Polyhedral computational geometry for averaging metric phylogenetic trees

Ezra Miller   Megan Owen   J. Scott Provan

February 18, 2014

Abstract

This paper investigates the computational geometry relevant to calculations of the Fréchet mean and variance for probability distributions on the phylogenetic tree space of Billera, Holmes and Vogtmann, using the theory of probability measures on spaces of nonpositive curvature developed by Sturm. We show that the combinatorics of geodesics with a specified fixed endpoint in tree space are determined by the location of the varying endpoint in a certain polyhedral subdivision of tree space. The variance function associated to a finite subset of tree space has a fixed C^∞ algebraic formula within each cell of the corresponding subdivision, and is continuously differentiable in the interior of each orthant of tree space. We use this subdivision to establish two iterative methods for producing sequences that converge to the Fréchet mean: one based on Sturm's Law of Large Numbers, and another based on descent algorithms for finding optima of smooth functions on convex polyhedra. We present properties and biological applications of Fréchet means and extend our main results to more general globally nonpositively curved spaces composed of Euclidean orthants.

Contents

Introduction

The development of statistical methods for studying phylogenetic trees, and in particular the search for meaningful notions of consensus tree for phylogenetic data, has been of considerable importance in biology for four decades. Starting with the problem as posed by Adams [1], a great deal of research has been done, and a myriad of definitions proposed, relating to consensus trees in phylogenetics; see [13] for an excellent overview. The problem has been confounded by the combinatorial nature of the trees themselves. According to Cranston and Rannala [18], "Phylogenetic inference has long been troubled by the difficulty of performing statistical analysis on tree topologies. The topologies are discrete, categorical, and non-nested hypotheses about the species relationships. They are not amenable to standard summary analyses such as the calculation of means and variances and cause difficulties for many traditional forms of hypothesis testing." Other papers share concerns about issues such as these [9, 26].

The introduction by Billera, Holmes, and Vogtmann of phylogenetic tree space [12] opened statistical analysis of tree-like data to a wide and computationally tractable variety of techniques [27]. Tree space, with its geodesic distance, is a *globally nonpositively curved* (abbreviated to *global NPC*) space, and as a result it has convexity properties that imply uniqueness of means as well as other important statistical and geometric objects, while also giving a

framework for effective computational methods to calculate these objects. One of the major uses of the convexity properties was the discovery by Owen and Provan [41] of a fast algorithm for computing geodesics in this space (see Section 1 for this algorithm as well as the background tree space geometry necessary to state it). Chakerian and Holmes [17] subsequently showed that phylogenetic tree space provides an excellent platform for implementing several distance-based statistical techniques, and Nye [37] has shown how this space can be used to perform principal component analysis on tree data.

Perhaps the two most fundamental concepts of interest in statistical analysis of data are that of *sample mean* (or *average*) and its associated *variance*. The basic goal of this paper is to demonstrate the computational effectiveness of certain notions of statistical mean and variance for probability distributions on tree space. The average that we use is the *Fréchet mean*, or *barycenter*: the point in tree space that minimizes its sum of squared geodesic distances to the sample points (Section 2). Our decision to use this definition is motivated by work of Sturm [48], who identified the Fréchet mean as a theoretically rich statistical object associated with sampling from a specified distribution on a global NPC space (see Theorem 2.4). Fréchet means in tree space and the algorithm for computing them that arises from Sturm's work (Algorithm 2.5) have been independently developed by Bačák [8].

Our principal theoretical contribution lies in the discovery of polyhedral structure governing the variation of geodesics in tree space as one endpoint varies (Section 3). To be more precise, if T is a fixed point in tree space, then in appropriate coordinates on tree space, the set of points whose geodesics to T share the same combinatorics comprise a convex polyhedral cone called a *vistal cell* (Theorem 3.25), and the vistal cells constitute a polyhedral subdivision of tree space (Theorem 3.30). This *metric combinatorics* also arises in single source shortest path queries (see [35] for a survey), and has direct roots in surprisingly similar statements for boundaries of convex polyhedra [34], the parallel being unexpected because boundaries of convex polyhedra are positively curved, in contrast to the negative curvature of tree space. However, polyhedrality of the subdivision is generally not encountered outside of the planar or positively curved cases, and thus is completely unexpected here; see Example 3.2 for a hint of the complexity that can occur even for global NPC cubical complexes.

Metric combinatorics of tree space, particularly its polyhedral nature, combines with generalities on nonlinear optimization in NPC spaces to give a second iterative method converging to the mean (Algorithm 4.4) via descent procedures. The crucial observations are that the variance function has a unique local minimum on tree space, is continuously differentiable on each Euclidean orthant in tree space, and has a simple algebraic formula within the interior of each vistal cell.

Means in tree space have subtle, sometimes peculiar properties that inform our particular motivations (Section 5), which come primarily from biological and medical applications, although we expect these observations to impact other fields where distributions of metric trees naturally appear. Evolutionary biology, for instance, considers actual phylogenetic trees, each representing a putative evolutionary history of a set species or genes (Example 5.5). In med-

ical imaging, trees can represent blood vessels in human brain scans [47] or lung airway trees [31], for example.

Some of the theory in Sections 1–4 extends to arbitrary global NPC spaces, and all of it extends to global NPC orthant spaces (Section 6). For the first iterative procedure (Algorithm 2.5) and the rest of Section 2, as well as for the shortest path combinatorics in Section 1, this means working in arbitrary global NPC spaces (Sections 6.1–6.2). For the second iterative procedure (Algorithm 4.4) and the rest of Section 4, as well as for the metric combinatorics in Section 3, this means working in piecewise Euclidean global NPC spaces that are formed by gluing orthants together by rules similar to — but substantially more general than — those defining tree space (Section 6.3). The extensions suggest exciting new research in applying both statistical methods and numerical nonlinear programming techniques to a wide variety of problems. **Important note:** readers interested in the generality of abstract orthant spaces or arbitrary NPC spaces are urged to begin with Section 6, which sets up the notation and concepts in Sections 1–4 from that perspective. Hence such readers can avoid checking the proofs in the earlier sections twice.

1 Tree space and the geodesic algorithm

In this section, we describe the space of phylogenetic trees introduced by Billera, Holmes, and Vogtmann [12], as well as a distance and characterization of geodesics in this space.

1.1 Phylogenetic tree space

A *phylogenetic n-tree* T, or simply an *n-tree*, is an acyclic graph T with edge set $\mathcal{E} = \mathcal{E}_T$ whose leaves (degree 1 nodes) are labeled with index set $L = \{0, 1, \dots, n\}$, and whose interior vertices have degree at least 3. (The label 0 is often referred to as the *root* of T, although that is not relevant in this paper.) The maximum number of edges in an n-tree is $2n - 1$. Each edge e of T is assigned a nonnegative *length* $|e|_T$, or $|e|$ in case the ambient tree is clear. Removal of any edge e from T determines a unique partition of the leaves of T into two subsets X_e and $\overline{X}_e$; the pair $X_e|\overline{X}_e$ is called the *split* associated with e. A key property of

splits in trees is that the splits $X_e|\overline{X}_e$ and $X_f|\overline{X}_f$ of any pair of edges e and f are *compatible*, that is, one of the sets $X_e \cap X_f$, $X_e \cap \overline{X}_f$, $\overline{X}_e \cap X_f$, or $\overline{X}_e \cap \overline{X}_f$ is empty. A set S of splits is called *compatible* if every pair of splits in S is compatible. It turns out [46, Theorem 3.1.4] that any compatible set of splits on L corresponds to a unique tree, and so from now on we identify a tree T by simply giving the splits and edge lengths for each edge in T.

A tree T can have an edge e whose associated length $|e|_T$ is 0. This corresponds to the edge e having been *contracted* in T. Denoting the set of edges of T with nonzero length by $\mathcal{E}_T^+$ allows the identification $T \sim T'$ between two trees T and T' whenever (i) $\mathcal{E}_T^+ = \mathcal{E}_{T'}^+$ and (ii) their nonzero edge lengths are equal.

Example 1.1. Two 5-trees are depicted in Figure 1. For simplicity, we only give the splits and edge lengths for the three internal edges in each tree. The six internal edges are distinct since they have different splits, and the splits within each tree are compatible. The only compatible pairs between the two trees, however, are $\{e_1, e_6\}$, $\{e_2, e_5\}$, and $\{e_3, e_4\}$.

$$
\begin{array}{ll}
& e_1 \colon \{0, 1, 2, 5\}\,|\,\{3, 4\} \\
\text{splits} & e_2 \colon \{0, 3, 4, 5\}\,|\,\{1, 2\} \\
& e_3 \colon \quad\ \ \{0, 5\}\,|\,\{1, 2, 3, 4\}
\end{array}
\qquad
\begin{array}{ll}
& e_4 \colon \{0, 1, 4, 5\}\,|\,\{2, 3\} \\
\text{splits} & e_5 \colon \{0, 1, 2, 3\}\,|\,\{4, 5\} \\
& e_6 \colon \quad\ \ \{0, 1\}\,|\,\{2, 3, 4, 5\}
\end{array}
$$
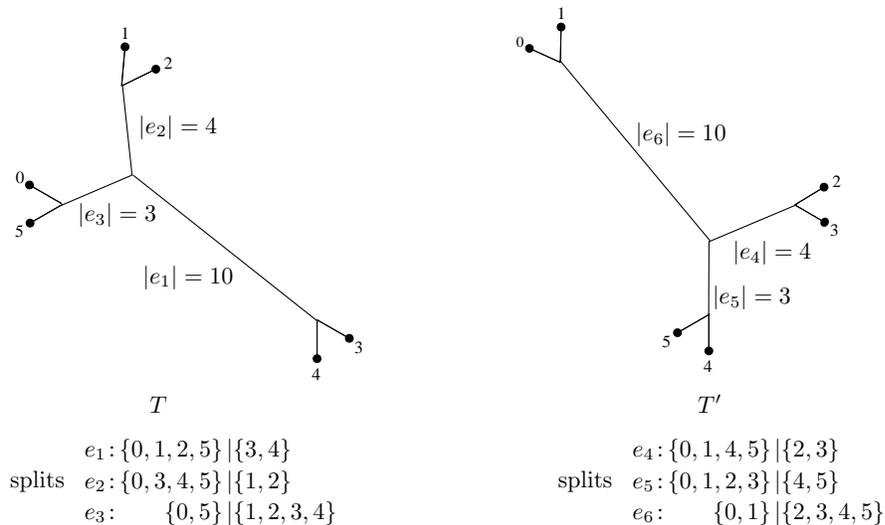
Figure 1: An example of two 5-trees.

The *tree space* $\mathcal{T}_n$ introduced by Billera, Holmes, and Vogtmann [12] is the space of all phylogenetic n-trees. It is obtained by representing each tree $T \in \mathcal{T}_n$ on edge set $\mathcal{E}$ by a vector in the Euclidean *orthant* $\mathcal{O}(T) = \mathcal{O}(\mathcal{E}) = \mathbb{R}_+^{\mathcal{E}}$, whose coordinate values are equal to the corresponding lengths of the edges of T. As above, trees T and T' are identified between orthants whenever the associated trees satisfy $T \sim T'$. This makes $\mathcal{T}_n$ a union of $(2n-1)$-dimensional orthants —called *maximal* orthants —whose interiors are disjoint and which are identified along their boundaries through the equivalence $\sim$ given above. A *path* in $\mathcal{T}_n$ is the image of a continuous map $\gamma : [0, 1] \to \mathcal{T}_n$. The *Euclidean length* of a path in $\mathcal{T}_n$ is the sum

of the Euclidean lengths of its restrictions to the maximal orthants. This length endows $\mathcal{T}_n$ with the metric d in which $d(T, T')$ is the infimum of the Euclidean lengths of the paths from T to T'. Note that $d(T, T') < \infty$, since the space $\mathcal{T}_n$ is path-connected: any two points can be joined by straight line segments through the origin.

1.2   Geodesics in tree space

Billera, Holmes, and Vogtmann [12] show that tree space is *globally non-positively curved* (a *global NPC* space), equivalently known in this context as *CAT(0)*. Among other things, this implies that shortest paths in tree space are unique, so they are unambiguously referred to as *geodesics*. This section summarizes the key results of [39] and [41], which investigate the structure of geodesics in tree space and provide an $O(n^4)$-algorithm — the *GTP algorithm* —to find shortest paths. For notation, if T is a tree with edge set $\mathcal{E}$ and $A \subseteq \mathcal{E}$, then we write

$$\|A\|_T = \sqrt{\sum_{e \in A} |e|_T^2}$$

and use $\|A\|$ if the tree T is clear. This means that $\|A\| = |e|$ whenever $A = \{e\}$.

We express a geodesic with endpoints X and T as a parameterized curve $\gamma : [0, 1] \to \mathcal{T}_n$ with $\gamma(0) = X$, $\gamma(1) = T$, and $d(\gamma(t), \gamma(t')) = |t - t'| \cdot d(X, T)$ for all $t, t' \in [0, 1]$. If an edge e lies in both X and T, then it lies in every tree on the path γ, with length uniformly changing between the two terminal values [12, Section 4.2]. We therefore focus first on the case when X and T have no internal edges in common, and ignore the lengths of the *pendant* edges (those containing leaves) in the distance computation.

Each geodesic in tree space is a sequence of straight line segments, called *legs*, because tree space is piecewise Euclidean. Each leg is contained within a single orthant $\mathcal{O}(E_i \cup F_i)$, where $E_i \subseteq \mathcal{E}_X$ and $F_i \subseteq \mathcal{E}_T$. The precise properties of the sets E_i and F_i making up these legs were determined in [39]. In particular, define the *support* $(\mathcal{A}, \mathcal{B}) = ((A_1, \ldots, A_k), (B_1, \ldots, B_k))$ of a geodesic γ to consist of a pair consisting of a partition $A_1 \cup \cdots \cup A_k$ of $\mathcal{E}_X$ and a partition $B_1 \cup \cdots \cup B_k$ of $\mathcal{E}_T$ such that the following property holds:

(P1) for each $i > j$, the union $A_i \cup B_j$ is compatible.

The geodesic γ has legs in $\mathcal{O}(E_i \cup F_i)$, where

$$E_i = A_{i+1} \cup \cdots \cup A_k$$
$$\text{and } F_i = B_1 \cup \cdots \cup B_i.$$

The individual pairs (A_i, B_i) are the *support pairs* for the geodesic.

Whether the shortest piecewise-linear path having these legs actually forms the geodesic between X and T is determined by the following two properties for $(\mathcal{A}, \mathcal{B})$.

(P2) $\dfrac{\|A_1\|}{\|B_1\|} \leq \dfrac{\|A_2\|}{\|B_2\|} \leq \cdots \leq \dfrac{\|A_k\|}{\|B_k\|}$. This is called the *ratio sequence* for $(\mathcal{A}, \mathcal{B})$.

(P3) For all (A_i, B_i) and partitions $I_1 \cup I_2$ of A_i and $J_1 \cup J_2$ of B_i such that $I_2 \cup J_1$ is compatible, the inequality $\frac{\|I_1\|}{\|J_1\|} \geq \frac{\|I_2\|}{\|J_2\|}$ holds.

The properties (P1)–(P3) determine the geodesic between X and T, as well as the algebraic description of this geodesic given in Theorem 2.4 in [41].

The case where X and T have a nonempty set C of common edges was addressed in [41, Section 4]: remove the common edges between X and T from each tree, and then find the paths between the remaining disjoint forests, matching trees by their leaf sets. The common edges are then placed into the path with the length of each such edge being

$$(1 - \lambda)|e|_X + \lambda|e|_T. \tag{1}$$

This also allows pendant edges to be taken into account.

To be able to work more easily with trees having common edges, we extend Theorem 2.4 in [41] to the case where X and T have common edges, and in the process simplify the description of the geodesic considerably. To do this, we use the following three important conventions.

(a) An edge is never compatible with itself; thus the pairs of identical edges in X and T must appear in the same support pair (A_i, B_i).

(b) $\|A_i\| = -\sqrt{\sum_{e \in A_i} |e|_T^2}$ for any set A_i of edges of X in common with T.

(c) We extend the notation for support pair by adding the additional sets

$$A_0 = B_0 = A_{k+1} = B_{k+1} = \varnothing$$

and define $\frac{\|A_0\|}{\|B_0\|} = -\infty$ and $\frac{\|A_{k+1}\|}{\|B_{k+1}\|} = \infty$.

With these conventions we can restate the unified result.

