Polyhedral Geometry of Phylogenetic Rogue Taxa

María Angélica Cueto · Frederick A. Matsen

Received: 28 January 2010 / Accepted: 28 May 2010 / Published online: 17 July 2010
© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract It is well known among phylogeneticists that adding an extra taxon (e.g. species) to a data set can alter the structure of the optimal phylogenetic tree in surprising ways. However, little is known about this “rogue taxon” effect. In this paper we characterize the behavior of balanced minimum evolution (BME) phylogenetics on data sets of this type using tools from polyhedral geometry. First we show that for any distance matrix there exist distances to a “rogue taxon” such that the BME-optimal tree for the data set with the new taxon does not contain any nontrivial splits (bipartitions) of the optimal tree for the original data. Second, we prove a theorem which restricts the topology of BME-optimal trees for data sets of this type, thus showing that a rogue taxon cannot have an arbitrary effect on the optimal tree. Third, we computationally construct polyhedral cones that give complete answers for BME rogue taxon behavior when our original data fits a tree on four, five, and six taxa. We use these cones to derive sufficient conditions for rogue taxon behavior for four taxa, and to understand the frequency of the rogue taxon effect via simulation.

Keywords Minimum evolution · Distance-based phylogenetic inference · Linear programming · Polytope · Normal fan

The first author was supported by a UC Berkeley Chancellor’s Fellowship. The second author was supported by the Miller Institute for Basic Research at UC Berkeley.

M.A. Cueto

Department of Mathematics, University of California, Berkeley, CA 94720, USA

e-mail: macueto@math.berkeley.edu

F.A. Matsen (✉)

Program in Computational Biology, Fred Hutchinson Cancer Research Center, 1100 Fairview Ave. N. M1-B514, P.O. Box 19024, Seattle, WA 98109-1024, USA

e-mail: matsen@fhcrc.org

1 Introduction

Ideally, phylogenetic data sets would have the property that the optimal tree for a subset X of taxa Y would be the same as the tree obtained by restricting the optimal tree on Y to the set X . However, practicing phylogeneticists are well aware that this is not the case; the extensive literature on “taxon sampling” reviewed below is evidence to the contrary. One can also find references to “rogue taxa” which, although not clearly defined or rigorously investigated, are taxa who do not fit into a tree and whose inclusion may disrupt the inference of evolutionary relationships of the other taxa. For example, Sullivan and Swofford (1997) state “... the hedgehog therefore appears to represent a ‘rogue’ taxon that cannot be placed reliably with these data and that possibly confounds attempts to estimate the relationships among the remaining taxa.” The “rogue” descriptor is also used by Baurain et al. (2007) to describe taxa with a “strong nonphylogenetic signal”; these authors describe the importance of finding and eliminating these taxa from phylogenetic studies.

Surprisingly, we were unable to find any mathematical or simulation-based analysis of the action of rogue taxa in phylogenetic trees. The closest studied subject is “taxon sampling.” This area of research is focused on the following question: If we are interested in the phylogenetic tree on a set of taxa Y , do we do better or worse by adding more taxa into the tree? If better, is the improvement more significant than would be gained by increasing the length of the sequences (by redirecting resources)?

The origins of the taxon sampling debate can be traced to the pioneering paper of Felsenstein (1978) that demonstrated mathematically the existence of “long-branch attraction,” where two pendant branches are artifactually placed close together by parsimony algorithms. This led to the question of whether parsimony long-branch problems could be dispensed with by adding new taxa to the data set to break up the long branches; Hendy and Penny (1989) have answered affirmatively under certain conditions. The investigation was continued by Kim (1996), who showed that the situation is subtle and that the new taxa must appear in specific regions of the tree in order to counter the long-branch attraction problem.

These mathematical investigations of parsimony were followed by a flood of simulation-based papers investigating maximum likelihood, parsimony, and distance methods for phylogenetics. Hillis (1996), Graybeal (1998), and Poe (1998) indicated that a larger number of taxa improved estimation, whereas the high-profile publication of Rosenberg and Kumar (2001) claimed the opposite. The Hillis group responded (Zwickl and Hillis 2002; Pollock et al. 2002; Hillis et al. 2003), which led Rosenberg and Kumar (2003) to somewhat moderating their position. The debate on taxon sampling has continued to the present day, with additional simulations (Poe 2003; DeBry 2005; Hedtke et al. 2006), review articles (Heath et al. 2008a), and studies to understand the impact of taxon sampling on the inference of macroevolutionary processes (Heath et al. 2008). The simulation literature in this area is considered important enough to even have a paper (Rannala et al. 1998) about methodology for taxon-sampling simulations.

There are two inherent difficulties with simulations of this type. First, the collection of possible parameter values for simulation is vast, and any simulation study must make choices about which parameters to use. This first problem alone may

be the source of the disagreement found in the taxon selection literature. Second, the simulations are done by simulating data with a single model on a tree, then reconstructing. This does not address the problem of what happens when considering unusual data sets, such as those obtained by major model misspecifications.

A mathematical approach can address these difficulties, although with certain caveats. Theorems can indicate that a phenomenon will always happen given certain criteria, and the construction of the complete spaces of examples or counterexamples gives very precise information about these questions. By exploring the complete space of data sets of a certain type, one is not limited to data sets which are within a certain class of models. The trade-off for the strength of these conclusions is that often the setting must be simplified to make the problem mathematically tractable.

In order to address taxon selection and the rogue taxon effect problem mathematically, we have chosen to use distance-based phylogenetics, specifically the Balanced Minimum Evolution (BME, described below) criterion. Because the optimality criterion is expressed in terms of the minimization of an inner product, we are able to harness the power of polyhedral geometry to answer the questions of interest with a high degree of precision. Although BME-based algorithms are not among the most popular in phylogenetics, implementations do exist which show good performance under simulation (Desper and Gascuel 2002b). The BME criterion is consistent (Desper and Gascuel 2004), as is FastBME which minimizes BME through tree rearrangements (Bordewich et al. 2009). Another motivation for studying BME is the close relationship between BME and the very popular Neighbor-Joining (NJ) algorithm (Saitou and Nei 1987). Specifically, NJ has been shown to be a heuristic BME minimizer (Desper and Gascuel 2005); the relationship between the two algorithms has been investigated by Eickmeyer et al. (2008).

After describing a bit of terminology, we will discuss the main results of the paper. Note that by *dissimilarity map* we simply mean a mapping D from unordered pairs of taxa to nonnegative numbers such that $D(x, x) = 0$ for all x . These are sometimes called “distance matrices” in the phylogenetics literature but we use dissimilarity map to emphasize that they need not satisfy the triangle inequality.

Definition 1.1 Let t be a phylogenetic tree equipped with branch lengths $\mathbf{b}$. The tree metric associated with t and $\mathbf{b}$ is the dissimilarity map obtained as follows: the distance between taxa i and j of t is given by the total length (i.e., sum of branch lengths) of the path from i to j in t with respect to $\mathbf{b}$.

Next we define some core objects of study for this paper.

Definition 1.2 Let D be a dissimilarity map on n taxa. A “lifting” $\tilde{D}$ of D is a dissimilarity map on $n + 1$ taxa obtained from D by adding distances from the first n taxa to an $(n + 1)$ st taxon.

Definition 1.3 Let D be a dissimilarity map on n taxa, and let $\tilde{D}$ be a lifting of D . The BME tree for D will be called the “lower tree,” while the BME tree for $\tilde{D}$ will be called the “upper tree.” The “restricted upper tree” will be the tree induced on the original n taxa by restricting the upper tree to this set.

Our primary goal is to understand topological differences between the upper and lower trees for various original dissimilarity maps D and various liftings $\tilde{D}$.

1.1 Overview of the Paper

The first section describes the effect of adding a new taxon when the original dissimilarity map D is arbitrary. Theorem 3.2 shows that for any D , there exists a lifting such that the intersection of the split sets for the restricted upper tree and the lower tree consists of the trivial pendant splits. In other words, we show that the restricted upper tree and the lower tree can be maximally distant in terms of the Robinson–Foulds metric (Robinson and Foulds 1981). However, the upper tree cannot deviate from the lower in an arbitrary way: Theorem 3.5 shows that certain combinations of lower and upper trees are not possible. We also note that the trees of Theorem 3.2 need not be maximally distant in terms of the quartet distance (Remark 3.4).

The second section addresses the case where the original dissimilarity map D is a tree metric for some tree t ; in this setting there is no question of what the optimal tree for the lower taxa “should” be. That is, if the upper tree does not contain the lower tree, the additional taxon is definitely a disrupting “rogue” taxon. When D is a tree metric, there exists a simplified formulation of the BME computations. This “reduced” formulation has a linear rather than a quadratic number of variables, and allows polyhedral computation directly over the parameters of interest. We study the associated “reduced polytope” and several of its combinatorial and geometric properties, including its dimension. Using this “reduced” formulation we are able to give sufficient conditions (Propositions 4.16 and 4.17) for the rogue taxon effect when the lower tree has four taxa, as well as a perspective on the frequency of the rogue effect through simulations for up to six lower taxa.

The computations in this paper were done with a combination of Gfan (Jensen 2009), Polymake (Gawrilow and Joswig 2000), and custom ocaml (Chailloux et al. 2000) code using GSL (Galassi et al. 2009), the GNU scientific library. For the interested reader, source code is available at

<http://github.com/matsen/roguebme>.

2 Polyhedral Geometry and BME Phylogenetics

In this section, we introduce the mathematical problem we wish to investigate and walk through the necessary background in polyhedral geometry. We start by defining the Balanced Minimum Evolution (BME) criterion for phylogenetic inference.

