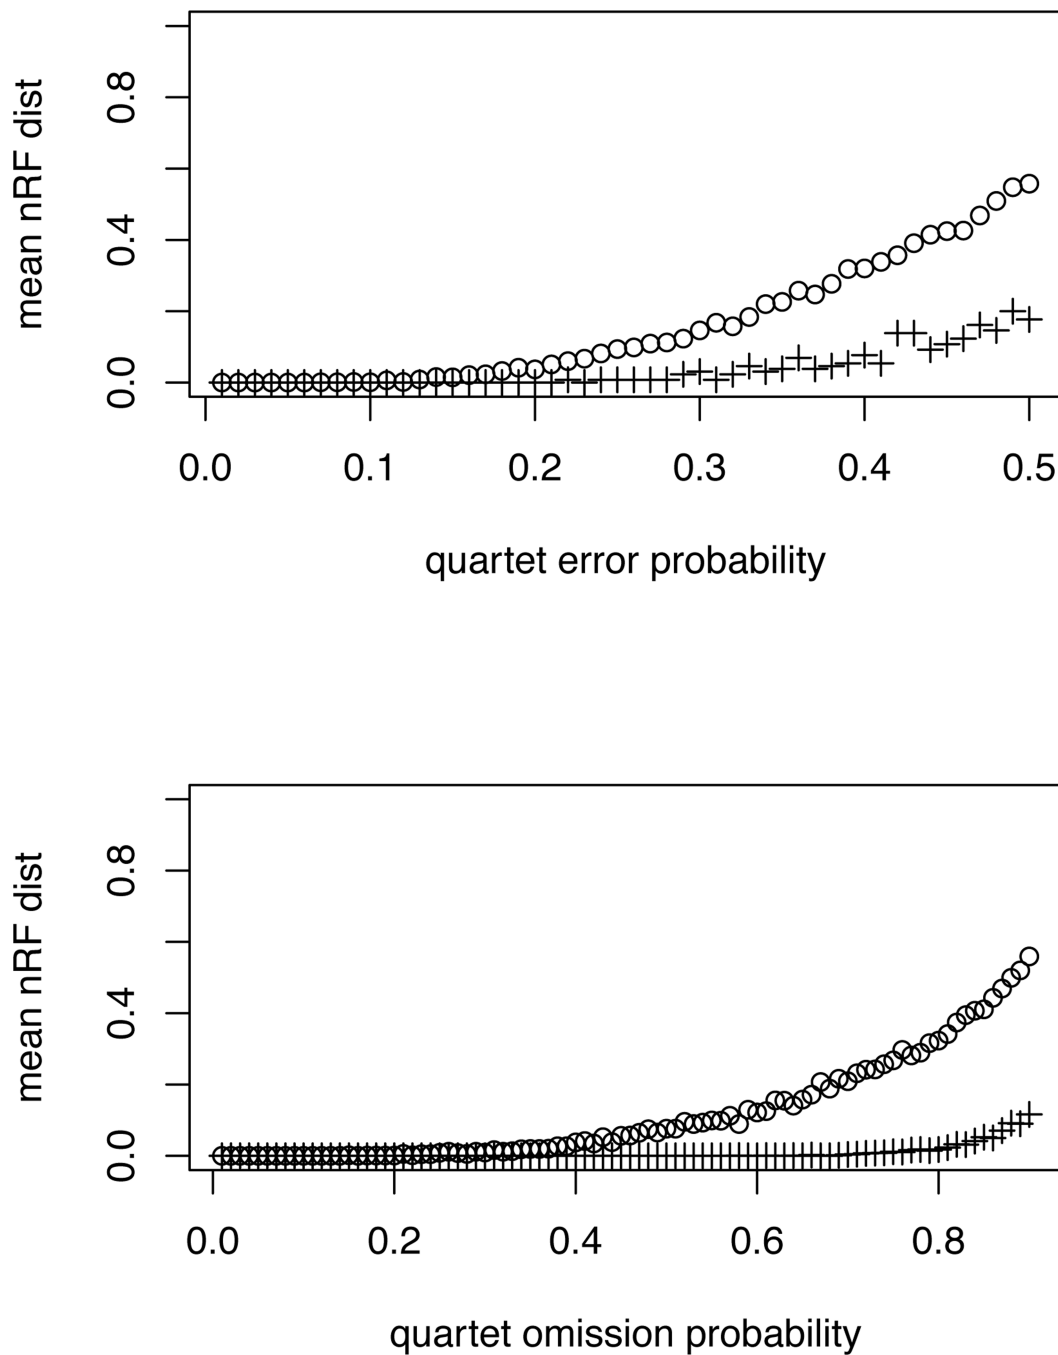


Fig. 2.