Theorem 1.2. *Let X and T be any two trees in $\mathcal{T}_n$ (not necessarily disjoint), and let $(\mathcal{A}, \mathcal{B})$ be a support for X and T satisfying (P2) and (P3). The unique geodesic $\gamma = \{\gamma(\lambda) : 0 \leq \lambda \leq 1\}$ from X to T has legs*

$$\gamma^i = \left\{ \gamma(\lambda) : \frac{\|A_i\|}{\|B_i\|} \leq \frac{\lambda}{1 - \lambda} < \frac{\|A_{i+1}\|}{\|B_{i+1}\|} \right\} \qquad for \ i = 0, \dots, k, \tag{2}$$

The points on each leg γ^i are associated with the tree T_i having edge set

$$B_1 \cup \cdots \cup B_i \cup A_{i+1} \cup \cdots \cup A_k$$

and edge lengths

$$|e|_{T_i} = \begin{cases} \dfrac{(1 - \lambda)\|A_j\| - \lambda\|B_j\|}{\|A_j\|}|e|_X & if \ e \in A_j \\[2mm] \dfrac{\lambda\|B_j\| - (1 - \lambda)\|A_j\|}{\|B_j\|}|e|_T & if \ e \in B_j. \end{cases} \tag{3}$$

7

The length of γ is

$$L(\gamma) = \left\| \left(\|A_1\| + \|B_1\|, \ldots, \|A_k\| + \|B_k\| \right) \right\|. \tag{4}$$

Proof. The presentation in this theorem matches that of the original Theorem 2.4 in [41] except for the treatment of the common edges of X and T. Consider any edge e common to X and T. The definition of a support ensures that e lies in both A_i and B_i for some i. Further, by convention the ratio $\|e\|_X/\|e\|_T$ is negative, so (P2) is never satisfied unless all of the common edges are placed at the front of the ratio sequence. This also means that for any $\lambda > 0$, each common edge is contained in some B_i for the computation of its edge length at that point along the geodesic. Furthermore, since the common edges are mutually compatible with each other, they are placed in different support pairs whenever the ratios $|e|_X/|e|_T$ differ. It follows that the common edges are always grouped in support pairs (A_i, B_i) having $\|A_i\|/\|B_i\| = -|e|_X/|e|_T$ for any e in that support pair.

Now consider the length of a common edge e in leg γ^i of the path. By (3),

$$|e|_{T_i} = \frac{\lambda\|B_j\| - (1-\lambda)\|A_j\|}{\|B_j\|}|e|_T = \left(\lambda - (1-\lambda)\frac{\|A_j\|}{\|B_j\|} \right)|e|_T$$

$$= \left(\lambda + (1-\lambda)\frac{|e|_X}{|e|_T} \right)|e|_T = \lambda|e|_T + (1-\lambda)|e|_X,$$

which matches (1).

Next look at the term in (4) corresponding to a support pair (A_i, B_i) of common edges:

$$(\|A_j\| + \|B_j\|)^2 = \left(\frac{\|A_j\|}{\|B_j\|} + 1 \right)^2 \|B_j\|^2 = \sum_{e \in B_j} \left(1 - \frac{|e|_X}{|e|_T} \right)^2 |e|_T^2$$

$$= \sum_{e \in B_j} \left(|e|_T - |e|_X \right)^2 .$$

Summing this over all such pairs (A_i, B_i) yields

$$\sum_{e \in C} \left(|e|_T - |e|_X \right)^2 ,$$

where C is the set of common edges. This matches the expression given in [41, Section 4].

Finally, take the case where an edge e lies in only one of the sets X and T, but is compatible with all edges in the other set. Intuitively, we can think of adding e to the other set with length 0, and treatin these as common edges. Formally, if e lies in X, then it appears in a support pair $(A_i, \varnothing)$ with A_i a set of edges compatible with all of T; and if e lies in T, then it appears in a support pair $(\varnothing, B_i)$ with B_i a set of edges compatible with all of X. Since the ratios of these pairs is either 0 or ∞, respectively (since $\|\varnothing\| = 0$), these pairs appear before and after any nontrivial pairs, respectively. Further, the edge component values and path length are as indicated in (3) and (4), respectively. This completes the proof of the theorem. $\square$

Example 1.3. Figure 2 shows an example of the geodesic γ between the trees T and T' in Figure 1 in Example 1.1. The associated support $(\mathcal{A}, \mathcal{B})$ for γ has $\mathcal{A} = \{\{e_2, e_3\}, \{e_1\}\}$ and $\mathcal{B} = \{\{e_6\}, \{e_4, e_5\}\}$, and the coordinates of seven equally spaced trees in γ are given in the table. The length of this path, as given by (4), is

$$L(\gamma) = \left\| \left(\|\{e_2, e_3\}\| + \|\{e_6\}\|, \|\{e_1\}\| + \|\{e_4, e_5\}\| \right) \right\| = 15\sqrt{2}.$$
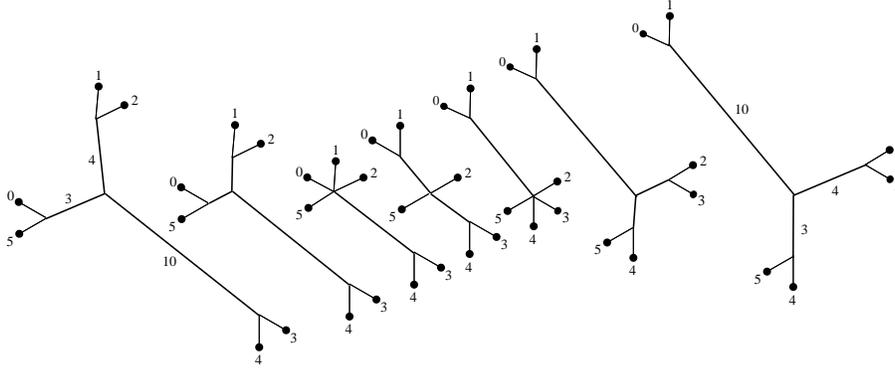
$\gamma(i/6)$

i	$\|e_1\|$	$\|e_2\|$	$\|e_3\|$	$\|e_4\|$	$\|e_5\|$	$\|e_6\|$
0	10	4	3	0	0	0
1	7.5	2	1.5	0	0	0
2	5	0	0	0	0	0
3	2.5	0	0	0	0	2.5
4	0	0	0	0	0	5
5	0	0	0	2	1.5	7.5
6	0	0	0	4	3	10

Figure 2: Seven trees in the geodesic γ between T and T', sampled at the points $\gamma(i/6)$ for $i \in \{0, 1, ..., 6\}$, The table gives the interior edge lengths for the trees, using the same edge labels as Figure 1.

We end the section by giving a canonical representation for any geodesic.

Lemma 1.4. *Any geodesic γ can be represented by unique support $(\mathcal{A}, \mathcal{B})$ satisfying*

$$\frac{\|A_1\|}{\|B_1\|} < \frac{\|A_2\|}{\|B_2\|} < \cdots < \frac{\|A_k\|}{\|B_k\|}. \tag{5}$$

9

This support is called the minimal support.

Proof. This is the content of the remark in [41, Section 2.3]. The basic argument is as follows. Any support $(\mathcal{A}', \mathcal{B}') \neq (\mathcal{A}, \mathcal{B})$ of form (5) results in a different geodesic, since by (2) they have different legs. On the other hand, for any representation of γ having equalities in the ratio sequence, combine the respective sets in every equality subsequence. The resulting support continues to satisfy (P2), and hence there is a shortest piecewise linear path from X to T through the prescribed orthants. Further, from (4) it follows that the length of this path equals that of γ, and hence defines the unique geodesic γ. $\qquad\square$

Remark 1.5. Theorem 1.2 positions the support pairs corresponding to edges compatible with both trees into (5) as follows.

(i) The set N_X of edges of X that are not in T but are compatible with all edges of T is the set A_0, with $B_0 = \varnothing$ and ratio $\frac{\|A_0\|}{\|B_0\|} = \frac{\|N_X\|}{\|\varnothing\|} = -\infty$.

(ii) The set N_T of edges of T that are not in X but are compatible with all edges of X is the set B_k with $A_k = \varnothing$, and so its ratio is $\frac{\|A_k\|}{\|B_k\|} = \frac{\|\varnothing\|}{\|N_T\|} = 0$.

(iii) Any edge e that lies in both X and T (and hence has positive length in both sets) appears in both sets of some support pair (A_i, B_i), and so the ratio is $-\infty < \frac{\|A_i\|}{\|B_i\|} < 0$.

(iv) All other support pairs have $\frac{\|A_i\|}{\|B_i\|} > 0$, so both sets in the support pair are nonempty.

The ordering of the support pairs in (i)–(iii) has no effect on the structure of the geodesic between X and T, so for the remainder of the paper we take the ratio sequence for a geodesic to represent only the *positive* ratios in the sequence.

2 The mean and variance in tree space

Given a finite point set $\boldsymbol{T} = \{T^1, \ldots, T^r\}$ of trees in $\mathcal{T}_n$, the *mean* of $\boldsymbol{T}$, alternatively known as the *Fréchet mean* or *barycenter*, is the tree $\overline{T} \in \mathcal{T}_n$ that minimizes the sum $S(X, \boldsymbol{T})$ of the squares of the distances from X to the points in $\boldsymbol{T}$. The *variance* of $\boldsymbol{T}$ is $S(X, \boldsymbol{T})/r$. Since r is constant throughout the following discussion, we abuse notation and henceforth refer to the variance as simply $S(X, \boldsymbol{T})$. The motivation for considering these notions of mean and variance as the appropriate statistical objects in tree space was given by Sturm [48], who established the mathematical foundations for probability theory on global NPC spaces. This section reviews the required basics of Sturm's geometric methods in the context of tree spaces. (Readers interested in Fréchet means of more general distributions, and those in arbitrary global NPC spaces, should read Section 6 now to put the material below in these more general settings.)

2.1 The variance function

Let $T \in \mathcal{T}_n$ be a fixed tree, and consider the geodesics from T to a variable tree $X \in \mathcal{T}_n$. The tree X can be thought of as a vector in $\mathbb{R}_+^{\mathcal{E}}$, whose coordinates are expressed using the corresponding lower-case letter x. If the geodesic from X to T has support pair $(\mathcal{A}, \mathcal{B})$ as in Theorem 1.2, then the squared distance $d(X, T)^2$ from X to T is expressed as the function

$$S_T(x) = \sum_{i=1}^{k} \big(\|x_{A_i}\| + \|B_i\| \big)^2 \tag{6}$$

in which x_{A_i} is the vector whose coordinates are restricted to edges in A_i. It follows that for a set $\boldsymbol{T} = \{T^1, \ldots, T^r\} \subseteq \mathcal{T}_n$ of trees, the variance function $S(X, \boldsymbol{T})$ can be written

$$S(x) := S(X, \boldsymbol{T}) = \sum_{\ell=1}^{r} S_{T^\ell}(x). \tag{7}$$

Thus the mean $\overline{T}$ can be thought of as the point x^* that minimizes $S(x)$ over $x \in \mathcal{T}_n$.

To state the next result, a real-valued function $f : \mathcal{T} \to \mathbb{R}$ on a metric space $\mathcal{T}$ is *strictly convex* if $f \circ \gamma$ is a strictly convex real-valued function on $\mathbb{R}$ for all geodesics γ; that is, if

$$f\big(\gamma(\lambda)\big) < (1 - \lambda)f\big(\gamma(0)\big) + \lambda f\big(\gamma(1)\big) \text{ whenever } 0 < \lambda < 1.$$

Proposition 2.1. *The variance function $S(x)$ is strictly convex as a function on $\mathcal{T}_n$. Consequently, the mean is the unique local minimum of $S(x)$ in $\mathcal{T}_n$.*

Proof. [48, Proposition 1.7]. See also Example 6.3. □

The differentiability of the variance function S is critical to the construction of gradient-descent methods for minimizing S. This in turn depends on the differentiability of the individual functions S_T in (7). Identifying the geodesic from X to T by its support $(\mathcal{A}, \mathcal{B})$, Eq. (6) yields the partial derivatives of S_T with respect to each of the coordinates x_e:

$$\frac{\partial S_T(X)}{\partial x_e} = 2\big(\|x_{A_i}\| + \|B_i\| \big) \frac{x_e}{\|x_{A_i}\|}$$

$$= 2x_e \Big(1 + \frac{\|B_i\|}{\|x_{A_i}\|} \Big), \tag{8}$$

where A_i is the set containing e. This is well-defined whenever x lies in the interior of its maximal orthant, although the functional form of (8) depends upon the combinatorial type of the geodesic, and in particular on $(\mathcal{A}, \mathcal{B})$. It turns out, however, that throughout the interior of any maximal orthant the function S is continuously differentiable.

Theorem 2.2. *The variance function $S(x)$ is continuously differentiable on the interior of every maximal orthant $\mathcal{O}$.*

11

Proof. From (7) it suffices to show that the function (8) is continuously differentiable on the interior of $\mathcal{O}$. By Lemma 1.4 the geodesic between X and T can be represented uniquely by support $(\mathcal{A}^0, \mathcal{B}^0)$ satisfying (5). Any other support $(\mathcal{A}, \mathcal{B})$ for this geodesic consists of a sequence of sets partitioning the A_i^0 and B_i^0 into equality subsequences, with the ratios $\|A_j\|/\|B_j\|$ in the equality subsequence equal to the corresponding ratio $\|A_i^0\|/\|B_i^0\|$ of the sets from which they were partitioned. But this means that the ratio $\|B_i\|/\|x_{A_i}\|$, and hence $\partial S_T(X)/\partial x_e$, is the same regardless of which representation we choose for the geodesic. It follows that the partial derivatives are continuous everywhere in the interior of $\mathcal{O}$. $\square$

2.2 Sturm's algorithm

The orthant structure of tree space $\mathcal{T}_n$ prevents the averaging of finite point sets using the standard Euclidean centroid. The following serves as an approximate replacement, introduced by Sturm [48, Definition 4.6] in the context of probability theory on arbitrary globally nonpositively curved spaces.

Definition 2.3. For a set $X^1, X^2, \ldots$ of points in $\mathcal{T}_n$ and an index k, the *inductive mean value* of $X^1, \ldots, X^k$ is the point μ_k defined by setting $\mu_1 = X^1$ and for $\ell = 2, \ldots, k$, letting μ_ℓ be the point $\gamma_{1/\ell} = \gamma(1/\ell)$ that is $1/\ell$ along the geodesic γ from $\mu_{\ell-1}$ to $\gamma_1 = X^\ell$.

Note that if all of the X^ℓ lie in the same orthant, then the inductive mean value of $X^1, \ldots, X^k$ in fact equals the standard centroid of $X^1, \ldots, X^k$. For general points in $\mathcal{T}_n$, though, the inductive mean may not be the Fréchet mean; it may in fact give different points for different orderings of the points $X^1, \ldots, X^k$ (see Example 5.3). Sturm goes on to prove [48, Theorem 4.7] the following *strong law of large numbers* for the Fréchet mean.

Theorem 2.4. *Fix a set $\{T^1, \ldots, T^r\} \subseteq \mathcal{T}_n$ of trees. If $X^1, X^2, \ldots$ is a sequence of points sampled uniformly and independently from $\{T^1, \ldots, T^r\}$, then with probability 1, the sequence of inductive mean values $\mu_1, \mu_2, \ldots$ approaches the mean $\overline{T}$ of $\{T^1, \ldots, T^r\}$.*

To be precise, Sturm shows that if we take the inductive mean μ_k as a random variable dependent on the sampling of the points X^ℓ, then the distance $d(\mu_k, \overline{T})$ from μ_k to the true mean $\overline{T}$ has expected value bounded above by $S(\boldsymbol{T}, \overline{T})/k$. This gives us a way of estimating the Fréchet mean $\overline{T}$ through a sequence of inductive means $\mu_1, \mu_2, \ldots$ obtained by randomly sampling trees from the set $\{T^1, \ldots, T^r\}$.

Algorithm 2.5 (Sturm's algorithm).

