

PRINCIPAL COMPONENT ANALYSIS IN PHYLOGENETIC TREE SPACE

Haojin Zhai

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Statistics and Operations Research.

Chapel Hill
2016

Approved by:

Shu Lu

J. S. Marron

Ezra Miller

Gabor Pataki

J. Scott Provan

©2016
Haojin Zhai
ALL RIGHTS RESERVED

ABSTRACT

Haojin Zhai : Principal Component Analysis in Phylogenetic Tree Space
(Under the direction of J. S. Marron and J. Scott Provan)

Complex data objects arise in many fields of modern science including drug discovery, psychology, dynamics of gene expression and anatomy. Object oriented data analysis describes the statistical analysis of a population of complex data objects. The specific case of tree-structured data objects is a large and promising research area with many interesting questions and challenging problems. This dissertation focuses on principal component analysis in the tree space introduced by Billera, Holmes, and Vogtmann.

Principal component analysis has been a widely used method in aiding visualization and reducing dimensions, and it is natural to extend this type of analysis into tree space. In this dissertation, we will discuss three interesting approaches to this extension. The first approach is multidimensional scaling, which focuses on better visualization of data in tree space, in particular, the out-of-sample embedding problem which inserts additional points into previously constructed multidimensional scaling configurations. It is shown that a better visualization can be achieved by choosing a higher dimensional embedding space and displaying only the first two dimensions. The other two approaches rely on our novel definitions of tree space line, and it is proven that there are only two types of such lines. The second approach is sample-limited geodesic which is an analog of the first type of line. This idea defines the first principal component for a set of trees by maximizing the data projection variance over geodesic segments connecting pairs of trees. Our study shows that the sample-limited geodesic is not an effective principal component object in terms of capturing data variation, due to the intrinsic geometry of the data used in this dissertation, and it is not natural to be generalized into higher-order principal component objects. The third approach is based on the principal ray set, which is a representative of the second type of line. We develop some heuristic searching algorithms for first order principal ray sets and higher order principal axis sets, which are special cases of principal ray sets. Principal ray sets are better summaries for less variable data, but gain very limited information for data with larger spread.

For Yuren and Jing, my beloved and caring parents.

TABLE OF CONTENTS

LIST OF FIGURES	vii
LIST OF TABLES	x
CHAPTER 1: INTRODUCTION	1
1.1 Motivation	1
1.2 Phylogenetic Tree Space	2
1.2.1 Phylogenetic n-trees	2
1.2.2 Construction of Phylogenetic Tree Space	3
1.3 Construction of Geodesics	5
1.4 The Fréchet Mean	7
1.5 The Combinatorics of Geodesics	8
1.6 PCA in Euclidean Space	11
1.7 PCA on Manifolds	12
1.8 PCA of Tree Structured Data Objects	13
1.9 Overview of Dissertation	14
CHAPTER 2: REAL AND SIMULATED DATA	16
2.1 Brain Artery Data	16
2.1.1 Magnetic Resonance Angiography	16
2.1.2 Cortical Correspondence	17
2.2 Generated Data Set I: Uniformly Random	18
2.3 Generated Data Set II: The Wright-Fisher Model	21
2.3.1 The Wright-Fisher Model	21
2.3.2 Data Generation	22
2.4 Exploratory Data Analysis	24
2.4.1 Angle-based Data Summaries	24
2.4.2 Distance-based Data Summaries	25
2.5 Stickiness and Landmark-reduced Brain Artery Data	26

CHAPTER 3: MULTIDIMENSIONAL SCALING IN TREE SPACE	28
3.1 Review of Multidimensional Scaling	28
3.2 MDS of Tree Data	32
3.3 MDS for Embedding Geodesics	35
3.4 Out-of-Sample Embedding	38
CHAPTER 4: NOTIONS OF “LINE” IN TREE SPACE	48
4.1 Tree-space Line	48
4.2 Type I Lines	49
4.3 Type II Lines	57
CHAPTER 5: SAMPLE-LIMITED GEODESICS	60
5.1 The First Sample-limited Geodesic	60
5.2 Higher Order Sample-limited Geodesic	66
CHAPTER 6: USING PRINCIPAL RAYS TO MODEL PCA IN TREE SPACE	67
6.1 Fixed Rays	67
6.2 Direction-variable Rays	68
6.3 The First Principal Ray Sets	77
6.4 Higher Order Principal Axis Set	79
REFERENCES	86

LIST OF FIGURES

1.1	(a) shows the brain blood vessel system of a person. (b) illustrates one of the brain artery trees used as data objects in the present research.	1
1.2	A portion of \mathcal{T}_4 with three trees T^1, T^2, T^3 and their projections onto the common edge. Orange represents common edge, while blue, red, yellow represent non-common edges. Black lines illustrate the contractions of non-common edges.	4
1.3	Three examples of shortest paths in \mathcal{T}_4 . All three paths are shortest paths between pairs of trees. In particular, the solid line represents a cone path.	4
1.4	Comparison triangle	5
2.5	(a) One slice of a Magnetic Resonance Angiography (MRA) image for one person. Bright regions indicate blood flow. These are tracked through MRA slices to recover artery tubes as shown in the figure on the right. (b) The data object is a reconstructed brain artery tree for one person. The goal of this research is PCA type statistical analysis of a sample of such data objects.	17
2.6	Illustrates steps for <i>cortical correspondence</i> . (a) Brain artery system (blue) with the cortical landmarks (red). (b) Find the closest point on the brain artery system to each landmark (red segments are landmark connections). Arteries which are not between the landmark projections and the base (cyan) are trimmed. (c) The result of this procedure is a cortical correspondence tree.	18
2.7	Schematic plots showing the three possibilities to insert a new leaf in the first iteration. . . .	20
2.8	Comparison between KDE of edge lengths in brain artery data (blue solid curve) and Gamma distribution with parameters 1.28 and 15.8 (red dashed curve) shows the Gamma distribution is an excellent fit for the edge lengths distribution of the brain artery data.	20
2.9	(a) A template phylogenetic tree with 4 leaves and edge lengths in boxes. (b) The corresponding population genealogical tree has the same topological structure as the template tree. The number of reproduction cycles in each subpopulation equals the corresponding edge length in the template tree and the size of subpopulation in each generation is the width parameter which equals 4.	22
2.10	Two possible sample gene trees generated within the same population genealogical tree, from Figure 2.9(b).	23
2.11	Two possible sample gene trees with different topologies can be generated within the same population genealogical tree, from Figure 2.9(b).	23
2.12	(a) Overlay of pairwise angle KDE plots shows the decreasing similarity ordering of the five data sets: WF2, WF10, WF40, brain artery data, and uniformly random data. (b) Overlay of logarithm of pairwise angle KDE plots indicates that the variability in pairwise angles is proportional to the magnitude of angles for WF2, WF10, WF40, and brain artery data, but not for the uniformly random data.	25
2.13	(a) Overlay of pairwise distance KDE plots shows the decreasing similarity ordering of the five data sets: WF2, WF10, WF40, brain artery data, and uniformly random data. (b) Overlay of the logarithm of pairwise distance KDE plots implies the proportionality between the spread and magnitude of pairwise distances.	26

3.14	The above scatter plot matrix gives a visualization of 3-D MDS configuration on 3 clusters from front, side, and top views. Both distance MDS and classical MDS produce this same 3-D configuration. Three data clusters are colored red, green, and blue.	31
3.15	Comparison across classical MDS, distance MDS, and PCA on 3-cluster (colored red, green, and blue) simulated data set. Shows low rank approximations of Euclidean data are different between classical MDS and distance MDS.	32
3.16	2-D MDS plots associated with 3 biological variables: age, gender and handedness. No clear correlates of these biological variables can be drawn.	33
3.17	3-D MDS plots in rainbow color scheme across ages along with top, side, and front views in separate sub-figures. Very limited information gained from adding the third dimension. . . .	33
3.18	3-D MDS plots with males in blue dots and females in green dots along with top, side, and front views in separate sub-figures. Very limited information gained from adding the third dimension.	34
3.19	Examples of embedding geodesics in 2-D MDS plots	36
3.20	2-D MDS plots with one triangle embedded	37
3.21	Examples of embedding geodesics as out-of-sample objects in 2-D MDS plots	41
3.22	HDOS embedding of median length geodesic with embedding dimension $p = 3, 5, 10, 66$	43
3.23	HDOS embedding of median length geodesic with embedding dimension $p = 3, 5, 10, 66$	44
3.24	HDOS embedding of median length geodesic with embedding dimension $p = 3, 5, 10, 66$	46
4.25	The solid path denotes $P(T^1, T^2)$ and $P(T^1, T^2) \neq \Gamma(T^1, T^2)$. The dashed ellipse represents the vial cell \mathcal{V} containing T^* , \tilde{T} and T' . The dotted path between T^1 and T' together with the solid segment between T' and \tilde{T} is $\Gamma(T^1, \tilde{T})$. We reach a contradiction that there are two distinct geodesics between T^1 and T^*	49
4.26	Examples of lines in \mathcal{T}_4	50
4.27	The above figure shows three geodesics in a portion of \mathcal{T}_4 : $\Gamma(T^1, T^2)$ is uniquely extendable and can be extended into a type I line; $\Gamma(T^1, T^6)$ is also uniquely extendable but can not be extended into a type I line; $\Gamma(T^2, T^6)$ is not even uniquely extendable.	51
4.28	The above figure shows extensions of $\Gamma(T^1, T^2)$ and $\Gamma(T^2, T^3)$ which are uniquely extendable geodesics for the three trees in Figure 4.27.	52
4.29	The above figure gives a detailed view of what happens when $\Gamma(T^1, T^2)$ crosses $\mathcal{O}_{i-1} \cap \mathcal{O}_i$ which is intuitively represented by the central vertical axis in bold.	56
4.30	The above figure gives a detailed view of what happens when e is contracted.	57
5.31	The above figures show the distributions of data projections along SLG1 across five data sets. The projections are only distributed within a relatively narrow range between the two end points of SLG1 for all five data sets.	65
6.32	A randomly generated small data set in a portion of \mathcal{T}_4 : 10 trees are represented by blue dots, and all the edge lengths are integers between 1 and 10.	70

6.33	The black dotted lines represent locally optimal rays, and the red circle denotes the Fréchet mean. This plot shows that the locally optimal rays and the Fréchet mean do not locate in the same orthant.	71
6.34	This figure shows the non-concavity of the objective function (6.12) as well as some other aspects of this toy example in \mathcal{T}_4 . The horizontal axis gives the angle travelled by the moving ray. The vertical axis represents the sum of squared projections from 10 data points onto the moving ray. The orthant boundaries are indicated by dashed vertical lines. The Fréchet mean is plotted as the green vertical line. The blue curve shows how the value of the objective function changes as the ray moves. The red horizontal line segments keep track of which data points project positively onto the ray.	72
6.35	Red dots represent the sorted proportions of the total variation in the brain artery data captured by 67 locally optimal rays which are returned by the steepest descent algorithm. Blue circles represent the sorted proportions obtained in the way that the squared norm of each tree is divided by the total squared norm of the brain artery data. The two plots are the same, which indicates that locally optimal rays go through trees in the data set.	75
6.36	Peterson graph illustrating how the steepest descent search algorithm works for the landmark-reduced brain artery data set with only 5 leaves. All 85 data trees are contained in 5 orthants represented by dashed line segments. The search algorithm will end at locally optimal rays in different orthants depending on where the search process starts.	76
6.37	Bar graph showing the comparison of proportions of total variation captured by the best locally optimal ray across 8 landmark-reduced brain data sets.	77
6.38	This plot shows that there will be a contradiction if tree T projects positively onto two antipodal rays \vec{r}_1 and \vec{r}_2	78
6.39	Shows the comparison of the four greedy approaches for three data sets: (a)Brain Artery data, (b)Uniformly Random data, (c)WF40, and (d)a zoom-in version of WF40. Suggests that the “Forward SS” is the best among four greedy approaches.	84

LIST OF TABLES

3.1	Row 1 gives the values for the embedding dimension p , and Row 2 shows the Stress of the 2-dimensional out-of-sample representation for each p . The Stress gets slightly larger as p increases, which indicates that the Stress of 2-dimensional representation is not a good measure of HDOS distortion.	44
3.2	Row 1 gives the values for the embedding dimension p , and Row 2 shows the Stress of the p -dimensional out-of-sample representation. The Stress decreases as p increases, which is compatible with what we saw in Figure 3.22.	44
5.3	Column 1 shows the performance of SLG1 for five data sets without counting end points. Column 2 shows the same statistics but counting end points. Column 3 lists the total variation with respect to the Fréchet mean in each data set. Overall, SLG1 explains a relatively small amount of the total variation in all five data sets.	63
5.4	Column 1, 2, and 3 show the total variation with respect to the origin, the total variation with respect to the Fréchet mean, and the effect of the Fréchet mean in each data set.	64
6.5	This table lists four summaries across five data sets and the proportion in the third column is used to measure the performance of the best locally optimal ray.	73
6.6	Comparison of the proportions of data variation captured by the 1 st PA set using 4 greedy approaches across 5 data sets. Suggests that the forward approaches are better than the backward approaches in terms of capturing data variation.	82
6.7	Comparison of the numbers of antipodal axes in the 1 st PA set using 4 greedy approaches across 5 data sets explains what we have seen in Table 6.6.	82
6.8	Comparison of the numbers of PA sets needed to capture at least 80% of the data variation by using the 4 greedy approaches across the 5 data sets. Reveals that the “Forward SS” use fewer principal components than the other 3 approaches.	85

CHAPTER 1: INTRODUCTION

1.1 Motivation

This dissertation focuses on Object Oriented Data Analysis in the context of tree-structured objects. The concept of “Object Oriented Data Analysis”(OODA) was first defined by [Wang and Marron, 2007] and has more recently been discussed in [Marron and Alonso, 2014]. Essentially OODA is the statistical analysis of a population of complex objects. In traditional statistical analysis, the atoms are generally either numbers or vectors. In functional data analysis, a currently active research area, each sample element is considered to be a function; see [Ramsay and Silverman, 2002, 2005] for a detailed review. Wang and Marron extended the idea of functional data analysis to even more complicated objects such as images, two-dimensional or three-dimensional shapes, and combinatorial structures such as graphs or trees.