For the purposes of this paper, all trees will be unrooted phylogenetic trees. We will use parenthetical “Newick format” to describe trees, such that $((a, b), (c, d), e)$ indicates a five-taxon tree with the pairs a, b and c, d being sister taxa (Felsenstein 2004). Sometimes we will write these unrooted trees in a rooted manner, as we feel that $((a, b), (c, d))$ is clearer than $(a, b, (c, d))$. The degree-two vertex of the rooted representation should be suppressed. *Trivalent* trees are trees such that all internal nodes have degree three.

Definition 2.1 Given a dissimilarity map D in $\mathbb{R}^{\binom{n}{2}}$, the “Balanced Minimum Evolution” (BME) length of a phylogenetic tree t with respect to a dissimilarity map D is the quantity

$$\lambda(t, D) := \sum_{1 \leq i < j \leq n} \omega_{ij}^t D_{ij}, \quad (1)$$

where $\omega_{ij}^t = \prod_{v \in p_{ij}^t} (\deg(v) - 1)^{-1}$, and p_{ij}^t denotes the internal vertices in t on the path between leaves i and j .

Remark 2.2 In the case of a trivalent tree t , the weight ω_{ij}^t equals $2^{-|p_{ij}^t|}$.

A BME tree for an $n \times n$ nonnegative matrix D will be a tree t minimizing $\lambda(t, D)$ over all n -taxon trees. The BME algorithm is *consistent* on trivalent trees: if D is tree metric with trivalent tree topology t , then the BME tree of D is t (Desper and Gascuel 2004).

Note that there is a volume-zero set of dissimilarity maps with multiple optimal BME trees, and therefore it is not quite right to speak of “the” BME tree. All of our statements are true by replacing “the BME tree” with “a BME tree”; however, we prefer stating the former. More precisely, given a dissimilarity map, we have two cases: either the set of possible BME trees of D consists of a single (trivalent) tree, or the set has size at least two, and it is closed under degenerations. That is, if a trivalent tree t contracts to a BME tree for D , then t is also a BME tree for D ; this claim will be clear from the polyhedral perspective described below.

There are several equivalent formulations of the BME length (Eickmeyer et al. 2008), although we prefer (1) because of its polyhedral interpretation.

Global BME minimization is known to be hard (Guillemot and Pardi 2009). The widely used Neighbor-Joining algorithm approaches the BME problem from a greedy perspective (Studier and Keppler 1988). The Fastme algorithm starts with a heuristically obtained tree and then refines it using Nearest-Neighbor Interchange (NNI) to attempt to find the BME minimal tree (Desper and Gascuel 2004). A better understanding of the BME polytope (defined below) could lead to better such algorithms (Desper and Gascuel 2002a), analogous to how understanding the traveling salesman polytope provides insight into the traveling salesman problem (Padberg and Grötschel 1985).

We now introduce the BME polytope, first investigated by Eickmeyer et al. (2008). A *polytope* in $\mathbb{R}^m$ is the convex hull of a finite number of points in $\mathbb{R}^m$. Fix a positive integer n . The *BME polytope* in $\mathbb{R}^{\binom{n}{2}}$ is the convex hull of the points $(\omega_{ij}^t)_{i,j}$, where t varies among all possible tree topologies on n taxa.

Using this polyhedral interpretation, the problem of finding the BME-optimal tree t on n taxa corresponds to picking a vertex ω^t of the BME polytope minimizing the Euclidean dot product of the vertex with a given dissimilarity map (considered as a vector in $\mathbb{R}^{\binom{n}{2}}$). The BME tree is the tree associated with this vertex.

We can characterize this optimization process by constructing the corresponding *inner normal fan*. The inner normal fan of a polytope $\mathcal{P} \subset \mathbb{R}^N$ is given as a finite collection of cones (i.e., a set closed under multiplication by positive scalars) as follows.

Each cone in the inner normal fan of $\mathcal{P}$ corresponds to a face $\mathcal{F}$ of the polytope $\mathcal{P}$ and is defined as

$$\mathcal{C}_{\mathcal{F}} := \{w \in \mathbb{R}^N : \langle w, v \rangle = \min\{\langle w, u \rangle : u \in \mathcal{P}\}, \forall v \in \mathcal{F}\}, \quad (2)$$

i.e., those vectors such that the minimum inner product is achieved at all points of the face $\mathcal{F}$.

By construction, each cone is polyhedral: it is the solution set of a system of linear inequalities. As such, it can be expressed as the positive span (i.e., using nonnegative scalars) of finitely many vectors, which we call *extremal rays*. In addition, the inner normal fan of $\mathcal{P}$ is a *polyhedral fan* because the family $\{\mathcal{C}_{\mathcal{F}} : \mathcal{F} \subset \mathcal{P} \text{ face}\}$ is closed under intersections. Moreover, this fan is *complete* (i.e., the union of all cones equals the ambient space $\mathbb{R}^N$), and each cone $\mathcal{C}_{\mathcal{F}}$ has dimension equal to $\text{codim } \mathcal{F} = N - \dim \mathcal{F}$, where $\dim \mathcal{F}$ denotes the dimension of the affine span of face $\mathcal{F}$. In particular, if $\mathcal{F}$ is a vertex, then $\mathcal{C}_{\mathcal{F}}$ is full-dimensional. We call these full-dimensional cones *chambers*. The inner normal fan of the BME polytope will be referred to as the *BME fan*. We refer the reader to Ewald (1996, Chap. 1) for a complete exposition of normal fans.

Remark 2.3 From the previous discussion we see that the BME criterion is equivalent to the membership of a dissimilarity map D to a *chamber* in the *BME fan*. Thus D belongs to the interior of a chamber in the BME fan *if and only if* the BME tree of D is unique. The boundary of these chambers is the volume-zero set having multiple BME trees (discussed earlier in this section).

Since the BME polytope encodes the problem of finding the BME tree of a dissimilarity map, it is worth understanding its structure. Some of its combinatorial properties have been studied for small number of taxa, although several questions remain open for $n \geq 6$. We investigate some of its features below, as described by Eickmeyer et al. (2008).

The vertices of the BME polytope correspond to the points $(\omega_{ij}^t)_{i,j}$, where t is a *trivalent* tree for a total of $(2n - 5)!!$ vertices (Pachter and Sturmfels 2005, Lemma 2.33). Here, $(2n - 5)!! = (2n - 5) \cdot (2n - 3) \cdots 3 \cdot 1$. In addition, the vector ω_{ij}^s associated with the star tree s (the tree with a single internal node) lies in the interior of the polytope, whereas all other points ω^t lie on its boundary (Eickmeyer et al. 2008, Lemma 2.1).

The *dimension* of the BME polytope (i.e., the dimension of the affine space spanned by this polytope) is $\binom{n}{2} - n$. The polytope is not full-dimensional because, after translation to the origin, the orthogonal complement of its affine span is spanned by the n *shift vectors* $\{h_{\mathbf{a}} : \mathbf{a} \in \{1, \dots, n\}\}$. Here, the shift vector $h_{\mathbf{a}}$ refers to a dissimilarity map in which leaf $\mathbf{a}$ is at distance 1 from all other leaves, while all other pairwise distances are 0.

The *f-vector* $\mathbf{f}(\mathcal{P}) \subset \mathbb{R}^N$ of an N -dimensional polytope $\mathcal{P}$ gives the number of faces of each dimension of $\mathcal{P}$. That is, $\mathbf{f}(\mathcal{P})_i = \#\{\text{faces of dimension } i - 1 \text{ of } \mathcal{P}\}$. The *f-vectors* of BME polytopes have been studied for up to seven taxa. In particular, for four and five taxa, these vectors have been completely described in Eickmeyer et al. (2008, Table 1), whereas for six and seven taxa, some of the entries of the *f-vector*

have remained unknown up to now. We were able to compute the complete f -vector for six taxa by methods of tropical geometry, using `Gfan`. The resulting f -vector is

$$(105, 5460, 105945, 635265, 1715455, 2373345, 1742445, 640140, 90262).$$

In particular, we see that the polytope has 90262 facets. It also has 105 vertices, labeled by all trivalent trees on six taxa.

As a corollary of these computations, it follows that the edge graph of the BME polytope for six taxa is the complete graph K_{105} (Eickmeyer et al. 2008). This says that any two vertices of the BME polytope can be connected by an edge. Similar behavior occurs for four and five taxa, but this is no longer true for seven or more taxa (Eickmeyer et al. 2008).

By construction, the BME polytope comes equipped with a natural symmetry given by the symmetric group $\mathbb{S}_n$ on n elements. Namely, relabeling the leaves of a trivalent tree t by a permutation $\sigma \in \mathbb{S}_n$ sends t to the relabeled trivalent tree σt and hence the vertex ω^t to $\omega^{\sigma t}$. In a similar way, higher-dimensional faces of the BME polytope will have this symmetry. Therefore, we can encode these symmetries in the f -vector, and record the number of faces of each dimension, up to the combinatorial action of $\mathbb{S}_n$ on all faces. In the case of six taxa, we get

$$(2, 20, 182, 982, 2492, 3489, 2626, 1032, 169).$$

We illustrate these constructions and their properties in the case of four taxa.

Example 2.4 (Eickmeyer et al. 2008) Fix $n = 4$. The points ω^t are:

$$\begin{aligned} \omega^{((1,2),(3,4))} &= \frac{1}{4}[2, 1, 1, 1, 1, 2]; & \omega^{((1,3),(2,4))} &= \frac{1}{4}[1, 2, 1, 1, 2, 1]; \\ \omega^{((1,4),(2,3))} &= \frac{1}{4}[1, 1, 2, 2, 1, 1]; & \omega^{\text{star}(4)} &= \frac{1}{3}[1, 1, 1, 1, 1, 1]. \end{aligned}$$

The BME polytope is a triangle in $\mathbb{R}^6$ with vertices $\omega^{((1,2),(3,4))}$, $\omega^{((1,3),(2,4))}$, and $\omega^{((1,4),(2,3))}$. It spans the two-dimensional space $\{(x_{12}, x_{13}, x_{14}, x_{23}, x_{24}, x_{34}) \in \mathbb{R}^6 : x_{12} + x_{13} + x_{14} = x_{12} + x_{23} + x_{24} = x_{13} + x_{23} + x_{34} = x_{14} + x_{24} + x_{34} = 1\}$.