      INPUT   a set $\{T^1, \ldots, T^r\}$ of trees in $\mathcal{T}_n$
                  positive integers K and N
                  positive real number ε

    OUTPUT $\mu_k = k^{\text{th}}$ approximation of the mean tree

INITIALIZE   choose a tree $T \in \{T^1, \ldots, T^r\}$ uniformly at random
                  set $\mu_1 := T$
                  set $k := 1$

      WHILE $k < K$ or pairwise distances $d(\mu_j, \mu_\ell)$ for $k - N < j, \ell \leq k$ are not all $\leq \varepsilon$

          DO   choose tree $T \in \{T^1, \ldots, T^r\}$ uniformly at random
                  set $\gamma :=$ the geodesic from T to μ_k
                  set $\mu_{k+1} := \gamma_{1/(k+1)}$
                  set $k := k + 1$

        END   WHILE-DO

    RETURN μ_k, the k^{th} approximation of the mean tree

Remark 2.6. The choice of stopping criterion involves two parts.

(i) Running the algorithm a specified initial number K of iterations guarantees an upper bound of $\frac{r}{K+1} S(\overline{T}, \boldsymbol{T})$ on the expected distance of the final tree μ_k to the mean $\overline{T}$. This is derived from the proof of Theorem 4.7 in [48].

(ii) Comparing the final N sample means serves as a proxy for testing that the sample means act like a Cauchy sequence for N steps.

Thus in principle, proper settings for K, N, and ε could be used to set confidence intervals on the distance $d(\mu_i, \overline{T})$ by using Sturm's result. This would involve a more sophisticated statistical analysis, which we did not undertake in this paper. In practice, we chose K, N, and ε to balance run-time with the desired precision, working under the rough assumption that if N is chosen large enough, then the approximate mean will be within ε of the mean tree. For Example 5.5, we chose $K = 1\,000\,000$, $N = 10$, and $\varepsilon = 10^{-4}$.

Remark 2.7. We have made software implementing this algorithm freely available [40].

3   The combinatorics of geodesics in tree space

This section investigates the combinatorial structure of geodesics in $\mathcal{T}_n$ and their relationship to the variance function. To be more precise, fix a source tree $T \in \mathcal{T}_n$. The shortest path from an arbitrary tree $X \in \mathcal{T}_n$ to T has a "combinatorial type", determined through Theorem 1.2 by the sequence of orthants that it passes through, or more specifically the support pair $(\mathcal{A}, \mathcal{B})$ associated with the geodesic. This combinatorial type can change, even when X has

the same topology, depending on the precise values of the lengths of the edges in X. We are interested in the partition of $\mathcal{T}_n$ — called[1] the *vistal subdivision* of $\mathcal{T}_n$—into regions for which the geodesics to the fixed tree T have the same combinatorial type.

We begin by describing a simple change of coordinates, the *squaring map*, and characterizing the faces of maximal dimension in the vistal subdivision (Section 3.1). In particular, Propositions 3.5 and 3.6 establish that after applying the squaring map the vistal facets are polyhedral regions that cover tree space but have disjoint interiors. Next we provide a simple description of the faces of lower dimension in these polyhedra (Section 3.2). Finally, we prove that the vistal facets constitute the maximal cells of a polyhedral subdivision of tree space, called the *vistal polyhedral subdivision* (Section 3.3).

Remark 3.1. The idea of studying the paths taken by geodesics emanating from a source point has been studied in computational geometry, in the areas of single source shortest path queries [35] and polyhedral unfolding [34]. Recently Chepoi and Maftuleac studied the single source shortest path problem for CAT(0) rectangular complexes, where each cell is a 2D rectangle. When the underlying space is intrinsically 2D, these shortest path subdivisions are often polyhedral. However, in general, we can not expect this in higher dimensions, and indeed, without the squaring map, the vistal subdivisions are not polyhedral, as illustrated in Example 3.32. The squaring map is only possible in $\mathcal{T}_n$ because all of the combinatorial complexity of the space happens about the origin.

Example 3.2. Vistal subdivisions are in general far from polyhedral, even after changes of coordinates such as squaring. A prerequisite for a squaring map to produce polyhedrality would be that (every component of the) the bounding hypersurface has degree at most 2. However, global NPC cubical complexes can have vistal cells bounded by hypersurfaces of degree greater than 2. We conjecture that the bounding hypersurface can have components of arbitrarily high degree.

For a specific example, consider first the arrangement
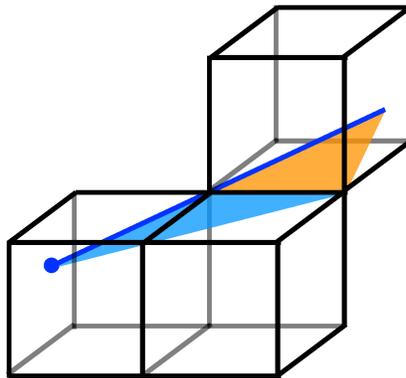
---

of three 3-dimensional cubes in which the bottom two cubes are joined along a common face, the top cube meets one bottom cube along an edge, and all three cubes meet at a central vertex. Consider the set $\mathcal{V}$ of points whose shortest path back to a fixed starting point (the dot in the left bottom cube) passes through the interior of the edge joining the top cube to the bottom cubes, as opposed to passing through the central vertex. The set $\mathcal{V}$ is obtained by rotating the top shaded triangle about the shared cube edge. The boundary of $\mathcal{V}$ is a cone and hence is described as the vanishing locus (in the top cube) of a polynomial of degree 2.

To get the desired example, glue a 4-cube to the right-hand 2-face of the top cube in the figure. If it were merely a 3-cube glued on, then the cone $\mathcal{V}$ would simply continue to expand into the added cube, creating a frustum in the added 3-cube. But once the new facet is a 4-cube, the frustum rotates freely around the shared 2-face. The boundary of such a rotated frustum is the locus of zeros (in the 4-cube) of a polynomial, but any such polynomial f has degree at least 4. Indeed, intersecting the hypersurface in question with a generic hyperplane yields a union of two cones, each of which has degree 2. Therefore f restricts to a polynomial of degree at least 4 on the hyperplane, whence the hypersurface itself has degree at least 4.

3.1   Vistal facets

Definition 3.3. Given a source tree $T \in \mathcal{T}_n$, a maximal orthant $\mathcal{O} \subseteq \mathcal{T}_n$, and a support $(\mathcal{A}, \mathcal{B})$, let $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ be the closure of the set of trees $X \in \mathcal{O}$ for which the geodesic joining X to T has support $(\mathcal{A}, \mathcal{B})$ satisfying (P2) and (P3) with strict inequalities. A *previstal facet* is any nonempty set $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ of this form.

The description of $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ becomes linear after a simple change of variables.

Definition 3.4. The *squaring map* $\mathcal{T}_n \to \mathcal{T}_n$ acts on $x \in \mathcal{T}_n \subseteq \mathbb{R}_+^E$ by squaring coordinates:

$$(x_e \mid e \in E) \mapsto (\xi_e \mid e \in E), \ \text{where } \xi_e = x_e^2.$$

Denote by $\mathcal{T}_n^2$ the image of this map, and let $\xi_e = x_e^2$ denote the coordinate indexed by $e \in E$. The image of an orthant in $\mathcal{T}_n$ is then the equivalent orthant in $\mathcal{T}_n^2$, and the image of a previstal facet $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ in $\mathcal{T}_n^2$ is a *vistal facet* denoted $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$. With this change of variables, $\|A\| = \sum_{e \in A} \xi_e$ for any set of splits A.

The squaring map induces on the variance function S a corresponding pullback function

$$S^2(\xi) = S(\sqrt{\xi}), \ \text{where } (\sqrt{\xi})_e = \sqrt{\xi_e}. \tag{9}$$

Since the variance function $S(x)$ is continuous on $\mathcal{T}_n$ with a uniquely attained minimum by Proposition 2.1, and continuously differentiable on the interior of each maximal orthant by Theorem 2.2, the same properties hold for S^2. Thus we can apply steepest descent methods after squaring just as we would beforehand. This is further explored in Section 4.

Theorem 1.2 implies a nice description of the vistal facets of $\mathcal{T}_n^2$.

Proposition 3.5. *The vistal facet $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ is a convex polyhedral cone in $\mathcal{T}_n^2$ defined by the following inequalities on $\xi \in \mathbb{R}^E$, where all norms $\|\cdot\|$ are to be interpreted as $\|\cdot\|_T$.*

(O) $\xi \in \mathcal{O}$; *that is,* $\xi_e \geq 0$ *for all* $e \in E$, *and* $\xi_e = 0$ *for* $e \notin \mathcal{E}$, *where* $\mathcal{O} = \mathbb{R}_+^{\mathcal{E}}$.

(P2) $\|B_{i+1}\|^2 \sum_{e \in A_i} \xi_e \leq \|B_i\|^2 \sum_{e \in A_{i+1}} \xi_e$ *for all* $i = 1, \dots, k-1$.

(P3) $\|B_i \setminus J\|^2 \sum_{e \in A_i \setminus I} \xi_e \geq \|J\|^2 \sum_{e \in I} \xi_e$ *for all* $i = 1, \dots, k$ *and subsets* $I \subseteq A_i$, $J \subseteq B_i$ *such that* $I \cup J$ *is compatible.*

Proof. For a vector $x = (x_e \mid e \in E) \in \mathcal{O}$ to lie in $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$, the tree X must be in an orthant of $\mathcal{T}_n$ and satisfy properties (P2) and (P3). The orthant condition immediately implies the nonnegativity conditions in (O). The inequalities corresponding to (P2) are

$$\frac{\|A_i\|}{\|B_i\|} \leq \frac{\|A_{i+1}\|}{\|B_{i+1}\|}.$$

Squaring, cross-multiplying, and substituting ξ_e for x_e^2 yields the corresponding linear inequality in $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ in $\mathcal{T}_n^2$. The inequalities for (P3) are obtained in the same manner. $\qquad\square$

Proposition 3.6. *The vistal facets are of dimension $2n - 1$, have pairwise disjoint interiors, and cover $\mathcal{T}_n^2$. A point $\xi \in \mathcal{T}_n^2$ lies interior to a vistal facet $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ if and only if the inequalities in (O), (P2), and (P3) are strict.*

Proof. The vistal facets cover $\mathcal{T}_n^2$ by definition: $T \in \mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ if the geodesic from $X \in \mathcal{O}$ to T has support $(\mathcal{A}, \mathcal{B})$. The second statement follows by definition and by standard properties of convex polyhedra presented as solutions to systems of linear inequalities. $\qquad\square$

3.2 Vistal cells

Henceforth in this section we focus our attention on the squared tree space $\mathcal{T}_n^2$ and its expression as a union of polyhedral vistal facets as given by Proposition 3.5. This subsection concerns faces of vistal facets, including compact characterizations thereof.

3.2.1 Signatures and vistal cells

Definition 3.7. Fix a source tree $T \in \mathcal{T}_n$, a (not necessarily maximal) orthant $\mathcal{O} \subseteq \mathcal{T}_n$, and a support $(\mathcal{A}, \mathcal{B})$. A *signature* associated with the support $(\mathcal{A}, \mathcal{B})$ is a length $k-1$ sequence $\mathcal{S} = (\sigma_1, \dots, \sigma_{k-1})$ of symbols $\sigma_i \in \{=, \leq\}$. The *previstal cell* defined by $\mathcal{O}$, $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{S}$ is the set $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ of points X in $\mathcal{O}$ for which the ratio sequence for $(\mathcal{A}, \mathcal{B})$ at the point X has the following specific form:

$$\frac{\|A_1\|}{\|B_1\|} \sigma_1 \frac{\|A_2\|}{\|B_2\|} \sigma_2 \cdots \sigma_{k-2} \frac{\|A_{k-1}\|}{\|B_{k-1}\|} \sigma_{k-1} \frac{\|A_k\|}{\|B_k\|}. \tag{10}$$

The *vistal cell* $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S}) \subseteq \mathcal{T}_n^2$ is the image of $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ under squaring.

Remark 3.8. Vistal cells are convex polyhedra that need not be bounded, and as such they might not be topological cells. However, the interior of a convex polyhedron is a topological cell, so every vistal cell is the closure of a topological cell.

Lemma 3.9. *The dimension of the vistal cell $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ is at most $\dim(\mathcal{O}) - m(\mathcal{S})$, where $m(\mathcal{S})$ is the number of "=" components in $\mathcal{S}$. The vistal cell $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ is full-dimensional if and only if there exists a point $X \in \mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ satisfying the following two properties.*

(V1) *For each $i = 1, \ldots, k-1$, $\frac{\|A_i\|}{\|B_i\|} = \frac{\|A_{i+1}\|}{\|B_{i+1}\|}$ if σ_i is "=" and $\frac{\|A_i\|}{\|B_i\|} < \frac{\|A_{i+1}\|}{\|B_{i+1}\|}$ if σ_i is "$\leq$".*

(V2) *The inequalities in (P3) are satisfied strictly.*

Proof. This follows from standard polyhedral theory, as treated in [50], for instance. $\square$

Proposition 3.5 implies that (i) vistal cells are faces of vistal facets, and that (ii) vistal facets are vistal cells for which $\mathcal{O}$ is maximal and the signature contains only "$\leq$" symbols. What we prove here is that *all* faces of vistal facets can be represented as vistal cells, and that under some simple conditions on $(\mathcal{A}, \mathcal{B})$, Definition 3.7 provides a canonical description of each vistal cell. We start by determining all supports and signatures associated with the geodesic γ from T to a particular point X. By Lemma 1.4, the geodesic γ can be represented by a unique minimal support $(\mathcal{A}, \mathcal{B})$ satisfying (5):

$$\frac{\|A_1\|}{\|B_1\|} < \frac{\|A_2\|}{\|B_2\|} < \cdots < \frac{\|A_k\|}{\|B_k\|}.$$

Any other support $(\mathcal{A}', \mathcal{B}')$ of γ corresponds to a ratio sequence in which at least one ratio $\|A_i\|/\|B_i\|$ is replaced by a *ratio subsequence* formed from a partition of A_i and B_i, with equalities between all terms. Any ratio subsequence for which X continues to satisfy (P3) together with equalities between terms of the ratio subsequences constitutes a valid support for γ. We next give a specific method for determining all such support sequences.

3.2.2 Incompatibility graphs and equality subsequences

In [41] it was shown how condition (P3) for support pair (A_i, B_i) can be rephrased in terms of conditions on a special node-weighted graph derived from the compatibility relations between X and T and their coordinate values. We summarize the technique here. Denote the coordinates of X and T by $X = (x_e \mid e \in \mathcal{E}_X)$ and $T = (t_e \mid e \in \mathcal{E}_T)$, and let $\xi_e = x_e^2$ and $\tau_e = t_e^2$ be their squared coordinates.

Definition 3.10. The *incompatibility graph* $G(A_i, B_i)$ between A_i and B_i is the weighted bipartite graph with vertex set $A_i \cup B_i$ and an edge from $a \in A_i$ to $b \in B_i$ whenever a and b are incompatible. The weight of each vertex $a \in X$ is $\tilde{\xi}_a = \xi_a / \sum_{e \in A_i} \xi_e$, and the weight of each vertex $b \in T$ is $\tilde{\tau}_b = \tau_b / \sum_{e \in B_i} \tau_e$. A *(vertex) cover* for $G(A_i, B_i)$ is a set $C \subset A_i \cup B_i$

having the property that every edge of $G(A_i, B_i)$ has at least one endpoint in C. The weight of C is the sum of the weights of its vertices.

Lemma 3.11 ([41, Section 3]). *Property (P3) holds for support pair (A_i, B_i) if and only if every cover of $G(A_i, B_i)$ has weight ≥ 1.* $\square$

By Lemma 3.11, testing a support pair (A_i, B_i) for property (P3) is equivalent to showing that the *min weight cover* for $G(A_i, B_i)$ has weight 1. The problem of finding the minimum cover in $G(A_i, B_i)$ in turn can be reduced to solving a max flow problem (see [2], Section 12.3) on a specially defined flow network $F(A_i, B_i)$. To construct $F(A_i, B_i)$, start with $G(A_i, B_i)$, attach a source $\bar{s}$ to the A_i-vertices of $G(A_i, B_i)$ and a sink $\bar{t}$ to the B_i-vertices of $G(A_i, B_i)$, and direct all edges from $\bar{s}$ toward $\bar{t}$. Set the capacity of each edge $(\bar{s}, a)$ to $\tilde{\xi}_a$, set the capacity of each edge $(b, \bar{t})$ to $\tilde{\tau}_b$, and set the capacities of edges in $G(A_i, B_i)$ to ∞. The Max-Flow-Min-Cut Theorem implies that the value of the maximum $(\bar{s}, \bar{t})$-flow f for $F(A_i, B_i)$ is equal to the capacity of a minimum capacity of an $(\bar{s}, \bar{t})$-cut K in $F(A_i, B_i)$, which in turn corresponds to a minimum weight cover C for $G(A_i, B_i)$. Thus the condition in Lemma 3.11 for $G(A_i, B_i)$ is equivalent to the property that the max flow in $F(A_i, B_i)$ is ≥ 1. The precise relationship between max flows in $F(A_i, B_i)$ and min covers in $G(A_i, B_i)$ is crucial to determining the possible ratio subsequences that can replace a term $\|A_i\|/\|B_i\|$ in (5), and we clarify this relationship below.