In [Wang and Marron, 2007], the authors modeled the human brain blood vessel systems as binary trees. Two types of information were taken into account: topological structures and attributes associated with nodes. Although both that work and the present research are based on a common set of human brain blood vessel systems, very different aspects of the data are studied here. Figure 1.1(a) gives an example of such a blood vessel system. Each blood vessel system studied in the present research is oriented with respect to 128

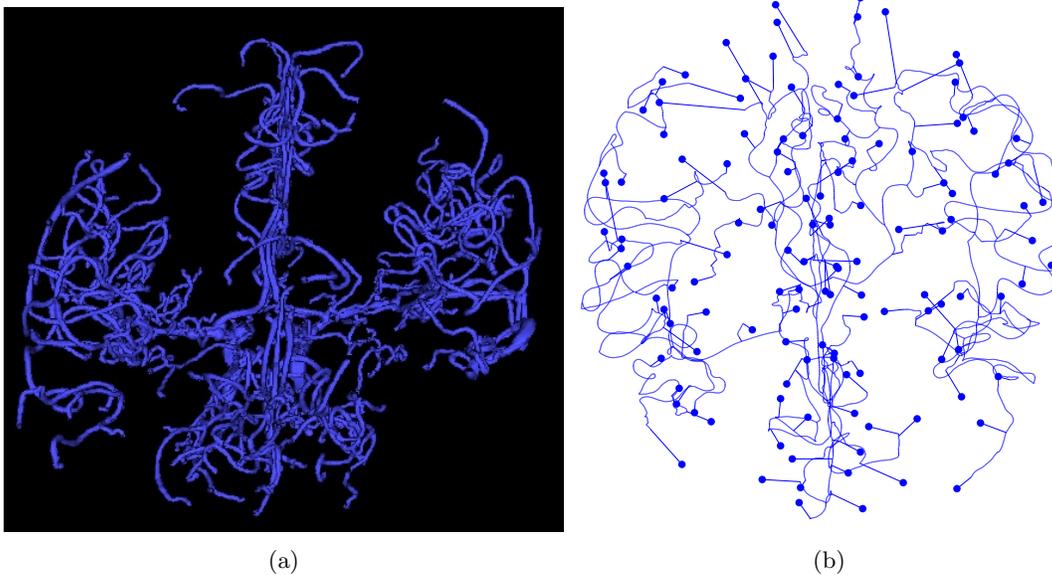


Figure 1.1: (a) shows the brain blood vessel system of a person. (b) illustrates one of the brain artery trees used as data objects in the present research.

carefully chosen (i.e. to correspond across patients) landmarks on the cortical surface. As in [Skwerer et al., 2014a], the cortical correspondence method of [Oguz et al., 2008] gave the landmarks. All 128 landmarks are used as non-root nodes and an artificial conceptual point is defined to be the root node. Corresponding to each landmark, a closest point on the brain artery is found by minimizing the 3-D Euclidean distance from the landmark. Due to the natural tree-like structure of the brain artery system, a brain artery tree is formed by connecting each landmark with this closest point, and one example is given in Figure 1.1(b). For the tree representation in the present research, the edge lengths are considered as important geometrical information. This is different from [Wang and Marron, 2007] in which the authors used nodal attributes containing other geometrical summaries.

Many applications of OODA come from the field of medical image analysis. In [Singh et al., 2010], researchers took brain shapes as objects and applied Partial Least Squares Regression to characterize the neuroanatomical variations observed in neurological disorders. In [Geneser et al., 2011], lung shapes were taken to be the objects and researchers modeled the changes in lung shapes as a function of chest wall amplitudes to calculate the resulting variability in radiation dose accumulation.

The present research is performed under the framework of *phylogenetic tree space* which was first built in [Billera et al., 2001]. In this space, there is a unique shortest path between each pair of trees, called a *geodesic*. Later [Owen and Provan, 2009] constructed a polynomial time algorithm to compute the geodesic distance which is implemented in [Skwerer, 2014]. These works enable us to find projection of one tree onto a subset of trees in phylogenetic tree space, and furthermore makes *Principal Component Analysis* (PCA) possible in this space.

1.2 Phylogenetic Tree Space

The foundation of this research was developed by [Billera et al., 2001], which gave a rigorous geometric definition of the space of rooted labeled trees. We will start this section with introducing some of the basic concepts in graph theory.

1.2.1 Phylogenetic n-trees

We are using standard graph theory terminology, such as given in [Ahuja et al., 1993; Bazaraa et al., 2010; Cormen et al., 2009]. A *tree* T is a connected graph with no cycles. A node of T is called a *leaf* if there is a unique edge connected to it, and the edge is called a *leaf edge*. $V = \{0, 1, \dots, n\}$ usually denotes the leaf set of T . A weighted tree is a tree with each edge e assigned a weight $|e|$, or $|e|_T$ if we want to emphasize the tree T to which e belongs. An edge e is *contracted* if $|e| = 0$. A *split* associated with edge e is defined as $\sigma_e = V_e | \bar{V}_e$, where V_e, \bar{V}_e are sets of leaves and $V_e | \bar{V}_e$ represents the partition of V resulting from deletion of the edge e from T . In order that a set Σ of splits actually form a tree, the corresponding edges must be pairwise *compatible*. Two edges e and f are compatible if their splits $V_e | \bar{V}_e$ and $V_f | \bar{V}_f$ satisfy the

property that one of the sets $V_e \cap V_f$, $V_e \cap \bar{V}_f$, $\bar{V}_e \cap V_f$, or $\bar{V}_e \cap \bar{V}_f$ is empty. This concept can be naturally extended to the compatibility of two sets of edges A and B : if for each edge $e \in A$ and each edge $f \in B$, e is compatible with f , then we say A and B are compatible. Here is some facts about compatibility:

- Each edge is identified uniquely by its split, henceforth edges from different trees with the same set of leaves are comparable.
- Each pair of edges in a tree are compatible.
- A tree is determined uniquely by its set of splits.

A *phylogenetic n-tree* (or simply *n-tree*) is a weighted tree $T = (V, \mathcal{E}, W, \Sigma)$, where $V = \{0, 1, \dots, n\}$ is a labeled set of leaves (with 0 arbitrarily denoting the root of T), \mathcal{E} is the set of interior (non-leaf) edges, $W = \{|e| : e \in \mathcal{E}\}$ is the set of edge weights which is also the set of edge lengths, and $\Sigma = \{\sigma_e : e \in \mathcal{E}\}$ is the set of splits for T .

In an n-tree, all the non-leaf nodes are assumed to have at least 3 adjacent edges. A *maximal n-tree* is an n-tree with the largest number $(2n - 1)$ of edges, or the largest number $(n - 2)$ of interior edges. Since all trees with the same set of leaves contain the same set of leaf edges, we will ignore them in this dissertation that follows. Notice that every n-tree can be represented topologically by contracting a set of edges from some maximal n-tree.

1.2.2 Construction of Phylogenetic Tree Space

Phylogenetic tree space, or *tree space* \mathcal{T}_n is a geometric space in which each point represents an n-tree and is placed in an $(n - 2)$ -dimensional *orthant* (copy of \mathbb{R}_+^{n-2}) with each orthant associated with some maximal *n-tree*. Orthants are attached to each other through common edges. Here we denote an orthant associated with tree $T = (V, \mathcal{E}, W, \Sigma)$ as $\mathcal{O}(\mathcal{E})$. Tree space is thus a union of $(2n - 3)!!$ orthants [Schröder, 1870], each of which corresponds to a distinct maximal tree topology. For any two trees T^1 and T^2 in \mathcal{T}_n , $d(T^1, T^2)$, the distance between T^1 and T^2 , will be defined as the length of the shortest path connecting them in \mathcal{T}_n . This extends the standard Euclidean metric, which is the standard distance between two vectors.

We will take \mathcal{T}_4 as an example to visualize some of the main features of a tree space. Note that each tree in \mathcal{T}_4 has only 2 interior edges, so each orthant of \mathcal{T}_4 will be 2-dimensional. Figure 1.2 shows three orthants of \mathcal{T}_4 , along with examples of tree topologies T^1 , T^2 , T^3 attached together along the common edge associated with the split $\{0, 1\}|\{2, 3, 4\}$ (green). The trees in these three orthants have this edge along with one additional edge with split $\{0, 1, 2\}|\{3, 4\}$ (blue), $\{0, 1, 3\}|\{2, 4\}$ (red) and $\{0, 1, 4\}|\{2, 3\}$ (yellow), respectively. Note that any one of the distinct splits $\{0, 1, 2\}|\{3, 4\}$, $\{0, 1, 3\}|\{2, 4\}$ and $\{0, 1, 4\}|\{2, 3\}$ is compatible with the vertical line $\{0, 1\}|\{2, 3, 4\}$, while no two of them are compatible with each other. Further, each tree T^i has

a projection P^i onto the vertical line by contracting its non-common edge as shown by black lines in Figure 1.2.

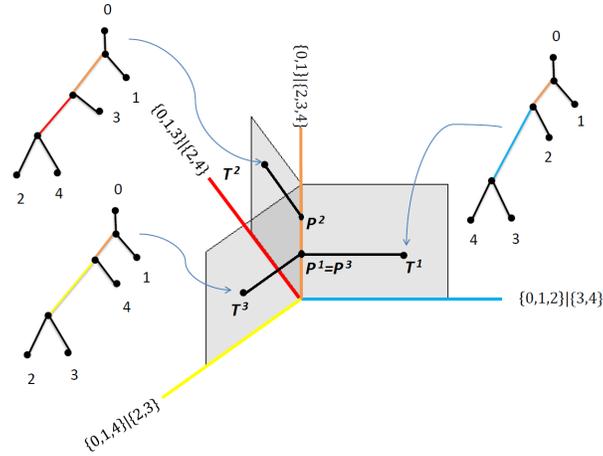


Figure 1.2: A portion of \mathcal{T}_4 with three trees T^1 , T^2 , T^3 and their projections onto the common edge. Orange represents common edge, while blue, red, yellow represent non-common edges. Black lines illustrate the contractions of non-common edges.

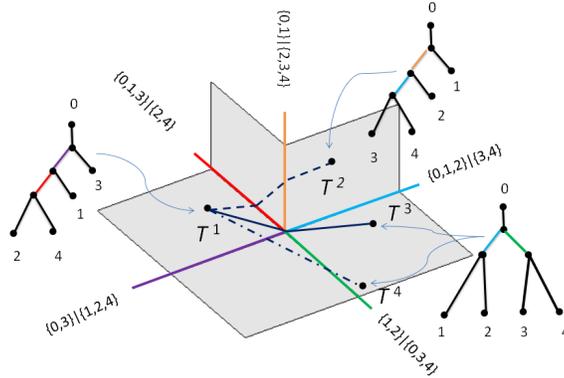


Figure 1.3: Three examples of shortest paths in \mathcal{T}_4 . All three paths are shortest paths between pairs of trees. In particular, the solid line represents a cone path.