The *lineality space* of a fan is defined as the maximal linear space contained in all cones of the fan. If this space is just the origin, we say that the fan is *pointed*. In the case of the BME fan, this linear subspace is n -dimensional with basis given by the n shift vectors h_a corresponding to the n leaves. Since the lineality space lies in all cones of the fan, we can mod out by this subspace (for example, by taking a projection to its orthogonal complement) and reduce our study to the case of pointed complete polyhedral fans in $\mathbb{R}^{\binom{n}{2}-n}$. We illustrate the construction of the BME fan and the associated pointed fan on four taxa.

Example 2.5 Let $n = 4$. We mod out by the lineality space $L = (h_1, h_2, h_3, h_4)$ via the *canonical* projection map $p: \mathbb{R}^{\binom{n}{2}} \rightarrow L^\perp \simeq \mathbb{R}^{\binom{n}{2}-n}$ to the orthogonal comple-

Fig. 1 Quartets minimizing the BME criterion for each dissimilarity map on four taxa

ment of the subspace L given by the matrix

$$\begin{pmatrix} 0 & 1 & -1 & -1 & 1 & 0 \\ 1 & 0 & -1 & -1 & 0 & 1 \end{pmatrix}.$$

We apply this projection to the BME fan, and we get a fan in $\mathbb{R}^2$, which we can plot. Alternatively, we project the BME polytope into 2-space, and we take the inner normal fan of the resulting polytope.

From Example 2.4 we know that the BME polytope is the triangle with vertices corresponding to the three quartet trees $((1, 2), (3, 4))$, $((1, 3), (2, 4))$, and $((1, 4), (2, 3))$. The projection p maps this triangle to the triangle with vertices $(-2, 4)$, $(4, 0)$ and $(-2, -2)$. Its inner normal fan consists of the rays spanned by $r_1 = (1, 0)$, $r_2 = (-1, -1)$ and $r_3 = (0, 1)$, plus the origin. Figure 1 shows the quartets corresponding to the relative interior of each chamber.

3 Behavior of BME Under the Addition of an Extra Taxon

The purpose of this section is to investigate the relationship between lower and upper trees for arbitrary D . Section 3.1 shows that for any D , there exists a lifting such that the upper tree is as different as possible from the lower tree in terms of splits. Section 3.2 provides a counterpoint by demonstrating that certain combinations of lower and upper trees are not possible, i.e., that a rogue taxon cannot affect a BME tree in arbitrary ways.

Notation 3.1 Throughout the remainder of the paper, we label our taxa by $[n] = \{1, \dots, n\}$. We write $\mathbb{R}_+$ for the set of nonnegative reals.

3.1 A Theorem Demonstrating the Existence of Unusual Upper Trees

We show that every lower tree has an upper tree whose restriction to the lower taxa is maximally different from it in terms of the *Robinson–Foulds metric* δ_{RF} on tree

topologies, although perhaps not in terms of *quartet distance*. The δ_{RF} metric on phylogenetic trees is defined in terms of bipartitions in the tree, also called “splits.” A split in a phylogenetic tree is simply the bipartition of the taxa induced by cutting that edge. For example, the split $\{1, 2\}, \{3, 4\}$ is induced by cutting the internal edge of the quartet $((1, 2), (3, 4))$. Let $\Sigma(t)$ denote the set of splits of tree t ; the distance $\delta_{\text{RF}}(s, t)$ is simply one half the size of the symmetric difference of $\Sigma(t)$ and $\Sigma(s)$ (Robinson and Foulds 1981). The quartet distance is analogous to the Robinson–Foulds distance but with the role of splits replaced by that of quartets (induced subtrees of size four) contained in a tree. The naive algorithm for computation is $O(n^4)$, although it can be computed in $O(n^2)$ via a simple algorithm (Bryant et al. 2000) and in $O(n \log n)$ via a more complex algorithm (Brodal et al. 2004). In this paper, s will have one more taxon than t ; we accommodate this difference for the Robinson–Foulds and quartet distances by simply taking the induced tree on s given by the set of lower taxa.

Theorem 3.2 *Let D be a dissimilarity map on n taxa with BME tree t . There exists a lifting $\tilde{D}$ whose upper tree s maximizes $\delta_{\text{RF}}(s, t)$ among all trees on n taxa.*

This theorem will follow easily from the following lemma.

Lemma 3.3 *Given an ordering of n taxa $z_1, \dots, z_n$ and any distance matrix D on taxa $\{z_i : 1 \leq i \leq n\}$, there exists a lifting $\tilde{D}$ such that the BME tree for $\tilde{D}$ restricted to $z_1, \dots, z_n$ is the caterpillar tree $(z_1, (z_2, \dots, (z_{n-1}, z_n) \dots))$.*

Proof Pick arbitrary numbers $1 < \alpha_1 < \dots < \alpha_n$. Let y denote the extra “rogue” taxon. We construct a family of liftings $\tilde{D}^c$ as an exponential function for a given base number c . Set $\tilde{D}^c(y, z_i) = c^{\alpha_i}$.

We write the BME length as

$$\lambda(s, \tilde{D}^c) = \sum_{1 \leq i < j \leq n} \omega_{i,j}^s D_{i,j} + \sum_{1 \leq i \leq n} \omega_{i,n+1}^s c^{\alpha_i}.$$

As c goes to infinity, the dominant term in the summation becomes $\omega_{n,n+1}^s c^{\alpha_n}$. For c greater than some c_n , the BME tree must be a caterpillar tree with y as far as possible from z_n . Indeed, any other topology would have a smaller coefficient for c^{α_n} . We can repeat the same argument replacing $n - 1$ for n , finding a c_{n-1} such that for $c \geq c_{n-1}$ the BME tree must be a caterpillar tree with y as far as possible from the subtree (z_{n-1}, z_n) . Continue in this way until a large enough lower bound on c is found such that the described caterpillar tree is the BME tree for $\tilde{D}^c$. $\square$

With this lemma, all that is needed to prove Theorem 3.2 is to show that there exists a caterpillar tree s such that the restriction of the caterpillar to the original taxa has maximal $\delta_{\text{RF}}(s, t)$.

Proof of Theorem 3.2 Color the taxa of t with black and white colors as follows: for every cherry (two-taxon subtree) of t , color one taxon white and the other black, and color the remaining taxa arbitrarily. Now order the taxa with all of the black taxa

first and all of the white taxa second. The caterpillar tree from Lemma 3.3 using this ordering will have the required maximal δ_{RF} . $\square$

Remark 3.4 The extension of Theorem 3.2 to quartet distances does not hold for more than seven taxa. Indeed, let t be $(1, (((((2, 3), 4), 5), 6), 7), 8))$. The maximally quartet-distant trees on 8 taxa (of quartet distance 61) are the following non-caterpillars:

$$\begin{aligned} &(1, (2, (((3, 8), 5), (4, 7)), 6))) \\ &(1, (2, (((3, 8), 6), (4, 7)), 5))) \\ &(1, (((((2, 8), 5), (4, 7)), 6), 3))) \\ &(1, (((((2, 8), 6), (4, 7)), 5), 3))). \end{aligned}$$

These trees were found by our code and distances were confirmed with the `qdist` program of Mailund and Pedersen (2004).

One could perform a similar analysis for the path distance metric of Steel and Penny (1993), although we have not done so.

3.2 A Theorem Restricting Topology of Upper Trees

The previous section shows that the lower and upper trees can be quite different. It is natural then to ask about the collection of possible upper trees for a given lower tree. That is, if we have a dissimilarity map D on n taxa with BME tree t , what are the possible BME trees s for liftings of D ? This question narrows the potential effect of rogue taxa.

We first gain intuition by investigating the case of four taxa. This setting is simple, as there is only one trivalent tree topology on five taxa (up to relabeling of its leaves).

Using `Polymake`, we can show that all but two tree topologies can be realized as upper trees for a lower quartet. The two trees not above $((1, 2), (3, 4))$ are shown in Fig. 2.

This example can be established analytically and generalized to the case of more taxa by replacing the leaves 1 through 4 with rooted subtrees a through d . In particular, we show that we can never obtain a tree where pairs of subtrees are exchanged “over” the extra taxon.

Let y denote the new leaf to be attached. The original tree t is the tree $((a, b), (c, d))$. Call s the tree $((a, c), (b, d))$ as in Fig. 3.

Theorem 3.5 *Let D be a dissimilarity map such the BME score of $t = ((a, b), (c, d))$ is strictly less than that of $s = ((a, c), (b, d))$ (Fig. 3). Then the BME score of $t_y :=$*

Fig. 2 The trees that do not sit above $((1, 2), (3, 4))$ for any lifting of a dissimilarity map D with BME tree $((1, 2), (3, 4))$

Fig. 3 The trees t , s , t_y , and s_y

$((a, b), y, (c, d))$ is strictly less than that of $s_y := ((a, c), y, (b, d))$ for any lifting $\tilde{D}$ of D . Consequently, if t is the BME tree for D , then s_y cannot be a BME tree for any lifting $\tilde{D}$.