Example 3.12. Figure 3 demonstrates this for a hypothetical support pair (A_i, B_i) with $A_i = \{x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8\}$ and $B_i = \{t_1, t_2, t_3, t_4, t_5, t_6, t_7\}$, compatibility graph $G(A_i, B_i)$, and values ξ_a, and τ_b as given in Figure 3(a). Figure 3(b) depicts the associated flow graph $F(A_i, B_i)$ and max flow. For simplicity, the weights are not normalized, so that all numbers are scaled by 23, the sum of the weights. This flow has value 23, which means that the pair (A_i, B_i) satisfies (P3).

3.2.3 Residual graphs and ratio subsequences

Now consider the problem of determining the possible ratio subsequences replacing a term $\|A_i\|/\|B_i\|$ in the ratio sequence of a minimal support for X and T. We use the optimal flow conditions on $F(A_i, B_i)$ to do this. Recall that here (A_i, B_i) also satisfies (P3), so that the max flow f on $F(A_i, B_i)$ has value 1. The associated minimum weight cover for $G(A_i, B_i)$ can then be obtained from this flow. To do this, we define another auxiliary graph.

Definition 3.13. The *residual graph* G_i^r with respect to f has

(a) all edges of $G(A_i, B_i)$, directed as in $F(A_i, B_i)$, and
(b) all edges e of $F(A_i, B_i)$ — but in the reverse direction — where $f_e > 0$.

An $(\bar{s}, \bar{t})$-*cut* in G_i^r is any partition $(H, \overline{H})$ of the nodes of G_i^r having the property that no edge of G_i^r goes from H to $\overline{H}$.

(a) Original incompatibility graph with (squared) edge weights

(b) Flow graph and associated max flow
(Flow given for intermediate edges only; edges with no numbers have 0 flow.)

(c) Residual graph G^r with contracted nodes circled
(thick edges are doubly-directed.)

(d) Final acyclic graph G^*.
$U = \{x_1, x_2, x_3, t_1\}$, $V = \{x_4, x_5, x_6, t_2, t_3, t_4, t_5\}$, $W = \{x_7, t_6\}$, $X = \{x_8, t_7\}$
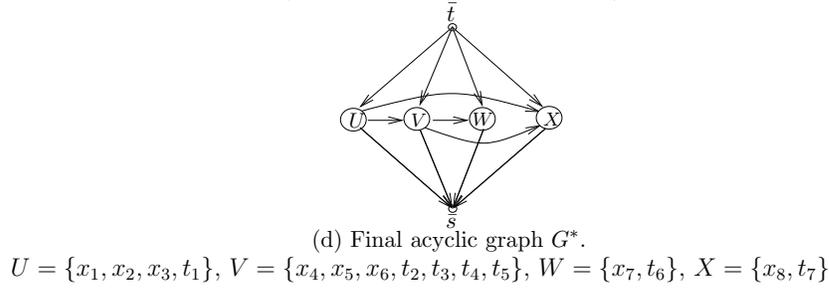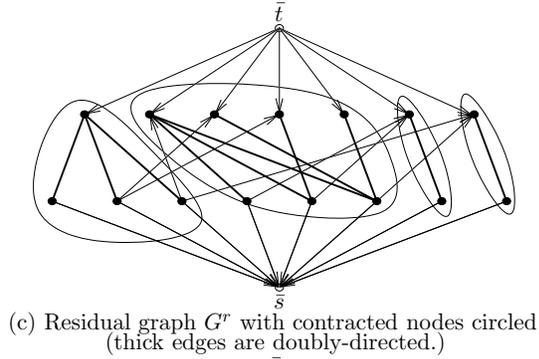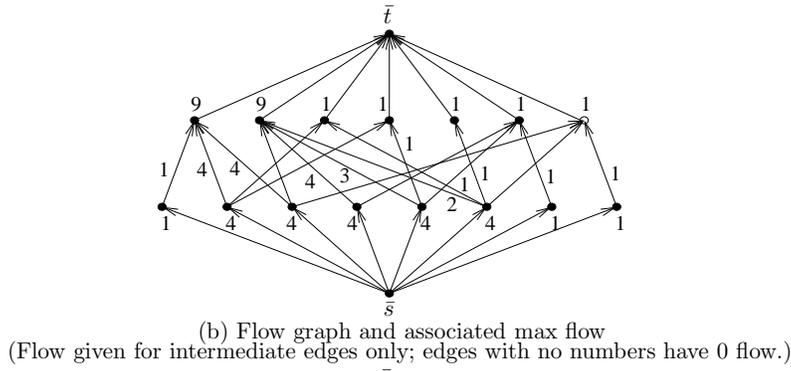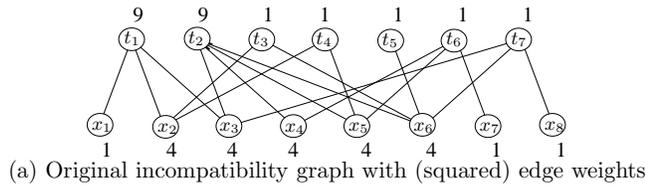
Figure 3: Characterizing ratio subsequences

It is easy to see that by this definition, H contains $\bar{s}$ and $\overline{H}$ contains $\bar{t}$. The definition of residual graph is based on the structure of $F(A_i, B_i)$ and the fact that the flow f saturates (is at capacity on) all edges adjacent to either $\bar{s}$ or $\bar{t}$. The Max-Flow-Min-Cut Theorem states that every $(\bar{s}, \bar{t})$-cut in G_i^r corresponds to a cut of capacity 1 in $F(A_i, B_i)$, which in turn corresponds to a cover of weight 1 in $G(A_i, B_i)$. This leads to the following result.

Lemma 3.14. *Let $(H, \overline{H})$ be a $(\bar{s}, \bar{t})$-cut in the residual graph G_i^r. Then the sets $I_1 = \overline{H} \cap A_i$, $J_1 = \overline{H} \cap B_i$, $I_2 = H \cap A_i$, and $J_2 = H \cap B_i$ have the property that $\frac{\|I_1\|}{\|J_1\|} = \frac{\|I_2\|}{\|J_2\|}$ can replace $\frac{\|A_i\|}{\|B_i\|}$ in (5) and the resulting sequence still satisfies (P2) and (P3).*

Proof. By Definition 3.13(a) all edges of $G(A_i, B_i)$ are in G_i^r, so in particular there can be no edge from any element in I_2 to any element in J_1. Thus $I_2 \cup J_1$ is compatible. Further, by Definition 3.13(b) there are no edges of G_i^r from H to $\overline{H}$, so the flow is conserved in H, and hence in $\overline{H}$. This implies $\|I_1\| = \|J_1\|$ and $\|I_2\| = \|J_2\|$, and thus the ratios are equal. Finally, since the flow f restricted to each of the subgraphs $F(I_1, J_1)$ and $F(I_2, J_2)$ continues to saturate the edges adjacent to $\bar{s}$ and $\bar{t}$, property (P3) continues to be satisfied on the replacement support pairs (I_1, J_1) and (I_2, J_2). □

Example 3.15 (continuation of Example 3.12). One min cut with respect to the flow in Figure 3(b) has $\overline{H} = \{x_1, x_2, x_3, t_1, \bar{t}\}$ and H its complement; this corresponds to the pairs $I_1 = \{x_1, x_2, x_3\}$, $J_1 = \{t_1\}$, $I_2 = \{x_4, x_5, x_6, x_7, x_8\}$, and $J_2 = \{t_2, t_3, t_4, t_5, t_6, t_7\}$, with squared ratios $\frac{9}{9} = \frac{14}{14}$.

Iteratively applying Lemma 3.14 to the resulting graphs $G(I_1, J_1)$ and $G(I_2, J_2)$ can produce various replacement subsequences for (A_i, B_i), depending upon the choice of min cuts and the number of times the lemma is applied. Picard and Queyranne [42] give a method to find all cuts for this flow problem, thereby allowing us to characterize all ratio subsequences associated with (A_i, B_i).

Definition 3.16. Write $G_i^*(X)$ for the result of modifying the residual graph G_i^r by contracting all edges contained in directed cycles.

The directed graph $G_i^*(X)$ is acyclic, is independent of the actual (max) flow f, and has nodes corresponding to a partition of the nodes of $A_i \cup B_i \cup \{\bar{s}, \bar{t}\}$. Furthermore, the nodes in any partition obtained by iteratively applying Lemma 3.14 must consist of unions of the sets corresponding to the nodes of G_i^*.

Definition 3.17. An *upper ideal* for $G_i^*(X)$ is any set I of nodes of $G_i^*(X)$ such that $v \in I$ whenever $u \in I$ and (u, v) is an edge of $G_i^*(X)$.

A partition $(H, \overline{H})$ is therefore a cut if and only if H is an upper ideal. Let $\mathcal{I}_i$ denote the set of upper ideals of $G_i^*(X)$, excluding the trivial ideal $\{\bar{s}\}$. The next corollary follows from this discussion.

Corollary 3.18. *The maximum size of any ratio subsequence that can replace (A_i, B_i) in (5) is equal to the number of vertices in $G_i^*(X) \setminus \{\bar{s}, \bar{t}\}$. Moreover, the ratio subsequences*

$$\frac{\|A_{i,1}'\|}{\|B_{i,1}'\|} = \frac{\|A_{i,2}'\|}{\|B_{i,2}'\|} = \cdots = \frac{\|A_{i,\ell}'\|}{\|B_{i,\ell}'\|}$$

are in bijection with nested sequences of sets in $\mathcal{I}_i$. ☐

This simplifies further. A *topological ordering* of $G_i^*(X)$ is any numbering of the vertices so that for every edge (u, v) of $G_i^*(X)$, vertex v is numbered lower than u. Every acyclic graph has at least one topological ordering.

Corollary 3.19. *The maximum-cardinality ratio subsequences of Corollary 3.18 are in bijection with the topological orderings of $G_i^*(X)$. In fact, any ratio subsequence for a particular pair (A_i, B_i) corresponds to a partition of the vertices of $G_i^*(X)$ according to one of these topological orderings.*

Example 3.20 (continuation of Example 3.12)**.** Applying Corollary 3.19 to the example in Figure 3, the only two acyclic orderings of G^* are (U, V, X, W) and (U, V, W, X), which results in the two maximal subsequences

$$(\mathcal{A}, \mathcal{B}) = \begin{cases} \big(\{x_1, x_2, x_3\}, \{t_1\}\big), \big(\{x_4, x_5, x_6\}, \{t_2, t_3, t_4, t_5\}\big), \big(\{x_7\}, \{t_6\}\big), \big(\{x_8\}, \{t_7\}\big) \\ \big(\{x_1, x_2, x_3\}, \{t_1\}\big), \big(\{x_4, x_5, x_6\}, \{t_2, t_3, t_4, t_5\}\big), \big(\{x_8\}, \{t_7\}\big), \big(\{x_7\}, \{t_6\}\big) \end{cases}$$

respectively, both of which have squared ratios of $\frac{9}{9} = \frac{12}{12} = \frac{1}{1} = \frac{1}{1} = 1$. The set of possible replacement subsequences for $(\mathcal{A}, \mathcal{B})$ corresponds to the twelve distinct contiguous partitions that can be formed from one of the above two sequences.

3.2.4 Valid support sequences

We next set up the combinatorial structure to give a canonical description of the vistal cell $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$.

Definition 3.21. Let (A_i, B_i) be a support pair for the minimal support $(\mathcal{A}, \mathcal{B})$. A *valid support sequence* for (A_i, B_i) is comprised of a set of pairs $(A_{i,1}', B_{i,1}'), \ldots, (A_{i,\ell}', B_{i,\ell}')$ with the following properties.

    (F1) The sets $A_{i,j}'$ and $B_{i,j}'$ are nonempty and partition A_i and B_i, respectively.
    (F2) The incompatibility graph $G(A_{i,j}', B_{i,j}')$ is connected for each $j = 1, \ldots, \ell$.
    (F3) Contracting the sets $A_{i,j}' \cup B_{i,j}'$ in $G(A_i, B_i)$ results in an acyclic graph.

21

Example 3.22 (continuation of Examples 3.12 and 3.20). Any support derived from the maximal supports in Example 3.20 is a valid support sequence, except for the two supports

$$\big(\{x_1, x_2, x_3\}, \{t_1\}\big), \big(\{x_4, x_5, x_6\}, \{t_2, t_3, t_4, t_5\}\big), \big(\{x_7, x_8\}, \{t_6, t_7\}\big)$$
$$\text{and } \big(\{x_1, x_2, x_3, x_4, x_5, x_6\}, \{t_1, t_2, t_3, t_4, t_5\}\big), \big(\{x_7, x_8\}, \{t_6, t_7\}\big),$$

whose final pairs do not correspond to connected subgraphs of the compatibility graph.

Lemma 3.23. *Let $X \in \mathcal{T}_n$ have associated (X,T)-geodesic with minimal support $(\mathcal{A}, \mathcal{B})$ satisfying (5), and for some index i let $(A'_{i,1}, B'_{i,1}), \ldots, (A'_{i,\ell}, B'_{i,\ell})$ be a valid support sequence for (A_i, B_i). There is an element $X' \in \mathcal{T}_n$ in the same orthant as X for which the geodesic between X' and T has support*

$$\mathcal{A}' = A_1, \ldots, A_{i-1}, A'_{i,1}, \ldots, A'_{i,\ell}, A_{i+1}, \ldots, A_k$$
$$\mathcal{B}' = B_1, \ldots, B_{i-1}, B'_{i,1}, \ldots, B'_{i,\ell}, B_{i+1}, \ldots, B_k$$

with

$$\frac{\|A_1\|}{\|B_1\|} < \cdots < \frac{\|A_{i-1}\|}{\|B_{i-1}\|} < \frac{\|A'_{i,1}\|}{\|B'_{i,1}\|} = \cdots = \frac{\|A'_{i,\ell}\|}{\|B'_{i,\ell}\|} < \frac{\|A_{i+1}\|}{\|B_{i+1}\|} < \cdots < \frac{\|A_k\|}{\|B_k\|}.$$

Further, for any pair $(A'_{i,j}, B'_{i,j})$ and any partition $I_1 \cup I_2$ of $A'_{i,j}$ and $J_1 \cup J_2$ of $B'_{i,j}$ in which $I_2 \cup J_1$ is compatible,

$$\frac{\|I_1\|}{\|J_1\|} > \frac{\|I_2\|}{\|J_2\|}.$$

Proof. For support pair (A_i, B_i), let $\tilde{\xi}$ and $\tilde{\tau}$ be the weights on the vertices of $G(A_i, B_i)$. Define X' by replacing the (squared) weights on X for each $a \in A'_{i,j}$ by

$$\tilde{\xi}'_a = \sum_{b \in E_j(a)} \frac{\tilde{\tau}_b}{\deg_j(b)},$$

where $E_j(a)$ is the set of vertices $b \in B'_{i,j}$ such that (a, b) is in the incompatibility graph, and $\deg_j(b)$ is the number of edges of the incompatibility graph from $A'_{i,j}$ to b. These values are all well-defined and positive by (F1) and (F2). Place the following flow f on the associated flow graph: for edge (a, b) where $a \in A'_{i,j}$ and $b \in B'_{i,j}$ for any $1 \leq j \leq l$, let the flow on that edge be $\tilde{\tau}_b / \deg_j(b)$; for all other edges, let the flow be 0. Then the flow into node b is exactly $\tilde{\tau}_b$ and the flow out of a is exactly $\tilde{\xi}_a$. Corollary 3.18 and property (F3) ensure that f is a max flow with respect to the flow graph, with flow value $\sum_{B_i} \tilde{\tau}_b = \sum_{A_i} \tilde{\xi}_a = 1$, and since flow is conserved between each $A'_{i,j}$ and $B'_{i,j}$, the original (un-normalized) weights satisfy

$$\frac{\|A'_{i,1}\|}{\|B'_{i,1}\|} = \cdots = \frac{\|A'_{i,\ell}\|}{\|B'_{i,\ell}\|} = \frac{\|A_i\|}{\|B_i\|}.$$

22

Finally, for a pair $(A'_{i,j}, B'_{i,j})$, let $I_1 \cap I_2$ and $J_1 \cap J_2$ be partitions of $A'_{i,j}$ and $B'_{i,j}$ respectively, in which $I_2 \cup J_1$ is compatible. This means that there are no edges of $G(A'_{i,j}, B'_{i,j})$ from I_2 to J_1, and since $G(A'_{i,j}, B'_{i,j})$ is connected there must be at least one edge going from I_1 to J_2. Since flow is positive on all edges of $G(A'_{i,j}, B'_{i,j})$, there is a net flow from I_1 away from J_1, and from the definition of ξ' it follows that $\frac{\|I_1\|}{\|J_1\|} > \frac{\|I_2\|}{\|J_2\|}$. $\qquad\square$

3.2.5 Canonical description of vistal cells

Finally, we extend Propositions 3.5 and 3.6 to describe all vistal cells associated with (X, T)-geodesics from points X in an orthant $\mathcal{O}$. Since a valid support sequence is determined by the combinatorics of the splits and not by their edge lengths, we can define the following.

