

L-INFINITY OPTIMIZATION TO BERGMAN FANS OF MATROIDS
WITH AN APPLICATION TO PHYLOGENETICS*

DANIEL IRVING BERNSTEIN†

Abstract. Given a dissimilarity map δ on a finite set X , the set of ultrametrics (equidistant tree metrics) which are l^∞ -nearest to δ is a tropical polytope. We give an internal description of this tropical polytope which we use to derive a polynomial-time checkable test for the condition that all ultrametrics l^∞ -nearest to δ have the same tree structure. It was shown by Ardila and Klivans [*J. Combin. Theory Ser. B*, 96 (2006), pp. 38–49] that the set of all ultrametrics on a finite set of size n is the Bergman fan associated with the matroid underlying the complete graph on n vertices. Therefore, we derive our results in the more general context of Bergman fans of matroids. This added generality allows our results to be used on dissimilarity maps where only subsets of the entries are known.

Key words. tropical polytopes, Bergman fans, phylogenetics, matroids

AMS subject classifications. 14T05, 05B35

DOI. 10.1137/18M1218741

1. Introduction. A fundamental problem in phylogenetics is to infer the evolutionary history among a set of genes or species from data. One approach is to use *distance-based methods*. The data required for such an approach are some measure of distance between each pair of species. If these distances are computed using some property that is expected to change in proportion to time elapsed, then one often assumes that the pairwise distances approximate an ultrametric. Finding a best-fit ultrametric to an arbitrary dissimilarity map is therefore an important computational problem. For background, see [21, Chapter 7].

A major source of difficulty in this endeavor stems from the fact that two of the most basic sets with which one would reason about distance-based phylogenetics, namely, the set of tree metrics and the set of ultrametrics, do not interact with Euclidean geometry in a clean way, thus making naive application of traditional statistical methods problematic. Beginning with the work of Billera, Holmes, and Vogtman [6], the past two decades have seen much research into developing and studying geometric theories that interact nicely with the sets of tree metrics and ultrametrics, with the hope that reinterpreting traditional statistical theory and methods in these new geometries will lead to something useful. Speyer and Sturmfels [22] and Ardila and Klivans [4] showed that the sets of tree metrics and ultrametrics are tropical varieties, thus giving the first indication that tropical geometry might offer useful tools for phylogenetics.

Since then, researchers have been exploring tropical geometry's potential as a fundamental theory on which to develop statistical methods designed specifically for phylogenetic applications. Tropical geometry is a geometric theory that one naturally obtains when redefining arithmetic over $\mathbb{R} \cup \{-\infty\}$ so that the sum of two numbers is their maximum and the product is their sum (in the usual sense). The natural

*Received by the editors October 3, 2018; accepted for publication (in revised form) December 20, 2019; published electronically March 12, 2020.

<https://doi.org/10.1137/18M1218741>

Funding: This work was partially supported by the U.S. National Science Foundation (grants DMS 0954865 and 1802902) and the David and Lucile Packard Foundation.

†Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, Massachusetts, 02139 (dibernst@mit.edu).

choice for a metric in tropical geometry is the l^∞ -metric. A recent preprint of Lin, Monod, and Yoshida [16] shows that the set of phylogenetic trees endowed with the tropically projectivized l^∞ -metric, which they call *palm tree space*, has many features of Euclidean space that enable classical statistical theory to work. In particular, palm tree space supports probability measures and a reasonable theory of linear algebra. In [17], Lin et al. compare tropical convexity to the convexity theory of Billera, Holmes, and Vogtmann [6] with regard to their potential as theoretical frameworks for developing algorithms to reduce the complexity of a dataset consisting of several ultrametrics on the same taxa. They show that in the convexity theory of [6], a triangle (i.e., the convex hull of three points) can have arbitrarily high dimension, whereas triangles are always two dimensional in the projective tropical setting [9]. In [23], Yoshida, Zhang, and Zhang develop a theory of tropical principal component analysis.

A recurring frustration one encounters when tropicalizing a classical object is that uniqueness guarantees may disappear. In order for tropical geometry to be considered a reasonable mathematical foundation for phylogenetic analysis, failures of uniqueness must be understood when they have potential to cause problems. Lin and Yoshida [18] studied nonuniqueness of the tropical Fermat–Weber point, which is analogous to the geometric mean from Euclidean geometry. They showed that the set of all tropical Fermat–Weber points is a (classical) polytope, and gave a necessary condition for uniqueness of the tropical Fermat–Weber point. In this paper, we provide analogous results for nonuniqueness of the ultrametric that is nearest to a given dissimilarity map in the l^∞ -metric. Colby Long and this author began a study of this, and other related phenomena, in [5]. The main mathematical results of [5] concern the nonuniqueness of the point in a (nontropical) linear subspace of \mathbb{R}^n that is l^∞ -nearest to a given $x \in \mathbb{R}^n$. In that paper, it is also shown that there exist dissimilarity maps in $\mathbb{R}^{\binom{n}{2}}$ whose set of l^∞ -nearest ultrametrics contains $\frac{1}{3} \cdot (2n - 3)!!$ different tree topologies.

This paper builds on some of these observations. In particular, Proposition 3.3 says that the set of ultrametrics l^∞ nearest to a given dissimilarity map is a tropical polytope, Theorem 3.6 provides an internal description, and Theorem 3.8 gives a polynomial-time checkable condition, telling us exactly when all nearest ultrametrics have the same tree structure. From a phylogenetics perspective, this is useful information since the tree structure describes the evolutionary relationship among the species being studied.

We derive our results in a more general context. Ardila and Klivans showed that the set of ultrametrics on n species is the Bergman fan associated with the matroid underlying the complete graph on n vertices [4]. Therefore we can view the problem of finding the set of l^∞ -nearest ultrametrics as a special case of the problem of finding the set of l^∞ -nearest points in the Bergman fan of a matroid. This latter set is also a tropical polytope (Proposition 5.2) and Theorem 5.9 provides an internal description. Feichtner and Sturmfels describe a refinement of the Bergman fan underlying a matroid [11] which can be used to generalize the concept of tree topology. In light of this, Theorem 5.9 is the straightforward generalization of Theorem 3.6.

The added generality of Bergman fans of matroids has a potential application in phylogenetics. Namely, if one wishes to reconstruct a phylogeny from partial distance data where observed distances correspond to the edges of some graph G , then one can begin by optimizing to the Bergman fan of G 's matroid which will give a partial ultrametric (see Proposition 5.12). This reconstruction problem is a special case of the *sandwich to ultrametric problem* studied by Farach, Kannan, and Warnow in [10]. The added generality is also interesting from a pure tropical geometry perspective. In

particular, given the Bergman fan $\tilde{\mathcal{B}}(\mathcal{M})$ of a matroid \mathcal{M} , the question of describing points in $\tilde{\mathcal{B}}(\mathcal{M})$ that are tropically nearest to a given $x \notin \tilde{\mathcal{B}}(\mathcal{M})$ is in some sense the tropical analog of finding the point of a (classical) linear space L that is Euclidean-nearest to a given $x \notin L$.

Just as with ordinary polytopes, tropical polytopes admit *external descriptions* as the intersection of tropical half-spaces, as well as internal descriptions [15]. Theorem 7.1 in [1] can be used to obtain an external description of the tropical polytopes we are interested in. However, an internal description is more advantageous for our purposes because it gives us a way to check whether all ultrametrics l^∞ -nearest to a given dissimilarity map have the same tree topology (see Theorem 3.8 and Proposition 5.11).

This paper is organized as follows. Section 2 gives the necessary background on tropical convexity. Section 3 contains Theorem 3.6, which is an internal description of the tropical polytope consisting of the ultrametrics that are l^∞ -nearest to a given dissimilarity map. A proof is deferred until section 5. Section 3 also states and proves Theorem 3.8, which provides a polynomial-time method for checking that all ultrametrics l^∞ -nearest to a given dissimilarity map have the same tree topology. Section 4 uses results of Feichtner and Sturmfels [11] to generalize the tree structure underlying an ultrametric to a similar combinatorial structure underlying an element of the Bergman fan of an arbitrary matroid. This combinatorial structure is used in section 5 to generalize Theorem 3.6 to get Theorem 5.9. Section 6 applies Theorem 3.6 to a biological dataset.

2. Preliminaries on tropical convexity. This section reviews the necessary concepts from tropical convexity. There are at least two different sets of basic definitions related to tropical convexity. One is used in [9], and the other in [2]. We adhere to the conventions of the latter, as their definition of tropical polytope is more natural in our context.

The *tropical semiring*, also known as the max-plus algebra, is the set $\mathbb{R} \cup \{-\infty\}$ together with the operations $a \oplus b := \max\{a, b\}$ and $a \odot b := a + b$. We denote this semiring by \mathbb{R}_{\max} . The additive identity of \mathbb{R}_{\max} is $-\infty$ and the multiplicative identity is 0. The set \mathbb{R}_{\max}^n is an \mathbb{R}_{\max} -semimodule, where for $x, y \in \mathbb{R}_{\max}^n$ and $\alpha \in \mathbb{R}_{\max}$, $(x \oplus y)_i := x_i \oplus y_i$ and $(\alpha \odot x)_i := \alpha + x_i$. If $A \in \mathbb{R}_{\max}^{m \times n}$ is a matrix and $x \in \mathbb{R}_{\max}^n$, then the product $A \odot x$ is the usual matrix product, but with multiplication and addition interpreted tropically. That is, if A has columns a_1, \dots, a_n , then

$$A \odot x := \bigoplus_{j=1}^n x_j \odot a_j.$$

Several notions from ordinary convexity theory have tropical analogs. We say that $P \subseteq \mathbb{R}_{\max}$ is a *tropical cone* if whenever $x, y \in P$ and $\lambda, \mu \in \mathbb{R}_{\max}$, $\lambda \odot x \oplus \mu \odot y \in P$. If this only holds with the restriction that $\lambda \oplus \mu = 0$, then we say that P is *tropically convex*. A *tropical polyhedron* is a set of the form

$$\{x \in \mathbb{R}_{\max}^n : A \odot x \oplus b \geq C \odot x \oplus d\},$$

where $A, C \in \mathbb{R}_{\max}^{m \times n}$ and $b, d \in \mathbb{R}_{\max}^m$. We denote this set $P(A, b, C, d)$. It follows from the discussion below that $P(A, b, C, d)$ is tropically convex. When $b = d = (-\infty, \dots, -\infty)^T$ then $P(A, b, C, d)$ is a tropical cone and we call it a *tropical polyhedral cone*. Bounded tropical polyhedra are called *tropical polytopes*. Given $V \subseteq \mathbb{R}_{\max}$, $\text{tconv}(V)$ is the *tropical convex hull* of V . That is,

$$\text{tconv}(V) := \{\lambda \odot x + \mu \odot y : x, y \in V, \lambda \oplus \mu = 0\}.$$

We define the *tropical conic hull* $\text{tcone}(V)$ similarly. Gaubert and Katz showed in [14] that any tropical polytope (cone) P can be expressed as the tropical convex (conic) hull of a finite set V . Conversely, Gaubert showed that if $V \subseteq \mathbb{R}^n$ is a finite set and $P = \text{tcone}(V)$, then P is a tropical polyhedral cone [13, Corollary 1.2.5]. The analogous result for $P = \text{tconv}(V)$ follows from results in [14]. There exists a minimal such V (see [7, Theorem 18] or [14, Theorem 3.1]) called the *tropical vertices (extreme rays) of P* . See also [15, Theorem 2].