Performance of QDS/NJ in simulation under two scenarios, erroneous quartets (top) and omitted quartets (bottom), as described in the text. Circles (O) denote the 16-taxon caterpillar tree and pluses (+) the 16-taxon balanced tree. The horizontal axes on the top plot gives the probability that a true quartet is replaced with an alternative on the same taxa in forming the quartet set. On the bottom plot, that axis gives the probability a quartet is omitted from the quartet set. Inference error is measured by the normalized Robinson-Foulds (nRF) distance between the correct and inferred tree. The vertical axes show the mean nRF

distance over 100 replicates. Note that two resolved 16 taxon trees differing by one NNI have nRF distance $2/2(16 - 3) \approx 0.077$.

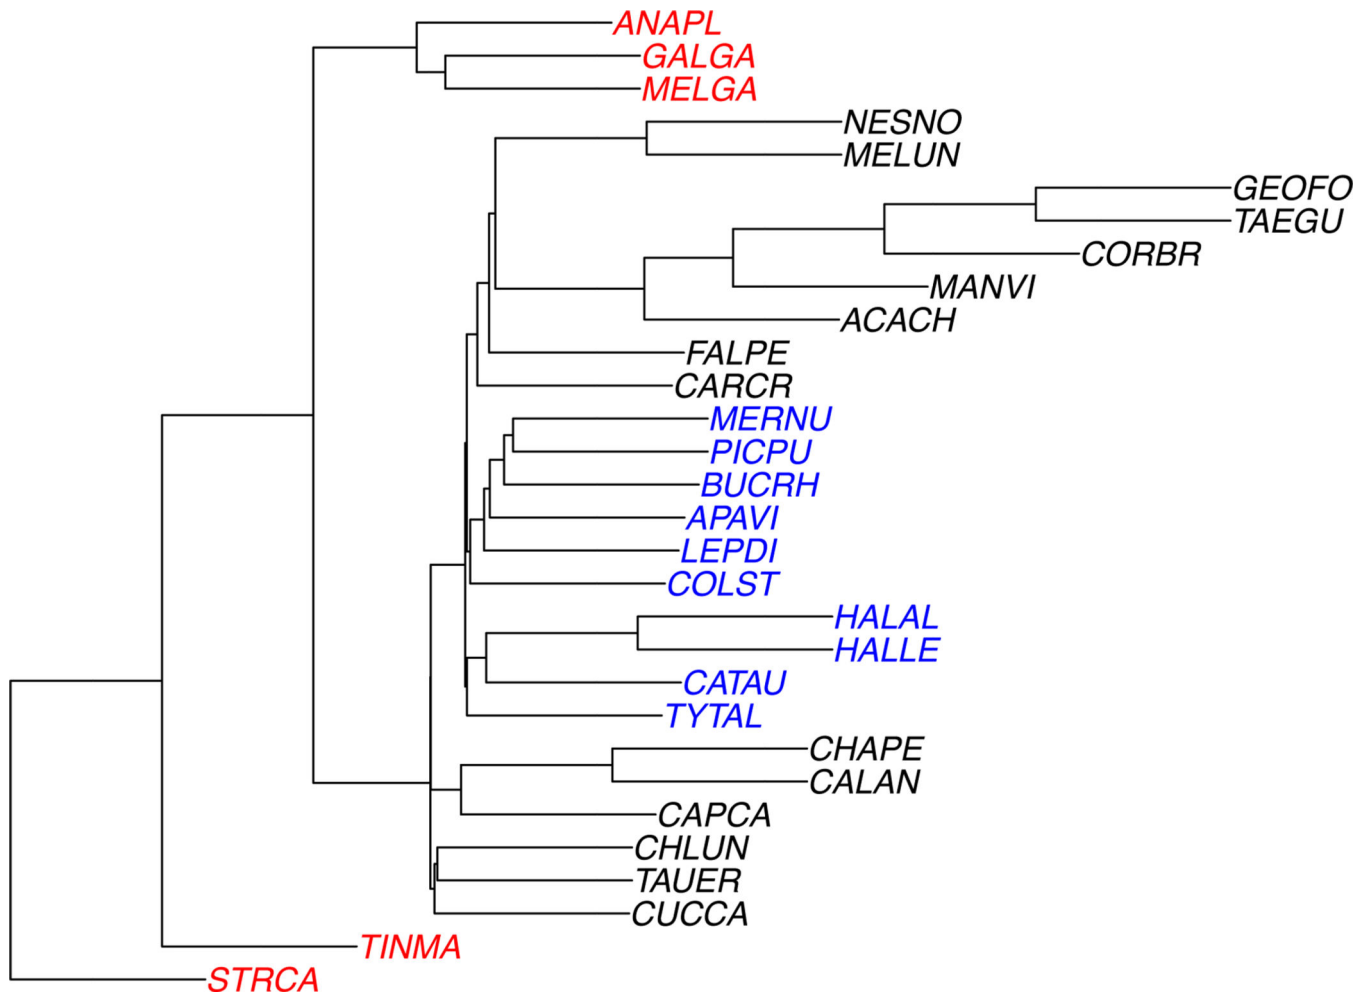


Fig. 3.

The 30-taxon species tree used for simulations. For trials involving missing taxa, a deletion probability p was chosen, and a group of taxa deleted with probability $2p$ from each gene tree. The deleted group was the taxa shown in red (anapl, galga, melga, tinma, strca) or blue (mervu, picpu, bucrh, apavi, lepd, colst, halal, halle, catau, tytal), with equal probability. Thus the the expected proportion of gene trees on all taxa is $1 - 2p$, on the the black and blue taxa is p , and on the black and red is p .