Definition 3.24. A *valid support sequence for* $(\mathcal{O}, T)$ is a support $(\mathcal{A}, \mathcal{B})$ for which each maximal equality subsequence

$$\frac{\|A_i\|}{\|B_i\|} = \frac{\|A_{i+1}\|}{\|B_{i+1}\|} = \cdots = \frac{\|A_j\|}{\|B_j\|} \tag{11}$$

satisfies properties (F1)–(F3) with respect to the pair $(\bigcup_{\ell=i}^{j} A_\ell, \bigcup_{\ell=i}^{j} B_\ell)$. Write $G(\mathcal{O}, T)$ for the corresponding incompatibility graph $G(\mathcal{A}, \mathcal{B})$.

Theorem 3.25. *Fix a tree* $T \in \mathcal{T}_n$.

1. *Vistal cells associated with geodesics to* T *are exactly those of the form* $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$, *where* $(\mathcal{A}, \mathcal{B})$ *is a valid support sequence for* $(\mathcal{O}, T)$ *and* $\mathcal{S}$ *is a signature on* $(\mathcal{A}, \mathcal{B})$.
2. *The dimension of the vistal cell* $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ *is* $\dim(\mathcal{O}) - m(\mathcal{S})$, *where* $m(\mathcal{S})$ *is the number of "=" components in* $\mathcal{S}$.
3. *The representation by a valid support sequence and signature is unique up to reordering the support sets within each equality subsequence of* $\mathcal{S}$.

Proof. Claim 1. Let $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ be a vistal cell containing an interior point ξ. The definition of support and the fact that ξ is positive implies that (F1) and (F3) hold for $G(\mathcal{O}, T)$. Now suppose that (F2) fails to hold; that is, some $G(A_i, B_i)$ has a partition into two disjoint subgraphs on vertex sets $I_1 \cup J_1$ and $I_2 \cup J_2$, respectively. Let f be the max flow in $G(A_i, B_i)$. Since (P3) is satisfied, f saturates all arcs adjacent to the source and sink. But since flow in each of the disjoint subgraphs $G(I_1, J_1)$ and $G(I_2, J_2)$ is self-contained, $\frac{\|I_1\|}{\|J_1\|} = \frac{\|I_2\|}{\|J_2\|} = \frac{\|A_i\|}{\|B_i\|}$. This means that the corresponding tree X satisfies one of its (P3) inequalities at equality, so ξ cannot be in the interior of $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$, a contradiction. Thus (F2) is also satisfied, so $(\mathcal{A}, \mathcal{B})$ is a valid support sequence with respect to $(\mathcal{O}, T)$.

Conversely, let $(\mathcal{A}, \mathcal{B})$ be a valid support sequence with respect to $(\mathcal{O}, T)$. Consider a ratio subsequence (11) with all terms equal. Since $(\mathcal{A}, \mathcal{B})$ is a valid support sequence, Lemma 3.23 constructs positive weights X^ℓ on the edges indexed by A_ℓ, for $\ell = i, \ldots, j$, so that (11) holds and all (P3) inequalities are strict inside each support pair. Now for each maximal-length

equal-ratio subsequence, scale the vectors of each term by the same positive multiplier λ_{ij} so that the sequence of multipliers λ_{ij} is increasing with the indices. The scaled x^ℓ vectors concatenate into a vector X in the interior of $\mathcal{O}$ having the correct signature indicated by $\mathcal{S}$, and for which the (P2) inequalities hold strictly between the equal-ratio subsequences. The squared point ξ corresponding to X therefore lies interior to $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$, and the desired result follows.

Claim 2. The vector ξ constructed in the proof of Claim 1 is positive in $\mathcal{O}$, satisfies all (P3) inequalities strictly, and satisfies all (P2) inequalities strictly for which the corresponding component of $\mathcal{S}$ is "$\leq$". Therefore the dimension of $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ is determined entirely by the set of equalities defined by the component of $\mathcal{S}$ that are "$=$". Since these are linearly independent, the dimension is as stated.

Claim 3. Let $F = \mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ and $F' = \mathcal{V}^2(T, \mathcal{O}'; \mathcal{A}', \mathcal{B}'\mathcal{S}')$ be two representations of vistal cells, defined by valid supports $(\mathcal{A}, \mathcal{B})$ and $(\mathcal{A}', \mathcal{B}')$ respectively. Any permutation of support pairs within an equality subsequence (11) results in the same set of equalities, so if the representations differ only by such a permutation, then $F = F'$. Conversely, suppose $F = F'$. Since all cell constraint inequalities other than those specified by $\mathcal{S}$ are satisfied strictly, the set of equalities dictated by $\mathcal{S}$ define the affine hulls of F and F'. This means that the two associated equality systems are row-equivalent. Now suppose that the supports $(\mathcal{A}, \mathcal{B})$ and $(\mathcal{A}', \mathcal{B}')$ do not comprise the same sets; that is, by symmetry the two sets A_i and A_j both have nonempty intersection with the same set A'_k. Since the variables of A'_k do not appear in any other A'_ℓ for $\ell \neq k$, no row transformation of the equality system for F' could possibly separate the variables in $A_i \cap A'_k$ from those in $A_j \cap A'_k$. Thus the two equality systems are not the same, a contradiction. □

Corollary 3.26. *Distinct vistal cells have disjoint relative interiors.*

Proof. Let ξ be an element in the relative interior of two faces in $\mathcal{T}_n^2$, given by valid representations. Then ξ satisfies (F2) and (F3) with respect to both faces, and by Theorem 3.25 the only way this could happen is if the faces coincide. □

3.3   Vistal subdivisions

Theorem 3.25 allows us a purely combinatorial way of describing vistal cells. This gives us the machinery to prove the principal result of the section, namely that the vistal cells are the faces of a polyhedral subdivision of tree space under the squaring map. To make this precise, we start with some definitions concerning polyhedra; see [50, Lecture 5] for further background.

Definition 3.27. A *polyhedral complex* Σ is a finite collection of polyhedra such that

(C1) every polyhedral face of every polyhedron in Σ is a polyhedron in Σ;

(C2) the intersection of any pair of polyhedra in Σ is a face of each.

The *dimension* of Σ is the largest dimension of a polyhedron in Σ. The *facets* of Σ are the maximal cells. The *underlying set* of Σ is the union $\bigcup_{V \in \Sigma} V$ of the polyhedra in Σ.

Example 3.28. Tree space $\mathcal{T}_n$ has a *natural polyhedral structure* as the underlying space of a polyhedral complex whose polyhedra are its orthants. This polyhedral structure is unchanged by the squaring map, and thus also found in T_n^2.

The relation between vistal cells and orthants is one of refinement, in the following sense.

Definition 3.29. Let Σ and Σ' be polyhedral complexes. Then Σ' is a *subdivision* of Σ (it is also said that Σ' *refines* Σ) if each polyhedron in Σ' is contained in a single polyhedron in Σ.

Theorem 3.30. *For tree space $\mathcal{T}_n$ and fixed source tree T, the vistal cells of $\mathcal{T}_n^2$ with respect to T refine the natural polyhedral structure of $\mathcal{T}_n^2$ to form a vistal polyhedral subdivision of $\mathcal{T}_n^2$.*

Proof. The vistal cells are polyhedra whose union is $\mathcal{T}_n^2$ by Propositions 3.5 and 3.6.

For (C1), we show that changing any of the inequalities defining a vistal cell to equality results in a set that can be described as a vistal cell. Let $V = \mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ be a vistal cell, so that by Lemma 3.23, $(\mathcal{A}, \mathcal{B})$ is a valid support sequence, and let F be a proper face of V obtained by setting one of its boundary inequalities to equality. There are three types of inequalities that define F: (P2) constraints, nonnegativity constraints, and (P3) constraints.

For the (P2) constraints, consider the inequality $\frac{\|A_i\|}{\|B_i\|} < \frac{\|A_{i+1}\|}{\|B_{i+1}\|}$, where the corresponding component of the signature $\mathcal{S}$ is "$\leq$". Let $\mathcal{S}'$ be obtained from $\mathcal{S}$ by setting this inequality to "$=$". Since neither $\mathcal{O}$ nor $(\mathcal{A}, \mathcal{B})$ has changed, this constitutes a valid support sequence, and $F = \mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S}')$.

For the nonnegativity constraints, consider the inequality $x_e > 0$, where e is a split indexing a coordinate of $\mathcal{O}$. Let A_i be the set containing e. Now remove e from $G(\mathcal{O}, T)$. This splits $G(A_i, B_i)$ into components corresponding to partitions $(A'_1, B'_1), \ldots, (A'_\ell, B'_\ell)$ of (A_i, B_i). Because these partitions correspond to separate components in $G(A_i, B_i)$, they can appear in any order in a valid support sequence for F. Thus every point in F must satisfy every (P2) inequality between the pairs (A'_i, B'_i) at equality, since otherwise the (A'_j, B'_j) sets could be interchanged so that some (P3) condition is violated. First consider the case where all of the A'_j are nonempty. Define the support $(\mathcal{A}', \mathcal{B}')$ by inserting $(A'_1, B'_1), \ldots, (A'_\ell, B'_\ell)$ in place of (A_i, B_i) in $(\mathcal{A}, \mathcal{B})$:

$$(\mathcal{A}', \mathcal{B}') = (A_1, B_1), \ldots, (A_{i-1}, B_{i-1}), (A'_1, B'_1), \ldots, (A'_\ell, B'_\ell),$$
$$(A_{i+1}, B_{i+1}), \ldots, (A_k, B_k)$$

and extend the signature $\mathcal{S}$ to $\mathcal{S}'$ by adding "$=$" signs between each of the sets in the primed subsequence. Then $(\mathcal{A}', \mathcal{B}')$ is valid, and $F = \mathcal{V}^2(T, \mathcal{O} \setminus \{e\}; \mathcal{A}', \mathcal{B}'; \mathcal{S}')$.

Now suppose that one of the support pairs (A'_j, B'_j) has $A'_j = \varnothing$. The associated ratio must be 0, which implies in turn that every ratio corresponding to the pairs $(A'_1, B'_1), \ldots, (A'_\ell, B'_\ell)$

25

is 0. Furthermore, the ratios are also 0 for any earlier support pairs. So $x_f = 0$ for every $f \in \overline{H}_i = A_1 \cup \cdots \cup A_i$. In this case set

$$\overline{\mathcal{O}}' = \mathcal{O} \setminus \overline{H}_i$$
$$(\overline{\mathcal{A}}', \overline{\mathcal{B}}') = (A_{i+1}, B_{i+1}), \ldots, (A_k, B_k)$$
$$\overline{\mathcal{S}}' = \mathcal{S} \text{ restricted to the last } k - i \text{ pairs of the sequence.}$$

By Remark 1.5 we have been ignoring the non-positive ratios; however, they still exist if there are common edges between X and T. In this case, the edges $B_1 \cup \cdots \cup B_i$ become common edges, and are added to the 0-valued ratio if it already exists, or form it anew, if it does not. Again $(\overline{\mathcal{A}}', \overline{\mathcal{B}}')$ is valid, and $F = \mathcal{V}^2(T, \overline{\mathcal{O}}'; \overline{\mathcal{A}}', \overline{\mathcal{B}}'; \overline{\mathcal{S}}')$.

Next consider the (P3) constraints. For some support pair (A_i, B_i) let $I_1 \cup I_2$ and $J_1 \cup J_2$ be partitions of A_i and B_i with $I_2 \cup J_1$ compatible, and consider the constraint

$$\frac{\|I_1\|}{\|J_1\|} > \frac{\|I_2\|}{\|J_2\|}.$$

Let $(A_1', B_1'), \ldots, (A_k', B_k')$ and $(A_{k+1}', B_{k+1}'), \ldots, (A_\ell', B_\ell')$ be pairs corresponding to the components of $G(I_1, J_1)$ and $G(I_2, J_2)$, respectively.

First consider the case where all of the A_j' and B_j' are nonempty. The same nonempty sets argument as above applies, and we obtain the the face $F = \mathcal{V}^2(T, \mathcal{O}; \mathcal{A}', \mathcal{B}'; \mathcal{S}')$ with $(\mathcal{A}', \mathcal{B}')$ and $\mathcal{S}'$ defined as in the nonempty-set case above.

Next suppose that one of the sets (A_j', B_j') has $A_j' = \varnothing$. As in the empty-set case above, this forces x_f to be 0 for every $f \in S_i = A_1 \cup \cdots \cup A_i$, and so $F = \mathcal{V}^2(T, \overline{\mathcal{O}}'; \overline{\mathcal{A}}', \overline{\mathcal{B}}'; \overline{\mathcal{S}}')$ with $\overline{\mathcal{O}}'$, $(\overline{\mathcal{A}}', \overline{\mathcal{B}}')$ and $\overline{\mathcal{S}}'$ defined as in the empty-set case above.

Now suppose that one of the sets (A_j', B_j') has $B_j' = \varnothing$. This forces the ratios for every pair in $(A_1', B_1'), \ldots, (A_\ell', B_\ell')$ to be ∞, which in turn means that $x_f = 0$ for every $f \in \tilde{S}_i' = B_{i+1} \cup \cdots \cup B_k$. Thus if we define

$$\widetilde{\mathcal{O}}' = \mathcal{O} \setminus \tilde{S}_i'$$
$$(\tilde{\mathcal{A}}', \tilde{\mathcal{B}}') = (A_1, B_1), \ldots, (A_{i-1}, B_{i-1})$$
$$\tilde{\mathcal{S}}' = \mathcal{S} \text{ restricted to the first } i - 1 \text{ pairs of the sequence,}$$

then again $(\tilde{\mathcal{A}}', \tilde{\mathcal{B}}')$ is a valid sequence, and so we obtain the face $F = \mathcal{V}^2(T, \tilde{\mathcal{O}}'; \tilde{\mathcal{A}}', \tilde{\mathcal{B}}'; \mathcal{S}')$. As before, the edges $A_i \cup \cdots \cup A_k$ become common edges, and hence be added to the ∞-valued ratio if it exists and otherwise form that ratio.

Finally, suppose that there are pairs $(A_{j'}', B_{j'}')$ and $(A_{j''}'', B_{j''}'')$ with $A_{j'}' = B_{j''}'' = \varnothing$. This forces all of the x_f where f is not a common edge to be 0, and we just get the face corresponding to the common edges.

For (C2), suppose that V and V' are vistal cells, so that $V \cap V'$ is a convex polyhedron. Let $F \subseteq V$ and $F' \subseteq V'$ be minimal faces of V and V', respectively, containing $V \cap V'$. Then

by (C1), F and F' are vistal cells, and since F and F' are minimal, then there must be a $p \in V \cap V'$ in the relative interior of F and a $p' \in V \cap V'$ in the relative interior of F'. It follows that the midpoint of the line segment joining p to p' must lie in the relative interiors of both F and F', and Corollary 3.26 then implies that $F = F'$. Thus $F = F' \subseteq V \cap V'$, whence $V \cap V' = F = F'$ and (C2) follows. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

3.4 Examples of vistal complexes

Example 3.31. To demonstrate Theorem 3.30, consider the incompatibility graph from Figure 3 and treat it as the incompatibility graph for two trees T and X. Take values on T as given in the figure, and consider the vistal cell $V = \mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ defined by

$$\mathcal{O} = \{x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8\}$$
$$(\mathcal{A}, \mathcal{B}) = \big(\{x_1, x_2, x_3\}, \{t_1\}\big), \big(\{x_4, x_5, x_6, x_7, x_8\}, \{t_2, t_3, t_4, t_5, t_6, t_7\}\big)$$
$$\mathcal{S} = (\leq).$$

This is a valid sequence, and in particular, using Lemma 3.23 we can assign weights as follows.