Figure 1.3 shows three examples of paths in \mathcal{T}_4 . These three paths are chosen to be the shortest path between T^1 and T^2 , T^3 , T^4 respectively, since shortest paths are the paths of interest in this dissertation. The shortest path between T^1 and T^2 looks like a piecewise linear path, while the shortest path between T^1 and T^4 is a straight line segment. However, if the three orthants containing the shortest path between T^1 and T^2 are flattened out, the piecewise linear path will become a straight line. The shortest path between T^1 and T^3 passes through the origin and is a special type of path called *cone path*. The cone path consists of only two line segments: one connecting T^1 with the origin and the other connecting the origin with T^3 .

Tree space has properties that are very useful for deriving and constructing statistical properties in the space. One important property noted in [Billera et al., 2001] is that \mathcal{T}_n is a *non-positively curved*, or *CAT(0)* space. Intuitively speaking, every triangle in a *CAT(0)* space is skinnier than a triangle with exactly the same lengths of sides in Euclidean space. More precisely, a metric space X is said to be *CAT(0)* if the following statement is true. Given any three points a , b and c , as illustrated in Figure 1.4, with distances $d_1 = d(b, c)$, $d_2 = d(a, c)$, and $d_3 = d(a, b)$, form a “comparison triangle” in the Euclidean plane. The “comparison triangle” has vertices a' , b' , and c' with side lengths $d_1 = d(b', c')$, $d_2 = d(a', c')$, and $d_3 = d(a', b')$. If x is a point on the geodesic from a to b , at distance d_4 from a , find the corresponding point x' on the straight line from a' to b' at distance d_4 from a' . Then $d(x, c) \leq d(x', c')$. This leads to many useful consequences as seen in the next section.

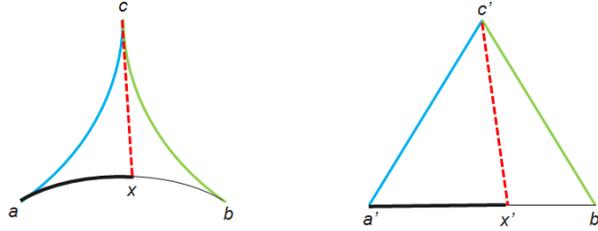


Figure 1.4: Comparison triangle

1.3 Construction of Geodesics

Because the tree space is *CAT(0)*, it follows by [Gromov, 1987] that there is a unique shortest path connection any two points $T = (V, \mathcal{E}, W, \Sigma)$ and $T' = (V, \mathcal{E}', W', \Sigma')$ of \mathcal{T}_n , called the *geodesic* $\Gamma(T, T')$ between T and T' , and its length $d(T, T')$ can be used to define a metric in tree space. This is the metric defined by [Billera et al., 2001]. This section reviews the topological structure of $\Gamma(T, T')$ and the algorithm given by [Owen and Provan, 2009] to compute $d(T, T')$. The topological structure of a geodesic was first given in [Billera et al., 2001] which showed that $\Gamma(T, T')$ is a piecewise linear path contained in a sequence of orthants. Each linear portion of $\Gamma(T, T')$ is called a *leg*, and the sequence of orthants is called a *path space* \mathcal{P} associated with $\Gamma(T, T')$. Assume there are $k + 1$ orthants $\mathcal{O}_0, \mathcal{O}_1, \dots, \mathcal{O}_k$ in \mathcal{P} , and let \mathcal{C} denote the set of common edges between T and T' . Now let $\mathcal{A} = (A_1, \dots, A_k)$ and $\mathcal{B} = (B_1, \dots, B_k)$ be partitions of $\mathcal{E} \setminus \mathcal{C}$ and $\mathcal{E}' \setminus \mathcal{C}$, respectively, such that A_i and B_j are compatible for each $i > j$. Then $(\mathcal{A}, \mathcal{B})$ is called the *support* of \mathcal{P} , and $A_1 \cup A_2 \cup \dots \cup A_k \cup \mathcal{C}$ is the edge set associated with \mathcal{O}_0 and $B_1 \cup \dots \cup B_i \cup A_{i+1} \cup \dots \cup A_k \cup \mathcal{C}$ is the edges set of \mathcal{O}_i for $1 \leq i \leq k$. Some properties about path space have been clarified further in [Owen, 2011]. A polynomial time algorithm for computing geodesic distances was developed in [Owen and Provan, 2009]. Here we summarize the relevant properties of the geodesic $\Gamma(T, T')$ and the associated path space \mathcal{P} from [Owen and Provan, 2009].

Since the Euclidean metric is preserved within each orthant of \mathcal{T}_n , the geodesic $\Gamma(T, T')$ will consist of a series of straight line segments through the orthants of the path space \mathcal{P} . The properties of $\Gamma(T, T')$ will be given in Theorem 1.3.1 which is the combination of Theorem 2.2, Theorem 2.3, Theorem 2.4 and Theorem 2.5 in [Owen and Provan, 2009]. For a set A of edges, we use $\|A\| = \sqrt{\sum_{e \in A} |e|^2}$ to denote the norm of the vector whose components are the lengths of the edges in A .

Theorem 1.3.1. *Let $T = (V, \mathcal{E}, W, \Sigma)$ and $T' = (V, \mathcal{E}', W', \Sigma')$ be two n -trees. Let \mathcal{C} be the set of edges which are common to both trees, and let \mathcal{P} be a path space containing T and T' with support $(\mathcal{A}, \mathcal{B})$ of all the non-common edges. Then $(\mathcal{A}, \mathcal{B})$ corresponds to a geodesic if and only if it satisfies the following three properties:*

(P1) *For each $i > j$, A_i and B_j are compatible.*

(P2) $\frac{\|A_1\|}{\|B_1\|} \leq \frac{\|A_2\|}{\|B_2\|} \leq \dots \leq \frac{\|A_k\|}{\|B_k\|}$.

(P3) *For each support pair (A_i, B_i) , there is no nontrivial partition $C_1 \cup C_2$ of A_i , and partition $D_1 \cup D_2$ of B_i , such that C_2 is compatible with D_1 and $\frac{\|C_1\|}{\|D_1\|} < \frac{\|C_2\|}{\|D_2\|}$.*

Further let Γ denote the geodesic between T and T' , and parameterize Γ as $\Gamma = (\gamma(\lambda) : 0 \leq \lambda \leq 1)$ where λ is the ratio of distance to T . In this way, Γ can be represented in \mathcal{T}_n with legs

$$\Gamma^i = \begin{cases} \left[\gamma(\lambda) : \frac{\lambda}{1-\lambda} \leq \frac{\|A_1\|}{\|B_1\|} \right], & i = 0 \\ \left[\gamma(\lambda) : \frac{\|A_i\|}{\|B_i\|} \leq \frac{\lambda}{1-\lambda} \leq \frac{\|A_{i+1}\|}{\|B_{i+1}\|} \right], & i = 1, \dots, k-1, \\ \left[\gamma(\lambda) : \frac{\lambda}{1-\lambda} \geq \frac{\|A_k\|}{\|B_k\|} \right], & i = k \end{cases}$$

where the points on each leg Γ^i are associated with the tree $T_i = (V, \mathcal{E}^i, W^i, \Sigma^i)$ having edge set

$$\mathcal{E}^i = B_1 \cup \dots \cup B_i \cup A_{i+1} \cup \dots \cup A_k \cup \mathcal{C}$$

and where the edge lengths in W^i are given by

$$|e|_{T_i} = \begin{cases} \frac{(1-\lambda)\|A_j\| - \lambda\|B_j\|}{\|A_j\|} |e|_T & e \in A_j \\ \frac{\lambda\|B_j\| - (1-\lambda)\|A_j\|}{\|B_j\|} |e|_{T'} & e \in B_j \\ (1-\lambda)|e|_T + \lambda|e|_{T'} & e \in \mathcal{C} \end{cases} .$$

The length of Γ is

$$L(\Gamma) = \left\| \left(\|A_1\| + \|B_1\|, \dots, \|A_k\| + \|B_k\|, \sqrt{\sum_{e \in \mathcal{C}} (|e|_T - |e|_{T'})^2} \right) \right\| \quad (1.1)$$

where \mathcal{C} is the set of common edges in the corresponding trees.

We note that geodesics can trivially be extended to include leaf edges, since these are always common to both trees and thus will be elements of \mathcal{C} .

In the above Theorem 1.3.1, the combination of (P1) and (P2) only gives us a necessary but not sufficient condition for a (T, T') -path with support $(\mathcal{A}, \mathcal{B})$ to be a geodesic, and (P3) focuses on characterizing when this path is guaranteed to be a geodesic by specifying when no local improvement can be made for the path. Based on this theorem, an $O(n^4)$ algorithm was developed in [Owen and Provan, 2009] and a Java implementation of the algorithm is available at <http://www.stat-or.unc.edu/webpace/miscellaneous/provan/treespace>.

1.4 The Fréchet Mean

One of the important applications of the geodesic algorithm is computing the Fréchet mean. For most data sets, notions of center of data, such as the mean or median provide useful descriptive statistics. For data sets in tree space though, it is challenging to define analogous objects, since trees can neither be ordered nor operated on as Euclidean points. If we only take into account the structure information of trees but no edge length, then identifying a single “best” representative for a set of trees is well-studied in phylogeny, and such trees are usually called “consensus trees”. Consensus trees are difficult to calculate, however, and do not tend to take edge lengths into account effectively. The *Fréchet mean* is a more promising candidate for measuring the center of a set of trees.

The Fréchet mean is characterized as the solution to a non-linear optimization problem. For a set of points $\mathbf{T} = \{T^1, T^2, \dots, T^r\}$ in \mathcal{T}_n the *Fréchet function* $F : \mathcal{T}_n \rightarrow \mathbb{R}_{\geq 0}$ is the mean square distance

$$F(T) = \frac{1}{r} \sum_{l=1}^r d(T, T^l)^2$$

where $d(\cdot, \cdot)$ is the geodesic distance, and the *Fréchet mean* is

$$\bar{T} = \operatorname{argmin}_{T \in \mathcal{T}_n} F(T).$$

The Fréchet mean \bar{T} is unique because \mathcal{T}_n is $CAT(0)$ and $d(T, T')$ is a convex function of T for fixed T' [Sturm, 2003]. The minimum value $F(\bar{T})$ of the Fréchet function is the *Fréchet variance*. The Fréchet mean is a useful notion of geometrically central tree of the data points in \mathbf{T} , and the Fréchet variance summarizes the variability of the dataset. The purpose of the Fréchet mean is to give a single point summary of an entire

dataset.

Since the Fréchet mean is such a useful summary statistic for tree data, it is worth some effort to review the *inductive mean algorithm* presented by [Sturm, 2003, Theorem 4.7], and further analyzed in [Bačák, 2012] and [Miller et al., 2015]. Given a set of trees $\mathbf{T} = \{T^1, T^2, \dots, T^r\}$, let $\{Y_1\}, \{Y_2\}, \dots$ denote a sequence of independently and identically distributed random trees chosen uniformly from \mathbf{T} . Then the following sequence of trees converge to the Fréchet mean of \mathbf{T} :

$$S_1 := Y_1,$$

and

$$S_k := \left(1 - \frac{1}{k}\right) S_{k-1} + \frac{1}{k} Y_k,$$

where the right hand side of the second equation denotes the point $\frac{1}{k}$ of the distance along the geodesic from S_{k-1} to Y_k . The point S_k is called the *kth inductive mean* of Y_1, \dots, Y_k . Theorem 4.7 of [Sturm, 2003] states that this inductive mean converges to the Fréchet mean for data sampled from any $CAT(0)$ space. In particular, the expected squared distance between S_k and \bar{T} is bounded by $F(\bar{T})/k$. This is also the theoretical foundation of the following inductive mean algorithm in [Bačák, 2012].

Inductive Mean Algorithm	
Input: $\{T^1, T^2, \dots, T^r\}$	
Step 1	$S_1 := T^1, i := 1$
Step 2	choose $k \in \{1, \dots, r\}$ at random
Step 3	$S_{i+1} := \frac{1}{i+1} T^k + \frac{i}{i+1} S_i$
Step 4	$i := i + 1$
Step 5	go to Step 2

1.5 The Combinatorics of Geodesics

Geodesics play an essential role in tree space data analysis and this section summarizes some useful results about the combinatorial structure of geodesics in \mathcal{T}_n from [Miller et al., 2015, Section 3]. Fix a source tree $T \in \mathcal{T}_n$, and consider the geodesic $\Gamma_X = \Gamma(T, X)$ from an arbitrary tree $X \in \mathcal{T}_n$ to T . Γ_X has a combinatorial structure specified by the support pair $(\mathcal{A}, \mathcal{B})$ associated with the geodesic. This support pair can change even when X stays in the same orthant, depending on the precise values of the edge lengths in X . Miller et al. constructed a partition of \mathcal{T}_n into regions for which all geodesics to the fixed tree T have the same combinatorial structure. This partition is called the *vistal subdivision* of \mathcal{T}_n and its major properties are summarized below.