3. Results for phylogenetics: l -infinity nearest ultrametrics. This section presents the results of section 5 in the context of our main motivation. In particular, Theorem 3.6 gives a combinatorial description of a finite set of ultrametrics whose tropical convex hull is the set of ultrametrics nearest in the l^∞ -norm to a given dissimilarity map. We also use Theorem 3.6 to derive Theorem 3.8, which gives a polynomial-time checkable condition guaranteeing that all ultrametrics l^∞ nearest to a given dissimilarity map have the same tree topology. We begin by reviewing the necessary background from [21] about ultrametrics, which are a special type of tree metric.

Let $X = \{x_1, \dots, x_n\}$ be a finite set. A *dissimilarity map* on X is a function $\delta : X \times X \rightarrow \mathbb{R}$ such that $\delta(x, x) = 0$ and $\delta(x, y) = \delta(y, x)$ for all $x, y \in X$. Note that we allow dissimilarity maps to take negative values. We can express a dissimilarity map d as a matrix D , where $D_{ij} = \delta(x_i, x_j)$. Note that D is symmetric with zeros along the diagonal. A *rooted X -tree* is a tree with leaf set X , where one interior (i.e., nonleaf) vertex has been designated the “root.” We use the notation $\text{root}(T)$ for the root of a rooted X -tree T . A *descendant* of a vertex v in a rooted tree T is a node $u \neq v$ such that the unique path from u to $\text{root}(T)$ contains v . Note that all nonroot vertices are descendants of $\text{root}(T)$. The set of descendants of a vertex v in a rooted tree T is denoted $\text{Des}_T(v)$. We let T° denote the set of interior vertices of T .

Let T be a rooted X -tree and let $\alpha : T^\circ \rightarrow \mathbb{R}$ be a weighting of the internal nodes of T . We say that α is *compatible with T* if $\alpha(u) \leq \alpha(v)$ whenever $u \in \text{Des}_T(v)$. The pair (T, α) gives rise to a dissimilarity map $\delta_{T, \alpha}$ on X defined by $\delta_{T, \alpha}(x_i, x_j) := \alpha(v)$, where $v \in T^\circ$ is the vertex nearest to $\text{root}(T)$ in the unique path from x_i to x_j . Given a dissimilarity map δ on X , if we can express δ as $\delta_{T, \alpha}$ for some X -tree T and compatible internal node weighting α , then δ is said to be an *ultrametric*. If we require that $\alpha(u) < \alpha(v)$ whenever $u \in \text{Des}_T(v)$, then the rooted X -tree T is unique and called the *(tree) topology* of δ . Some readers from other areas of mathematics take issue with this use of the word “topology,” but it is standard in the phylogenetics literature [21]. Figure 1 shows an ultrametric along with an interior-vertex-weighted tree displaying it.

Some readers may be familiar with a seemingly different definition of ultrametric which says that $\delta : X \times X \rightarrow \mathbb{R}$ is an ultrametric if and only if for every triple $x, y, z \in X$ of distinct elements, the maximum of $\delta(x, y), \delta(x, z), \delta(y, z)$ is attained twice. This is equivalent to the definition given above. Sometimes the requirement that x, y, z be distinct is relaxed. This gives the more restricted class of ultrametrics, consisting only of ultrametrics representable as $\delta_{T, \alpha}$ for nonnegative α compatible with T . See [21, Chapter 7] for details. We use the more inclusive definition of an ultrametric because it simplifies connections with tropical geometry.

A *polytomy* of a rooted tree is either a nonroot internal node of degree at least four, or the root node if it has degree at least three. We say that a rooted tree is *binary* if it does not have any polytomy. A *resolution* of a tree T is a binary tree T' such that T can be obtained from T' via a (possibly empty) series of edge contractions [21]. Note that if the topology underlying δ is not binary, then there will be multiple

$$\begin{array}{cccc} & A & B & C & D \\ A & \left(\begin{array}{cccc} 0 & 5 & 7 & 9 \\ 5 & 0 & 7 & 9 \\ 7 & 7 & 0 & 9 \\ 9 & 9 & 9 & 0 \end{array} \right) \\ B \\ C \\ D \end{array}$$

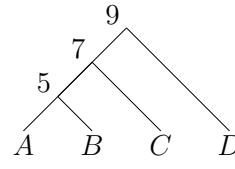


FIG. 1. An ultrametric on $\{A, B, C, D\}$ and its representation on a rooted tree.

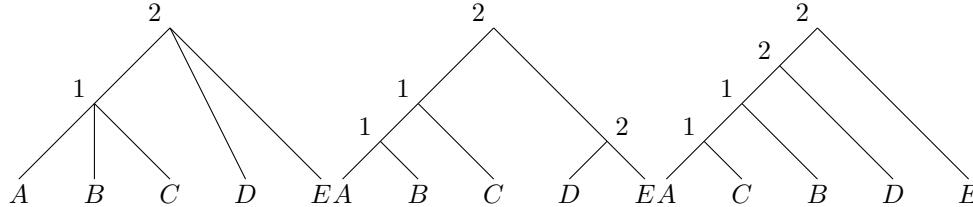


FIG. 2. An ultrametric whose topology has two polytomies. Above, we see it represented on its topology and on two different resolutions.

resolutions of the topology of δ . Figure 2 illustrates these concepts by representing a single ultrametric in three ways—on its topology and on two different resolutions.

Given two dissimilarity maps δ_1, δ_2 on X with associated matrices D_1, D_2 , we define the l^∞ distance between δ_1 and δ_2 , denoted $\|\delta_1 - \delta_2\|_\infty$, to be the greatest absolute value among entries in $D_1 - D_2$. An important question that comes up in phylogenetics is then, given a dissimilarity map δ , which ultrametrics are nearest to δ in the l^∞ metric? Chepoi and Fichet [8] give an algorithm for producing a single ultrametric l^∞ -nearest to a given dissimilarity map which we now describe. We denote the all-ones vector or dissimilarity map by $\mathbf{1}$.

THEOREM 3.1 (see [8, Corollary 1 and discussion on p. 607]). *Let δ be a dissimilarity map on a finite set X . Then the following algorithm produces an ultrametric on X that is nearest to δ in the l^∞ norm.*

1. Draw the complete graph on vertex set X .
2. Label the edge between x and y by $\delta(x, y)$.
3. Define $\delta_u : X \times X \rightarrow \mathbb{R}$ so that for each $x, y \in X \times X$,

$$\delta_u(x, y) := \min_P \left(\max_{\text{edges } (i, j) \text{ of } P} \delta(i, j) \right),$$

where the minimum is taken over all paths P from x to y .

4. Define $d := \|\delta_u - \delta\|_\infty$. Then $\delta_u + \frac{d}{2}\mathbf{1}$ is an ultrametric that is l^∞ -nearest to δ .

Although the algorithm given by Theorem 3.1 produces only one ultrametric, there can be multiple ultrametrics that are l^∞ -nearest to a given dissimilarity map. Figure 3 shows a dissimilarity map alongside two l^∞ -nearest ultrametrics with differing topologies.

DEFINITION 3.2. *We call the ultrametric given by Theorem 3.1 the maximal closest ultrametric to δ and denote it symbolically as δ_m .*

That δ_m is coordinatewise-maximal among ultrametrics nearest to δ is shown in [8], and also follows from Lemma 5.1.3.

$$\begin{array}{cccc} A & B & C & D \\ A & \left(\begin{array}{cccc} 0 & 2 & 4 & 6 \\ 2 & 0 & 8 & 10 \\ 4 & 8 & 0 & 12 \\ 6 & 10 & 12 & 0 \end{array} \right) & & \end{array}$$

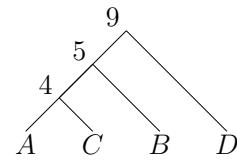
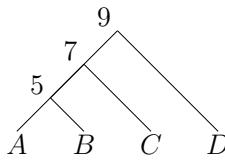


FIG. 3. A dissimilarity map and two l^∞ -nearest ultrametrics with different topologies.

PROPOSITION 3.3. *Let δ be a dissimilarity map on a finite set X . The set of ultrametrics that are nearest to δ in the l^∞ -norm is a tropical polytope.*

We will prove Proposition 3.3 in a more general setting later (see Proposition 5.2). Given a dissimilarity map $\delta : X \times X \rightarrow \mathbb{R}$, Theorem 3.6 describes a finite set of ultrametrics whose tropical convex hull is the set of ultrametrics l^∞ -nearest to δ . The statement of Theorem 3.6 requires the following definition.

DEFINITION 3.4. *Let $\delta : X \times X \rightarrow \mathbb{R}$ be a dissimilarity map and let u be an ultrametric that is closest to δ in the l^∞ -norm. Let T be a resolution of the topology of u and let $\alpha : T^\circ \rightarrow \mathbb{R}$ be a compatible weighting of T 's internal nodes such that $\delta_{T,\alpha} = u$. An internal node v of T is said to be mobile if there exists an ultrametric $\hat{u} \neq u$, expressible as $\hat{u} = \delta_{T,\hat{\alpha}}$ for $\hat{\alpha} : T^\circ \rightarrow \mathbb{R}$ such that*

1. \hat{u} is also nearest to δ in the l^∞ -norm,
2. $\hat{\alpha}(x) = \alpha(x)$ for all internal nodes $x \neq v$, and
3. $\hat{\alpha}(v) < \alpha(v)$.

In this case, we say that \hat{u} is obtained from u by sliding v down. If moreover v is no longer mobile in $\delta_{T,\hat{\alpha}}$, i.e., if $\hat{\alpha}(v) = \max\{\hat{\alpha}(y) : y \in \text{Des}_T(v)\}$, or $\hat{\alpha}(v)$ is the minimum value such that $\delta_{T,\hat{\alpha}}$ is nearest to δ in the l^∞ -norm, then we say that \hat{u} is obtained from u by sliding v all the way down.

Example 3.5. Let δ be the dissimilarity map shown on the left in Figure 3 and consider the l^∞ -nearest ultrametrics u_1, u_2 , and u_3 shown in Figure 4. Note that u_2 is obtained from u_1 by sliding the node with weight 7 all the way down, and u_3 is obtained from u_1 by sliding the node with weight 5 all the way down.