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
2	2	2	$3\frac{1}{3}$	$4\frac{1}{3}$	$5\frac{1}{2}$	$\frac{1}{3}$	$\frac{1}{2}$

(The first three weights have additionally been scaled so that (P2) is satisfied strictly.) Here are examples of the three types of faces of V.

- Setting the single (P2) constraint to equality: this gives the face corresponding to the numbers in Figure 3.
- Setting $x_j = 0$: for $j \neq 5, 6$ the face has the same structure as the cell V, except that x_j is removed from the corresponding sets. For $j = 5, 6$, removal of x_j disconnects (A_2, B_2) by isolating t_4 or $\{t_3, t_5\}$, respectively, and thus setting x_5 or x_6 to 0 collapses the face to the single origin point.
- Setting the (P3) constraint with $I_1 = \{x_4, x_5, x_6\}$, $J_1 = \{t_2, t_3, t_4, t_5\}$, $I_2 = \{x_7, x_8\}$, and $J_2 = \{t_6, t_7\}$ to equality: here

$$\frac{\|I_1\|^2}{\|J_1\|^2} = \frac{79}{66} > \frac{5}{12} = \frac{\|I_2\|^2}{\|J_2\|^2}.$$

  Now $G(I_2, J_2)$ is not connected, and has nontrivial components on vertex sets $\{x_7, t_6\}$ and $\{x_8, t_7\}$. Thus the face obtained by setting the above inequality to equality is $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}', \mathcal{B}'; \mathcal{S}')$, where

$$(\mathcal{A}', \mathcal{B}') = \big(\{x_1, x_2, x_3\}, \{t_1\}\big), \big(\{x_4, x_5, x_6\}, \{t_2, t_3, t_4, t_5\}\big), \big(\{x_7\}, \{t_6\}\big),$$
$$\big(\{x_8\}, \{t_7\}\big)$$
$$\mathcal{S}' = (\leq, =, =).$$

Example 3.32. Figure 4 gives the restriction of a vistal polyhedral subdivision to a maximal orthant in $\mathcal{T}_5$. The trees are depicted in Figure 4(a), with $t_1 = t_2 = t_3 = 1$. Figures 4(b) and 4(c) depict the vistal cells in orthant $\mathcal{O}(\{x_1, x_2, x_3\})$ before and under the squaring map, respectively, as they intersect with the sets $x_1 + x_2 + x_3 = 1$ and $\xi_1 + \xi_2 + \xi_3 = 1$, respectively. The vistal cells are labeled by the corresponding ratio sequences, using "=" or "<" to indicate the behavior of points in the interior of the cell. We also label the six cells of lower dimension that are the intersections of the vistal facets. Some of the vistal cells before the squaring map are not polyhedral, because the boundary equations are those of circular cones.

3.5 Multivistal complexes

It is a straightforward matter to extend vistal cells to the case where there is a collection $\boldsymbol{T} = \{T^1, \ldots, T^r\}$ of source trees in $\mathcal{T}_n$, and we are interested in the set of points $X \in \mathcal{T}_n$ for which the geodesic to each tree in $\boldsymbol{T}$ has a specified combinatorial structure.

Definition 3.33. A *premultivistal cell* for a collection $\boldsymbol{T}$ of trees is a set of the form

$$\mathcal{V}(\boldsymbol{T}; \mathcal{O}; \mathcal{A}^{\boldsymbol{T}}, \mathcal{B}^{\boldsymbol{T}}) = \bigcap_{\ell=1}^{r} \mathcal{V}(T^\ell, \mathcal{O}; \mathcal{A}^\ell, \mathcal{B}^\ell),$$

where $\mathcal{V}(T^\ell, \mathcal{O}; \mathcal{A}^\ell, \mathcal{B}^\ell)$ are previstal cells, $\mathcal{O} \subseteq \mathcal{T}_n$ is an orthant, and

$$(\mathcal{A}^{\boldsymbol{T}}, \mathcal{B}^{\boldsymbol{T}}) = \{(\mathcal{A}^1, \mathcal{B}^1) \ldots, (\mathcal{A}^r, \mathcal{B}^r)\}$$

is a collection of support pairs for (T^i, X)-geodesics. A *multivistal cell* is the image in $\mathcal{T}_n^2$ of a premultivistal cell.

Corollary 3.34. *The multivistal cells of tree space $\mathcal{T}_n$ for any fixed set source trees refine the natural polyhedral structure of $\mathcal{T}_n$ to form a* multivistal polyhedral subdivision *of $\mathcal{T}_n$.*

Proof. The common refinement of any finite collection of polyhedral subdivisions of a given polyhedral complex is a polyhedral subdivision of the same polyhedral complex, and so the result follows from Theorem 3.30. $\square$

Remark 3.35. The previstal cells for any fixed source tree form a subdivision of $\mathcal{T}_n$, called a *premultivistal complex*, that is the image of the corresponding multivistal polyhedral subdivision of $\mathcal{T}_n^2$ under the inverse $\xi \to \sqrt{\xi}$ of the squaring map, which is a homeomorphism. However, the cells in this subdivision are not polyhedral. One might hope that the premultivistal complex is a *CW complex*, in the standard topological sense (see [36], for example), but it is not, for the same reason that multivistal polyhedral subdivisions are not CW complexes: the closed cells are not images of closed balls under continuous maps (a cone of positive dimension fails to be compact). The situation can be remedied by considering the *link* L_n

(a) Trees X and T.

$$\frac{\|\{x_1\}\|}{\|\{t_1\}\|} < \frac{\|\{x_2,x_3\}\|}{\|\{t_2,t_3\}\|}$$

$$\frac{\|\{x_1\}\|}{\|\{t_1\}\|} = \frac{\|\{x_2,x_3\}\|}{\|\{t_2,t_3\}\|}$$

$$\frac{\|\{x_1\}\|}{\|\{t_1\}\|} < \frac{\|\{x_2\}\|}{\|\{t_2\}\|} = \frac{\|\{x_3\}\|}{\|\{t_3\}\|}$$

$$\frac{\|\{x_1,x_2,x_3\}\|}{\|\{t_1,t_2,t_3\}\|}$$
(cone path)

$$\frac{\|\{x_2\}\|}{\|\{t_2\}\|} = \frac{\|\{x_1,x_3\}\|}{\|\{t_1,t_3\}\|}$$

$$\frac{\|\{x_1\}\|}{\|\{t_1\}\|} < \frac{\|\{x_2\}\|}{\|\{t_2\}\|} < \frac{\|\{x_3\}\|}{\|\{t_3\}\|}$$

$$\frac{\|\{x_1\}\|}{\|\{t_1\}\|} = \frac{\|\{x_2\}\|}{\|\{t_2\}\|} = \frac{\|\{x_3\}\|}{\|\{t_3\}\|}$$

$$\frac{\|\{x_2\}\|}{\|\{t_2\}\|} < \frac{\|\{x_1,x_3\}\|}{\|\{t_1,t_3\}\|}$$

$$\frac{\|\{x_1\}\|}{\|\{t_1\}\|} = \frac{\|\{x_2\}\|}{\|\{t_2\}\|} < \frac{\|\{x_3\}\|}{\|\{t_3\}\|}$$

$$\frac{\|\{x_2\}\|}{\|\{t_2\}\|} < \frac{\|\{x_1\}\|}{\|\{t_1\}\|} < \frac{\|\{x_3\}\|}{\|\{t_3\}\|}$$
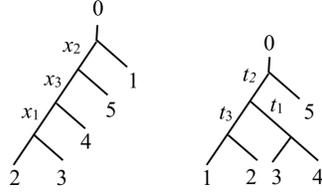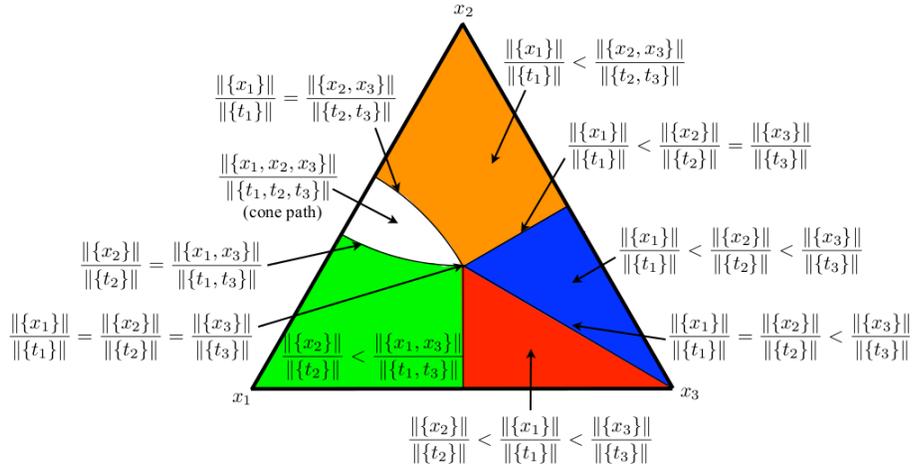
(b) A cross-section of the orthant corresponding to tree topology X before the squaring map.
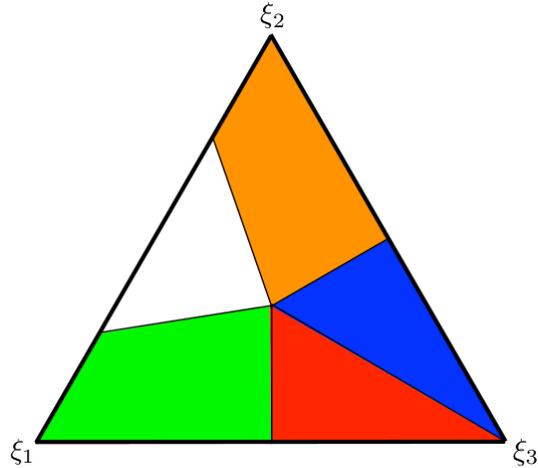
(c) A cross-section of the orthant corresponding to tree topology X under the squaring map. Vistal cells are labelled as in Figure 4(b)

Figure 4: The vistal polyhedral subdivision between variable tree X and fixed tree T in $\mathcal{T}_5$.

of the origin in $\mathcal{T}_n^2$, namely the set of trees whose edge lengths sum to 1. Intersecting L_n with any multivistal polyhedral subdivision yields a polyhedral CW complex whose preimage under the squaring map is a (non-polyhedral) CW complex. Thus a premultivistal complex is essentially a (noncompact, unbounded) cone over a CW complex.

Remark 3.36. In general, the number of vistal facets is exponential in n, even within a single orthant [39]. Thus an efficient method to move through the vistal facets – or prune the list of relevant ones – would likely improve calculation time of the mean.

4 Computing the mean in tree space

Although the algorithm to calculate the mean in Section 2.2 is simple and seems to perform well for small data sets, Remark 2.6 indicates that its convergence rate is sublinear, so in theory it is a poor iterative method. This section outlines a general framework for a descent method to find the mean of a set $T^1, \ldots, T^r$ of n-trees. It promises to accelerate the convergence considerably by generalizing powerful nonlinear programming techniques to apply to optimization in tree space.

4.1 Optimality criteria

We start by analyzing the variance function $S(x)$ of a variable point $X \in \mathcal{T}_n$ whose components are represented by the variable vector x. For $\ell = 1, \ldots, r$, let γ_ℓ be the geodesic from X to T_ℓ, with associated support pair $(\mathcal{A}^\ell, \mathcal{B}^\ell)$. By summing the lengths $L(\gamma_\ell)$ of these geodesics as given by Eq. (4), write the variance in $\mathcal{T}_n$ as

$$S(x) = \sum_{\ell=1}^{r} (L(\gamma_\ell))^2 = \sum_{\ell=1}^{r} \left[\sum_{i=1}^{k_\ell} \left(\|A_i^\ell\| + \|B_i^\ell\| \right)^2 \right]$$

with its derivative given by Eq. (8). Consider this function in its $\mathcal{T}_n^2$-version $S^2(\xi)$ as given in Definition 3.4. Using the notation

$$\bar{\xi}_i^\ell = \sum_{e \in A_i^\ell} \xi_e$$

and

$$\delta_i^\ell = \begin{cases} +1 & \text{if } A_i^l \text{ and } B_i^l \text{ are disjoint} \\ -1 & \text{if } A_i^l = B_i^l \text{ are made up of common edges,} \end{cases}$$

then the corresponding pullback function for $\xi \in \mathcal{T}_n^2$ can be derived from Eqs. (6) and (7):

$$S^2(\xi) = \sum_{\ell=1}^{r} \sum_{i=1}^{k_\ell} \left(\delta_i^\ell \sqrt{\bar{\xi}_i^\ell} + \|B_i^\ell\| \right)^2. \tag{12}$$

If $i(e, \ell)$ denotes the index of the set A_i^ℓ containing e, then the gradient of S^2 can be obtained from Eq. (12):

$$\frac{\partial S^2(\xi)}{\partial \xi_e} = \sum_{\ell=1}^{r} \left(1 + \delta_{i(e,\ell)}^\ell \frac{||B_{i(e,\ell)}^\ell||}{\sqrt{\bar{\xi}_{i(e,\ell)}^l}} \right).$$

The differentiability of S transfers to S^2, as well.

Corollary 4.1. *The function $S^2(\xi)$ is continuously differentiable on the interior of every maximal orthant $\mathcal{O}$.*

Proof. The inverse of the squaring map is continuously differentiable on the interior of $\mathcal{O}$. Now apply Theorem 2.2. □

The function $S^2(\xi)$ is not necessarily convex on $\mathcal{T}_n^2$. By Proposition 2.1, however, it does have a unique local minimum, which is therefore the mean. Consequently, optimality conditions for the function $S^2(\xi)$ on $\mathcal{T}_n^2$ can be based on its behavior in *any one* of the multivistal facets in which ξ lies. In particular, we have the following important result.

Corollary 4.2. *The squared image $\overline{\mathcal{X}}$ of the Fréchet mean $\overline{X}$ must satisfy $\nabla S^2(\overline{\mathcal{X}}) = 0$ on its orthant $\mathcal{O}(\overline{\mathcal{X}})$. If $\overline{\mathcal{X}}$ lies interior to a maximal orthant $\mathcal{O}$, then $\overline{\mathcal{X}}$ is the squared image of the mean if and only if the gradient satisfies $\nabla S^2(\overline{\mathcal{X}}) = 0$. These statements are true regardless on which multivistal facet of $\mathcal{O}$ the variance function is derived.*

Proof. Since by Corollary 4.1 the gradient is independent of which vistal facet it is calculated from, the gradient $\nabla S^2(\overline{\mathcal{X}})$ must be zero on any of them in order $\overline{\mathcal{X}}$ to be optimal. Conversely, since S^2 attains a unique minimum on $\mathcal{T}_n^2$, it follows that X must be the mean whenever $\nabla S^2(\overline{\mathcal{X}}) = 0$ on an entire maximal orthant. □

Remark 4.3. When a point $\mathcal{X}$ lies on the boundary of a maximal orthant, the gradient $\nabla S^2(\mathcal{X})$ may be zero on $\mathcal{O}(\mathcal{X})$ even if $\mathcal{X}$ is not the squared mean, since there may be a maximal orthant $\mathcal{O} \supset \mathcal{O}(\mathcal{X})$ having a point with smaller variance than $\mathcal{X}$. Finding $\mathcal{O}$ from $\mathcal{X}$ can be quite difficult, since of $\partial S^2(\mathcal{X})/\xi_e$ may be undefined or infinite for $e \notin \mathcal{O}(\mathcal{X})$. Furthermore, directional derivatives may fail to be continuous along orthant boundaries. This issue presents serious optimization difficulties in locating sample means, since there is ample evidence that reasonably evenly distributed data in tree space yield means that are likely to occur on orthant boundaries, or indeed, even to lie at the origin; see Section 5.3. Thus optimality conditions for the mean when it occurs on orthant boundaries is an important topic of further research.