Proof We denote with sans-serif font the elements in each subtree, so a denotes a leaf in subtree a , etc. For simplicity, we abbreviate ω^t by ω . By definition, we get

$$\begin{aligned} \omega_{ab}^{s_y} &= \omega_{ab}/4; & \omega_{ac}^{s_y} &= 2\omega_{ac}; & \omega_{ad}^{s_y} &= \omega_{ad}/2; & \omega_{bc}^{s_y} &= \omega_{bc}/2; \\ \omega_{bd}^{s_y} &= 2\omega_{bd}; & \omega_{cd}^{s_y} &= \omega_{cd}/4; & \omega_{ab}^{t_y} &= \omega_{ab}; & \omega_{ac}^{t_y} &= \omega_{ac}/2; \\ \omega_{ad}^{t_y} &= \omega_{ad}/2; & \omega_{bc}^{t_y} &= \omega_{bc}/2; & \omega_{bd}^{t_y} &= \omega_{bd}/2; & \omega_{cd}^{t_y} &= \omega_{cd}. \end{aligned}$$

Similarly,

$$\begin{aligned} \omega_{ab}^s &= \omega_{ab}/2; & \omega_{ac}^s &= 2\omega_{ac}; & \omega_{ad}^s &= \omega_{ad}; & \omega_{bc}^s &= \omega_{bc}; \\ \omega_{bd}^s &= 2\omega_{bd}; & \omega_{cd}^s &= \omega_{cd}/2. \end{aligned}$$

Since we are interested in the difference between the two scores, we do not compute the weights w.r.t. leaf y nor weights within a cluster, since both trees have the same weight in these two cases. Then for any given lifting $\tilde{D}$, we have by subtraction

$$\lambda(s_y, \tilde{D}) - \lambda(t_y, \tilde{D}) = 3/2(\lambda(s, D) - \lambda(t, D)).$$

The term on the right-hand side is positive by hypothesis. $\square$

4 Liftings of Tree Metrics

In the previous section, we analyzed the relationship between the lower and upper trees for liftings of a general dissimilarity map D . For a practicing phylogeneticist, however, this provides limited useful information. Indeed, the basic assumption of phylogenetic inference is that the data evolve in a primarily tree-like manner. Namely, in distance-based inference, the assumption is that the given dissimilarity map is “close” to a tree metric. In the rogue setting, we are interested in n taxa which evolve in a tree-like manner and one, the rogue, that does not.

In this section we formalize these notions by assuming that D is a tree metric with respect to the tree topology t . By the consistency of BME inference, the lower tree will be t . With this assumption, our primary interest will be in understanding how the upper tree can differ from t in the sorts of situations more likely to be encountered in phylogenetics. Although Theorem 3.2 provides an interesting theoretical result in this vein, the required lifting is quite unlikely to appear in data. By reformulating the problem below directly in terms of the branch lengths of the tree metric, we are able to obtain more precise and relevant information about the action of rogue taxa.

4.1 Preliminaries

Notation 4.1 Given a positive integer n , we define $\mathcal{D}_n$ to be the cone of dissimilarity maps on n taxa. We identify $\mathcal{D}_n$ with $\mathbb{R}_+^{\binom{n}{2}}$. Similarly, we define $\mathcal{T}_n \subset \mathcal{D}_n$ to be the space of tree metrics on n taxa. We omit the subscript n whenever it is clear from the context. Finally, given a tree topology t , we denote by $\mathcal{T}_t \subset \mathcal{T}_n$ the set of tree metrics with underlying tree topology t .

Notation 4.2 Given a trivalent tree t , the BME cone $\mathcal{C}_{\omega^t}$ associated with t will be denoted by $\mathcal{C}_t$. Moreover, we call $\mathcal{C}_t^+ = \mathcal{C}_t \cap \mathbb{R}_+^{\binom{n}{2}}$ the *positive BME cone* of t , also known as the BME cone of dissimilarity maps associated with t .

Notation 4.3 In what follows, we write $\mathcal{P}_n$ for the BME polytope on n taxa. If the number of taxa is understood, we omit the subscript.

Given a tree topology t on n taxa, let $\pi_t: \mathbb{R}_+^{\binom{n}{2}} \rightarrow \mathbb{R}^{2n-3}$ denote a map generalizing the branch-length map for tree metrics as follows. The coordinates of this map are indexed by the branches of the tree t , and each coordinate is a linear function on the metric cone whose value on tree metrics with topology t is precisely the length of the corresponding edge. Note that this linear function is not unique, and it is positive on tree metrics with topology t . An expression defining the coordinate e of the map π_t (that is, the branch length of e) can be obtained by the four-point condition equations (Pachter and Sturmfels 2005, Theorem 2.36) characterizing the tree topology t . For example, let $t = ((1, 2), (3, 4))$, let e_i be the edge adjacent to leaf i , let e be the internal edge, and let b_{e_i}, b_e be their corresponding lengths. Then $\pi_t(D) := (b_{e_1}(D), b_{e_2}(D), b_{e_3}(D), b_{e_4}(D), b_e(D))$, where $b_{e_1}(D) = (D_{31} - D_{32} + D_{12})/2$, $b_{e_2}(D) = (D_{32} - D_{31} + D_{12})/2$, $b_{e_3}(D) = (D_{23} - D_{24} + D_{34})/2$, $b_{e_4}(D) = (D_{24} - D_{23} + D_{34})/2$, and $b_e(D) = (D_{13} + D_{24} - D_{12} - D_{34})/2$. The map π_t has the property that it identifies the cone of tree metrics realizing t with $\mathbb{R}_+^{2n-3}$.

Our goal for this subsection is to understand the interplay between the branch lengths of a tree metric $D \in \mathcal{T}_t$ and the possible upper trees one can obtain by lifting this metric. In particular, we wish to characterize the branch lengths of lower trees admitting a prescribed upper tree s . It is clear that if we start from a tree metric $D = d_t$ and its corresponding branch-length vector $\pi_t(D)$, we can easily lift D to a tree metric $\tilde{D}$ whose underlying tree s contains t as a subtree. Hence, the union of the

sets $\{\pi_t(D) : D \text{ s.t. } \exists \tilde{D} \in \mathcal{C}_s^+\}$ as s varies among possible upper BME trees equals the set $\mathbb{R}_+^{2n-3}$. We want to understand each one of these sets. In particular, we want to answer the following challenge:

Problem 4.4 *Given a tree topology t on n taxa and $s \in \mathcal{T}_{n+1}$, describe the cone of dissimilarity maps on $n + 1$ taxa whose BME tree equals s and whose restriction to the first n taxa is a tree metric of combinatorial type t .*

For each upper tree s , the elements of the corresponding set in Problem 4.4 can be thought of as vectors in $\mathbb{R}_+^{3n-3}$, where the first $2n - 3$ entries encode the branch lengths of the lower tree t , and the remaining ones refer to distances to the new taxon. That is,

$$X_s(t) := \{(\pi_t(D), \tilde{D}_{1,n+1}, \dots, \tilde{D}_{n,n+1}) : D \in \mathcal{T}_t, \tilde{D} \in \mathcal{C}_s^+\}. \quad (3)$$

By construction, these sets are polyhedral cones, and they partition the set $\mathbb{R}_+^{3n-3}$:

Proposition 4.5 *$X_s(t)$ is a rational (possibly empty) polyhedral cone for every s and t . It is described by two types of homogeneous linear constraints:*

- All entries $\tilde{D}_{ij} \geq 0$, and $\pi_t(D) \geq 0$.
- Inequalities describing $\mathcal{C}_s$: they correspond to the directions $\omega^s - \omega^u$ for all trivalent trees u on $n + 1$ taxa, and all constants are zero. That is: $\langle \omega^s - \omega^u, \tilde{D} \rangle \geq 0$ for all trivalent trees u .

Proof $X_s(t)$ is a polyhedral cone because it is the image of the linear map $\tilde{D} \mapsto (\pi_t(\tilde{D}|_{[n]}), \tilde{D}_{1,n+1}, \dots, \tilde{D}_{n,n+1})$, where $\tilde{D} \in \mathcal{C}_s^+ \cap (\mathcal{T}_t \times \mathbb{R}_+^n)$. The inequalities describing $X_s(t)$ follow by construction. The entries of $\tilde{D}|_{[n]}$ are expressed as linear combinations of the entries $\pi_t(\tilde{D}|_{[n]})$. The second group of inequalities include facet inequalities of the cone $\mathcal{C}_s$: whose directions are given by the edges containing vertex ω^s . To simplify the construction, we add the inequalities coming from differences between ω^s and all other vertices of $\mathcal{P}$ and not only of vertices ω^u adjacent to ω^s . Adding these inequalities makes no harm, and it simplifies the problem by avoiding the computation of the edges adjacent to ω^s , which can be hard if the number of taxa is too big. $\square$

4.2 The Reduced BME Polytope

We now present an equivalent approach to our lifting task in the setting of this section, i.e., when D is a tree metric on n taxa with (trivalent) tree t and branch lengths $\mathbf{b}_e$. As shown below, all that is needed to study the restricted BME problem is a change of order of summation followed by a grouping of appropriate terms. This small modification reduces the problem from having a quadratic number of free variables to a linear number, as well as simplifying the constraints. After introducing the reduced polytope, we show that it has dimension $2n - 4$ by characterizing its affine hull.