THEOREM 3.6. *Let $\delta : X \times X \rightarrow \mathbb{R}$ be a dissimilarity map. Let $S_0 = \{\delta_m\}$, and for each $i \geq 1$ define S_i to be the set of ultrametrics obtained from some $u \in S_{i-1}$ by sliding a mobile internal node of a resolution of the topology of u all the way down. Then*

1. $\bigcup_i S_i$ is a finite set, and
2. the tropical convex hull of $\bigcup_i S_i$ is the set of ultrametrics l^∞ -nearest to δ , and
3. every vertex of this tropical polytope has at most one mobile internal node.

Theorem 3.6 is a special case of Theorem 5.9, which will be proven later. We now illustrate Theorem 3.6 on an example.

Example 3.7. Let δ be the dissimilarity map given in Figure 3 on the left. We will make reference to ultrametrics u_1, \dots, u_5 which are shown in Figure 4. Using Theorem 3.1, we can see that $\delta_m = u_1$. Let v_1 be the internal node of u_1 's topology with weight 5. Then v_1 is mobile and sliding it all the way down yields u_3 . Let v_2 be the internal node of u_1 's topology with weight 7. Then v_2 is mobile and sliding it all the way down yields u_2 . The topology of u_4 is a resolution of the topology of u_2 . Letting v_3 be the internal node of u_4 's topology with weight 1, we can see that

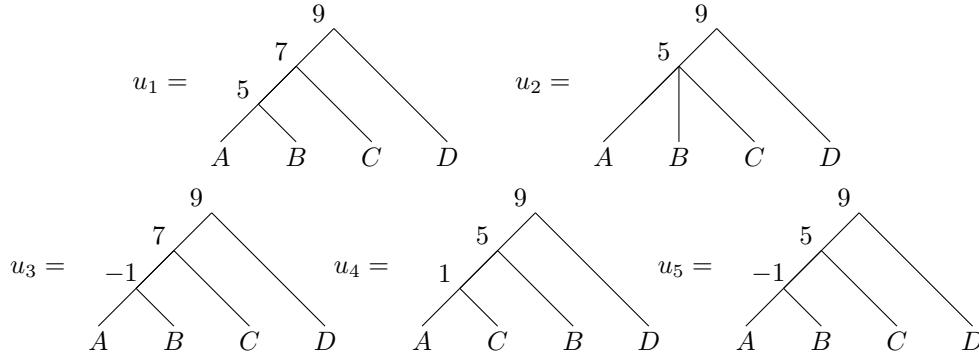


FIG. 4. By Theorem 3.6, the tropical convex hull of the ultrametrics above is the set of ultrametrics l^∞ -nearest to the dissimilarity map given on the left side of Figure 3.

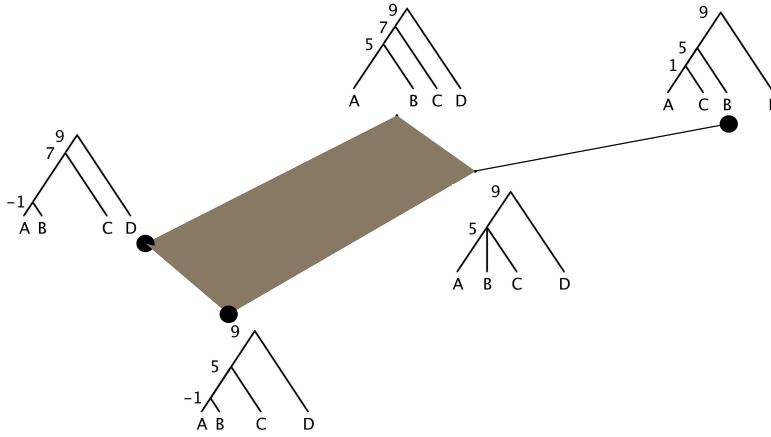


FIG. 5. The tropical polytope consisting of ultrametrics that are l^∞ -nearest to δ . The large points are the tropical vertices.

u_4 is obtained from u_2 by sliding v_3 all the way down. The topology of u_5 is also a resolution of the topology of u_2 . Letting v_4 be the internal node of u_5 's topology with weight -1 , we can see that u_5 is obtained from u_2 by sliding v_4 all the way down. Beyond v_3 and v_4 , no internal nodes of any resolution of the topology of u_2 are mobile. The only mobile node of u_3 is the node labeled 7 ; denote this v_5 . Then sliding v_5 all the way down gives us u_5 once again.

Using the notation of Theorem 3.6, we have $S_0 = \{u_1\}$, $S_1 = \{u_2, u_3\}$, and $S_2 = \{u_4, u_5\}$. Note that no internal nodes of u_4 and u_5 are mobile. Hence S_i is empty for all $i \geq 3$. Since u_1 and u_2 each have two mobile internal nodes, Theorem 3.6 implies that the set of ultrametrics l^∞ -nearest to δ is the tropical convex hull of $\{u_3, u_4, u_5\}$. This tropical polytope is contained in the three-dimensional affine subspace $\{\tilde{\delta} \in \mathbb{R}^{[4]} : \tilde{\delta}(1, 4) = \tilde{\delta}(2, 4) = \tilde{\delta}(3, 4) = 9\} \subset \mathbb{R}^{[4]}$. Therefore, we can visualize it as in Figure 5.

Theorem 3.6 implies that the elements of $\bigcup_i S_i$ that have at most one mobile internal vertex are a superset of the vertex set of the tropical polytope consisting of the ultrametrics l^∞ -nearest to a given dissimilarity map. A recent preprint of Yu

shows that this containment can be strict for dissimilarity maps with at least four elements [24]. A complete characterization of the vertices of this tropical polytope is still open.

We now describe a polynomial-time checkable condition that is equivalent to the condition that all ultrametrics l^∞ -nearest to a given dissimilarity map have the same topology.

THEOREM 3.8. *All ultrametrics l^∞ -nearest to a given dissimilarity map δ on n elements have the same topology if and only if the ultrametrics in $S_0 \cup S_1$ from Theorem 3.6 all have the same topology. This condition can be checked in $O(n^2)$ time.*

Proof. The second claim follows from the fact that Chepoi and Fichet's algorithm in Theorem 3.1 runs in $O(n^2)$ time (cf. [8]), and that δ_m has at most $n - 1$ internal vertices.

If we were to replace $S_0 \cup S_1$ with $\bigcup_i S_i$ in the statement of the theorem, then it would immediately follow from Proposition 5.11. So it suffices to show that if all trees in $S_0 \cup S_1$ have the same topology, then all trees in $\bigcup_i S_i$ do as well.

Assume that all ultrametrics in $S_0 \cup S_1$ have the same topology T . For the sake of contradiction, let $i \geq 2$ be minimal such that there exists $u \in S_i$ such that the topology of u is not T . This means that there exists some $u' \in S_{i-1}$ with topology T such that an internal node v of T is mobile in u' and sliding v all the way down in u' yields u . Let $\alpha_m, \alpha, \alpha' : T^\circ \rightarrow \mathbb{R}$ be internal edge weightings of T expressing δ_m, u , and u' , respectively (i.e., $\delta_m = \delta_{T, \alpha_m}$, $u = \delta_{T, \alpha}$, and $u' = \delta_{T, \alpha'}$). Since u is obtained from u' by sliding v all the way down, $\alpha(y) = \alpha'(y)$ unless $y = v$, in which case $\alpha(v) = \max\{\alpha'(y) : y \in \text{Dest}_T(v)\}$. Since $\max\{\alpha_m(y) : y \in \text{Dest}_T(v)\} \geq \max\{\alpha'(y) : y \in \text{Dest}_T(v)\}$, δ_m and u' both have topology T , and all internal nodes of T that are mobile for u' are also mobile for δ_m , we can slide v all the way down in δ_m to get an element of S_1 with the topology of u , contradicting that all elements of S_1 have topology T . \square

4. Bergman fans and nested sets. The goal of this section is to generalize the notion of tree topology for ultrametrics to elements of Bergman fans of arbitrary matroids. Nested sets of matroids, as described in [11], will play the role of rooted trees in this more general context. Familiarity with matroid connectivity is assumed; for this we refer the reader to [19, Chapter 4]. We begin by defining the Bergman fan of a matroid. Equivalent cryptomorphic definitions exist. The one we provide is due to Ardila (see [3, Proposition 2]).

DEFINITION 4.1. *Let \mathcal{M} be a matroid on ground set E . A vector $w \in \mathbb{R}^E$ is said to be an \mathcal{M} -ultrametric if for each circuit C of \mathcal{M} , the cardinality of $\{x \in C : w_x = \max_{y \in C} w_y\}$ is at least two. The set of \mathcal{M} -ultrametrics, denoted $\tilde{\mathcal{B}}(\mathcal{M})$, is called the Bergman fan of \mathcal{M} .*

As the name suggests, \mathcal{M} -ultrametrics generalize the ultrametrics discussed in section 3. In particular, letting K_n denote the complete graph on n vertices and $\mathcal{M}(G)$ denote the matroid underlying a graph G , the following theorem of Ardila and Klivans tells us that ultrametrics are $\mathcal{M}(K_n)$ -ultrametrics.

THEOREM 4.2 (see [4, Theorem 3]). *A dissimilarity map on the set $\{1, \dots, n\}$ is an ultrametric if and only if it is an $\mathcal{M}(K_n)$ -ultrametric.*

We would like to generalize Theorem 3.6, i.e., describe a generating set of the tropical polytope consisting of the \mathcal{M} -ultrametrics that are l^∞ -nearest to a given $x \in \mathbb{R}^E$. To do this, we need to generalize the notion of tree topology for arbitrary \mathcal{M} -

ultrametrics. Definition 4.3 below provides the desired generalization. It is essentially the special case of Definition 3.2 in [11] where the required lattice is the lattice of flats of a connected matroid \mathcal{M} and the required building set is the set of connected flats of \mathcal{M} .

DEFINITION 4.3. *Given a connected matroid \mathcal{M} on ground set E , a nested set of \mathcal{M} is a set \mathcal{S} of connected nonempty flats of \mathcal{M} such that $E \in \mathcal{S}$ and whenever $F_1, \dots, F_k \in \mathcal{S}$ are pairwise incomparable with respect to the containment order, the closure of $F_1 \cup \dots \cup F_k$ is disconnected. If \mathcal{M} is disconnected with connected components $\mathcal{M}_1, \dots, \mathcal{M}_k$, then a nested set of \mathcal{M} is the union of nested sets $\mathcal{S}_1, \dots, \mathcal{S}_k$ of $\mathcal{M}_1, \dots, \mathcal{M}_k$.*

Example 4.4. Let \mathcal{M} be the uniform matroid of rank three on ground set $\{a, b, c, d\}$. The nested sets of \mathcal{M} are the sets of any of the following forms:

$$\{\{a, b, c, d\}\} \quad \{\{x\}, \{a, b, c, d\}\} \quad \{\{x\}, \{y\}, \{a, b, c, d\}\},$$

where $x, y \in \{a, b, c, d\}$. If \mathcal{N} is the uniform matroid of rank two on ground set $\{e, f, g\}$, then the nested sets of the direct sum $\mathcal{M} \oplus \mathcal{N}$ are sets of the form

$$S \cup \{\{e, f, g\}\}, \quad S \cup \{\{x\}, \{e, f, g\}\},$$

where S is a nested set of \mathcal{M} and $x \in \{e, f, g\}$.