Definition 1.5.1. [Miller et al., 2015, Definition 3.1] Given a source tree $T \in \mathcal{T}_n$, a maximal orthant $\mathcal{O} \subset \mathcal{T}_n$, 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 Γ_X joining each 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. For convenience in notations, the tree $X = (V, \mathcal{E}, W, \Sigma)$ can be thought of as a vector in $\mathbb{R}_+^{\mathcal{E}}$, whose coordinates are expressed using the corresponding lower-case letter x .

Definition 1.5.2. [Miller et al., 2015, Definition. 3.2] The *squaring map* $\mathcal{T}_n \rightarrow \mathcal{T}_n$ acts on $X \in \mathcal{T}_n \subset \mathbb{R}_+^{\mathcal{E}}$ by squaring the coordinates of X :

$$(x_e | e \in \mathcal{E}) \rightarrow (\xi_e | e \in \mathcal{E}), \text{ where } \xi_e = x_e^2$$

Denote by \mathcal{T}_n^2 the image of this map, and let $\xi_e = x_e^2$ denote the coordinate indexed by $e \in \mathcal{E}$. 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 by $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$. With this change of variables, given any edge set $S \subset \mathcal{E}$, $\|S\| = \sum_{e \in S} \xi_e$.

The squaring map induces on the Fréchet function F a corresponding pullback function

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

Since the Fréchet function $F(T)$ is continuous on \mathcal{T}_n with a uniquely attained minimum by convexity, and continuously differentiable on the interior of every maximal orthant, the same properties hold for F^2 . Thus descent methods apply after squaring just as beforehand.

Proposition 1.5.1. [Miller et al., 2015, Proposition 3.3] 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 $\xi \in \mathbb{R}^{n-2}$, where all norm $\|\cdot\|$ are to be interpreted as $\|\cdot\|_T$.

(O) $\xi \in \mathcal{O}$; that is, $\xi_e \geq 0$ for all $e \in \mathcal{E}$, and $\xi_e = 0$ for $e \notin \mathcal{E}$, where $\mathcal{O} = \mathbb{R}_{\geq 0}^{n-2}$.

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

Remark: (P2) and (P3) here are actually the same as (P2) and (P3) in Theorem 1.3.1, but written in the form of multiplication instead of fraction. (O) is just a nonnegativity constraint.

Proposition 1.5.2. [Miller et al., 2015, Proposition 3.4] *The vial 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 vial facet $\mathcal{V}^2(T, \mathcal{O}; \mathcal{A}, \mathcal{B})$ if and only if the inequalities in (O), (P2), and (P3) are strict.*

Definition 1.5.3. [Miller et al., 2015, Definition 3.5] Fix a source tree $T \in \mathcal{T}_n$, a (not necessarily maximal) orthant $\mathcal{O} \subset \mathcal{T}_n$, and a support $(\mathcal{A}, \mathcal{B})$. A *signature* associated with the support $(\mathcal{A}, \mathcal{B})$ is a length $k - 1$ sequence $\mathcal{S} = (s_1, \dots, s_{k-1})$ of symbols $s_i \in \{=, \leq\}$. The *previal 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 each point X has the following specific form:

$$\frac{\|A_1\|}{\|B_1\|} s_1 \frac{\|A_2\|}{\|B_2\|} s_2 \dots s_{k-2} \frac{\|A_{k-1}\|}{\|B_{k-1}\|} s_{k-1} \frac{\|A_k\|}{\|B_k\|}.$$

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

Note that vial cells are convex polyhedra. A canonical description of vial cells is given in [Miller et al., 2015, Theorem 3.23].

Theorem 1.5.3. *Fix a tree $T \in \mathcal{T}_n$.*

1. *Vial 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 $(\mathcal{A}, \mathcal{B})$. Here a signature is a list of "=", "<", and " \leq " symbols in (P2).*
2. *The dimension of a vial cell $\mathcal{V}(T, \mathcal{O}; \mathcal{A}, \mathcal{B}; \mathcal{S})$ 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} .*

Definition 1.5.4. [Miller et al., 2015, Definition 3.31] A *premultivial cell* for a collection \mathbf{T} of trees is a set of the form

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

where $\mathcal{V}(T^l, \mathcal{O}; \mathcal{A}^l, \mathcal{B}^l)$ are previal cells, and $\mathcal{O} \subset \mathcal{T}_n$ is an orthant, and

$$(\mathcal{A}^{\mathbf{T}}, \mathcal{B}^{\mathbf{T}}) = \{(\mathcal{A}^1, \mathcal{B}^1), \dots, (\mathcal{A}^r, \mathcal{B}^r)\}$$

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

1.6 PCA in Euclidean Space

Principal component analysis (PCA) has been a workhorse method for understanding population structure of a data set in Euclidean space. A good overview and discussion of many aspects of PCA can be found in [Jolliffe, 2005]. A common goal of PCA is *dimension reduction*: finding principal components out of the original set of variables and making sure the number of principal components is less than the number of original variables. There are *population* principal components and *sample* principal components. Since we mainly deal with a sample of trees rather than the entire tree population in this research proposal, only the sample PCA (see [Mardia et al., 1979]) will be studied here. Let $X = (x_1, \dots, x_n)'$ be an $n \times p$ sample data matrix, i.e. the rows of X are the data objects. Let a be a standardized p -vector, i.e. $a'a = 1$. Since variation is reasonably defined around a center point, we work with a mean centered version of the data, $X - \mathbf{1}\bar{x}'$, where $\mathbf{1}$ is an $n \times 1$ vector with all entries equal to 1 and \bar{x}' is a $1 \times p$ vector with the i^{th} entry being the mean for the column i in X . Then the *standard linear combination* (SLC) Xa gives n observations on a new variable defined as a weighted sum of the columns of X . The sample variance of this new variable is $a'Sa$, where S is the sample covariance matrix of X . The first result is that the SLC with largest variance is the first principal component defined by

$$y_{(1)} = (X - \mathbf{1}\bar{x}')g_{(1)} \quad (1.2)$$

where $g_{(1)}$ is the standardized eigenvector corresponding to the largest eigenvalue of S (i.e. $S = G\Lambda G'$). Similarly, for $i = 2, 3, \dots$, the i^{th} sample principal component is defined as

$$y_{(i)} = (X - \mathbf{1}\bar{x}')g_{(i)} \quad (1.3)$$

where $g_{(i)}$ is the standardized eigenvector corresponding to the i^{th} largest eigenvalue of S . And we have the following properties for principal components:

- The principal components are uncorrelated and the variance of $y_{(i)}$ is λ_i , the i^{th} largest eigenvalue of S .
- Among all the SLCs of columns of X which are uncorrelated with the first k principal components of X , the $(k + 1)$ th principal component has largest variance.
- The sum of the first k eigenvalues of S divided by the sum of all the eigenvalues of S , $(\lambda_1 + \dots + \lambda_k)/(\lambda_1 + \dots + \lambda_p)$, represents the “proportion of total variation” explained by the first k principal components.

- If S has rank $r < p$, then the total variation of X can be entirely explained by the first r principal components.

1.7 PCA on Manifolds

As noted in [Marron and Alonso, 2014], an interesting extension of PCA is to the “mildly non-Euclidean” setting, such as data objects on a manifold. The value of thinking about data on a curved manifold was first motivated in the area of directional data (i.e. the data objects are angles). Such data objects arise in many contexts, such as wind directions, magnetic directions, etc. Some good examples can be found in [Fisher, 1995; Fisher et al., 1993; Mardia, 1972]. To illustrate the motivation of thinking about directional data as lying on a manifold, we will take the following example from [Marron and Alonso, 2014]. Consider a data set consisting of four angles, 1° , 2° , 358° , 359° , and consider their average. By the fact that there are just four numbers, it is natural to compute the classical arithmetic mean. The result here of 180° is typically not satisfactory if we view the angles as points on the unit circle. From this perspective a mean of 0° makes more sense. The idea behind this toy example is: we think of paths between data points on a manifold as “moving along a surface using geodesics”.

As there are quite a few notions of data center in Euclidean space (such as mean, median, trimmed means, and so on), there are also many ways of defining center on manifolds. On a curved manifold, the Fréchet Mean (see section 1.4) is one of the commonly used descriptions of data center. To measure the variation about center within a data set on a manifold, there are plenty of analogs of PCA. We will start with a simple one in the sense that it only requires a metric, called *Multi-Dimensional Scaling* (MDS) [Borg and Groenen, 2005; Cox and Cox, 2001; Lingoes et al., 1979; Young and Hamer, 1987]. MDS takes a matrix of pairwise distances between data points and finds a lower dimensional representation of the data so as to optimize the relationships indicated by the distance matrix. The main advantage of MDS is that it works not only for manifold data but also for general metric space data. One major disadvantage of MDS is that it can be difficult to interpret. That is also why other generalizations of PCA for manifolds have been presented in the literature.

Since a manifold is a metric space where *there is an approximating tangent space at each point*, many of the refined generalizations of PCA rely on this property. The idea behind the *Principal Geodesic Analysis* (PGA) in [Fletcher et al., 2004] is exactly the tangent plane characterization. Similar to the idea of PCA, starting from the sample mean and then finding directions of maximal projected variance, PGA starts at the Fréchet mean and finds the geodesic along which the projections of the data have maximal variance. The nice thing about manifolds is that the geodesics passing through the Fréchet mean could be characterized by directions in the tangent plane. Hence we can first project all the data points onto a proper tangent plane

and get the PCA there, and then map the results back into the manifold. One concern about PGA is that the principal geodesics are constrained to go through the Fréchet mean. It was shown in [Huckmann et al., 2010] that this was a serious constraint by using an example of uniformly distributed data points along the equator on the sphere S^2 . The Fréchet mean of the data are the north and south poles, so PGA could only find a line of longitude which gives a poor representation of the data. To address this issue [Huckmann et al., 2010] proposed *Geodesic Principal Components* which finds the best fit over all geodesics.

On the sphere S^2 , by relaxing the constraint of searching within only geodesics, [Jung et al., 2011] proposed *Principal Arc Analysis* (PAA). PAA was motivated by the data following closely to a *small circle* (meaning not a great circle). In this case, the above modifications of PCA need to find at least two components to explain the variation. However, the nature of the data is still one dimensional, so having more than one component is unsatisfying. PAA addresses this by finding the best fit of any small circle to the data. Later on, [Jung et al., 2012] extended the idea of PAA to data lying on higher dimensional spheres, i.e. S^d , where $d > 2$, which is called *Principal Nested Spheres* (PNS). The main idea behind PNS is iterative dimension reduction. For addressing the problem that data objects lie in the Cartesian product of many spheres, [Pizer et al., 2012] proposed the method of *Composite Principal Nested Spheres* (CPNS). The approach of CPNS is to do a sphere by sphere decomposition of the data using PNS, then to combine the scores into a large Euclidean vector, and finally to apply PCA to a collection of such large vectors.

1.8 PCA of Tree Structured Data Objects

[Marron and Alonso, 2014] pointed out that a more challenging extension of PCA is to “strongly non-Euclidean” settings, such as tree-structured or graph-structured data objects. Several attempts at extending PCA into tree structured data have already been made in literature. An early approach to the tree structured data PCA was in [Wang and Marron, 2007] where the statistical use of the term OODA first appeared. In that approach, the focus was on topological structure. Given a set of trees, a two-step procedure was performed: first, an optimal nested sequence of trees, called *principal structure treeline*, was obtained; second, for each tree in the principal structure treeline, an optimal *principal attribute treeline* was calculated (attributes being only on nodes). This line of research was continued in [Aydin et al., 2009a]. Aydin et al. used the term of principal structure treeline from [Wang and Marron, 2007] as their foundation, and one of their main contributions is a linear time computational method for a production scale data set of trees.

In another recently published work [Shen et al., 2013], authors applied the *Dyck path representation* to transform trees into curves, and then used the power of PCA in Functional Data Analysis to explore statistical properties of tree data. However, all of the above approaches experience a common limitation which was correspondence between tree branches. This was the result of embedding 3-d trees into 2-d space.