The set of edges of t will be denoted by $E(t)$. Pick any lifting $\tilde{D}$ of D , and any tree s with $n + 1$ leaves. The BME length of s with respect to $\tilde{D}$ can be calculated as follows:

$$\lambda(s, \tilde{D}) = \langle \omega^s, \tilde{D} \rangle = \sum_{i, j \neq n+1} \omega_{ij}^s D_{i,j} + \sum_{i=1}^n \omega_{i,n+1}^s \tilde{D}_{i,n+1}.$$

Now we simply substitute the definition of the dissimilarity map D :

$$D_{i,j} = \sum_{e \in t(i \leftrightarrow j)} \mathbf{b}_e,$$

where $e \in t(i \leftrightarrow j)$ indicates that edge $e \in E(t)$ lies in the path between leaves i and j in tree t . Exchanging order of summation and regrouping, we have

$$\langle \omega^s, \tilde{D} \rangle = \sum_{e \in E(t)} \left(\sum_{\substack{i, j \neq n+1 \\ e \in t(i \leftrightarrow j)}} \omega_{ij}^s \right) \mathbf{b}_e + \sum_{i=1}^n \omega_{i,n+1}^s \tilde{D}_{i,n+1}, \quad (4)$$

which is again a simple inner product with a rational vector. For a tree s on $n + 1$ taxa, define $(v^s) \in \mathbb{R}^{3n-3}$ by

$$\begin{cases} (v^s)_e = \sum_{\substack{i, j \neq n+1 \\ e \in t(i \leftrightarrow j)}} \omega_{ij}^s, & e \text{ edge of lower tree,} \\ (v^s)_i = \omega_{i,n+1}^s, & 1 \leq i \leq n. \end{cases} \quad (5)$$

Note that this definition depends on the fixed tree t , but we do not incorporate it to the notation, as we will typically be fixing a lower tree.

To find the BME tree for a tree metric $(t, \{\mathbf{b}_e\}_{e \in E(t)})$, we build a vector $v^s \in \mathbb{R}^{3n-3}$ for each tree $s \in \mathcal{T}_{n+1}$. Each vector has entries indexed by the $2n - 3$ edges of t and the n distances $\{\tilde{D}_{i,n+1} : i = 1, \dots, n\}$. Our goal is to find s minimizing the quantity (4). As in the case of the BME problem, we build a polytope $\mathcal{B}^t$ (here in $(3n - 3)$ -space) which is the convex hull of the points v^s and study its properties.

Definition 4.6 Fix a tree t on n taxa and consider the points $(v^s)_{e,i}$ as in (5). The convex hull of these points is called the “reduced BME polytope,” and we denote it by $\mathcal{B}^t$. It only depends on the combinatorial type of the tree t , and it is symmetric with respect to the group of symmetries of the tree t . The points $\{v^s : s \in \mathcal{T}_{n+1}\}$ are called “reduced weights.” The *inner* normal fan of $\mathcal{B}^t$ is called the “reduced fan.” Cones in this fan are called “reduced cones” and their intersections with the positive orthant are to be called “positive reduced cones.”

From the previous construction it is clear that the BME polytope and the reduced BME polytope are closely related. We now explain this connection. The linear map $\alpha_t : \mathbb{R}^{\binom{n+1}{2}} \rightarrow \mathbb{R}^{3n-3}$ assigning the reduced weight v^s to the BME weight ω^s sends the polytope $\mathcal{P}$ surjectively onto the polytope $\mathcal{B}^t$. That is, the reduced polytope is a linear *projection* of the BME polytope. On the dual side, the dual of the linear map α_t will inject the dual space of the polytope $\mathcal{B}^t$ into the dual space of the polytope $\mathcal{P}$,

Table 1 Reduced weights for trivalent trees on five taxa, starting from the lower tree $t = ((1, 2), (3, 4))$, up to symmetry of the lower tree t . The column labels show the quantity for which the entry is the corresponding coefficient in the reduced weight vector: e.g., the first entry of the table shows that $7/8$ is the coefficient of $\mathbf{b}_1$ for topology $((1, 2), (3, (4, 5)))$

Upper tree	$\mathbf{b}_1$	$\mathbf{b}_2$	$\mathbf{b}_3$	$\mathbf{b}_4$	$\mathbf{b}_0$	x_1	x_2	x_3	x_4
$((1, 2), (3, (4, 5)))$	$7/8$	$7/8$	$6/8$	$4/8$	$6/8$	$1/8$	$1/8$	$2/8$	$4/8$
$((1, 2), (5, (3, 4)))$	$6/8$	$6/8$	$6/8$	$6/8$	$4/8$	$2/8$	$2/8$	$2/8$	$2/8$
$((1, 3), (2, (4, 5)))$	$7/8$	$6/8$	$7/8$	$4/8$	$9/8$	$1/8$	$2/8$	$1/8$	$4/8$
$((1, 3), (5, (2, 4)))$	$6/8$	$6/8$	$6/8$	$6/8$	$10/8$	$2/8$	$2/8$	$2/8$	$2/8$
$((1, 3), (4, (2, 5)))$	$7/8$	$4/8$	$7/8$	$6/8$	$9/8$	$1/8$	$4/8$	$1/8$	$2/8$
$((1, 5), (2, (3, 4)))$	$4/8$	$6/8$	$7/8$	$7/8$	$6/8$	$4/8$	$2/8$	$1/8$	$1/8$

and in this case the linear spaces of both polytopes are identified by the map α_t (Proposition 4.9). We refer the interested reader to Ziegler (2006, Sect. 7.2) for more information about projections of polytopes.

Example 4.7 We illustrate the previous construction in the case of liftings of the quartet tree $t = ((1, 2), (3, 4))$, describing the reduced weights v^s for six trivalent trees s in Table 1. The remaining reduced weights can be obtained by relabelings of s that respect the combinatorial type of t . The table is organized as follows. The first five columns encode the branch lengths of the lower tree: $\mathbf{b}_0$ for the internal edge of t , and $\mathbf{b}_i$ for the edge pendant to taxon i . The rest, x_1 through x_4 , are the four distances to the new taxon. The polytope $\mathcal{B}^{((1,2),(3,4))} \subset \mathbb{R}^9$ is four-dimensional, has 14 vertices and f -vector $(14, 46, 52, 20)$. The vertices of $\mathcal{P}_5$ corresponding to the trees $((1, 3), (5, (2, 4)))$ and $((1, 4), (5, (2, 3)))$ project to the same vertex of $\mathcal{B}^t$. Among all 14 vertices, only 5 correspond to upper BME trees: the reduced weight corresponding to the tree $s = ((2, 5), (3, (1, 4)))$ and its five relabelings that fix t . The affine hull of $\mathcal{B}^t$ has five defining linear equations $x_1 + x_2 + x_3 + x_4 = 1$ and $\mathbf{b}_i + x_i$ for $i = 1, 2, 3, 4$. Analogous equations will define the affine hull for all reduced BME polytopes, as we show in Proposition 4.9.

One can compute the dimension, number of vertices, and f -vector of the reduced polytope $\mathcal{B}^t$ as we did in the case of the BME polytope. We can also study the behavior of the vertices of the BME polytope under the projection map, and see how many of its vertices collapse to a single vertex in $\mathcal{B}^t$, how many lie in the interior and how many lie in proper faces of positive dimension. We now show that the reduced polytope has dimension $2n - 4$ by characterizing its affine hull. First we state a technical lemma. Questions involving vertices and their behavior under the projection map will be deferred to the next section.

Lemma 4.8 *Given a tree t on n taxa, let ω denote the BME weight for t . Then*

$$\sum_{j \neq i} \omega_{ij} = 1 \quad \forall 1 \leq i \leq n.$$

Proof If a nonbacktracking random walk starts at i , then w_{ij} is the probability of that walk ending at j . $\square$

Proposition 4.9 *The affine hull of $\mathcal{B}^t$ is characterized by $n + 1$ linearly independent linear equations. More precisely, they are given by $Ax = \mathbf{1} \in \mathbb{R}^{n+1}$, where*

$$A := \left(\begin{array}{c|c|c} I_n & \mathbf{0} & I_n \\ \hline \mathbf{0} & \mathbf{0} & \mathbf{1} \end{array} \right) \in \mathbb{Z}^{(n+1) \times (3n-3)},$$

and the columns of A and points in $\mathbb{R}^{3n-3}$ are labeled by partitioning the coordinates as $(\mathbf{b}_{e_1}, \dots, \mathbf{b}_{e_n} \mid \mathbf{b}_e : e \text{ interior edges of } t \mid \tilde{D}_{1,n+1}, \dots, \tilde{D}_{n,n+1})$. Here, e_i denotes the edge pendant to the leaf i in tree t . In particular, $\dim \mathcal{B}^t = 2n - 4$, and the $(n + 1)$ -dimensional lineality space of the reduced fan coincides with the row span of A .

Proof First, we rewrite the equations in terms of the coordinates of reduced weights and then apply Lemma 4.8. Fix an upper tree s and write v and ω for v^s and ω^s , respectively. The following equalities hold:

$$\sum_{j=1}^n v_j = \sum_{j \neq n+1} \omega_{i,n+1} = 1,$$

$$v_{e_i} + v_i = \sum_{j \neq i} \omega_{ij} = 1 \quad \forall 1 \leq i \leq n.$$

These are precisely the linear equations described by matrix A .

We now prove that these equations characterize the space. To simplify notation, let ψ be the surjective map $\psi(p) = (\pi_t(p|_{[n]}), p_{1,n+1}, \dots, p_{n,n+1})$ for any lifting p of a tree metric with tree t . We proceed by dimensionality arguments. We know that $rk(A) = n + 1$, so $\dim \mathcal{B}^t \leq 3n - 3 - (n + 1) = 2n - 4$. Our goal is to show that equality holds. It will suffice to show that the dimension of the lineality space of the “reduced fan” equals $n + 1$.

By construction, the shift vectors $\{h_a : 1 \leq a \leq n + 1\}$ represent tree metrics associated with a degeneration of the trivalent tree t with two nodes and one edge: a leaf labeled a , and the other leaf labeled by the set $\{1, \dots, \hat{a}, \dots, n + 1\}$. Hence, these tree metrics can be expressed as points $\tilde{h}_a = \psi(h_a)$ in $\mathbb{R}^{3n-3}$, and they generate an $(n + 1)$ -dimensional vector space. These points are precisely the rows of A as described in the statement. Hence, it suffices to show that these vectors span the lineality space of the “reduced fan.”