We remind the reader that $\mathbf{1}$ denotes the all-ones vector.

DEFINITION 4.5. *Let \mathcal{M} be a connected matroid on ground set E and let \mathcal{S} be a nested set of \mathcal{M} . For each $F \in \mathcal{S}$, let $v_F \in \mathbb{R}^E$ denote -1 times the characteristic vector of F . Define $K_{\mathcal{S}}$ to be the cone spanned by the v_F and $\pm \mathbf{1}$. The nested set fan of \mathcal{M} , denoted $\tilde{N}(\mathcal{M})$, is the polyhedral fan consisting of all the polyhedral cones $K_{\mathcal{S}}$ as \mathcal{S} ranges over all nested sets of \mathcal{M} . When \mathcal{M} is disconnected, we define its nested set fan to be the Cartesian product of the nested set fans of its connected components.*

Note that $\tilde{N}(\mathcal{M})$ is indeed a polyhedral fan since $K_{\mathcal{S}}$ is simplicial, and $K_{\mathcal{S}} \cap K_{\mathcal{S}'} = K_{\mathcal{S} \cap \mathcal{S}'}$. Also note that the lineality space of $\tilde{N}(\mathcal{M})$ is spanned by the characteristic vectors of the connected components of \mathcal{M} .

Definition 4.3 is slightly more restrictive than Definition 3.2 of [11]. Namely, a nested set in the sense of [11] does not require that each connected component of a matroid be present, nor that the entire ground set of a disconnected matroid not be present. For example, using \mathcal{M} and \mathcal{N} as in Example 4.4, Definition 3.2 of [11] would allow us to remove $\{a, b, c, d\}$ from any nested set of \mathcal{M} , or add $\{a, b, c, d, e, f, g\}$ to any nested set of $\mathcal{M} \oplus \mathcal{N}$. However, this is not an issue because these differences in definitions do not affect the nested set fan. Under the less restrictive definition, if $E \in \mathcal{S}$ for some nested set \mathcal{S} , then $K_{\mathcal{S}} = K_{\mathcal{S} \setminus \{E\}}$. We use this more restrictive definition to avoid this ambiguity when indexing cones of $\tilde{N}(\mathcal{M})$.

PROPOSITION 4.6. *The nested set fan $\tilde{N}(\mathcal{M})$ is a refinement of the Bergman fan $\tilde{\mathcal{B}}(\mathcal{M})$.*

Proof. When \mathcal{M} is connected, this follows from Theorem 4.1 in [11]. The rest of the proposition follows by noting that the Bergman fan of a disconnected matroid is the Cartesian product of the Bergman fans of its connected components. \square

In light of Proposition 4.6, we can make the following definition.

DEFINITION 4.7. Let w be an \mathcal{M} -ultrametric. Let $\mathcal{T}(w)$ denote the unique nested set of \mathcal{M} such that w lies in the relative interior of $K_{\mathcal{T}(w)}$. We call $\mathcal{T}(w)$ the topology of w .

Definition 4.7 might be unsettling to some readers since it appears to have nothing to do with topology in the usual sense. We use it because, as Proposition 4.8 below shows, it generalizes the notion of tree topology of an ultrametric in the phylogenetics sense.

PROPOSITION 4.8 (see [11, Remark 5.4]). Let w, u be $\mathcal{M}(K_n)$ -ultrametrics. Then the tree topologies of w, u are equal if and only if $\mathcal{T}(w) = \mathcal{T}(u)$.

The following proposition tells us that the topology of \mathcal{M} -ultrametrics is well-behaved with respect to tropical convexity.

PROPOSITION 4.9. The set of \mathcal{M} -ultrametrics that have a particular topology \mathcal{S} is tropically convex.

Proof. The lineality space of $\tilde{N}(\mathcal{M})$ contains $\mathbf{1}$ so the topology of an \mathcal{M} -ultrametric is preserved under tropical scalar multiplication. We now show that topology is preserved under tropical sums. To this end, let u, w be \mathcal{M} -ultrametrics that lie in the relative interior of the same cone $K_{\mathcal{S}}$. Modulo the lineality space of $\tilde{N}(\mathcal{M})$, $u = \sum_{F \in \mathcal{S}} \lambda_F^u v_F$ and $w = \sum_{F \in \mathcal{S}} \lambda_F^w v_F$, where the sums are taken over the flats in \mathcal{S} that are not connected components of \mathcal{M} , and λ_F^u, λ_F^w are all strictly positive. Then $(u \oplus w) = \sum_F (\max\{\lambda_F^u, \lambda_F^w\}) v_F$. So $(u \oplus w)$ also lies in the relative interior of $K_{\mathcal{S}}$ and so $\mathcal{T}(u \oplus w) = \mathcal{S}$. \square

Lemma 4.1 below implies that the Hasse diagram of the containment partial ordering on a nested set of a matroid \mathcal{M} is a forest with a tree for each connected component of \mathcal{M} . Proposition 4.11 implies that each \mathcal{M} -ultrametric can be displayed on this forest in the same way that an ultrametric can be displayed on its tree topology.

LEMMA 4.1. Let \mathcal{S} be a nested set of a matroid \mathcal{M} . Then for any pair $F, G \in \mathcal{S}$, $F \subseteq G$ or $G \subseteq F$ or $G \cap F = \emptyset$.

Proof. Assume F and G are connected flats of \mathcal{M} and that $F \cap G \neq \emptyset$. We will show that the closure K of $F \cup G$ is connected. It will then follow from the definition of a nested set that either $F \subseteq G$ or $G \subseteq F$. Let \sim be the relation on K where $a \sim b$ if and only if there exists a circuit $C \subseteq K$ containing both a and b . It suffices to show that there is only one equivalence class of K under \sim [19, Chapter 4.1]. Both F and G are connected, so each must lie entirely within one equivalence class. Moreover, their intersection is nontrivial so $F \cup G$ lies in a single equivalence class. Since K is the closure of $F \cup G$, each $e \in K \setminus (F \cup G)$ must also lie in this equivalence class. \square

Note that Lemma 4.1 implies that if \mathcal{S} is a nested set of a matroid \mathcal{M} , then for each e in the ground set of \mathcal{M} , there is a unique minimal flat in \mathcal{S} that contains e .

DEFINITION 4.10. Let \mathcal{M} be a matroid on ground set E and let \mathcal{S} be a nested set of \mathcal{M} . A function $\alpha : \mathcal{S} \rightarrow \mathbb{R}$ is said to be compatible with \mathcal{S} if $F \subseteq G$ implies $\alpha(F) \leq \alpha(G)$ for all $F, G \in \mathcal{S}$. For α compatible with \mathcal{S} , define $w^{\mathcal{S}, \alpha} \in \mathbb{R}^E$ by $w_e^{\mathcal{S}, \alpha} = \alpha(F)$, where F is the minimal flat in \mathcal{S} that contains e . If $w = w^{\mathcal{S}, \alpha}$, then we call the pair (\mathcal{S}, α) a nested set representation of w on \mathcal{S} .

PROPOSITION 4.11. Let \mathcal{M} and \mathcal{S} be as in Definition 4.10 and let $\alpha : \mathcal{S} \rightarrow \mathbb{R}$ be compatible with \mathcal{S} . Then $w^{\mathcal{S}, \alpha}$ is an \mathcal{M} -ultrametric. Every \mathcal{M} -ultrametric w has a unique nested set representation $w = w^{\mathcal{T}(w), \alpha}$ on its topology.

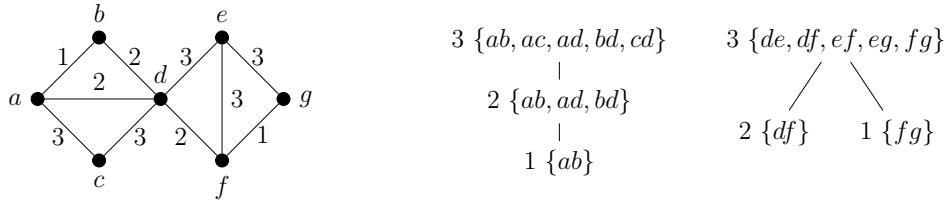


FIG. 6. An $\mathcal{M}(G)$ -ultrametric w , displayed as an edge weighting of G and using the $\alpha : \mathcal{T}(w) \rightarrow \mathbb{R}$ as described in Proposition 4.11.

Proof. It is sufficient to prove the proposition in the case where \mathcal{M} is connected, so assume \mathcal{M} is connected. We first show that $w^{\mathcal{S}, \alpha}$ is indeed an \mathcal{M} -ultrametric. Define $\lambda_E := -\alpha(E)$ and for each $F \in \mathcal{S} \setminus E$, define $\lambda_F := -\alpha(F) + \alpha(G)$, where G is the minimal element of \mathcal{S} strictly containing F (Lemma 4.1 implies that a unique such G exists). For each $F \in \mathcal{S}$, let v_F be as in Definition 4.5. Then $w^{\mathcal{S}, \alpha} = \sum_{F \in \mathcal{S}} \lambda_F v_F$. Since α is compatible with \mathcal{S} , $F \neq E$ implies that λ_F is nonnegative. This shows that $w^{\mathcal{S}, \alpha}$ is in the nested set fan. Proposition 4.6 then implies that $w^{\mathcal{S}, \alpha}$ is an \mathcal{M} -ultrametric.

Now let w be an arbitrary \mathcal{M} -ultrametric. By Proposition 4.6 and Definition 4.7, $w = \sum_{F \in \mathcal{T}(w)} \lambda_F v_F$ for some choice of coefficients λ_F satisfying $\lambda_F > 0$ when $F \neq E$. Set $\alpha(E) := -\lambda_E$, and for each $F \in \mathcal{T}(w) \setminus \{E\}$ inductively set $\alpha(F) := -\lambda_F + \alpha(G)$, where G is the minimal element of $\mathcal{T}(w)$ containing F . Note that $\alpha(F) < \alpha(F')$ whenever $F \subsetneq F'$ and that $w = w^{\mathcal{T}(w), \alpha}$. Uniqueness of α follows from the fact that this map from the λ_F 's to the $\alpha(F)$'s is invertible and that $\{v_F : F \in \mathcal{T}(w)\}$ is a linearly independent set. \square

Proposition 4.11 gives us a way to display an \mathcal{M} -ultrametric that generalizes the way we can display an ultrametric on its tree topology. Namely, if w is an \mathcal{M} -ultrametric and $\alpha : \mathcal{T}(w) \rightarrow \mathbb{R}$ is such that $w = w^{\mathcal{T}(w), \alpha}$, we can specify w by drawing the Hasse diagram for $\mathcal{T}(w)$ (which is a forest by Lemma 4.1) and labeling each $F \in \mathcal{T}(w)$ with $\alpha(F)$. We now show this in an example.