Some work has been done in the area of PCA in phylogenetic tree space. [Nye, 2011] explicitly studied

the topic of the *first principal component* in tree space. The basic idea of defining the first PC is closely related to the idea of [Wang and Marron, 2007]:

- 1 Given a set of trees $\{T^1, T^2, \dots, T^r\}$, construct a center T^0 .
- 2 Given a geodesic π through T^0 , project $\{T^1, T^2, \dots, T^r\}$ onto π by finding the closest point y_i in π to T^i for $i = 1, \dots, n$.
- 3 Find the geodesic π such that the points y_i have the greatest variance along the geodesic.

In step(1), Nye chooses the center T^0 to be the *majority consensus tree* [BARTHÉLÉMY, 1986]. The majority consensus topology consists of splits which are found in strictly more than half the trees in the data set. Due to the highly complex combinatorial features of tree space, it will be very computationally intense to directly follow the above procedure since it is not possible to try all the geodesics passing through the center. A natural idea is to come up with some good heuristics with acceptable computational complexity. The major contribution of [Nye, 2011] is to propose a greedy algorithm “ Φ PCA” which computes an approximated first principal component. For the detail of the algorithm, see [Nye, 2011, Section 2]. The main idea of Φ PCA is that the principal component is constructed greedily by adding one coordinate in each step instead of all at once. Intuitively speaking, Nye reduces the problem of finding the direction with largest projection variance in an n -dimensional space into a series of 2-dimensional subproblems. This fundamental work of Nye is pioneering in tree space, but this area needs more comprehensive investigation due to interesting and complex topological structure.

1.9 Overview of Dissertation

In this dissertation, we develop three approaches of Principal Component Analysis in tree space, in order to better understand the data structure in tree-like data sets. See [Jolliffe, 2005] for a thorough review of PCA. The main body of the dissertation is organized as follows. In Chapter 2, we introduce several data sets: Brain Artery data, Uniformly Random data, Wright-Fisher data, and Reduced Brain Artery data, which are going to be used as testbeds for three PCA type approaches throughout the rest of the dissertation. In Chapter 3, we focus on a Euclidean-based approach, namely *Multi-Dimensional Scaling* (MDS), since MDS approximates the tree space data in an appropriate Euclidean space. We explore the geometry of the Brain Artery data by using MDS, and in particular we discuss the *out-of-sample embedding* in aiding MDS visualization for the same data set. See [Borg and Groenen, 2005; Cox and Cox, 2001; Lingoes et al., 1979; Young and Hamer, 1987] for a detailed review of MDS. The other two approaches are based on analog of lines in tree space. A challenge is that the concept of line is not naturally defined in tree space as in Euclidean space. This motivates the characterization of notions of lines in tree space, developed in Chapter 4. We

show that these lines naturally split into type I lines and type II lines. Intuitively speaking, a type I line is the extension of a sufficiently long geodesic with no bifurcation, while a type II line is the extension of the union of a collection of short geodesics and it has many bifurcations. Directly modeling PCs as general type I lines can be computationally hard, thus, we study a practical analog of the type I line in Chapter 5, called *Sample-limited Geodesics*, which are the geodesics connecting a pair of data points. In Chapter 6, we continue to discuss methods to search for a restricted version of type II lines, called *principal ray sets*, of which all rays emanate from the origin. In particular, we investigate a special case of principal ray sets, namely *principal axis sets*, in more detail.

CHAPTER 2: REAL AND SIMULATED DATA

There are several data sets which will be used through out the rest of this thesis to test and contrast the methods being developed. The purpose of this chapter is to highlight relevant aspects of several test bed data sets in order to enhance intuitive understanding of the numerical studies. In Section 2.1 the source of the brain artery data and the generation of the corresponding tree data set is discussed. Section 2.2 describes the generation of random tree data sets. In Section 2.3 the Wright-Fisher model, which is a useful tool to generate topologically similar trees, is considered. Useful contrasts between these example data sets are given in Section 2.4, using some exploratory data analyses. At the end of this chapter, a method of generating reduced brain artery data sets will be proposed.

2.1 Brain Artery Data

One of the original motivations in this thesis is to look for latent correlates of biological variables such as sex, age, and handedness from a set of human brain artery trees. Those brain artery trees were constructed from a raw data set of Magnetic Resonance Angiography (MRA) brain images collected by the CASILab at the University of North Carolina at Chapel Hill. The data set is publicly available and can be downloaded following the link on the MIDAG website [Bullitt et al., 2008]. See [Bullitt et al., 2005] for a simple summary based analysis of this image database. The full data set contains images of the brains of 109 healthy subjects and each image is tagged with subject features of age, sex, handedness and self-identified race.

2.1.1 Magnetic Resonance Angiography

Magnetic Resonance Angiography (MRA) is an important non-invasive technique in medical imaging to visualize arteries. MRA uses the fact that blood flowing in the arteries has a distinct magnetic signature. And MRA can only detect arteries with blood in relatively rapid motion, but not veins because blood moves more slowly there². Data are stored as a 3-D image of intensities at voxel location, where brightness indicates the motion of blood flowing through arteries. This can be visualized using 2-D slices. Figure 2.5(a) gives one slice of a MRA image of a human brain. The outside dark region indicates the space out of the head. The outer layer mainly in grey represents skin and skull and the scattered bright pieces might be the blood vessels in the skin. The next dark layer is cerebrospinal fluid. The inside large region is a piece of the cerebral cortex. Four darker grey parts near the center are ventricles. The bright pieces are our focus, namely brain arteries, which can be either dots or segments. Bright dots are arteries going through the MRA slice and bright segments are those going across the MRA slice. Figure 2.5(b) gives a detailed view of

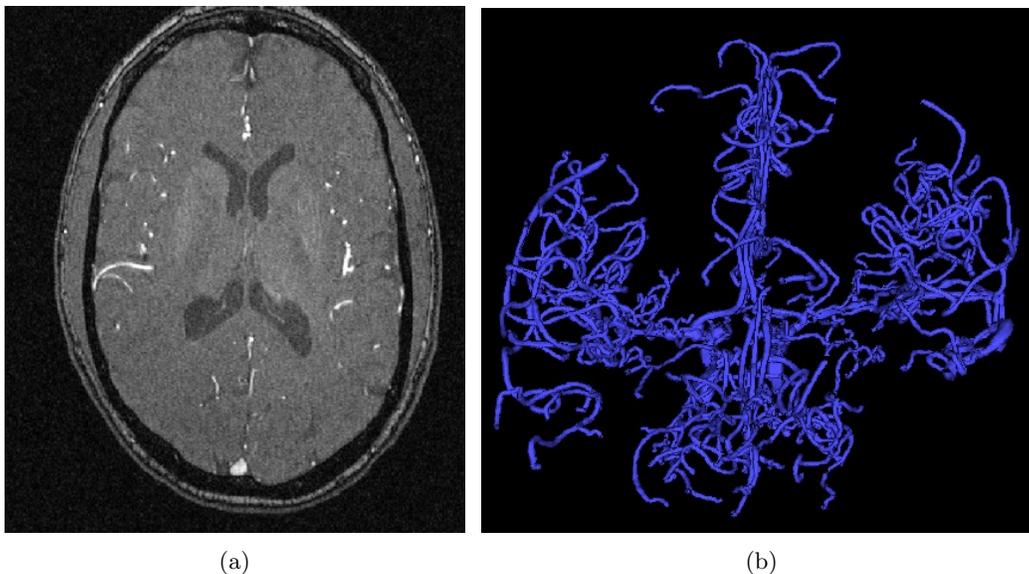


Figure 2.5: (a) One slice of a Magnetic Resonance Angiography (MRA) image for one person. Bright regions indicate blood flow. These are tracked through MRA slices to recover artery tubes as shown in the figure on the right. (b) The data object is a reconstructed brain artery tree for one person. The goal of this research is PCA type statistical analysis of a sample of such data objects.

one reconstructed 3-D artery system. An important point is that MRA has a resolution threshold of about 1mm. Consequently many small arteries, between 1mm in diameter down to capillaries, are too small for detection.

A challenge in this study is the starting point of the artery tree. For uniformity we work with a set of subtrees, where each subtree starts at the Circle of Willis, see [Aydin et al., 2009b]. The major part of this research develops PCA type statistical methods for the tree data set by using the phylogenetic tree space as a mathematical foundation. Recall from Section 1.2.2, a set of phylogenetic trees must have a common leaf set. However, the arteries detected by MRA do not even have the same number of branches across all the subjects. To address this issue, a common leaf set is artificially introduced by determining points on the cortical surface that correspond across different subjects. The next section describes the details of representing brain artery systems as points in phylogenetic tree space.

2.1.2 Cortical Correspondence

The work described in this section heavily relies on the elegant work done by [Oguz et al., 2008]. After reconstructing the brain artery system as the blue lines in Figure 2.6(a), we are ready to map it into a phylogenetic tree space by using cortical correspondence. A *cortical correspondence* is a method of determining landmarks on the cortical surface which correspond across different subjects. A group-wise shape correspondence algorithm based on spatial locations, which is proposed in [Oguz et al., 2008], is used here to place landmarks on the cortical surface. In this research 128 landmarks are selected, 64 for each hemisphere,

which are the red dots in Figure 2.6(a). For each landmark, the closest point on the artery tree system, called the *landmark projection*, is found. Each landmark is then connected to its landmark projection by a red segment in Figure 2.6(b). Each landmark and the line segment to its projection become part of the data object. Each phylogenetic tree data object is finally formed by tracing the parts of the tree that are between the base and the landmarks. Any parts of the original brain artery system that are not between a landmark projection and the base, found as cyan in Figure 2.6(b), are trimmed. The resulting tree is in Figure 2.6(c).

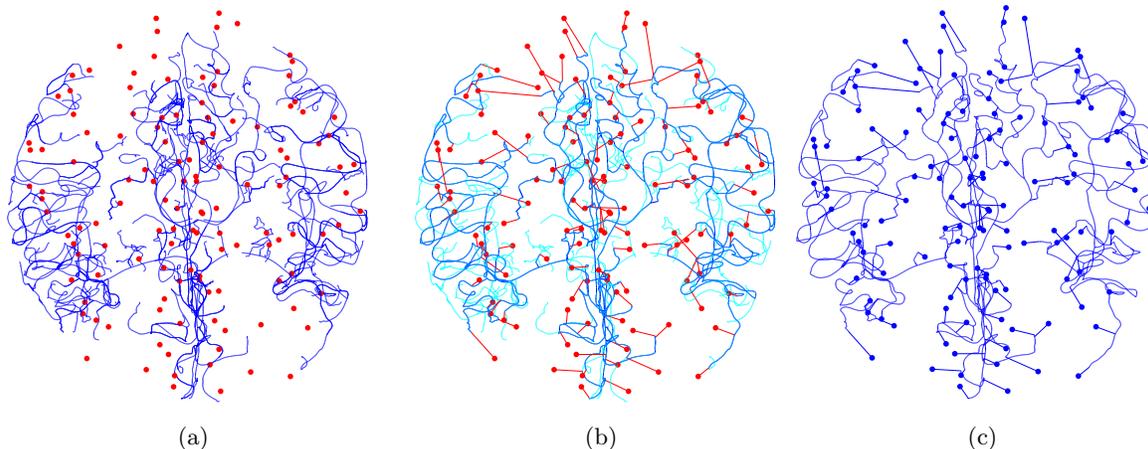


Figure 2.6: Illustrates steps for *cortical correspondence*. (a) Brain artery system (blue) with the cortical landmarks (red). (b) Find the closest point on the brain artery system to each landmark (red segments are landmark connections). Arteries which are not between the landmark projections and the base (cyan) are trimmed. (c) The result of this procedure is a cortical correspondence tree.

To form one tree instead of a few subtrees, an artificial leaf, called the *root*, is added to the base of the reconstructed data object to substitute for the Circle of Willis. The root and the 128 landmarks together add up to a common set of 129 leaves. Recall from Section 1.2.2, each edge of a reconstructed phylogenetic tree is associated with a positive edge length. The edge length for each interior edge is the corresponding arc length in the original brain artery system. The pendant for each landmark has length equal to the projection distance plus the original artery length from the projection point to the nearest artery branch. The pendant length for the root of the tree is zero.

Recall that the raw data set contains images of the brains of 109 healthy subjects. However, after performing cortical correspondence only 67 images produced eligible phylogenetic trees, hence the real data set will contain 67 brain artery trees throughout this thesis.