Fix any trivalent tree s_0 on $n + 1$ taxa. Given $p \in \mathbb{R}^{3n-3}$ in the lineality space of the reduced fan, by definition we have $\langle p, v^s \rangle = \langle p, v^{s_0} \rangle$ for all trees s . By construction, p lies in the image of ψ , so fix q with $p = \psi(q)$. Thus, $\langle q, \omega^s \rangle = \langle p, v^s \rangle$ for all s by (4), and so $\langle q, \omega^s \rangle = \langle q, \omega^{s_0} \rangle$ for all s . By definition, we have that q is in the lineality space of the BME fan, and so it is a linear combination of the shift vectors. After applying the map ψ , the same holds for p and the vectors $\tilde{h}_a$, and the result follows. $\square$

4.3 Analysis of the Reduced BME Polytope

In this section we focus on combinatorial properties of the reduced BME polytope and the behavior of the vertices of the BME polytope under the projection map α_t , as t varies along the set of combinatorial types of trees on n taxa. In particular, we give a complete description of the vertices for up to six taxa (see Table 2). As we mentioned earlier, two tree topologies on $n + 1$ taxa can give the same vertex in the polytope $\mathcal{B}^t$, and vertices of the BME polytope can map to interior points in $\mathcal{B}^t$ under the projection map. As Example 4.7 shows, for four taxa, there exists a pair of tree topologies with the same associated reduced weight, but all fourteen reduced weights are still vertices of $\mathcal{B}_4^t$. Similarly, in the case of five taxa, a `Polymake` computation shows that all 94 possible (out of 105) reduced weights $\{v^s : s \in \mathcal{T}_6\}$ are vertices. This is no longer true for six taxa.

By construction, the polytope $\mathcal{B}^t$ encodes an optimization problem where we restrict our ambient space $\mathbb{R}^{\binom{n+1}{2}}$ to the space of extensions of tree metrics with associated tree t . In terms of the BME fan, this means cutting out the fan with the $(2n - 3)$ -dimensional cone $\mathbb{R}_+ \mathcal{T}_t \subset \mathbb{R}^{\binom{n+1}{2}}$. Note that by intersecting the BME chambers with this cone, we may get a cone with dimension less than $2n - 3$. Moreover, it could very well happen that this intersection is just the lineality space $\mathbb{R}(\alpha_t(h_a) : 1 \leq a \leq n + 1)$ of the cone. This would imply that the point v^s lies in the interior of the polytope. This is indeed what happens for six taxa, as we have found through computation:

Proposition 4.10 *Let $t = ((1, 2), (3, 4), (5, 6))$ be the snowflake tree. Then the reduced polytope $\mathcal{B}_6^t$ is generated by the 792 reduced weights (out of the possible 945 reduced trivalent points) and it has 780 vertices and 83 227 facets. The remaining twelve reduced trivalent weights v^s that are not vertices of $\mathcal{B}_6^t$ lie in the interior of the polytope. They are associated with pairs of trivalent trees with topologies:*

$(1, (((2, 3), (4, 6)), 7), 5))$ $(1, (((2, 4), (3, 6)), 7), 5))$
 $(1, (((2, 3), 7), (4, 6)), 5))$ $(1, (((2, 3), 7), (4, 5)), 6))$
 $(1, ((2, 3), ((4, 6), 7)), 5))$ $(1, ((2, 5), ((4, 6), 7)), 3))$
 $(1, (((2, 5), (3, 6)), 7), 4))$ $(1, (((2, 6), (3, 5)), 7), 4))$
 $(1, (((2, 5), 7), (3, 6)), 4))$ $(1, (((2, 5), 7), (4, 6)), 3))$
 $(1, (((2, 5), ((3, 6), 7)), 4))$ $(1, (((2, 4), ((3, 6), 7)), 5))$
 $(1, (((2, 6), 7), (3, 5)), 4))$ $(1, (((2, 6), 7), (4, 5)), 3))$
 $(1, (((2, 6), ((3, 5), 7)), 4))$ $(1, (((2, 4), ((3, 5), 7)), 6))$
 $(1, (((2, 3), (4, 5)), 7), 6))$ $(1, (((2, 4), (3, 5)), 7), 6))$
 $(1, (((2, 3), ((4, 5), 7)), 6))$ $(1, (((2, 6), ((4, 5), 7)), 3))$
 $(1, (((2, 4), 7), (3, 6)), 5))$ $(1, (((2, 4), 7), (3, 5)), 6))$
 $(1, (((2, 5), (4, 6)), 7), 3))$ $(1, (((2, 6), (4, 5)), 7), 3))$.

Similarly if t is the lower tree $(1, (((3, 4), 6), 5), 2)$ (the caterpillar tree), then the polytope $\mathcal{B}_6^t$ has 804 distinct reduced weights, 800 vertices, and 116 701 facets. In this case, all four reduced trivalent weights v^s that are not vertices of $\mathcal{B}_6^t$ lie in the interior. In this case, each point corresponds to a single topology, and they are:

$(1, (((2, (3, 5)), 7), 4), 6))$
 $(1, (((2, 6), 3), 7), (4, 5)))$
 $(1, (((2, (4, 5)), 7), 3), 6))$
 $(1, (((2, 6), 4), 7), (3, 5)))$.

From the previous examples we see that in the case of four and five taxa, all reduced points are vertices. And for six taxa, reduced points are either vertices or *interior* points (Proposition 4.10). Thus, it is natural to ask whether these are the only two possibilities:

Question 4.11 For $n \geq 7$ and any tree $t \in \mathcal{T}_n$, are all reduced trivalent points either vertices or interior points of the reduced polytope $\mathcal{B}^t$?

We expect the answer to be positive, provided that the projection map α_t is generic.

We now switch gears and focus on the number of upper BME trees we can obtain from a lifting of a given tree metric with topology t . This study will highlight the behavior of “rogue taxa.” Equivalently, we want to know how many positive reduced cones $\mathcal{C}_s^+(\mathcal{B}^t)$ (for s trivalent trees on $n + 1$ taxa) are nonempty. We provide a complete answer for up to six taxa in Table 2 below.

The next natural question to ask is what are the asymptotics (or provide an upper bound) of the number of such nonempty positive reduced cones. As a first attempt, we give some insight about which topologies can be ruled out for upper BME trees. In other words, which are the *blocking* topologies for upper trees.

Definition 4.12 Fix $t \in \mathcal{T}_n$ and let v^s be the reduced weight for a trivalent tree $s \in \mathcal{T}_{n+1}$. We define a *partial order* on the set $\{v^s : s \in \mathcal{T}_{n+1}\}$ as follows: $v^s > v^{s'}$ if and only if $(v^s)_l \leq (v^{s'})_l$ for all $1 \leq l \leq 3n - 3$. We say that s *blocks* s' if $v^s > v^{s'}$.

Lemma 4.13 Let $t \in \mathcal{T}_n$, and let $s, s' \in \mathcal{T}_{n+1}$ be such that s blocks s' . Then, s' cannot be a BME tree for any lifting $\tilde{D}$ of $D \in \mathcal{C}_t^+$.

Table 2 A comparison between the BME and reduced BME polytopes for up to six taxa. In the case of six taxa, we have more than one combinatorial type for the lower tree t . Each vector in the last column gives the number of reduced BME positive cones classified by dimension, starting from dimension $n + 1$ and up to dimension $3n - 3$. The lowest-dimensional ones correspond to reduced weights of forbidden upper BME trees, since they lie in the linear space spanned by the shift vectors. The discrepancy between the first entry of these vectors and the entry of the column indicating the number of voided upper trees reflects that several of these void trees have equal reduced weights

n	dim.		# vertices		# void upper trees for t	f -vector of reduced BME positive cones
	BME	red.	BME	red.		
3	2	2	3	3	0	(0,0,3)
4	5	4	15	14	2	(1, 0, 0, 0, 13)
5	9	6	105	94	20	(16, 1, 6, 0, 0, 0, 71)
6a	14	8	945	800	208	(160, 32, 98, 10, 39, 0, 0, 461)
6b	14	8	945	780	154	(123, 0, 144, 9, 39, 0, 0, 0, 465)

Fig. 4 The blocking relations (up to symmetry) for trees on six taxa. Pairs of trees in a column are a single blocking relation, with the tree in the second row blocking the corresponding tree in the first row. Note that these blocking relations do not come from Theorem 3.5

Proof It suffices to show that for any $\tilde{D}$, $\lambda(s, \tilde{D}) \leq \lambda(s', \tilde{D})$, and this follows because $\tilde{D}$ has nonnegative entries. $\square$

We illustrate with examples on five taxa.

Example 4.14 Let $t = (1, ((3, 4), 5), 2)$. Out of all possible 94 vertices in $\mathcal{B}^t$, there are 19 reduced vertices that are blocked by other vertices, out of 20 empty positive reduced cones. The blocking relation is described in Fig. 4 and gives 26 blocking upper tree topologies. We simplify the picture by reducing the relation modulo relabeling of all leaves involved in each chain and that fix the lower tree t .

In particular we see that out of the 94 possible BME reduced vertices for t , we can rule out 19 of these vertices for upper trees by “blocking” relations.

Unfortunately, this partial order set is not a sufficient criterion to determine if a tree on $n + 1$ taxa can be an upper tree or not. In particular, it cannot explain the obstruction to exchange subtrees “over” the new pendant edge (Theorem 3.5), except in the case of quartet trees. However, understanding the blocking relation can give an upper bound for the asymptotics of the upper BME trees.

We end this section with a table describing the relation between the BME and reduced BME polytopes for up to six taxa. In the case of six taxa, we have two combinatorial types of lower trees, and each one will label a row in our table. The row starting with “6a” indicates the caterpillar tree on six taxa, whereas “6b” refers to the snowflake tree (see Proposition 4.10).

We conclude with an interesting computationally challenging question:

Question 4.15 What are the asymptotics of the number of vertices of the $\mathcal{B}^t$ and of the number of upper BME trees and upper BME reduced trees for different combinatorial types of lower trees t ?