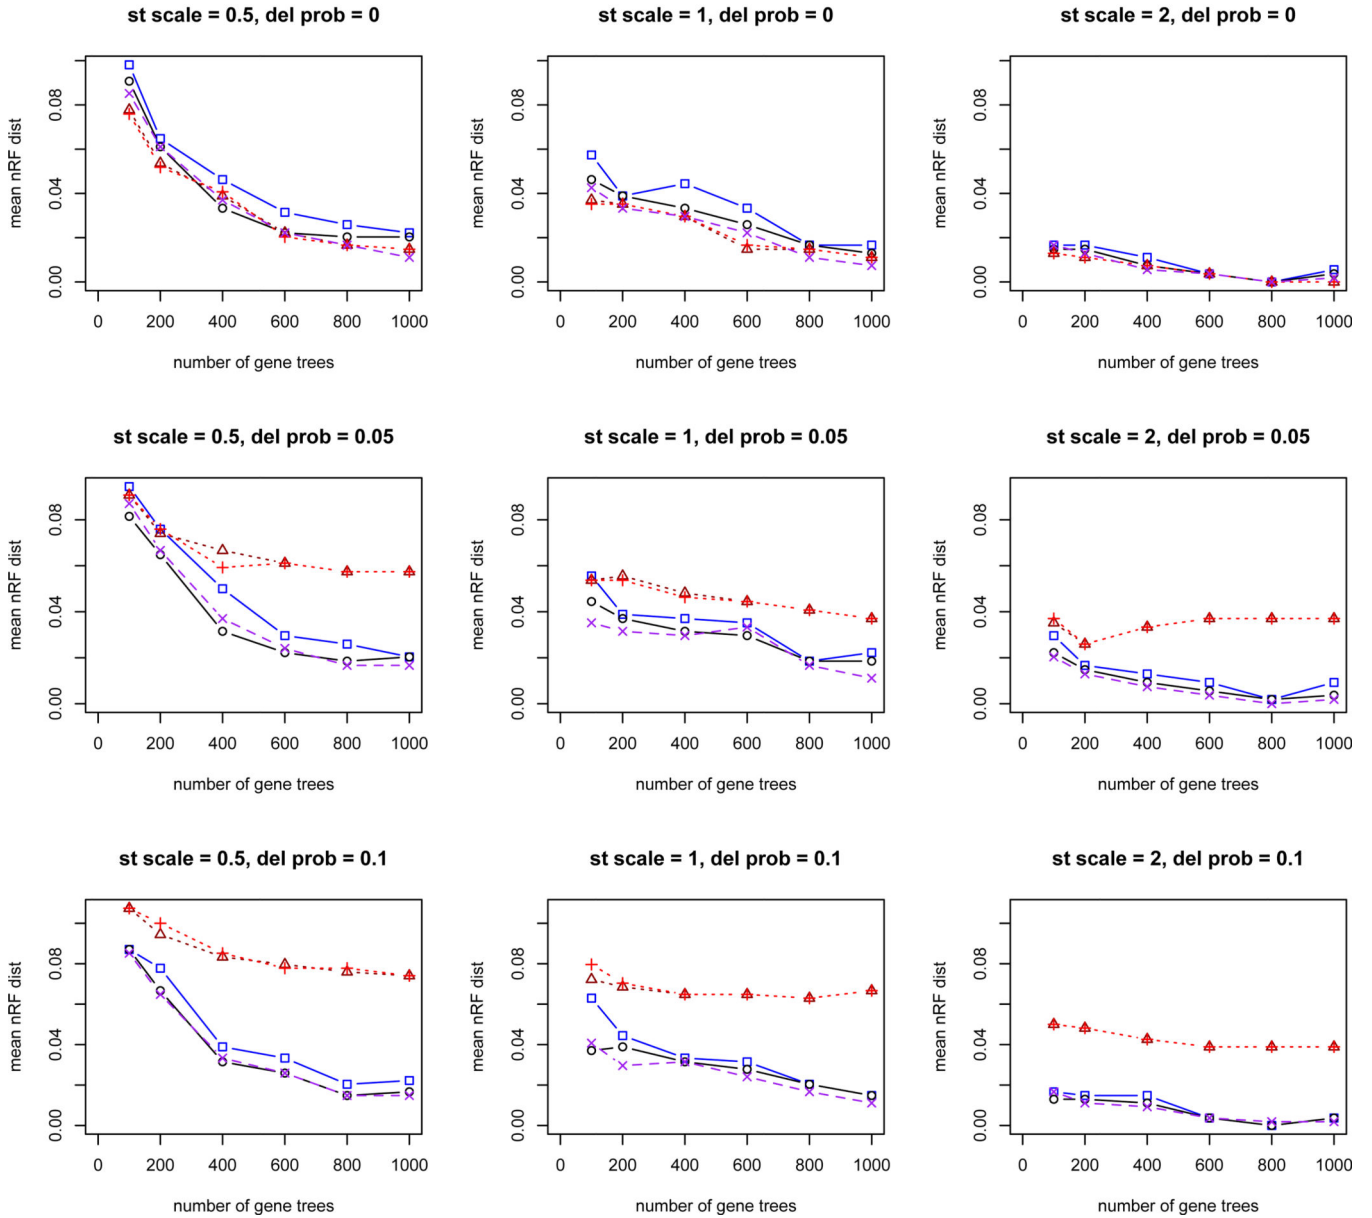



Fig. 4. Simulation results based on 20 replicates for each simulation condition, using gene trees sampled from the multispecies coalescent model (lacking estimation error). From left to right, “st scale” is the species tree scaling factor of .5,1, or 2, indicating decreasing amounts of ILS. From top to bottom, “del prob” controls the probability of missing taxa on gene trees, with values 0, .05, or .1 indicating increasing numbers of gene trees with missing taxa. On individual plots, increasing numbers of gene trees, 100, 200, 400, 600, 800, and 1000, were analyzed for species tree inference. The mean over the replicates of the normalized

Robinson-Foulds (nRF) distance from the species tree is used to measure accuracy. For 30 taxa, 2 trees differing by a single NNI move have $nRF = 2/2(30 - 3) \approx 0.037$.