2.2 Generated Data Set I: Uniformly Random

In this section generation of a *uniformly random tree data set* is discussed. This could be interpreted in several ways, but here it means that the generated trees are uniformly randomly distributed over all possible tree topologies. To make the uniformly random tree set comparable with the brain artery tree set, 67 random

phylogenetic trees with 129 leaves will be generated. Recall from Section 1.2.2 the fact that the number of all possible topologies for a tree with 129 leaves is $253!! \approx 10^{250}$.

Most of the algorithms for generating randomness heavily rely on some sort of random number generator. However, just to generate a random integer between 1 and 253!! is infeasible in double-precision floating-point numerical computing environments such as MATLAB. Therefore, in personal correspondence, Sean Skwerer proposed another approach which will be used to generate a random tree data set. Instead of randomly pulling one from a pool of topologies, let the tree randomly *grow* from a base-tree with only three leaves. The whole growth process is an iterative procedure. At each iteration, one of the existing edges is uniformly randomly selected and then cut in the middle. An interior node and a new leaf will be inserted into this cut. Iterations continue until the resulting tree has 129 leaves. Now we show that we indeed get uniformly random tree topologies on n leaves by using this approach. First, we need to show that all possible tree topologies on n leaves can be generated by *random growth*. For any tree topology on n leaves, we can sequentially delete leaves and their associated leaf edges in descending order of leaf label, until only leaves with label 1, 2, and 3 left. Then the reverse process is a random growth. Second, we show that each possible topology has the same probability to be constructed by using induction. Starting from the base-tree with only 3 leaves, it is straightforward to see that the 3 possible topologies on 4 leaves have equal probability to appear when inserting the leaf labeled 4. Now suppose all the topologies generated by random growth on k leaves have equal probability to appear. In the next iteration of random growth, since each of the $2k - 3$ edges has equal probability to be chosen and cut in middle, all the topologies generated by random growth on $k + 1$ leaves will also have equal probability to be constructed. Therefore, we can obtain uniformly random tree topologies by using random growth. See Figure 2.7 for an illustration of all three possibilities in the first iteration. Figure 2.7(a) shows the starting tree with only three leaves: 0, 1, and 2. If the red solid edge is cut and a new node 3 is inserted, the tree in (b) is obtained. If the blue dotted edge is cut and a new node 3 is inserted, the tree in (c) is obtained. If the green dashed edge is cut and a new node 3 is inserted, the tree in (d) is obtained. Starting from the tree in (a), there is $1/3$ chance of getting each tree in (b), (c), and (d).

After a random topology is generated, it remains to generate edge lengths. Random edge lengths create additional data variation. To make this uniformly random data comparable with the brain artery data, the distribution of edge lengths is taken to be the empirical distribution of edge lengths of the brain artery trees. Figure 2.8 shows the overlay of a *kernel density estimate* (KDE) of all the edge lengths in the brain artery data and a Gamma distribution with parameters 1.28 and 15.8. A kernel density estimate (KDE) generates an empirical density curve for a data sample, see [Wand and Jones, 1995] for more details. Intuitively speaking, a KDE is a smooth histogram. By comparing these two curves, the edge lengths distribution of brain artery data can be reasonably approximated by the Gamma distribution. Thus, in our simulations,

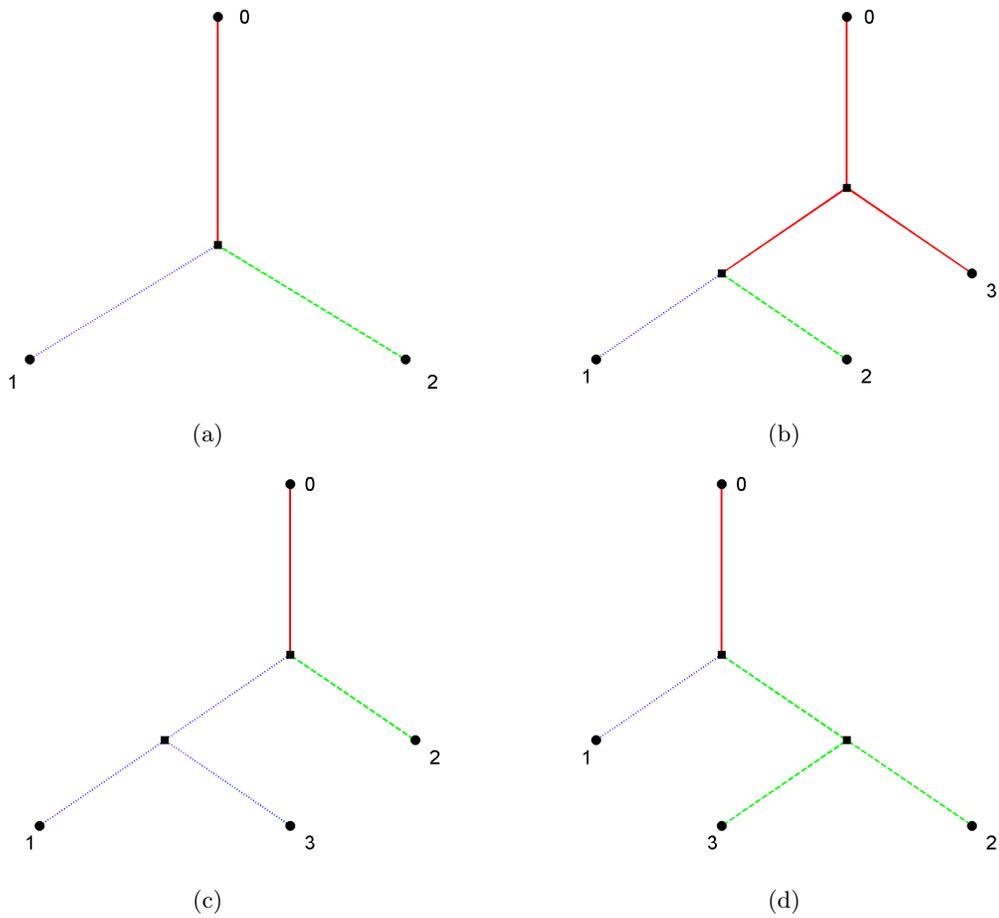


Figure 2.7: Schematic plots showing the three possibilities to insert a new leaf in the first iteration.

the edge lengths of uniformly random trees will be generated from this Gamma distribution.

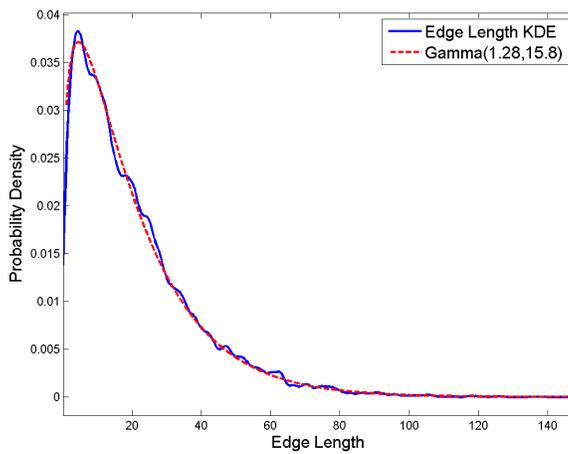


Figure 2.8: Comparison between KDE of edge lengths in brain artery data (blue solid curve) and Gamma distribution with parameters 1.28 and 15.8 (red dashed curve) shows the Gamma distribution is an excellent fit for the edge lengths distribution of the brain artery data.

2.3 Generated Data Set II: The Wright-Fisher Model

Besides the random tree data set, some simulated data sets with controlled similarity of topology are also considered. Using the relationship between phylogenetic trees and genealogy, the Wright-Fisher model, introduced by [Wright, 1931] and [Fisher, 1930], will also be used to generate some test bed data sets in this thesis.

2.3.1 The Wright-Fisher Model

The basic concepts of the Wright-Fisher model are summarized in [Hein et al., 2005, Section 1.4]. The Wright-Fisher model is a simplified mathematical model of populations describing genealogical relationships. This model of reproduction provides a dynamic description of the evolution of an idealised population and the transmission of genes from one generation to the next. The following assumptions are made in the Wright-Fisher model:

1. Generations are discrete and non-overlapping.
2. The population is made of haploid organisms, that is, the genes making up the present generation are drawn randomly with replacement from the parental generation.
3. The population size is constant.
4. All individuals have equal reproductive ability.
5. The population has no geographical or social structure.
6. The genes (or sequences) in the population do not recombine.

The original Wright-Fisher model does not consider growth and splitting of populations. However, to generate tree data, some extensions of the model need to be considered. [Wilson et al., 2003] proposed a more complex model based on the Wright-Fisher model. This model assumes that the population starts to grow exponentially at some point instead of being constant all the time. Furthermore, at random time points the population *bifurcates*, that is, splits into two separate subpopulations with equal size. We will adjust Wilson's model to fit our data generation scheme. In our model, population growth and bifurcation are assumed to occur at the same time. Further, the bifurcation points can be controlled to approximate a specified sample genealogical tree, as will be shown below. However, instead of exponential growth, the population grows according to a scalar factor which is the number of subpopulations in the next generation. Specifically, whenever the population splits, both of the subpopulations will have the same size as the ancestral population. The advantage of this assumption is that the subpopulations of all branches in the

genetic tree have the same population size and this single parameter is easy to control. From now on this parameter will be called the *width parameter*.

2.3.2 Data Generation

The actual data generating process starts from a template standard phylogenetic tree, with edge lengths assumed to be integer. (This may require the original lengths to be scaled or approximated.) This template tree is turned into a *population genealogical tree* which has the same topology, but additional attributes. A population genealogical tree preserves the edge lengths from the template tree by setting the number of reproduction cycles on each subpopulation to be the corresponding edge length in the template tree. To add population size as an additional uniform attribute on all edges, a population genealogical tree is constructed by choosing the population size equal to the predetermined width parameter. Figure 2.9(a) shows a template tree with 4 leaves and edge lengths as marked in boxes. Figure 2.9(b) gives the corresponding population genealogical tree with width parameter equal to 4. The population genealogical tree has the same topological structure as the template tree. Each horizontal line of dots represent a generation and the number of dots in each line equals the width parameter. The numbers of empty spaces between generations in each subpopulation reflect the edge lengths in the template tree. Each non-root leaf of the template tree becomes an individual from the latest generation of the corresponding subpopulation, and the root becomes the common ancestor of the whole population. Since the reproduction from generation to generation is random,

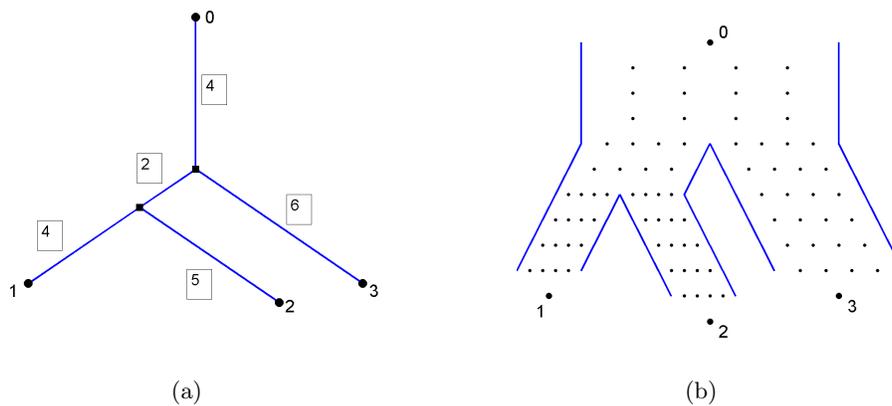


Figure 2.9: (a) A template phylogenetic tree with 4 leaves and edge lengths in boxes. (b) The corresponding population genealogical tree has the same topological structure as the template tree. The number of reproduction cycles in each subpopulation equals the corresponding edge length in the template tree and the size of subpopulation in each generation is the width parameter which equals 4.

a *sample gene tree* will be generated within the population genealogical tree. If the width parameter is set to be 1, then all the sample gene trees generated from this model will have exactly the same topology as the template tree. As the width parameter increases, a greater variety of topologies will appear among the generated trees. Figure 2.10 shows two possible sample gene trees generated within the same population

genealogical tree displayed in Figure 2.9(b). In both figures, three individuals labeled as 1, 2, and 3 have been randomly sampled in the most present generation of each subpopulation. Edges back in time tracking the ancestors of these three individuals are highlighted as green dashed lines. In (a) individuals 1 and 2 find a common ancestor first, then they find the common ancestor with 3. In (b) individuals 2 and 3 find a common ancestor first, then they find the common ancestor with 1.