Table 3 The extremal rays of the polyhedral cone $X_s(t)$ for four lower taxa for $t = ((1, 2), (3, 4))$ and $s = (((1, 5), 3), (2, 4))$. The rows represent the rays. Labeling conventions for rows and columns are described in the text

	$\mathbf{b}_1$	$\mathbf{b}_2$	$\mathbf{b}_3$	$\mathbf{b}_4$	$\mathbf{b}_0$	x_1	x_2	x_3	x_4
c_1	4	0	3	3	1	0	0	0	0
c_2	3	0	3	0	1	0	0	0	0
c_3	1	0	0	0	1	0	3	0	0
c_4	0	0	0	0	1	0	4	1	1
c_5	0	0	0	0	1	0	3	0	3
e_1	1	1	1	1	0	0	0	0	0
e_2	1	1	1	0	0	0	0	0	0
e_3	1	0	1	1	0	0	0	0	0
e_4	1	0	1	0	0	0	0	0	0
e_5	1	0	0	0	0	0	0	0	0
f_1	0	0	0	0	0	1	1	1	1
f_2	0	0	0	0	0	0	1	1	1
f_3	0	0	0	0	0	0	1	0	0
f_4	0	0	0	0	0	0	0	0	1
h_1	1	0	0	0	0	1	0	0	0
h_2	0	1	0	0	0	0	1	0	0
h_3	0	0	1	0	0	0	0	1	0
h_4	0	0	0	1	0	0	0	0	1

4.4 The Rogue Taxon Effect for Four Taxa

The extremal rays of each reduced cone can be interpreted to give precise information on the rogue taxon effect. In this section, we explore the reduced polyhedral cone associated with the lower tree $((1, 2), (3, 4))$ and the upper tree $(((1, 5), 3), (2, 4))$. Up to symmetry, this is the only lower/upper combination for this number of taxa such that the new taxon has “rogue” behavior. By understanding the extremal rays of the polyhedral cone, we establish Propositions 4.16 and 4.17.

Table 3 gives the extremal rays of the cone $X_s(t)$. We follow the notation of Example 4.7 to label the columns. The rows label the extremal rays of the cone, and are divided into sections. In the first section, labeled with c are the rays which give branch length/extra taxon distances with a nontrivial internal branch length for the lower tree. This is visible because of the 1 in the $\mathbf{b}_0$ column. These rays are interesting as they represent the “minimal” rogue taxon examples. We analyze these c_i in more detail below.

The second section, labeled with e , f , and h , shows how the pendant (leading to a leaf) branch lengths of the lower tree and the distances to the new taxon can be modified without changing the upper tree. That is, any positive multiple of these vectors can be added to a point in the cone while staying in the same polyhedral cone. For instance, e_4 says that we can increase the branch lengths $\mathbf{b}_1$ and $\mathbf{b}_3$ simultaneously while maintaining the same upper tree. The ray f_3 (which is all-zero except for the

x_2 column), for example, says that we can increase the distance of the new taxon to the second original taxon without changing the upper tree. The h_i are simply the *shift vectors* corresponding to the pendant branches. Thus h_i means that we can increase the i th pendant branch length while increasing the distance of the new taxon to the i th original taxon without changing the upper tree.

These extremal rays can give some sufficient conditions for rogue taxon behavior. We specify branch lengths of quartets by a vector giving branch lengths in the order $(\mathbf{b}_0, \dots, \mathbf{b}_4)$. We say that a vector $\mathbf{x}$ is a rogue vector for a branch-length vector $\mathbf{b}$ if the BME tree for the combined data as in Table 3 is the tree $((1, 5), 3), (2, 4))$. We will call the cone given by positive linear combinations of the set

$$\{(0, 1, 1, 1, 1), (0, 1, 1, 1, 0), (0, 1, 0, 1, 1), (0, 1, 0, 1, 0), (0, 1, 0, 0, 0)\}$$

the *extension cone*. Any element from this cone can be added to a branch-length set without changing the polyhedral cone; this can be seen by looking at the e_i vectors above.

Note that any vector satisfying $0 \leq x_1 \leq x_3 \leq \min(x_2, x_4)$ sits in the cone generated by the f_i restricted to their last four coordinates. Therefore we conclude:

Proposition 4.16 *Any vector satisfying $0 \leq x_1 \leq x_3 \leq \min(x_2, x_4)$ is a rogue vector for any tree with branch-length vector given by either $(1, 4, 0, 3, 3)$ or $(1, 3, 0, 3, 0)$ plus any element of the extension cone.*

The next proposition gives rogue criteria for a quartet tree with arbitrary internal branch length. The proof is simple: just look at c_5 in Table 3, which shows that $(0, 3, 0, 3)$ is a rogue vector for the quartet with trivial pendant branch lengths and internal branch length 1.

Proposition 4.17 *Any quartet tree has a rogue vector with an entry greater than or equal to three times the internal branch length of the lower tree.*

Although the above propositions do give some conditions on when the rogue taxon effect appears for four taxa, they do not specify how likely are we to end up in a rogue taxon situation. They also give no information about trees on larger number of taxa. In the next section, we gain some intuition about these questions via simulation.

4.5 Simulations

Here we describe simulations performed to better understand the rogue taxon effect as it might appear in biological data. These simulations show that, at least for small numbers of taxa, the rogue taxon effect is common when the extra distances are chosen without reference to the original tree. They also suggest that the effect gets worse as the number of taxa increases.

We assume a random distribution for the branch lengths and distances to the new taxon. Such simulations are not the only way to address these sorts of questions. Volume computations of, e.g., spheres intersected with our polyhedral cones are in principle possible, but they do not seem to admit a closed-form solution. Thus our

Table 4 Simulation results for 10^7 exponentially distributed branch lengths and distances to rogue taxa. The columns are labeled by the topology of the lower tree. The numbers in the table represent the fraction of time that the corresponding Robinson–Foulds distance between the upper and lower trees appeared via the rogue taxon effect

δ_{RF}	((1, 2), (3, 4))	((1, 2), 3, (4, 5))	((((1, 2), 5), ((3, 4), 6))	((1, 2), (3, 4), (5, 6))
0	0.705071	0.502925	0.380863	0.381869
1	0.294929	0.364874	0.367523	0.363955
2	–	0.132201	0.195223	0.209066
3	–	–	0.0563907	0.04511

understanding of such volumes still depends strongly on Monte Carlo simulations (Eickmeyer et al. 2008). Furthermore, such a volume may give less practical information than simulation using a reasonable model of branch lengths.

To better understand the frequency with which the rogue taxon phenomenon can occur, we simulate using the exponential distribution. Although a simple arbitrary choice, the exponential distribution is realistic enough to be a branch length prior for Bayesian phylogenetic inference (Ronquist et al. 2005). For a given lower tree, we generate branch lengths for that tree according to the mean one exponential distribution, then generate distances to the extra taxon via the exponential distribution with mean equal to the expected pairwise distance between tips of the tree. Then, we find the upper tree (i.e., the BME tree for the original data set plus the rogue taxon) and check to see how many bipartitions of the upper tree (restricted to the lower taxa) are not contained in the lower tree. This number is the Robinson–Foulds distance between the upper and lower trees used in Sect. 3.1.

The results of 10^7 exponentially drawn branch lengths are shown in Table 4; it shows that a taxon added with random data can substantially alter the structure of the phylogenetic tree. Indeed, almost 30% of the lifted four taxon trees do not contain the original topology, growing to almost 50% for five taxa, then almost 62% for the six taxon topologies.

We emphasize that such simulations do not paint an accurate picture of the rogue taxon effect for real data. Indeed, even the worst data does not have completely random distances: even “random” sequence data will not have random distances to the rest of the tree. Nevertheless, we believe that these results indicate that this area merits further investigation and that the effective volume of these “rogue” polyhedral cones is not small.

In the reduced BME setting it can happen that multiple bifurcating upper trees are associated with a cone of the reduced normal fan for a given lower tree. That is, the trees all have the same BME length for given lower tree branch lengths and rogue taxon distances. We have observed in the example presented here that when there are these multiple trees, the Robinson–Foulds distances between the lower tree and these multiple upper trees (restricted to the lower taxa) for a given cone are equal. It would be interesting to know if this is true in the general case.

The equivalent fact for the quartet distance is not true. In the case of the lower tree being $((((1, 2), 5), ((3, 4), 6)), 3)$, there is a cone of the reduced normal fan associated with both $(1, (((((2, 3), (4, 7)), 6), 5)))$ and $(1, (((((2, 6), (4, 7)), 3), 5)))$. Restricting to

the lower taxa, these trees are $(1, (((((2, 3), 4), 6), 5))$ and $(1, (((((2, 6), 4), 3), 5))$, which have quartet distances 10 and 11, respectively, to the lower tree.

5 Conclusions and Future Directions

We have investigated the effect of adding an extra “rogue” taxon into a phylogenetic data set for BME phylogenetic inference. We have shown that rogue taxa can have significant though not arbitrary effects on the tree. For a small number of taxa, we can delineate the domain of the rogue taxon effect. Simulations show that the rogue taxon effect is very significant when the data for the rogue taxon is chosen randomly without reference to the topology of the original tree.

The results presented here may have algorithmic consequences for phylogenetic inference. It is common for inference programs to start with a tree on three taxa and then build a tree by adding taxa sequentially. Software packages using sequential taxon addition, such as PHYLIP (Felsenstein 1995) and fastDNAmI (Olsen et al. 1994), do optimize the tree after addition using rearrangements; the question of strict sequential addition performance is still important in order to determine the amount of post-addition optimization required. Furthermore, “evolutionary placement algorithms” for large amounts of sequence data have been proposed whereby a “query” sequences are inserted into a fixed “reference tree” (Von Mering et al. 2007; Berger and Stamatakis 2009). The accuracy of such algorithms compared to traditional phylogenetics algorithms can be seen as an aspect of the rogue taxon problem.