Example 4.12. The left side of Figure 6 displays an $\mathcal{M}(G)$ -ultrametric w as an edge weighting of the graph G . On its right is $\mathcal{T}(w)$ where each flat $F \in \mathcal{T}(w)$ is labeled by $\alpha(F)$, where $\alpha : \mathcal{T}(w) \rightarrow \mathbb{R}$ satisfies $w = w^{\mathcal{T}(w), \alpha}$. Since the graph G is not biconnected, the matroid $\mathcal{M}(G)$ is disconnected and so $\mathcal{T}(w)$ is disconnected.

We now generalize the concepts of polytomy and resolution from rooted trees representing ultrametrics to nested sets representing \mathcal{M} -ultrametrics.

DEFINITION 4.13. Let \mathcal{M} be a matroid on ground set E and let \mathcal{S} be a nested set of \mathcal{M} . A polytomy of \mathcal{S} is an element $F \in \mathcal{S}$ such that $\text{rank}(F / \bigcup_{G \in \mathcal{S}} G) > 1$, where the union is taken over all $G \in \mathcal{S}$ such that $G \subsetneq F$. A resolution of \mathcal{S} is another nested set \mathcal{S}' without polytomies such that $\mathcal{S} \subseteq \mathcal{S}'$.

If $\mathcal{T}(w)$ has a polytomy, then the nested set representation of w is not unique. In particular, w can be represented on any nested set \mathcal{S} that is a resolution of $\mathcal{T}(w)$.

Example 4.14. On the left side of Figure 7, we see a nested set \mathcal{S}_1 of the matroid $\mathcal{M}(K_4)$ underlying the complete graph on vertex set $\{a, b, c, d\}$. Since $\mathcal{M}(K_4)/\{ab\}$ is a matroid of rank 2, the set $\{ab, ac, ad, bc, bd, cd\}$ is a polytomy of \mathcal{S}_1 . To its right are the two possible resolutions \mathcal{S}_2 and \mathcal{S}_3 . Each \mathcal{S}_i is shown with a compatible $\alpha_i : \mathcal{S}_i \rightarrow \mathbb{R}$, thus giving us the $\mathcal{M}(K_4)$ -ultrametrics $w^{\mathcal{S}_i, \alpha_i}$. Note that $w^{\mathcal{S}_1, \alpha_1} = w^{\mathcal{S}_2, \alpha_2} = w^{\mathcal{S}_3, \alpha_3}$ and that the topology of this $\mathcal{M}(K_4)$ -ultrametric is \mathcal{S}_1 .

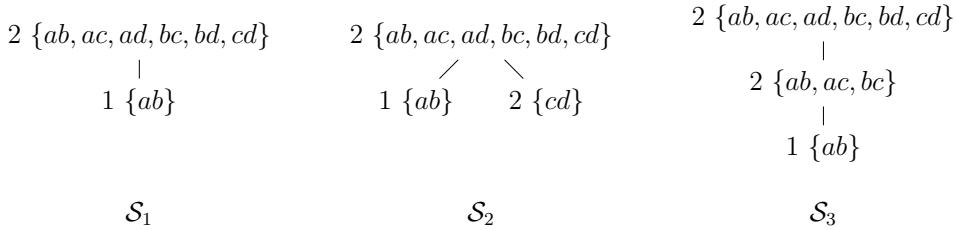


FIG. 7. A nested set \mathcal{S}_1 of the complete graph on vertex set $\{a, b, c, d\}$ with a polytomy and its two resolutions \mathcal{S}_2 and \mathcal{S}_3 . The weightings on each nested set all give rise to the same $\mathcal{M}(K_4)$ -ultrametric.

5. L-infinity optimization to Bergman fans of matroids. The first important result of this section is Proposition 5.2, which says that the subset of a Bergman fan $\tilde{\mathcal{B}}(\mathcal{M}) \subseteq \mathbb{R}^E$ consisting of all points l^∞ -nearest to a given $x \in \mathbb{R}^E$ is a tropical polytope. The main result of this section is Theorem 5.9, which describes a generating set of this tropical polytope. In light of Proposition 5.12, Theorem 5.9 is applicable to ultrametric reconstruction in cases where the data consist only of a subset of all pairwise distances. We begin by recalling a result of Ardila, establishing a connection between ultrametric reconstruction and tropical convexity.

PROPOSITION 5.1 (see [3, Proposition 4.1]). *The Bergman fan $\tilde{\mathcal{B}}(\mathcal{M})$ is a tropical polyhedral cone.*

We introduce some notation. Given points $x, y \in \mathbb{R}^E$ and a set $S \subseteq \mathbb{R}^E$, we denote the l^∞ -distance between x and y by $d(x, y)$ and $\inf_{y \in S} d(x, y)$ by $d(x, S)$. Given some $x \in \mathbb{R}^E$, we define the subset of $\tilde{\mathcal{B}}(\mathcal{M})$ consisting of the \mathcal{M} -ultrametrics that are l^∞ -nearest to x by $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. That is, $C(x, \tilde{\mathcal{B}}(\mathcal{M})) = \{w \in \tilde{\mathcal{B}}(\mathcal{M}) : d(x, w) = d(x, \tilde{\mathcal{B}}(\mathcal{M}))\}$. The next proposition says that this set is a tropical polytope.

PROPOSITION 5.2. *If \mathcal{M} is a matroid on ground set E and $x \in \mathbb{R}^E$, then the subset of the Bergman fan of \mathcal{M} consisting of elements l^∞ -nearest to x is a tropical polytope.*

Proof. Let C denote the cube of side length $d(x, \tilde{\mathcal{B}}(\mathcal{M}))$ centered at x . Therefore we can express $C(x, \tilde{\mathcal{B}}(\mathcal{M})) = \tilde{\mathcal{B}}(\mathcal{M}) \cap C$. Proposition 5.1 tells us that $\tilde{\mathcal{B}}(\mathcal{M})$ is a tropical polyhedron and C is clearly a tropical polytope. Their intersection is again a tropical polyhedron. Since it is bounded it is by definition a tropical polytope. \square

Much of the remainder of this section is devoted to describing the set of tropical vertices of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. Now we recall the concept of a subdominant \mathcal{M} -ultrametric, the existence of which was proven by Ardila in [3].

DEFINITION 5.3 (see [3]). *Let \mathcal{M} be a matroid on ground set E and let $x \in \mathbb{R}^E$. Let $x^{\mathcal{M}}$ denote the unique coordinatewise maximum \mathcal{M} -ultrametric which is coordinatewise at most x . We call $x^{\mathcal{M}}$ the subdominant \mathcal{M} -ultrametric of x .*

Given some $x \in \mathbb{R}^E$, Ardila shows how the first three steps of the algorithm from Theorem 3.1 can be extended to compute the subdominant \mathcal{M} -ultrametric of x . Then the subdominant ultrametric can be shifted to obtain an l^∞ -nearest ultrametric that is coordinatewise maximal among all l^∞ -nearest ultrametrics.

LEMMA 5.1. *Let \mathcal{M} be a matroid on ground set E , $x \in \mathbb{R}^E$, and $\delta = \frac{1}{2}d(x, x^{\mathcal{M}})$. Then*

1. *The l^∞ -distance from x to $\tilde{\mathcal{B}}(\mathcal{M})$ is δ ,*

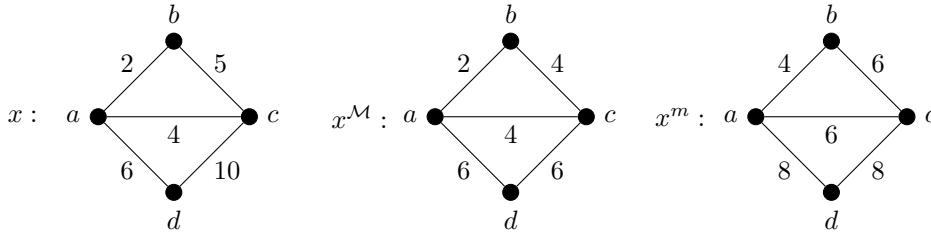


FIG. 8. An element $x \in \mathbb{R}^E$ alongside its subdominant $\mathcal{M}(G)$ -ultrametric x^M and the l^∞ -nearest \mathcal{M} -ultrametric x^m .

2. $x^M + \delta \cdot \mathbf{1}$ is an \mathcal{M} -ultrametric, l^∞ -nearest to x ,
3. $x^M + \delta \cdot \mathbf{1}$ is maximal among \mathcal{M} -ultrametrics l^∞ -nearest to x .

Proof. The existence of $x^M + \delta \cdot \mathbf{1}$ shows that $d(x, \tilde{\mathcal{B}}(\mathcal{M})) \leq \delta$. Suppose there exists $w \in \tilde{\mathcal{B}}(\mathcal{M})$ such that $d(x, w) < \delta$. Then $w - d(x, w) \cdot \mathbf{1}$ is coordinatewise at most x . There exists $e \in E$ such that $x_e - x_e^M = 2\delta$ and so $x_e^M < w_e - d(x, w)$. Thus, $w - d(x, w) \cdot \mathbf{1}$ is an ultrametric coordinatewise at most x but not coordinatewise at most x^M , contradicting that x^M is the subdominant \mathcal{M} -ultrametric. So (1) is proven and (2) immediately follows.

If (3) were false and there existed some \mathcal{M} -ultrametric $y \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ such that $y \geq x^M + \delta \cdot \mathbf{1}$ with inequality somewhere, then $y - \delta \cdot \mathbf{1}$ would not be coordinatewise at most x^M . However, it would be coordinatewise at most x , thus contradicting that x^M is the subdominant \mathcal{M} -ultrametric. \square

DEFINITION 5.4. Given $x \in \mathbb{R}^E$, we denote by x^m the l^∞ -nearest ultrametric $x^M + d(x, \tilde{\mathcal{B}}(\mathcal{M})) \cdot \mathbf{1}$ and call it the maximal closest \mathcal{M} -ultrametric to x .

Example 5.5. Let G be the graph displayed in Figure 8 and denote its edge set by E . Let $x \in \mathbb{R}^E$ be as on the left of Figure 8. Then the subdominant $\mathcal{M}(G)$ -ultrametric x^M and its translation giving the l^∞ -nearest $\mathcal{M}(G)$ -ultrametric x^m are shown to the right.

Definition 5.6 below introduces a way to decrease certain coordinates of an \mathcal{M} -ultrametric w that is l^∞ -nearest to a given $x \in \mathbb{R}^E$ to produce another \mathcal{M} -ultrametric l^∞ -nearest to x . The coordinates of w that can be decreased are determined by what we will call *mobile flats*. We call the process of decreasing these coordinates *sliding mobile flats (all the way) down*. Theorem 5.9 uses these concepts to describe a generating set of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$.