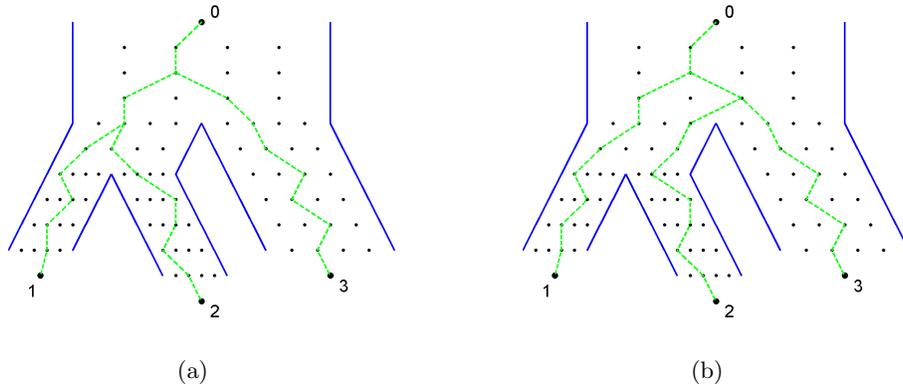


Figure 2.10: Two possible sample gene trees generated within the same population genealogical tree, from Figure 2.9(b).

To better visualize these two sample gene trees, we represent them in the form of phylogenetic trees in Figure 2.11 below. The phylogenetic tree in (a) corresponds to the sample gene tree in Figure 2.10(a), and it has the same topology as the template tree in 2.9(a). The phylogenetic tree in (b) corresponds to the sample gene tree in Figure 2.10(b), and it has a topology different from the template tree. Although the Wright-Fisher model can generate trees with topologies quite different from the template trees, they will tend to be “centered” around the original template tree. It is not known in what way the trees generated by the Wright-Fisher model statistically represent the original tree.

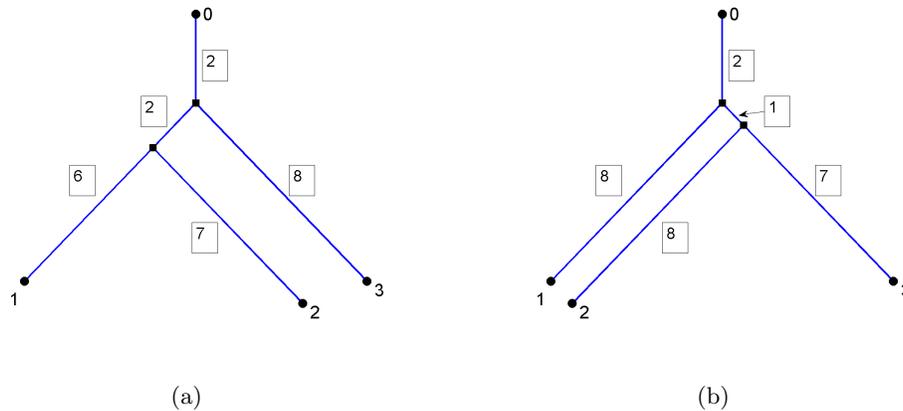


Figure 2.11: Two possible sample gene trees with different topologies can be generated within the same population genealogical tree, from Figure 2.9(b).

2.4 Exploratory Data Analysis

In sections 2.1, 2.2, and 2.3, a variety of example data sets that will be used in later chapters are introduced. These are expected to exhibit different levels of similarity. We now focus on 5 specific cases, in order of decreasing similarity:

- WF2 (Wright-Fisher data with width parameter = 2) has a high level of similarity.
- WF10 has similarity level lower than WF2 from Section 2.3.
- WF40 has even less similarity from Section 2.3.
- Brain artery data will be seen in Sections 2.4.1 and 2.4.2 to have less similarity than WF40.
- Uniformly random data will also be seen in Sections 2.4.1 and 2.4.2 to have the least similarity among all 5 cases.

2.4.1 Angle-based Data Summaries

One way to measure the similarity of tree data topologies is to study the distribution of angles, with vertex at the origin, between each pair of trees (called *pairwise angle*). Given two trees T_1 and T_2 , denote the pairwise angle between these two trees as θ , then we can define θ by the cosine law:

$$\cos\theta = \frac{\|T^1\|^2 + \|T^2\|^2 - L(\Gamma(T^1, T^2))^2}{2\|T^1\|\|T^2\|}. \quad (2.4)$$

It can be shown that θ does not depend on either $\|T^1\|$ or $\|T^2\|$. Under this definition, if $\Gamma(T^1, T^2)$ is a cone path, then $\theta = 180^\circ$, otherwise $\theta < 180^\circ$. A good general definition of angle in any metric space is the *Alexandrov angle* [Alexandrov, 1951]. In the special case of phylogenetic tree space, the definition of angle by the cosine law in (2.4) coincides with Alexandrov angle.

The distributions of pairwise angles for the five data sets are visualized using kernel density estimation (KDE), as in Section 2.2. Figure 2.12(a) shows the overlay of KDEs for the five example data sets which allow direct comparison of these populations. The red curve corresponds to WF2 and shows that all of these pairwise angles are smaller than 10° . The magenta curve represents WF10, indicating almost all the pairwise angles are between 30° and 50° . These much larger angles for WF10 are very consistent with the greater spread of WF10 data across tree space. The green curve corresponds to WF40, showing most of the pairwise angles are between 80° and 120° , again consistent with more spread for WF40. The blue curve represents the brain artery data, with all the pairwise angles between 120° and 170° . This shows that the spread of the brain artery data is more than even the diverse WF40 distribution. The black curve on the far right corresponds to the uniformly random data, showing most of the pairwise angles are greater than

160° and a big proportion of angles are 180°, showing the brain artery data set is not purely random. The overall comparison of these five distributions is consistent with the similarity ordering of these five data sets mentioned in the bullet points just before Section 2.4.1. Very often the spread of a data set is proportional to its mean, to investigate this issue, Figure 2.12(b) presents the overlay of the logarithms of the pairwise angles. Except for the uniformly random data, all other four data sets have similar spread, which indicates the variability in pairwise angles is proportional to the magnitude of angles.

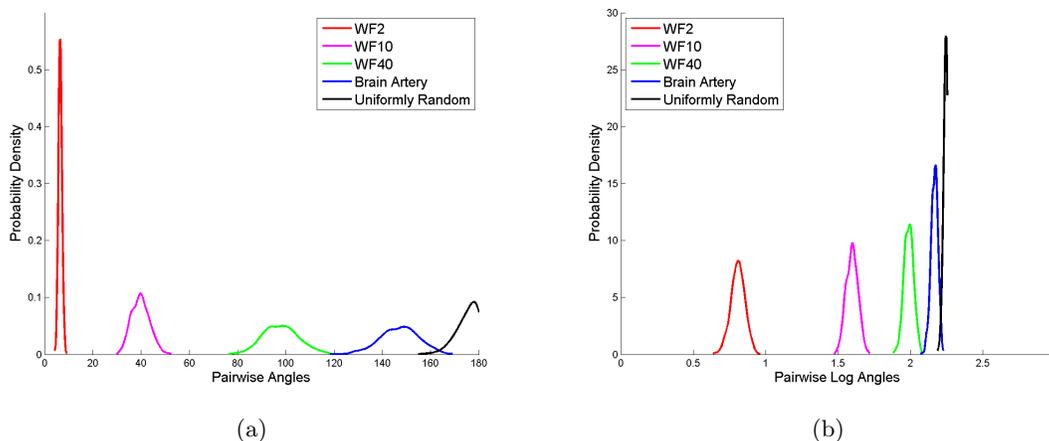


Figure 2.12: (a) Overlay of pairwise angle KDE plots shows the decreasing similarity ordering of the five data sets: WF2, WF10, WF40, brain artery data, and uniformly random data. (b) Overlay of logarithm of pairwise angle KDE plots indicates that the variability in pairwise angles is proportional to the magnitude of angles for WF2, WF10, WF40, and brain artery data, but not for the uniformly random data.

2.4.2 Distance-based Data Summaries

Another way to examine the similarity of a set of trees is to study the distances between each pair of trees (called *pairwise distance*). This is defined as the length of the geodesic between the pair of trees. It is intuitive that larger pairwise distances correspond to less similar data sets. Figure 2.13(a) shows the overlay of pairwise distance KDE plots for the same five data sets which is another useful comparison. Again, the red curve corresponds to WF2 and shows that all the pairwise distances are within a narrow range smaller than 50. The magenta curve represents WF10, indicating almost all the pairwise distances are between 150 and 250, which shows that WF10 has more spread than WF2. The green curve is consistent with even larger spread for WF40, showing most of the pairwise distances are between 400 and 600. The blue curve representing the brain artery data and the black curve representing the uniformly random data on the far right overlap heavily. Both indicate pairwise distances distributing from 450 to 750, again consistent with the fact that these two data sets have the largest spread. The comparison of distributions of pairwise distances for these five data sets are quite similar to that of pairwise angles, except that the blue and black curves are relatively separated for pairwise angles but overlapped heavily for pairwise distances. This indicates that

the distances between each tree and the origin differ for the brain artery data and the uniformly random data, probably because the distribution of edge lengths used in the uniform generation is different from the true distribution of the edge lengths in the brain artery data. This shows that it is worth looking at both pairwise angles and pairwise distances. Figure 2.13(b) presents the overlay of the logarithms of the pairwise distances, which again gives a good indication of proportionality between the spread and magnitude of pairwise distances.

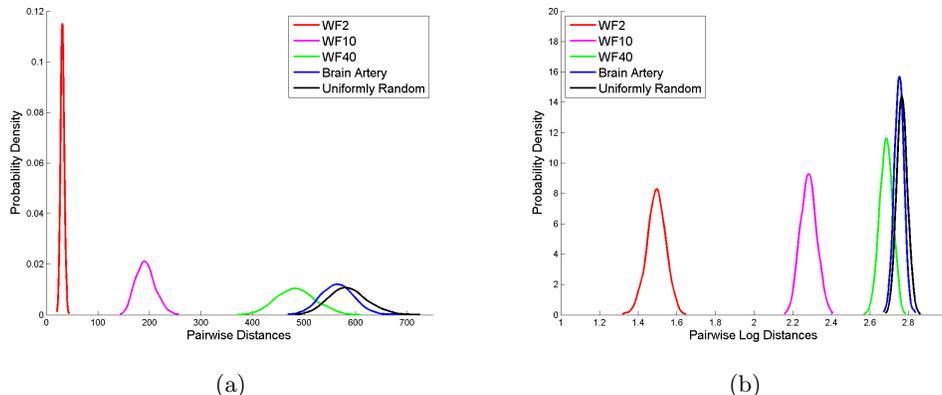


Figure 2.13: (a) Overlay of pairwise distance KDE plots shows the decreasing similarity ordering of the five data sets: WF2, WF10, WF40, brain artery data, and uniformly random data. (b) Overlay of the logarithm of pairwise distance KDE plots implies the proportionality between the spread and magnitude of pairwise distances.

2.5 Stickiness and Landmark-reduced Brain Artery Data

An important property of the data in tree space is the behavior of the Fréchet mean. In Euclidean space, the Fréchet mean, which is also the conventional entry-wise sample mean, of a finite set of data points will move slightly when one of the data points is perturbed. In contrast, we can find data sets in some particular tree spaces, whose Fréchet mean stays still when one of the data points moves within a certain range. This interesting phenomenon is called a *sticky* Fréchet mean. In [Hotz et al., 2013], the *stickiness* of the Fréchet mean was defined precisely on an “open book”, which is a metric space constructed by gluing a disjoint union of half-spaces along their boundary hyperplanes. An “open book” can also be viewed as a subset of an appropriate tree space. In [Basrak, 2010], the stickiness property was independently and simultaneously proven for arbitrary binary trees. As a continuation of the work initiated in [Hotz et al., 2013], [Barden et al., 2013] further investigated the stickiness in \mathcal{T}_4 , the second simplest tree space. From the results in [Hotz et al., 2013] and [Barden et al., 2013], a reasonable speculation is that the stickiness is likely to occur for a largely spread data set, even in any general tree space. From the discussion in the previous section, we saw that the spread of each of the five data sets is quite different from each other. The two data sets with large spread, namely the Brain Artery data and the Uniformly Random data, are likely to have a sticky

Fréchet mean.