An interesting next direction would be to consider situations where rogue taxa do not have arbitrary data but appear via misspecified evolutionary models. This will hopefully give a clearer understanding of the actual impact of rogue taxa. It would also be interesting to see if some of the results presented here also extend to other inference criteria, such as parsimony or maximum likelihood. Some results, such as the simulation results presented above, will certainly be different in this new setting, but others may correspond well. Maximum likelihood and parsimony are considerably more difficult to analyze, but hopefully the results presented here can act as a guide.

Acknowledgements We thank Tracy Heath for directing us to the taxon sampling debate, Mike Steel for simplifying the proof of Theorem 3.2, and Bernd Sturmfels, Lior Pachter and Rudy Yoshida for fruitful discussions. We are grateful to the two anonymous reviewers for their careful reading of the manuscript.

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

References

- Baurain, D., Brinkmann, H., & Philippe, H. (2007). Lack of resolution in the animal phylogeny: Closely spaced cladogeneses or undetected systematic errors? *Mol. Biol. Evol.*, 24(1), 6.
- Berger, S. A., & Stamatakis, A. (2009). Evolutionary placement of short sequence reads. <http://arxiv.org/abs/0911.2852>.
- Bordewich, M., Gascuel, O., Huber, K. T., & Moulton, V. (2009). Consistency of topological moves based on the balanced minimum evolution principle of phylogenetic inference. *IEEE/ACM Trans. Comput. Biol. Bioinfo.*, 6(1), 110–117.

- Brodal, G. S., Fagerberg, R., & Pedersen, C. N. S. (2004). Computing the quartet distance between evolutionary trees in time $O(n \log n)$. *Algorithmica*, 38(2), 377–395.
- Bryant, D., Tsang, J., Kearney, P., & Li, M. (2000). Computing the quartet distance between evolutionary trees. In *Proceedings of the eleventh annual ACM-SIAM symposium on discrete algorithms* (p. 286). Society for Industrial and Applied Mathematics, Philadelphia.
- Chailloux, E., Manoury, P., & Pagano, B. (2000). Developing applications with objective caml. <http://caml.inria.fr/ocaml/index.en.html>.
- DeBry, R. W. (2005). The systematic component of phylogenetic error as a function of taxonomic sampling under parsimony. *Syst. Biol.*, 54(3), 432.
- Desper, R., & Gascuel, O. (2002a). Fast and accurate phylogeny reconstruction algorithms based on the minimum-evolution principle. In *Workshop on algorithms in bioinformatics (WABI)* (pp. 357–374).
- Desper, R., & Gascuel, O. (2002b). Fast and accurate phylogeny reconstruction algorithms based on the minimum-evolution principle. *J. Comput. Biol.*, 9(5), 687–705.
- Desper, R., & Gascuel, O. (2004). Theoretical foundation of the balanced minimum evolution method of phylogenetic inference and its relationship to weighted least-squares tree fitting. *Mol. Biol. Evol.*, 21(3), 587–598.
- Desper, R., & Gascuel, O. (2005). The minimum evolution distance-based approach to phylogenetic inference. In O. Gascuel (Ed.), *Mathematics of evolution & phylogeny* (pp. 1–32). Oxford: Oxford University Press.
- Eickmeyer, K., Huggins, P., Pachter, L., & Yoshida, R. (2008). On the optimality of the neighbor-joining algorithm. *Algorithms Mol. Biol.*, 3(5).
- Ewald, G. (1996). *Graduate texts in mathematics: Vol. 168. Combinatorial convexity and algebraic geometry*. New York: Springer.
- Felsenstein, J. (1978). Cases in which parsimony or compatibility methods will be positively misleading. *Syst. Zool.*, 27(4), 401–410.
- Felsenstein, J. (1995). PHYLIP (phylogeny inference package), version 3.57 c. Department of Genetics, University of Washington, Seattle.
- Felsenstein, J. (2004). *Inferring Phylogenies*. Sunderland, MA: Sinauer Press.
- Galassi, M., Davies, J., Theiler, J., Gough, B., Jungman, G., Alken, P., Booth, M., & Rossi, F. (2009). *GNU scientific library reference manual* (3rd ed.). Network Theory Ltd. <http://www.gnu.org/software/gsl/>.
- Gawrilow, E., & Joswig, M. (2000). Polymake: A framework for analyzing convex polytopes. In G. Kalai, & G. M. Ziegler (Eds.), *Polytopes—combinatorics and computation* (pp. 43–74). Basel: Birkhäuser.
- Graybeal, A. (1998). Is it better to add taxa or characters to a difficult phylogenetic problem? *Syst. Biol.*, 47(1), 9.
- Guillemot, S. & Pardi, F. (2009). Personal communication.
- Heath, T. A., Hedtke, S. M., & Hillis, D. M. (2008a). Taxon sampling and the accuracy of phylogenetic analyses. *J. Syst. Evol.*, 46(3), 239–257.
- Heath, T. A., Zwickl, D. J., Kim, J., & Hillis, D. M. (2008b). Taxon sampling affects inferences of macroevolutionary processes from phylogenetic trees. *Syst. Biol.*, 57(1), 160.
- Hedtke, S. M., Townsend, T. M., & Hillis, D. M. (2006). Resolution of phylogenetic conflict in large data sets by increased taxon sampling. *Syst. Biol.*, 55(3), 522.
- Hendy, M. D., & Penny, D. (1989). A framework for the quantitative study of evolutionary trees. *Syst. Zool.*, 38(4), 297–309.
- Hillis, D. M. (1996). Inferring complex phylogenies. *Nature*, 383(6596), 130.
- Hillis, D. M., Pollock, D. D., McGuire, J. A., & Zwickl, D. J. (2003). Is sparse taxon sampling a problem for phylogenetic inference? *Syst. Biol.*, 52(1), 124–126.
- Jensen, A. N. (2009). Gfan—A software system for Gröbner fans. Available at <http://www.math.tu-berlin.de/~jensen/software/gfan/gfan.html>.
- Kim, J. (1996). General inconsistency conditions for maximum parsimony: Effects of branch lengths and increasing numbers of taxa. *Syst. Biol.*, 45(3), 363.
- Mailund, T., & Pedersen, C.N.S. (2004). QDist—Quartet distance between evolutionary trees. *Bioinformatics*, 971.
- Olsen, G. J., Matsuda, H., Hagstrom, R., & Overbeek, R. (1994). fastDNAm1: A tool for construction of phylogenetic trees of DNA sequences using maximum likelihood. *Bioinformatics*, 10(1), 41.
- Pachter, L., & Sturmfels, B. (Eds.) (2005). *Algebraic statistics for computational biology* (p. 69). Cambridge: Cambridge University Press. Chap. II.
- Padberg, M. W. & Grötschel, M. (1985). Polyhedral computations. In *Wiley-intersci. ser. discrete math. The traveling salesman problem* (pp. 307–360). Chichester: Wiley.

- Poe, S. (1998). Sensitivity of phylogeny estimation to taxonomic sampling. *Syst. Biol.*, 47(1), 18.
- Poe, S. (2003). Evaluation of the strategy of long-branch subdivision to improve the accuracy of phylogenetic methods. *Syst. Biol.*, 52(3), 423–428.
- Pollock, D. D., Zwickl, D. J., McGuire, J. A., & Hillis, D. M. (2002). Increased taxon sampling is advantageous for phylogenetic inference. *Syst. Biol.*, 51(4), 664–671.
- Rannala, B., Huelsenbeck, J. P., Yang, Z., & Nielsen, R. (1998). Taxon sampling and the accuracy of large phylogenies. *Syst. Biol.*, 47(4), 702–710.
- Robinson, D. F., & Foulds, L. R. (1981). Comparison of phylogenetic trees. *Math. Biosci.*, 53(1–2), 131–147.
- Ronquist, F., Huelsenbeck, J. P., & van der Mark, P. (2005). MrBayes 3.1 manual. http://mrbayes.csit.fsu.edu/mb3.1_manual.pdf.
- Rosenberg, M. S., & Kumar, S. (2001). Incomplete taxon sampling is not a problem for phylogenetic inference. *Proc. Natl. Acad. Sci.*, 98(19), 10751.
- Rosenberg, M. S., & Kumar, S. (2003). Taxon sampling, bioinformatics, and phylogenomics. *Syst. Biol.*, 52(1), 119–124.
- Saitou, N., & Nei, M. (1987). The neighbor-joining method: A new method for reconstructing phylogenetic trees. *Mol. Biol. Evol.*, 4(4), 406–425.
- Steel, M. A., & Penny, D. (1993). Distributions of tree comparison metrics—Some new results. *Syst. Biol.*, 42(2), 126.
- Studier, J. A., & Keppler, K. J. (1988). A note on the neighbor-joining method of Saitou and Nei. *Mol. Biol. Evol.*, 5(6), 729–731.
- Sullivan, J., & Swofford, D. L. (1997). Are guinea pigs rodents? The importance of adequate models in molecular phylogenetics. *J. Mammal. Evol.*, 4(2), 77–86.
- Von Mering, C., Hugenholtz, P., Raes, J., Tringe, S. G., Doerks, T., Jensen, L. J., Ward, N., & Bork, P. (2007). Quantitative phylogenetic assessment of microbial communities in diverse environments. *Science*, 315(5815), 1126.
- Ziegler, G. M. (2006). *Graduate texts in mathematics: Vol. 152. Lectures on polytopes*. Berlin: Springer.
- Zwickl, D. J. & Hillis, D. M. (2002). Increased taxon sampling greatly reduces phylogenetic error. *Syst. Biol.*, 51(4), 588.