4.2 A descent method to compute the mean

In spite of Remark 4.3, we can suggest a basic method for finding the mean in tree space. The general idea is to start with some feasible tree, and construct a sequence of trees whose variance function is decreasing, until arriving at the mean tree.

Algorithm 4.4 (Descent method for computing the mean).

INPUT	Trees $T^1, \ldots, T^r$ in $\mathcal{T}_n$
OUTPUT	The mean tree for $T^1, \ldots, T^r$
INITIALIZE	Choose some good starting point $\xi^0 \in \mathcal{T}_n^2$, for example, by running Sturm's algorithm for a predetermined number of iterations.
WHILE	the mean has not been found:

   DO 1. Find the set $\mathcal{M}$ of all maximal orthants containing ξ^t.

        2. For each $\mathcal{V} \in \mathcal{M}$, choose a point $u_\mathcal{V}^0$ in the interior of $\mathcal{V}$.

        3. Use a nonlinear interior point/penalty function method to find a local minimum $u_\mathcal{V}^*$ of S^2 in $\mathcal{V}$.

        4. If $u_\mathcal{V}^* \neq \xi^t$ for any $\mathcal{V} \in \mathcal{M}$, then choose the $u_\mathcal{V}^*$ with minimum $S^2(u_\mathcal{V}^*)$, and set $\xi^{t+1} = u_\mathcal{V}^*$.

   END WHILE-DO

  RETURN ξ^t

The local minimum search in Step 3 should be both straightforward and reasonably fast, and the accuracy of the points ξ^k as representing the true local minimum of course depends upon the method used to find it. Since the function S^2 is continuously differentiable on all $\mathcal{V} \in \mathcal{M}$, the search in fact finds a local minimum on the orthant $\mathcal{V}$. Since all neighboring orthants are searched from ξ^t, it follows that whenever all of these local searches converge back to ξ^t then ξ^t must necessarily be the mean. Finally, the algorithm terminates after a finite number of iterations, since no two ξ^t in the sequence can lie in the same orthant. The number of iterations depends both on the number of iterations t and also the size of $\mathcal{M}$, each of which may be exponentially large. Thus it is important for the implementation that a good starting point ξ^0 be found, and that a good method be used to determine descent directions in the set of maximal orthants adjacent to the point ξ^t. In general, better local search techniques and starting solutions, perhaps through a hybrid of Sturm's Algorithm and descent methods, could improve the accuracy and reliability of procedures to calculate the mean.

5   Properties and applications of the mean

This section contains a series of remarks, results, and computational studies related to the Fréchet mean in tree space.

All synthetic examples in this section will be given using the trees in Figure 5. This figure depicts three adjacent orthants in $\mathcal{T}_4$, "flattened out" into the plane, to make the visualization easier. The edges e_1' and e_2 are not compatible, so the (e_1', e_2)-orthant (shaded in Figure 5) is not part of $\mathcal{T}_4$. For tree T^1 (respectively, trees T^2 and T^3), we specify its interior edge lengths by a pair of coordinates (e_1, e_2) (respectively, (e_1, e_2') and (e_1', e_2')). The geodesic between any pair of these trees is a straight line, unless it would cross the shaded region, in which case the geodesic is the *cone path*, consisting of the two legs joining the given points to the origin.

Likewise, $\overline{T}$ is the Euclidean barycenter unless it lies in the shaded region, in which case $\overline{T}$ is the point on the boundary of the shaded region that minimizes the sum of the squared geodesic distances to the three trees.
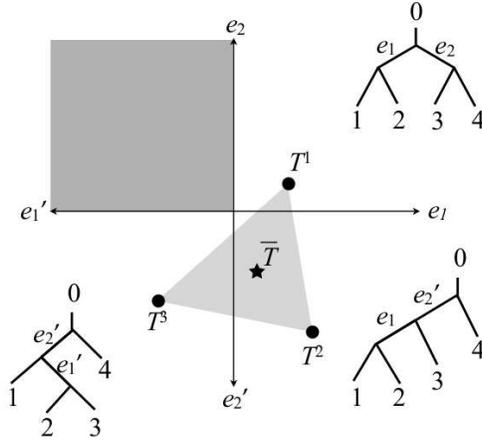
Figure 5: Example for the remarks.

5.1 Composition of the mean tree

The topology of the mean tree depends on both the topologies and the edge lengths of the sample trees. Consider, for example, trees $T^1 = (3,1)$ and $T^3 = (1,3)$. The mean between these two trees is the midpoint $\overline{T}^2 = (1,1)$ of the segment joining them. Changing both edge lengths of T^1 to 5, however, yields a midpoint $\overline{T}^1 = (1,2)$; similarly, by symmetry, changing both edge lengths of T^3 to 5 yield the midpoint $\overline{T}^3 = (1,2)$. That said, in general we can give some indication of what edges lie in the mean tree.

Lemma 5.1. *Every edge of the mean tree is an edge of some sample tree. Furthermore, if an edge appears in all sample trees, then it must also appear in the mean tree.*

Proof. If a tree contains an edge not in any sample tree, then contracting this edge gives a tree with a smaller variance function. Now suppose that the edge e is contained in all sample trees, and thus is compatible with all other edges in the sample trees. Since the mean contains only edges from the sample trees, edge e is also compatible with all edges in the mean tree. Thus if the mean tree does not contain e, we can add in e with length equal to the minimum of its lengths in the sample trees, yielding a tree that is closer to all the sample trees, which is a contradiction. $\square$

5.2 Other notions of consensus tree

Several authors have proposed notions of "center" for a set $\boldsymbol{T} = \{T^1, \dots, T^r\}$ of points in $\mathcal{T}_n$. The Euclidean or combinatorial properties of these centers make them useful for representing consensus trees. Here we compare three such centers with the Fréchet mean $\overline{T}$ of $\boldsymbol{T}$. All of these centers agree when $\boldsymbol{T}$ lies entirely in a single orthant of $\mathcal{T}_n$, but fail to agree for more globally distributed samples from tree space.

Example 5.2 (**The majority-rule consensus (MRC) tree**). First introduced by Margush and McMorris [33], this is the tree whose edge set is comprised of those edges that appear in at least half of the trees in $\boldsymbol{T}$. It, or an variation, is widely used in the phylogenetics literature. The topology of $\overline{T}$ is not a refinement of the MRC tree, unlike many other consensus methods [13]. For example, consider the trees in Figure 5 with coordinates $T^1 = (1,1)$, $T^2 = (1,1)$, and $T^3 = (5,6)$. The mean of these trees is the Euclidean barycenter $\overline{T}^3 = (1,2)$, while the MRC tree has the topology of tree T^2, so neither tree is a refinement of the other.

Example 5.3 (**Sturm's inductive mean**). The inductive mean (Definition 2.3) of the set $\boldsymbol{T}$, for some ordering of $\boldsymbol{T}$, does not coincide with $\overline{T}$, and it can differ depending upon the ordering. Consider the trees in Figure 5 with coordinates $T^1 = (3,10)$, $T^2 = (3,3)$, and $T^3 = (10,3)$. Either order having T^1 and T^3 first yields the inductive mean $\tilde{T}^2 = (1,1)$. Either order having T^1 and T^2 first yields the inductive mean $\tilde{T}^3 = (0.390, 0.117)$, and either order having T^2 and T^3 first yields the inductive mean $\tilde{T}^1 = (0.117, 0.390)$. These have different topologies, and none of them equals $\overline{T}$, which has all edges 0.

Example 5.4 (**The BHV centroid**). Billera, Holmes, and Vogtmann [12] define the *centroid* of $\boldsymbol{T} = \{T^1, \dots, T^r\}$ inductively on r. For $r = 2$, the centroid is the midpoint of the two trees. For $r > 2$, the centroid is obtained as follows: set $\boldsymbol{T}^1 = \boldsymbol{T}$ and inductively find the centroid of each subset of $r - 1$ trees in $\boldsymbol{T}^1$ to obtain a new set $\boldsymbol{T}^2$ of r trees. Repeat this process on the new set, creating a sequence $\boldsymbol{T}^1, \boldsymbol{T}^2, \dots$ of r-sets of trees. The BHV centroid of $\boldsymbol{T}$ is the limit $\hat{T}$ of any sequence of points chosen from each of the sets $\boldsymbol{T}^i$. This process converges in a general global NPC space [12, Theorem 4.1].

Billera, Holmes, and Vogtmann note that in Euclidean space, the centroid and Fréchet mean coincide. This is not generally the case in tree space. Consider, for example, the trees in Figure 5 with coordinates $T^1 = (2,4)$, $T^2 = (2,2)$, and $T^3 = (4,2)$. Then $\overline{T}$ is again the origin, while it is easy to see that the BHV centroid lies off the origin.

5.3 Stickiness of the mean

Sullivant [49] noticed the tendency of the Fréchet mean to be *sticky*, which in this context means that perturbing one or more of the trees in the set $\boldsymbol{T}$ does not necessarily change any of the coordinates of $\overline{T}$. Take, for example, the points $T^1 = (3,10)$, $T^2 = (3,3)$, and $T^3 = (10,3)$. The mean $\overline{T}$ lies at the origin, and remains there even if the coordinates of any

of the three trees T^i are perturbed even up to a full unit. Sticky means occur exclusively on orthants of lower dimension, underscoring the importance of closely investigating properties of mean trees that lie on orthant boundaries.

The notion of stickiness has been quantified via a Central Limit Theorem for means of probability distributions on certain NPC spaces [10, 28].

5.4 Application to biological data

Statistical applications of this research are important in several areas of mathematics, biology, and medicine. Here, we consider a well-studied data set in phylogenetics with respect to the Fréchet mean. For applications of the Fréchet mean to medical imaging, see [47] and [31].

Example 5.5 (**Gene trees vs. species trees**). A *gene tree* is a phylogenetic tree representing the evolutionary history of a particular gene found in some set of species. In contrast, a *species tree* is a phylogenetic tree representing the evolutionary history of the species themselves: the history of population bifurcations leading to divergence. Due to natural processes such as incomplete lineage sorting, gene trees for different genes can have different topologies, even when sampled from the same set of individuals—let alone the same set of species—and hence a gene tree need not share its topology with the species tree (see [32], for example). Furthermore, the most likely gene tree topology need not agree with the species tree topology [20]. However, species trees are usually reconstructed from gene trees, and a major open question is how best to accomplish this.

We examined the yeast data set of Rokas et al. [45]. For eight species of yeast, they identified 106 genes and reconstructed the corresponding gene tree with edge lengths for each using a maximum likelihood approach. In these 106 gene trees, there were 21 different topologies. We used Sturm's algorithm to compute the Fréchet mean of these gene trees. This mean tree had the same topology as the agreed-upon species tree [22]. In general, the mean gene tree does not necessarily identify the species tree, as a consequence of stickiness, when branch lengths are taken into account [49]; that is, two finite samples of gene trees can yield the same mean tree but have different species trees. However, we conjecture that the topology of the species tree is a refinement of the topology of the Fréchet mean of the gene trees. That is, stickiness of the Fréchet mean forces some edges to have zero length but should not add any extraneous edges to this mean.

6 Globally nonpositively curved spaces

Virtually all of our treatment of tree spaces extends to more general global NPC spaces. This section reframes the concepts and notation of the paper in the context of global NPC spaces, particularly orthant spaces, and shows how the results of the paper generalize to these spaces.

6.1 The geometry of nonpositively curved spaces

Fix a metric space $\mathcal{T} = (\mathcal{T}, d)$. A *path* in $\mathcal{T}$ is the image of a continuous map $\gamma : [0, 1] \to \mathcal{T}$. Write $\gamma_\lambda = \gamma(\lambda)$ for $0 \leq \lambda \leq 1$. The *length* of γ is the supremum of all sums

$$d(\gamma_{x_0}, \gamma_{x_1}) + d(\gamma_{x_1}, \gamma_{x_2}) + \cdots + d(\gamma_{x_{k-1}}, \gamma_{x_k})$$

such that $0 \leq x_0 \leq \cdots \leq x_k \leq 1$. A path is a *(global) geodesic* if the distance $d(\gamma_x, \gamma_y)$ between any pair of points on γ equals the length of that portion of γ between them. A *geodesic space* is a complete metric space such that every pair $\{x, y\}$ of points is joined by a path γ whose length is the distance $d(x, y)$ between x and y.

Definition 6.1. A metric space $(\mathcal{T}, d)$ is *globally nonpositively curved*, also known as *global NPC* or *CAT(0)*, if for every triple of points $a, b, c \in \mathcal{T}$, any point x on a geodesic joining a to b, and any *reference triangle* $a'b'c'$ in Euclidean space with edge lengths $d(a, b)$, $d(b, c)$, and $d(a, c)$, the unique point x' on $a'b'$ at distance $d(a, x)$ from a' satisfies $d(x, c) \leq ||x' - c'||$.

The definition essentially says that triangles created by joining points by geodesics in a global NPC space are "skinnier" than their counterparts in Euclidean space.

Lemma 6.2 ([48, Proposition 2.3]). *In a global NPC space every pair of points is joined by a unique geodesic.* $\square$

A real-valued function $f : \mathcal{T} \to \mathbb{R}$ is *convex* if $f \circ \gamma$ is convex for all geodesics γ; that is, if

$$f(\gamma_\lambda) \leq (1 - \lambda)f(\gamma_0) + \lambda f(\gamma_1) \tag{13}$$

for all geodesics $\gamma : [0, 1] \to \mathcal{T}$.

Example 6.3. For any point $t \in \mathcal{T}$, the distance $d_t(x) = d(x, t)$ from a point $x \in \mathcal{T}$ to t is a convex function of x [48, Corollary 2.5 and subsequent Remark (i)].

A real-valued function $f : \mathcal{T} \to \mathbb{R}$ is *strictly convex* if Eq. (13) holds strictly for $0 < \lambda < 1$.

Lemma 6.4 ([48, Proposition 1.7 and Remark 1.8]). *Any strictly convex continuous function on a global NPC space attains a unique minimum.* $\square$

Corollary 6.5. *If $\boldsymbol{T} = \{t^1, \ldots, t^r\}$ is a set of points in $\mathcal{T}$, and $f : \mathbb{R}_+^r \mapsto \mathbb{R}$ is any (strictly) convex function, then the function $F : \mathcal{T} \mapsto \mathbb{R}$ defined by*

$$F(x) = f(d_{t^1}(x), \ldots, d_{t^r}(x))$$

is a (strictly) convex function.

In particular, the variance function for a set $\boldsymbol{T}$ of points is a convex function and hence attains a unique minimum at the Fréchet mean.

6.2 Means and variances in global NPC spaces

This subsection generalizes the notion of mean and variance to general probability measures in global NPC spaces. The results follow from those of Sturm [48] in this area. Let $\mathcal{P}(\mathcal{T})$ be the set of probability measures on a global NPC space $\mathcal{T}$. If $\rho \in \mathcal{P}(\mathcal{T})$ is such a measure, then its *variance* is

$$\text{var}(\rho) = \inf_{x \in \mathcal{T}} \int_{\mathcal{T}} d^2(x, y) \rho(dy).$$

The variance can be infinite in general, but not in the case of most interest to us, when ρ has finite support, meaning that there is a set $\boldsymbol{T} = \{t^1, \ldots, t^r\}$ of points in $\mathcal{T}$, along with nonnegative weights $\omega_1, \ldots, \omega_r$ satisfying $\omega_1 + \cdots + \omega_r = 1$, such that the point t_i has mass $\rho(t_i) = w_i$ for $i = 1, \ldots, r$. Let $\mathcal{P}^2(\mathcal{T})$ be the set of measures in $\mathcal{P}(\mathcal{T})$ having finite variance.

Proposition 6.6 ([48, Proposition 4.3]). *For a global NPC space $\mathcal{T}$ and probability measure $\rho \in \mathcal{P}^2(\mathcal{T})$, there is a unique point $\overline{\rho} \in \mathcal{T}$ such that $\text{var}(\rho) = \int_{\mathcal{T}} d^2(\overline{\rho}, y) p(dy)$.*

The point $\overline{\rho}$ is referred to as the *Fréchet mean* or *barycenter* in this context as well, and when ρ has finite support with $\omega_i = \frac{1}{r}$ for all r, it is a direct generalization of the definition of mean given in Section 2. The notion of inductive mean given by Definition 2.3 extends easily to an arbitrary global NPC space $\mathcal{T}$, and the following result generalizes Theorem 2.4.