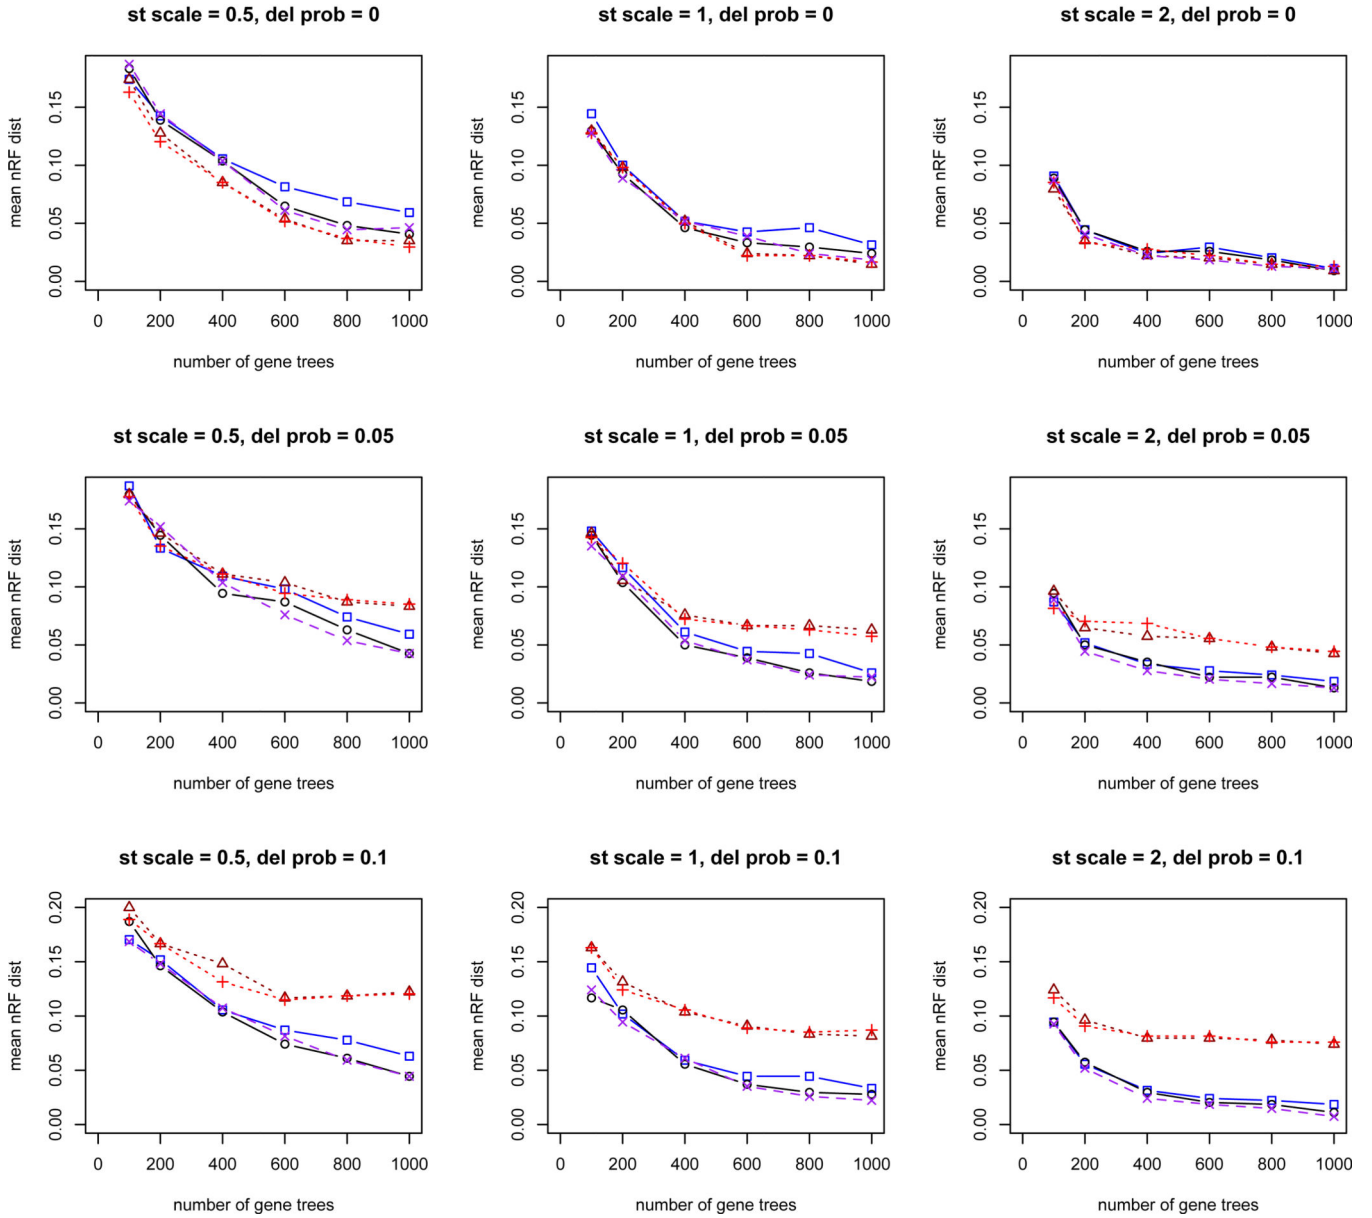



Fig. 5. Simulation results based on 20 replicates for each simulation condition, using gene trees estimated from sequences simulated on gene trees sampled from the multispecies coalescent model. From left to right, “st scale” is the species tree scaling factor of .5, 1, or 2, indicating decreasing amounts of ILS. From top to bottom, “del prob” controls the probability of missing taxa on gene trees, with values 0, .05, or .1 indicating increasing numbers of gene trees with missing taxa. On individual plots, increasing numbers of gene trees, 100, 200, 400, 600, 800, and 1000, were analyzed for species tree inference. The mean over the replicates of the normalized Robinson-Foulds (nRF) distance from the species tree is used to

measure accuracy. For 30 taxa, 2 trees differing by a single NNI move have $nRF = 2/2(30 - 3) \approx 0.037$.