DEFINITION 5.6. Let \mathcal{M} be a matroid on ground set E . Let $x \in \mathbb{R}^E$ and let $w \in \tilde{\mathcal{B}}(\mathcal{M})$ be l^∞ -nearest to x . Let \mathcal{S} be a resolution of $\mathcal{T}(w)$ and $\alpha : \mathcal{S} \rightarrow \mathbb{R}$ be compatible with \mathcal{S} satisfying $w = w^{\mathcal{S}, \alpha}$. We say that $F \in \mathcal{S}$ is mobile if there exists an \mathcal{M} -ultrametric $\hat{w} \neq w$ expressible as $\hat{w} = w^{\mathcal{S}, \hat{\alpha}}$ with $\hat{\alpha}$ compatible with \mathcal{S} such that

1. \hat{w} is also nearest to x in the l^∞ -norm,
2. $\hat{\alpha}(G) = \alpha(G)$ for all $G \neq F$, and
3. $\hat{\alpha}(F) < \alpha(F)$.

In this case, we say that \hat{w} is obtained from w by sliding F down. If moreover F is no longer mobile in $w^{\mathcal{S}, \hat{\alpha}}$, i.e., if $\hat{\alpha}(F) = \max\{\alpha(G) : G \in \mathcal{S} \text{ and } G \subsetneq F\}$ or $\hat{\alpha}(F)$ is the minimum value such that $w^{\mathcal{S}, \hat{\alpha}}$ is l^∞ -nearest to x , then we say that \hat{w} is obtained from w by sliding F all the way down.

Remark 5.7. Given some $x \in \mathbb{R}^E$ and some $w^{S,\alpha}$ that is l^∞ -nearest to x , one can determine that a given $F \in \mathcal{S}$ is mobile by decreasing $\alpha(F)$ by some small $\varepsilon > 0$ and seeing that the resulting \mathcal{M} -ultrametric is still l^∞ -nearest to x .

Remark 5.8. If \mathcal{S} is a resolution of $\mathcal{T}(w)$ and $F \in \mathcal{S} \setminus \mathcal{T}(w)$ is mobile, then F is contained in a polytomy of $\mathcal{T}(w)$ and all elements of \mathcal{S} covered by F are also in $\mathcal{T}(w)$.

THEOREM 5.9. *Let \mathcal{M} be a matroid on ground set E and let $x \in \mathbb{R}^E$. Define $S_0 := \{x^m\}$ and for each $i \geq 1$, define S_i to be the set of \mathcal{M} -ultrametrics obtained from some $w \in S_{i-1}$ by sliding a mobile flat in a resolution of $\mathcal{T}(w)$ all the way down. Then*

1. $\bigcup_i S_i$ is a finite set,
2. the tropical convex hull of $\bigcup_i S_i$ is $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$, and
3. each tropical vertex v of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ has at most one mobile flat across all resolutions of $\mathcal{T}(v)$.

Proof. We first prove that $\bigcup_i S_i$ is a finite set. Let $w \in S_i$ for some $i \geq 0$. Then each coordinate w_e is either x_f^m or $x_f^m - d(x, \tilde{\mathcal{B}}(\mathcal{M}))$ for some $f \in E$, not necessarily equal to e . So as w ranges over $\bigcup_i S_i$, there are only finitely many values that each w_e can take and so $\bigcup_i S_i$ is a finite set.

We now prove that each tropical vertex has at most one mobile flat. Let $v \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. Let α be such that $v = w^{\mathcal{T}(v), \alpha}$ (recall Definition 4.10). If \mathcal{S}_1 and \mathcal{S}_2 are resolutions of $\mathcal{T}(v)$ and $F_i \in \mathcal{S}_i$ is mobile, then there exist $\alpha_i : \mathcal{S}_i \rightarrow \mathbb{R}$ compatible with \mathcal{S}_i such that $w^{S_i, \alpha_i} \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$, and $w_e^{S_i, \alpha_i} = v_e - \varepsilon$ for a fixed small $\varepsilon > 0$ whenever $e \in F_i \setminus \bigcup_F F$ where the union is taken over all $F \in \mathcal{S}_i$ such that $F \subsetneq F_i$, and $w_e^{S_i, \alpha_i} = v_e$ for all other $e \in E$. We claim that $w_e^{S_1, \alpha_1} \neq v_e$ implies $w_e^{S_2, \alpha_2} = v_e$. When F_1 and F_2 are disjoint, the claim is obvious. When F_1 and F_2 are not disjoint, they must be subsets of the same polytomy $F \in \mathcal{T}(v)$. Let U be the union of all the flats covered by F in $\mathcal{T}(v)$. Then $U \subseteq F_1 \cap F_2$. Moreover, $U = F_1 \cap F_2$ because if $e \in F_1 \cap F_2 \setminus U$, then $\text{rank}(F_1/U) > \text{rank}((F_1 \cap F_2)/U) \geq \text{rank}((U \cup \{e\})/U) = 1$. In light of Remark 5.8, this is a contradiction because then F_1 would be a polytomy in \mathcal{S}_1 . The claim then follows because $w_e^{S_1, \alpha_1} \neq v_e$ if and only if $e \in F_i \setminus U$. Now we have $v = w^{S_1, \alpha_1} \oplus w^{S_2, \alpha_2}$ and so v is not a tropical vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$.

Now we prove that the tropical convex hull of $\bigcup_i S_i$ is $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ by showing that each vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ is a member of some S_i . So let v be a tropical vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. We construct a sequence $x^m = w^0 \geq w^1 \geq \dots \geq v$ such that $w^i \in S_i$ and $w^i \neq w^{i+1}$. Since $\bigcup_i S_i$ is finite, this sequence must eventually terminate and so the final w^i is equal to v . Assuming w^i has been constructed and satisfies $w^i \geq v$ and $w^i \neq v$, we show how to construct w^{i+1} satisfying $w^i \geq w^{i+1} \geq v$ and $w^{i+1} \neq w^i$.

First assume $\mathcal{T}(w^i) \subseteq \mathcal{T}(v)$. Let \mathcal{S} be a resolution of $\mathcal{T}(v)$. Then \mathcal{S} is also a resolution of $\mathcal{T}(w^i)$. Let α_{w^i}, α_v be such that $w^i = w^{S, \alpha_{w^i}}$ and $v = w^{S, \alpha_v}$. Let $F \in \mathcal{S}$ be a minimal element such that $\alpha_v(F) < \alpha_{w^i}(F)$. We can choose such an F to be non-mobile in v . Otherwise, the unique mobile flat in \mathcal{S} of w^i would be F , which would also be the unique mobile flat of \mathcal{S} in v and so for all $G \in \mathcal{S} \setminus \{F\}$, $\alpha_{w^i}(G) = \alpha_v(G)$. Since F is mobile in v , there exists some $\alpha : \mathcal{S} \rightarrow \mathbb{R}$ compatible with \mathcal{S} such that $\alpha(G) = \alpha_v(G)$ for $G \neq F$ but $\alpha(F) < \alpha_v(F)$ and $w^{S, \alpha} \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. This contradicts v being a vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ because $v = (\alpha_v(F) - \alpha_{w^i}(F)) \odot w^i \oplus w^{S, \alpha}$. So we can choose F to be mobile in w^i and not in v . Define $\alpha_{w^{i+1}} : \mathcal{S} \rightarrow \mathbb{R}$ by $\alpha_{w^{i+1}}(G) = \alpha_{w^i}(G)$ when $G \neq F$ and $\alpha_{w^{i+1}}(F) = \alpha_v(F)$. Define $w^{i+1} := w^{S, \alpha_{w^{i+1}}}$. Then $w^i \geq w^{i+1} \geq v$ and w^{i+1} is obtained from w^i by sliding F down. Since F was chosen to be minimal such that $\alpha_v(F) < \alpha_{w^i}(F)$ and $\alpha_{w^{i+1}}(G) = \alpha_{w^i}(G)$ when $G \neq F$, nonmobility of F

in v implies nonmobility of F in w^{i+1} . Hence w^{i+1} is obtained from w^i by sliding F all the way down and so $w^{i+1} \in S_{i+1}$.

Now assume $\mathcal{T}(w^i) \not\subseteq \mathcal{T}(v)$. Denote $v^t := (t \odot w^i) \oplus v$. Since $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ is tropically convex, $v^t \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ whenever $t < 0$. Let $t_0 < 0$ maximum such that $\mathcal{T}(w^i) \setminus \mathcal{T}(v^{t_0})$ is nonempty and let $G \in \mathcal{T}(w^i) \setminus \mathcal{T}(v^{t_0})$ be maximal. Note that for small $\varepsilon > 0$, $\mathcal{T}(w^i) \subseteq \mathcal{T}(v^{t_0+\varepsilon})$ and the minimal $H \in \mathcal{T}(v^{t_0+\varepsilon})$ that strictly contains G is also a member of $\mathcal{T}(v^{t_0})$. Let \mathcal{S} be a resolution of $\mathcal{T}(v^{t_0+\varepsilon})$ and therefore also a resolution of $\mathcal{T}(w^i)$. Choose $K \in \mathcal{S}$ such that $G \subsetneq K \subseteq H$ and let w^{i+1} be the result of sliding K all the way down in w^i . Then $w^{i+1} \in S_{i+1}$ and $w^i \geq w^{i+1} \geq v^t \geq v$. \square

As with Theorem 3.6, the set of \mathcal{M} -ultrametrics specified by Theorem 5.9.3 is, in general, a strict superset of the set of tropical vertices; see [24].

Example 5.10. Let G be the graph from Figure 8 and let x be the edge-weighting displayed. We now describe how to use Theorem 5.9 to obtain a generating set of the tropical polytope consisting of the $\mathcal{M}(G)$ -ultrametrics that are l^∞ -nearest to x . Figure 9 shows the $\mathcal{M}(G)$ -ultrametrics in each nonempty S_i , displayed on their topologies. The mobile flats of the unique element x^m of S_0 are $\{ab, ac, bc\}$ and $\{ab\}$. Sliding $\{ab\}$ all the way down yields the element of S_1 shown on the left, and sliding $\{ab, ac, bc\}$ all the way down yields the element of S_1 shown on the right. The only mobile flat of the element of S_1 shown on the left is $\{ab, ac, bc\}$. Sliding this all the way down yields the leftmost element displayed in S_2 . The element of S_1 shown on the right has $\{ab, ac, bc\}$ as a polytomy. There are three possible resolutions, the first obtained by adding the flat $\{ab\}$, the second by adding $\{ac\}$, and the third by adding $\{bc\}$. Each such flat is mobile, and the elements of S_2 obtained by sliding each all the way down are shown second, third, and fourth from the left in S_2 . Continuing in this way yields the elements shown in S_3 and S_4 . Note that there are no mobile flats in any element of S_4 so S_i is empty for $i \geq 5$. The leftmost element of S_2 also appears in S_3 and S_4 . A subset of $\bigcup_i S_i$ whose tropical convex hull is $C(x, \tilde{\mathcal{B}}(\mathcal{M}(G)))$ is shown in red. Note that we've omitted elements with two or more mobile flats, as well as repeated elements.