Theorem 6.7 ([48, Theorem 4.7]). *For a global NPC space $\mathcal{T}$ and probability measure $\rho \in \mathcal{P}^2(\mathcal{T})$, let $X^1, X^2, \ldots$ be a sequence of independent and identically distributed random variables drawn from ρ. Then with probability 1, the sequence of inductive mean values $\mu_1, \mu_2, \ldots$ approaches the mean $\overline{\rho}$ of ρ.*

Corollary 6.8. *The convergence properties of the sequence of inductive means given by Algorithm 2.5 continue to hold on any probability distribution $\rho \in \mathcal{P}^2(\mathcal{T})$, by sampling the points of $\mathcal{T}$ according to the specified distribution.*

Note that Corollary 6.8 was independently observed by Bačák [8].

Remark 6.9. As an application of Corollary 6.8, Markov chain Monte Carlo (MCMC) simulations produce phylogenetic trees sampled independently from a fixed finite-variance distribution on the entire space $\mathcal{T}_n$. Calculating inductive mean values of repeated samples from this distribution results in a method to approximate the mean of the distribution.

6.3 NPC orthant spaces

Definition 6.10. The *orthant space* $\mathcal{O}(\mathcal{E}, \Omega)$ consists of a set $\mathcal{E}$ of *axes* together with a simplicial complex $\Omega \subseteq 2^{\mathcal{E}}$, called the *scaffold complex*. Two elements of $\mathcal{E}$ are *compatible* if they appear in some face of Ω. Each face $F \in \Omega$ is associated with a copy $\mathcal{O}_F$ of $\mathbb{R}_+^F$, the *orthant* associated with F. The orthant space $\mathcal{O}(\mathcal{E}, \Omega)$ is the union of the orthants $\mathcal{O}_F$ for $F \in \Omega$, with points identified whenever their nonzero coordinates agree on all elements of $\mathcal{E}$.

An orthant space can be thought of as constructed by gluing together orthants according to instructions laid out by the scaffold complex, and in fact the scaffold complex is (homeomorphic to) the *link* of the origin in the orthant space $\mathcal{O}(\mathcal{E}, \Omega)$.

Example 6.11. Tree space $\mathcal{T}_n$ is an orthant space: $\mathcal{E}$ corresponds to the set of splits on $\{0, \ldots, , n\}$, and Ω corresponds to the collection of sets of splits that are compatible in the sense of Section 1.1.

A *path* in an orthant space $\mathcal{O}(\mathcal{E}, \Omega)$ is defined as in Section 6.1. A locally length-minimizing path is a *geodesics*, which always consists of a finite number of linear legs through intermediate orthants of $\mathcal{T}$, as in the case of tree space (Section 1.2). As with tree space, $\mathcal{O}(\mathcal{E}, \Omega)$ is always path-connected.

Although any orthant space is geodesic, it may not be global NPC.

Example 6.12. The space $\mathcal{T} = \mathcal{O}(\mathcal{E}, \Omega)$, where $\mathcal{E}$ is indexed by $\{1, 2, 3\}$ and the scaffold complex Ω has facets $\{1, 2\}$, $\{1, 3\}$, and $\{2, 3\}$ is not global NPC. Indeed, the two points $x = (1, 0, 0)$ and $y = (0, 1, 1)$ in $\mathcal{T}$ have a pair of geodesics between them, namely $[(1, 0, 0), (0, 1/2, 0)] \cup [(0, 1/2, 0), (0, 1, 1)]$ and $[(1, 0, 0), (0, 0, 1/2)] \cup [(0, 0, 1/2), (0, 1, 1)]$. By Lemma 6.2, $\mathcal{T}$ cannot be global NPC.

M. Gromov [24] determined conditions on Ω that characterize when $\mathcal{O}(\mathcal{E}, \Omega)$ has non-postive curvature (in fact, Gromov worked with arbitrary cubical complexes), based on the following standard notion from geometric combinatorics.

Definition 6.13. The simplicial complex Ω is *flag* if $F \in \Omega$ whenever all pairs of elements in F are compatible.

Proposition 6.14 ([24]). *An orthant space $\mathcal{O}(\mathcal{E}, \Omega)$ is global NPC if and only if Ω is flag.*

In particular, tree space $\mathcal{T}_n$ is global NPC, since its scaffold complex is defined precisely by the pairwise compatibility between its splits (this is the proof given in [12]). Generally, any global NPC orthant space can be defined entirely by its set of compatible elements.

Definition 6.15. The *scaffold graph* $\mathcal{G}(\mathcal{E}, \Omega)$ of an orthant space $\mathcal{O}(\mathcal{E}, \Omega)$ is the graph with vertex set $\mathcal{E}$ whose edges are the pairs of compatible elements of $\mathcal{E}$.

Lemma 6.16. *The orthants of a global NPC orthant space $\mathcal{O}(\mathcal{E}, \Omega)$ are precisely the* clique *sets (sets of mutually compatible edges) of the scaffold graph $\mathcal{G}(\mathcal{E}, \Omega)$.* □

Thus there is a one-to-one correspondence between orthant spaces and graphs. A general global NPC orthant space $\mathcal{O}(\mathcal{E}, \Omega)$ need not have all of its maximal orthants the same dimension, since maximal orthants correspond to the maximal cliques in $\mathcal{G}(\mathcal{E}, \Omega)$. The dimension of the maximal orthants, however, is not relevant to any of the results in the previous sections, except when the dimension is given explicitly.

Example 6.17. The space of trees in which each split is associated with an m-dimensional vector instead of a single length is an NPC orthant space. In this case, the scaffold graph is the scaffold graph of tree space $\mathcal{T}_n$, with each vertex replaced by K_m, the complete graph on m vertices. Our software implementation also computes geodesics and means in this space.

Example 6.18. The generality of scaffold graphs to define any global NPC orthant space provides an opportunity to extend the statistical structures of this paper to a wider range of applications. As one example, consider a computer network specified by its computational devices and the graph $\mathcal{G}$ denoting those pairs of computers that are compatible with each other. A *local area network (LAN)* for this system is a set C of mutually compatible computers—that is, a clique of $\mathcal{G}$. A *local network configuration (LNC)* is a LAN C together with a measure w_e of participation of each computer $e \in C$ in the LAN C. Some important areas of analysis of the network $\mathcal{G}$ might be the relationship between the LNCs associated with $\mathcal{G}$, in terms of the number and participation weight of common computers and the relative compatibility of the noncommon computers (although it does *not* model chaining-related measures such as the number of nodes in a communications path). The global NPC orthant space generated by $\mathcal{G}$ would be a good framework for answering questions like this associated with the LANs of the network.

The combinatorics of geodesics in Section 1 generalizes immediately to global NPC orthant spaces, using the generalized notation in this section. None of the proofs in Sections 1–4 rely on particulars of tree space except the flag property. Thus we have the following.

Corollary 6.19. *The results in Sections 1–4 (except for statements specifying dimension) extend to arbitrary global NPC orthant spaces, using the definitions in this section. In particular, the GTP algorithm [41] for finding geodesics in tree space, Sturm's Algorithm (Algorithm 2.5), and the Descent Method (Algorithm 4.4) apply in the more general setting of global NPC orthant spaces.*

Remark 6.20. The results here extend even further. For example, Ardila, Owen, and Sullivant [4] extend the global NPC theory, and in particular the GTP algorithm, to the case of *cubical complexes*, where orthants are replaced by Euclidean cubes. Sturm's algorithm extends to these global NPC cubical complexes, and there is every reason to believe that the idea of vistal cells and the Descent Method can be extended as well, although without the polyhedrality. Furthermore, the results extend to negative edge lengths, as described in Remark 6.21.

Remark 6.21. Negative values can also be allowed for the coordinates in a global NPC orthant space. In this case, it remains a global NPC space, and the results listed in Corollary 6.19 for NPC orthant spaces also hold here, with the following modification. If a negative value appears in a common 'split', then that negative value is used in the geodesic calculations. However, if it appears in a split that is not in common, then its absolute value is used in the geodesic calculations.

References

[1] E.N. Adams, Consensus techniques and the comparison of taxonomic trees, Syst. Zool. **21** (1972), 390–397.

[2] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin, Network Flows: Theory, Algorithms, and Applications, Prentice Hall, Upper Saddle River, NJ, 1993.

[3] B.L. Allen and M. Steel, Subtree transfer operations and their induced metrics on evolutionary trees, Ann. Comb. **5** (2001), 1–15.

[4] F. Ardila, M. Owen, and S. Sullivant, Geodesics in CAT(0) cubical complexes, Advances in Applied Mathematics **48** (2012), 142-163

[5] Elissaveta Arnaoudova, David Haws, Peter Huggins, Jerzy W. Jaromczyk, Neil Moore, Christopher Schardl, and Ruriko Yoshida, Statistical phylogenetic tree analysis using differences of means, Frontiers in Neuroscience **4** (2012), 47.

[6] B. Aydın, G. Pataki, H. Wang, A. Ladha, E. Bullitt, and J.S. Marron, Visualizing the structure of large trees, Electron. J. Statist. **5** (2011), 405-420.

[7] S. Aylward and E. Bullitt, Initialization, noise, singlularities, and scale in height ridge traversal for tubular object centerline extraction, IEEE Transactions on Medical Imaging **21** (2002), 61–75.

[8] M. Bačák, Computing medians and means in Hadamard space, preprint, arXiv:math.MG/ 1210.2145 (2012)

[9] M. Barrett, M. Donoghue, and E. Sober, Against consensus, Syst. Zool. **40** (1991), 486–493.

[10] Bojan Basrak, Limit theorems for the inductive mean on metric trees, J. Appl. Prob. **47** (2010), 1136–1149.

[11] Martin Bridson and André Haefliger, Metric Spaces of Non-positive Curvature, Springer-Verlag, 1999.

[12] L. Billera, S. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees, Adv. in Appl. Math. **27** (2001), 733–767.

[13] David Bryant, A classification of consensus methods for phylogenetics, DIMACS Series in Discrete Mathematics and Theoretical Computer Science **61** (2003), 163–183.

[14] E. Bullitt, D. Zeng, G. Gerig, S. Aylward, S. Joshi, J.K. Smith, W. Lin, M.G. Ewend, Vessel tortuosity and brain tumor malignancy: A blinded study, Academic Radiology **12** (2005), 1232-1240.

[15] E. Bullitt, N.U. Lin, J.K. Smith, D. Zeng, E.P. Winer, L.A. Carey, W. Lin, M.G. Ewend, Blood vessel morphologic changes depicted with MR Angiography during treatment of brain metastases: a feasibility study, Radiology **245** (2007), 824–830.

[16] E. Bullitt, S.R. Aylward, T. Van Dyke, W. Lin. Computer-assisted measurement of vessel shape from 3T magnetic resonance angiography of mouse brain, Methods **43** (2007), 29-34.

[17] J. Chakerian, S. Holmes, Computational tools for evaluating phylogenetic and hierarchical clustering trees. Journal of Computational and Graphical Statistics 21 (2012) 581-599.

[18] Karen A. Cranston and Bruce Rannala, Summarizing a posterior distribution of trees using agreement subtrees, Systematic Biology **56** (2007), no. 4, 578–590.

[19] Elizabeth Allman, James Degnan, and John Rhodes, Identifying the rooted species tree from the distribution of unrooted gene trees under the coalescent, Journal of Mathematical Biology **62** (2011), 833–862.

[20] James Degnan and Noah Rosenberg, Discordance of species trees with their most likely gene trees, PLoS Genetics **3** (2006), 762–768.

[21] James Degnan and Laura Salter, *Gene tree distributions under the coalescent process*, Evolution **59** (2005), 24–37.

[22] Scott V. Edwards, Liang Liu, and Dennis K. Pearl, High-resolution species trees without concatenation, Proceedings of the National Academy of Sciences **104** (2007), 5936–5941.

[23] C.R. Finden and A.D. Gordon, Obtaining common pruned trees, Journal of Classification **2** (1985), 255–276.

[24] Mikhail Gromov, Hyperbolic groups, in Essays in Group Theory, pp. 75–263, Springer-Verlag, New York, 1987.

[25] J. Hein, Reconstructing evolution of sequences subject to recombination using parsimony, Math. Biosci. **98** (1990), 185–200.

[26] S. Holmes, Statistics for phylogenetic trees, Theoretical population biology, **63** (2003), 17–32.

[27] S. Holmes, Statistical approach to tests Involving phylogenetics, Proc. Math. of Evolution and Phylogeny. Oxford University Press, 2005.

[28] Thomas Hotz, Stephan Huckemann, Huiling Le, J. Stephen Marron, Jonathan C. Mattingly, Ezra Miller, James Nolen, Megan Owen, Sean Skwerer, and Victor Patrangenaru,

Sticky central limit theorems on open books, Annals of Applied Probability, 23 (2013), 2238-2258.

[29] Bret Larget, Satish Kotha, Colin Dewey, and Cécile Ané, BUCKy: Gene tree/species tree reconciliation with Bayesian concordance analysis, Bioinformatics **26** (2010), 2910-2911.

[30] Liang Liu, BEST: Bayesian estimation of species trees under the coalescent model, Bioinformatics **24** (2008), 2542–2543.

[31] A. Feragen, M. Owen, J. Petersen, M.M.W. Wille, L.H. Thomsen, A. Dirksen, and M. de Bruijne. Tree-space statistics and approximations for large-scale analysis of anatomical trees In Proceedings of the 23rd biennial International Conference on Information Processing in Medical Imaging (IPMI), (2013) 74-85.

[32] Wayne P. Maddison, Gene trees in species trees, Systematic Biology **46** (1997), 523–536.

[33] T. Margush and F.R. McMorris, Consensus n-trees, Bulletin of Mathematical Biology **43** (1981), 239–244.

[34] Ezra Miller and Igor Pak, Metric combinatorics of convex polyhedra: cut loci and nonoverlapping unfoldings, Discrete and Computational Geometry **39** (2008), no. 1–3, 339–388.

[35] J.S.B. Mitchell Geometric shortest paths and network optimization. In: J.-R. Sack, J. Urrutia (Eds.), Handbook of Computational Geometry, Elsevier, Amsterdam, (2000) 633–701

[36] James R. Munkres, Elements of algebraic topology, Addison–Wesley, Menlo Park, CA, 1984.

[37] Tom M. W. Nye, Principal components analysis in the space of phylogenetic trees, Annals of Statistics **39** (2011), 2716-2739.

[38] İpek Oğuz, Groupwise Shape Correspondence with Local Features. Ph. D. dissertation, Dept. Computer Science, Univ. North Carolina at Chapel Hill, 2009.

[39] M. Owen, Computing geodesic distances in tree space, SIAM Journal on Discrete Mathematics **25** (2011), 1506–1529.

[40] M. Owen, Sturm algorithm implementation. http://comet.lehman.cuny.edu/owen/code.html

[41] M. Owen and S. Provan, A fast algorithm for computing geodesic distance in tree space, ACM/IEEE Transactions on Computational Biology and Bioinformatics **8** (2011), 2–13.

[42] J.-C. Picard and M. Queyranne, On the structure of all minimum cuts in a network and applications, Mathematical Programming Study **13** (1980), 8–16.

[43] D.F. Robinson, Comparison of labeled trees with valency three, J. Combinatorial Theory, **11** (1971), 105–119.

[44] D.F. Robinson and L.R. Foulds. Comparison of phylogenetic trees, Math. Biosci **53** (1981), 131–147.

[45] Antonis Rokas, Barry L. Williams, Nicole King, and Sean B. Carroll, Genome-scale approaches to resolving incongruence in molecular phylogenies, Nature **425** (2003), 798–804.

[46] C. Semple and M. Steel, Phylogenetics, Oxford University Press, Oxford, 2003.

[47] S. Skwerer, E. Bullitt, S. Huckemann, E. Miller, I. Oguz, M. Owen, V. Patrangenaru, S. Provan, and J.S. Marron Tree-oriented analysis of brain artery structure To appear Journal of Mathematical Imaging and Vision (2013).

[48] Karl-Theodor Sturm, Probability measures on metric spaces of nonpositive curvature, in Heat kernels and analysis on manifolds, graphs, and metric spaces: lecture notes from a quarter program on heat kernels, random walks, and analysis on manifolds and graphs, Contemporary Mathematics **338** (2003), 357–390.

[49] Seth Sullivant, personal communication, 2010.

[50] Günter M. Ziegler, Lectures on polytopes, Graduate Texts in Mathematics, Vol. 152, Springer, New York, 1995.

[51] H. Wang, J.S. Marron, Object oriented data analysis: Sets of trees, Ann. Statist. **35** (2007), 1849-1873.