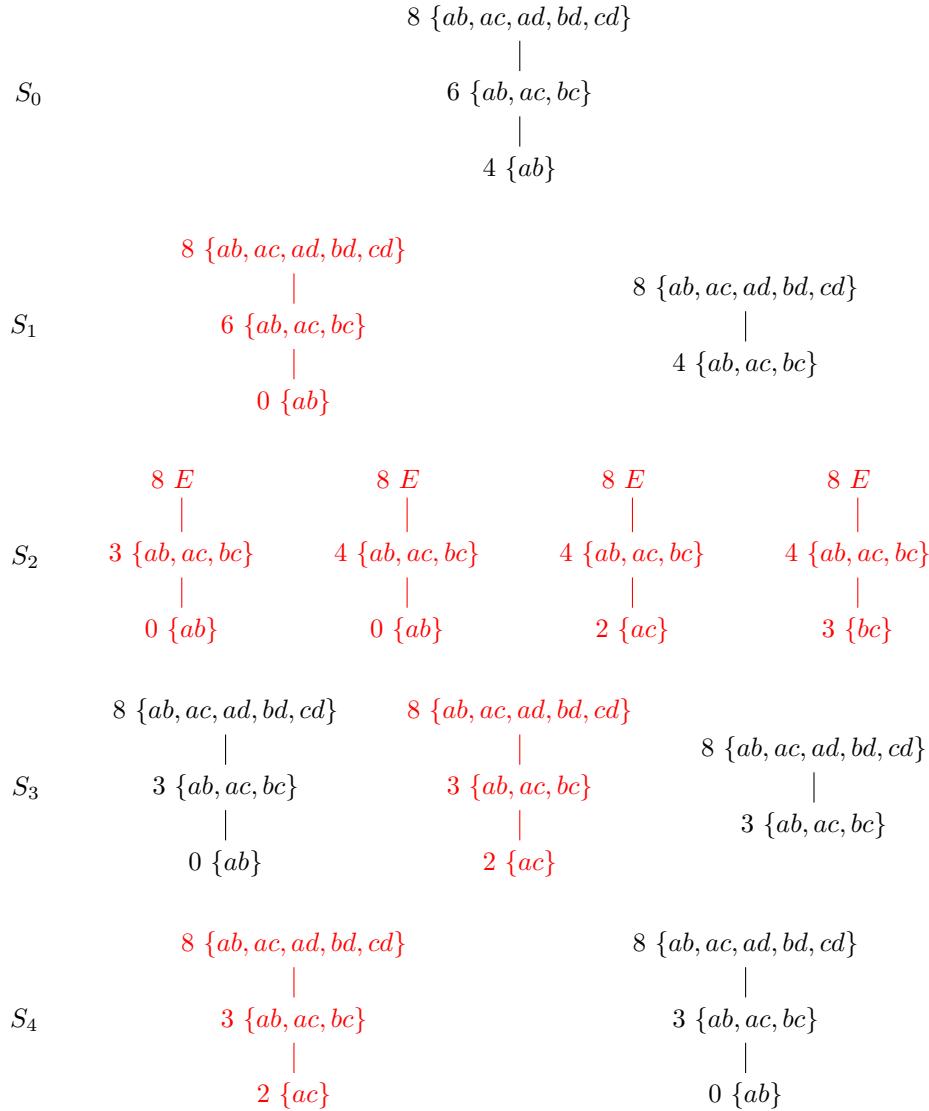
The following proposition tells us that if all the \mathcal{M} -ultrametrics in the generating set of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ indicated by Theorem 5.9 have the same topology, then all elements of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ have the same topology.

PROPOSITION 5.11. *Let \mathcal{M} be a matroid on ground set E and let $x \in \mathbb{R}^E$. Then the set of all \mathcal{M} -ultrametrics that are l^∞ -nearest to x have the same topology if and only if all tropical vertices of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ have the same topology.*

Proof. This follows immediately from Proposition 4.9. \square

When $\mathcal{M} := \mathcal{M}(G)$ is the matroid underlying some graph G , then Theorem 5.9 has potential use for phylogenetics even when G is not the complete graph. In particular, it sometimes happens that only a subset of the pairwise distances between n species can be computed within a reasonable budget. Then one may ask the question of which partial ultrametrics are l^∞ -nearest to the observed distances. Assuming that the observed distances correspond to the edge set E of a graph G , the following proposition tells us that the above question is equivalent to, given some partial dissimilarity map $x \in \mathbb{R}^E$, which $\mathcal{M}(G)$ -ultrametrics are l^∞ -nearest to x ?

PROPOSITION 5.12. *Let $E \subsetneq \binom{[n]}{2}$, let G be the graph with vertex set $[n]$ and edge set E , and let $x \in \mathbb{R}^E$. Then we may extend x to some ultrametric $x' \in \tilde{\mathcal{B}}(\mathcal{M}(K_n))$ if and only if x is an $\mathcal{M}(G)$ -ultrametric.*

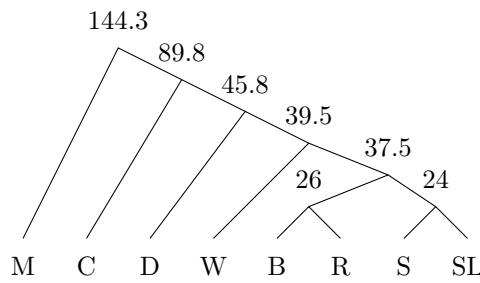
FIG. 9. The nonempty S_i 's from Theorem 5.9 for the edge weighting of the graph G in Figure 8.

Proof. First let x be an $\mathcal{M}(G)$ -ultrametric. Let $e \in \binom{[n]}{2} \setminus E$. Let G' be the graph obtained by adding e to G . We can extend x to an $\mathcal{M}(G')$ -ultrametric x' by setting x'_e to be the maximum of all the minimum edge weights appearing in some cocircuit of $\mathcal{M}(G')$. That this is indeed an $\mathcal{M}(G')$ -ultrametric follows from Ardila's characterization of \mathcal{M} -ultrametrics in terms of \mathcal{M} 's cocircuits [3]. By induction it follows that x may be completed to an $\mathcal{M}(K_n)$ -ultrametric.

Now let $x \in \mathbb{R}^E$ and assume that there exists some $x' \in \tilde{\mathcal{B}}(\mathcal{M}(K_n))$ such that $x_e = x'_e$ for each $e \in E$. Since x' is an $\mathcal{M}(K_n)$ -ultrametric, each $e \in E$ appears in some x' -minimal basis of $\mathcal{M}(K_n)$. As $x'_e = x_e$ for each $e \in E$, it follows that each $e \in E$ appears in some x -minimal basis of $\mathcal{M}(G)$. Therefore x is an $\mathcal{M}(G)$ -ultrametric. \square

TABLE 1
Pairwise immunological distances between eight species.

	Dog	Bear	Raccoon	Weasel	Seal	Sea Lion	Cat	Monkey
Dog	0	32	48	51	50	48	98	148
Bear	32	0	26	34	29	33	84	136
Raccoon	48	26	0	42	44	44	92	152
Weasel	51	34	42	0	44	38	86	142
Seal	50	29	44	44	0	24	89	142
Sea lion	48	33	44	38	24	0	90	142
Cat	98	84	92	86	89	90	0	148
Monkey	148	136	152	142	142	142	148	0

FIG. 10. *Ultrametric returned by the UPGMA algorithm.*

6. Example on a biological dataset. Now that we understand how uniqueness of the l^∞ -nearest ultrametric can fail to be unique, one might wonder if this is likely to happen for a dissimilarity map not explicitly constructed to break uniqueness. To this end, we now apply Theorem 3.6 to the dataset displayed in Table 1. It consists of pairwise immunological distances between the species dog, bear, raccoon, weasel, seal, sea lion, cat, and monkey that were obtained by Sarich in [20]. It is used in the textbook [12] to illustrate the UPGMA and neighbor joining algorithms, which are two other distance-based methods for phylogenetic reconstruction.

Theorem 5.9 suggests an algorithm for computing a generating set of the set of ultrametrics l^∞ -nearest to a given dissimilarity map. This consists of computing all nonempty S_i 's and removing all ultrametrics that have more than one mobile internal node. Applying this to the dataset in Table 1 gives us the twenty ultrametrics displayed in Figure 11. Four different tree topologies appear; for example, note that the topologies of the first, second, eighth, and fourteenth ultrametrics in the row-major order of Figure 11 are distinct.

The UPGMA algorithm always returns an ultrametric. Figure 10 shows the ultrametric computed by the UPGMA algorithm when applied to the dataset given in Table 1 (see [12, pp. 162–166]). No ultrametric sharing the topology of the ultrametric shown in Figure 10 will be l^∞ -nearest to the data. To see this, note that among the ultrametrics displayed in Figure 11, the distance between weasel and seal is 42 or 43, and that the distance between dog and seal is always 41. Since the set of l^∞ -nearest ultrametrics is tropically convex, any ultrametric l^∞ -nearest to the data will have the distance between weasel and seal strictly greater than the distance between dog and seal. However, the opposite relation will be true in any ultrametric whose topology is the tree displayed in Figure 10.

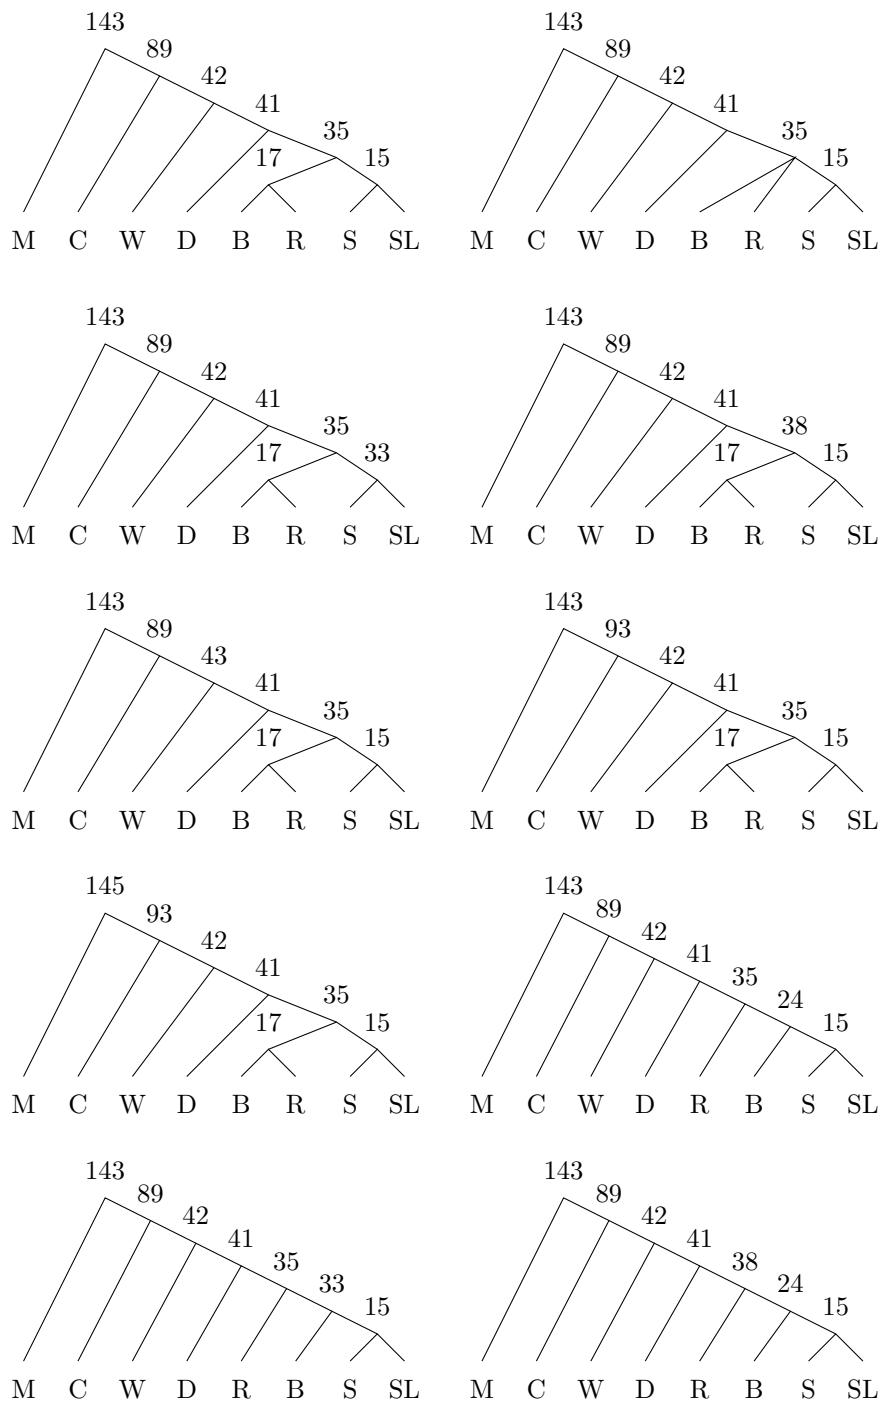
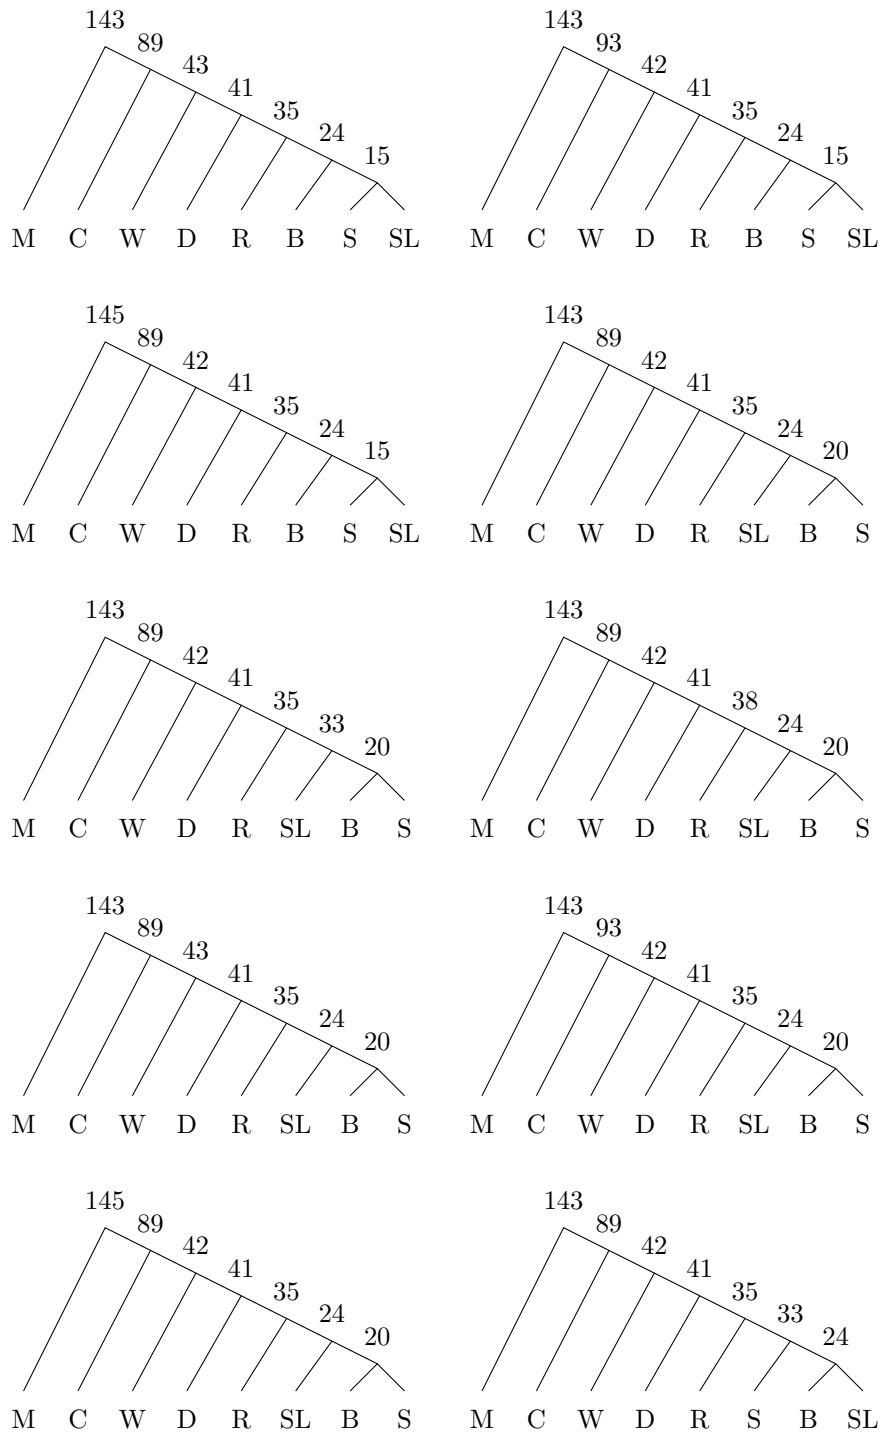


FIG. 11. A set of ultrametrics whose tropical convex hull is the set of ultrametrics l^∞ -nearest to the dataset in Table 1. Continued on next page.

FIG. 11. *Continued from previous page.*

Acknowledgments. The author is grateful to Colby Long and Seth Sullivant for many helpful conversations and for feedback on early drafts, and to several anonymous referees who provided thoughtful feedback that greatly improved this manuscript.

REFERENCES

- [1] M. AKIAN, S. GAUBERT, V. NIȚICĂ, AND I. SINGER, *Best approximation in max-plus semimodules*, Linear Algebra Appl., 435 (2011), pp. 3261–3296.
- [2] X. ALLAMIGEON, S. GAUBERT, AND É. GOUBAULT, *The tropical double description method*, in STACS 2010: 27th International Symposium on Theoretical Aspects of Computer Science, Vol. 5, Schloss Dagstuhl, Wadern, Germany, 2010, pp. 47–58.
- [3] F. ARDILA, *Subdominant matroid ultrametrics*, Ann. Comb., 8 (2004), pp. 379–389.
- [4] F. ARDILA AND C. J. KLIVANS, *The Bergman complex of a matroid and phylogenetic trees*, J. Combin. Theory Ser. B, 96 (2006), pp. 38–49.
- [5] D. I. BERNSTEIN AND C. LONG, *L -infinity optimization to linear spaces and phylogenetic trees*, SIAM J. Discrete Math., 31 (2017), pp. 875–889.
- [6] L. J. BILLERA, S. P. HOLMES, AND K. VOGTMANN, *Geometry of the space of phylogenetic trees*, Adv. Appl. Math., 27 (2001), pp. 733–767.
- [7] P. BUTKOVIČ, H. SCHNEIDER, AND S. SERGEEVC, *Generators, extremals and bases of max cones*, Linear Algebra Appl., 421 (2007), pp. 394–406.
- [8] V. CHEPOI AND B. FICHET, *l_∞ -approximation via subdominants*, J. Math. Psych., 44 (2000), pp. 600–616.
- [9] M. DEVELIN AND B. STURMFELS, *Tropical convexity*, Doc. Math., 9 (2004), pp. 1–27.
- [10] M. FARACH, S. KANNAN, AND T. WARNOW, *A robust model for finding optimal evolutionary trees*, in Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, ACM, New York, 1993, pp. 137–145.
- [11] E. M. FEICHTNER AND B. STURMFELS, *Matroid polytopes, nested sets and Bergman fans*, Port. Math., 62 (2005), pp. 437–468.
- [12] J. FELSENSTEIN, *Inferring Phylogenies*, Vol. 2, Sinauer Associates, Sunderland, MA, 2004.
- [13] S. GAUBERT, *Théorie des systèmes linéaires dans les dioïdes*, Ph.D thesis, ENMP, Paris, 1992.
- [14] S. GAUBERT AND R. D. KATZ, *The Minkowski theorem for max-plus convex sets*, Linear Algebra Appl., 421 (2007), pp. 356–369.
- [15] S. GAUBERT AND R. D. KATZ, *Minimal half-spaces and external representation of tropical polyhedra*, J. Algebraic Combin., 33 (2011), pp. 325–348.
- [16] B. LIN, A. MONOD, AND R. YOSHIDA, *Tropical Foundations for Probability & Statistics on Phylogenetic Tree Space*, preprint, <https://arxiv.org/abs/1805.12400>, 2018.
- [17] B. LIN, B. STURMFELS, X. TANG, AND R. YOSHIDA, *Convexity in tree spaces*, SIAM J. Discrete Math., 31 (2017), pp. 2015–2038.
- [18] B. LIN AND R. YOSHIDA, *Tropical Fermat–Weber points*, SIAM J. Discrete Math., 32 (2018), pp. 1229–1245.
- [19] J. OXLEY, *Matroid Theory*, 2nd ed., Oxford University Press, Oxford, 2011.
- [20] V. M. SARICH, *Pinniped phylogeny*, System. Biol., 18 (1969), pp. 416–422.
- [21] C. SEMPLE AND M. STEEL, *Phylogenetics*, Oxford University Press, Oxford, 2003.
- [22] D. SPEYER AND B. STURMFELS, *The tropical Grassmannian*, Adv. Geom., 4 (2004), pp. 389–411.
- [23] R. YOSHIDA, L. ZHANG, AND X. ZHANG, *Tropical principal component analysis and its application to phylogenetics*, Bull. Math. Biol., 81 (2019), pp. 568–597.
- [24] L. YU, *Extreme rays of the l^∞ -nearest ultrametric tropical polytope*, Linear Algebra Appl., 587 (2020), pp. 23–44.