One goal of this dissertation research is to understand the effect of different aspects of a data set on the performance of our tree space principal components. Stickiness is certainly a potentially interesting aspect of tree space data. Since the suspected strong stickiness of the Brain Artery data and the Uniformly Random data is associated with their large data spread, ideally we want a series of data sets with different levels of spread. Fortunately, a group of landmark-reduced Brain Artery data sets with different levels of spread have been created and thoroughly studied in [Skwerer, 2014]. These data sets come from a slightly different data source from our Brain Artery data. Each of them contains 85 subjects instead of 67, but 64 of the 67 subjects in our Brain Artery data are included in those 85 subjects. In addition, the original set of 128 landmarks are chosen differently from our Brain Artery data. The landmarks of the 85-subject data are chosen by only considering the locations on the cortical surface, but the landmarks of the 67-subject data are chosen by taking into account both locations and curvatures. However, from the results in [Skwerer, 2014, Section 3.2], we clearly see different levels of data spread across these data sets, which is our focus here.

Specifically, the landmark-reduced Brain Artery data sets are obtained based on a sequence of nested subsets of the original landmarks. The subset for each data set is chosen by maximizing the total data variation within the set. In Chapter 6, we will use some representative landmark-reduced Brain Artery data sets to investigate an important tree space version of principal components.

CHAPTER 3: MULTIDIMENSIONAL SCALING IN TREE SPACE

PCA type visualization is very helpful when dealing with high-dimensional complex data sets. A challenge of visualizing tree space is that it is strongly non-Euclidean. *Multidimensional Scaling* (MDS) gives one approach to addressing that. MDS is an analog of PCA, which is applicable in any metric space, such as tree space. It searches for a lower dimensional representation of a data set based only on pairwise distances defined between each pair of data objects. In Section 3.1, a brief review of MDS is given. In Section 3.2, 2-D and 3-D MDS are applied to the brain artery data set introduced in Chapter 2. In Section 3.3, it is seen that embedding one or more geodesics into a data set create a major distortion in the MDS, allowing study of how the geodesics behave in tree space. In Section 3.4, it is seen that an out-of-sample approach to MDS can mitigate this distortion.

3.1 Review of Multidimensional Scaling

In this section, we give a short introduction of MDS [Torgerson, 1952, 1958; Gower, 1966]. MDS finds an appropriate configuration in Euclidean space for a set of objects in any complex space as long as the dissimilarities between pairs of objects can be defined. A *dissimilarity matrix*, the most common input form of MDS, consists of dissimilarity data for each pair of objects. In the case of a metric space, the dissimilarity is taken to be distance, although MDS works in more general spaces. In this work, since we focus on tree space with a metric, the word “distance” will be used in most of this chapter. If the objects are labeled $i = 1, \dots, N$, the distances are given by $D_{i,j}$ and the distance matrix is given by D . MDS gives a good understanding of the relationships between the N data objects, by visually representing them as a set of “configuration points”, $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^k$, whose pairwise distances $\|\mathbf{x}_i - \mathbf{x}_j\|$ approximate each corresponding $D_{i,j}$. The dimension k of the configuration space is arbitrary in theory, but $k = 2, 3, 4$ are usual for the purpose of visualization.

Much of the early development of MDS was in the behavioural sciences and *Psychometrika* published many papers in this area. The theoretical foundation of MDS was built in the 1930s. [Eckart and Young, 1936] and [Young and Householder, 1938] proved when the input is a Euclidean distance matrix, the MDS representation is essentially the same as PCA coordinates. Based on this result, [Torgerson, 1952, 1958] and [Gower, 1966] made the first practical method available for MDS by replacing the Euclidean distance matrix with the more general dissimilarity matrix, which was the oldest version of non-Euclidean MDS, called “*classical MDS*”. [Kruskal, 1964a,b] offered a later version of MDS based on a direct measure of lack

of fit, called “*distance MDS*”. More recently, many variations of these ideas have been developed for a variety of applications, see [Borg and Groenen, 2005; Cox and Cox, 2001] for a good overview of this literature.

During decades of development, MDS has become a rich field in the literature. [Buja et al., 2008] suggests two dichotomies which give a clearer view of the whole topic. The first one is “*metric MDS* versus *nonmetric MDS*”: Metric MDS uses the actual values of the dissimilarities, while nonmetric MDS uses only their ranks [Shepard, 1962; Kruskal, 1964a]. Nonmetric MDS estimates an optimal configuration simultaneously with an optimal monotone transformation $f(D_{i,j})$ of the dissimilarities. Because the effect of edge lengths is taken into account in the tree space, the present research focuses on the metric MDS. The second dichotomy categorizes metric MDS into either *classical metric MDS* or *distance metric MDS*. Since we are going to focus mainly on metric MDS in this research, the terms “distance MDS” and “classical MDS” will be used instead of “distance metric MDS” and “classical metric MDS” in the rest of this chapter. The main difference between classical MDS and distance MDS is the loss function used in optimizing the MDS configuration. A loss function is a commonly used tool to measure the lack of fit between dissimilarities $D_{i,j}$ and fitted distances $\|\mathbf{x}_i - \mathbf{x}_j\|$. For metric MDS, *Stress* and *Strain* are two most frequently used loss functions.

Distance MDS uses Stress as the loss function which is a direct measure of disparities between dissimilarities $D_{i,j}$ and corresponding fitted distances $\|\mathbf{x}_i - \mathbf{x}_j\|$. In the simplest case, Stress is a residual sum of squares:

$$\text{Stress}_D(\mathbf{x}_1, \dots, \mathbf{x}_N) = \left(\sum_{i \neq j=1 \dots N} (D_{i,j} - \|\mathbf{x}_i - \mathbf{x}_j\|)^2 \right)^{1/2} \quad (3.5)$$

Then distance MDS will minimize Stress over all possible configurations $(\mathbf{x}_1, \dots, \mathbf{x}_N)^T$. The minimization can be carried out by applying standard gradient descent to Stress_D , which can be viewed as a function on \mathbb{R}^{Nk} .

Classical MDS uses Strain as the loss function and the idea of Strain originates from recovering a set of points in Euclidean space [Torgerson, 1952, 1958; Gower, 1966]. Some basic results about classical MDS are summarized here from [Borg and Groenen, 2005; Cox and Cox, 2001]. Given an $N \times N$ matrix of squared Euclidean distances D_E^2 , we want to find the MDS coordinate matrix X of $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^k$ up to a rotation or reflection. The matrix X has N rows, corresponding to the N objects. The number of columns of X can be taken to be any $k \leq n - 1$. Since distances do not depend on locations, we can assume that X has column means equal to zero to prevent arbitrary translation. In Euclidean space, there is an identity relating the distance $\|\mathbf{x}_i - \mathbf{x}_j\|$ and the inner product $\langle \mathbf{x}_i, \mathbf{x}_j \rangle$:

$$\|\mathbf{x}_i - \mathbf{x}_j\|^2 = \langle \mathbf{x}_i, \mathbf{x}_i \rangle - 2\langle \mathbf{x}_i, \mathbf{x}_j \rangle + \langle \mathbf{x}_j, \mathbf{x}_j \rangle.$$

Let I denote the $N \times N$ identity matrix, let $e = (1, \dots, 1)^T \in \mathbb{R}^N$, and denote the *centering matrix* as

$J = I - \frac{1}{N}ee^T$. From the above identity, the inner product matrix $B_E = -\frac{1}{2}JD_E^2J$ and this operation is called *double centering*. Note that matrix multiplication by J on the left or right removes the mean from each column or row respectively. To find the coordinate matrix X , we factor B_E by spectral decomposition, $B_E = Q\Lambda Q^T = (Q\Lambda^{1/2})(Q\Lambda^{1/2})^T = XX^T$. Classical MDS only differs from the above procedure by replacing the matrix of squared Euclidean distances D_E^2 with the matrix of the more general squared dissimilarities D^2 . We first obtain the inner product matrix $B = -\frac{1}{2}JD^2J$ by double centering D^2 . Then we compute the spectral decomposition of $B = Q\Lambda Q^T$, but this decomposition does not give us the coordinate matrix X directly. Instead, the coordinate matrix X is given by $X = Q_+\Lambda_+^{1/2}$, where Λ_+ is the diagonal matrix of the first k largest eigenvalues *greater than zero*.

It needs to be pointed out that both distance MDS and classical MDS will get the same configuration up to a rotation or shift of origin if the given distances are Euclidean and the dimension of the MDS space is the Euclidean rank of the data. However, if the distances are not Euclidean, these two versions of MDS will behave very differently. Even when the dissimilarities are Euclidean, but the dimension of the MDS space is less than the Euclidean rank of the data, these two versions of MDS still perform differently in the following example shown in Figure 3.14. In 3-D Euclidean space, consider three multivariate normal distributions: $N_3((5, 0, 0)^T, \mathbf{I})$, $N_3((0, 5, 0)^T, \mathbf{I})$, and $N_3((0, 0, 5)^T, \mathbf{I})$. From each of these three distributions, a sample of 10 data points is randomly generated. Now we have three clusters of points in 3-D Euclidean space, colored red, green, and blue. Then an identical 3 dimensional MDS configuration is achieved by both distance MDS and classical MDS. For a better visualization, this common 3 dimensional MDS configuration is displayed in the form of a *scatter plot matrix* in Figure 3.14. The scatter plot matrix puts 1-D projections on the diagonals and corresponding pairwise projection scatter plots off the diagonals based on the three MDS directions. On the diagonals, there are three smoothed histograms of the data projections onto each MDS direction, of which the horizontal axes give the projection scores onto the three MDS directions and the vertical axes give the probability density of the projection scores. The first diagonal plot for MDS direction 1 shows a clear separation of three colors. The second diagonal plot for MDS direction 2 displays only a separation of red from green and blue, since it is orthogonal to direction 1. The third diagonal plot for MDS direction 3 shows no separation at all, because it is orthogonal to both directions 1 and 2. Consequently, the scatter plot of MDS directions 1 and 2 shows a perfect separation of the three original clusters. From the scatter plot of MDS directions 1 and 3, we can see that the three clusters are still kind of separated, but the separation is not as clear as for directions 1 and 2. Finally, the scatter plot of MDS directions 2 and 3 only separates red from green and blue.

This section is concluded with discussion of the connection between two versions of metric MDS and PCA. [Cox and Cox, 2001, Section 2.2.7] proved: under the condition that the dissimilarities are given by

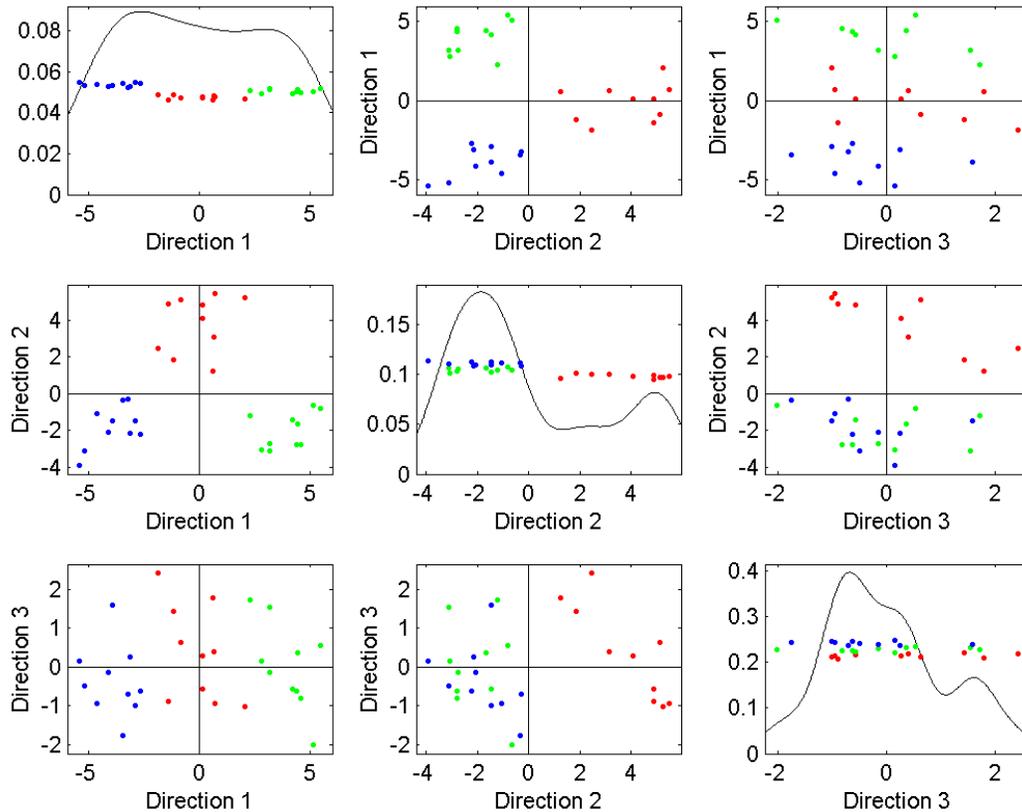
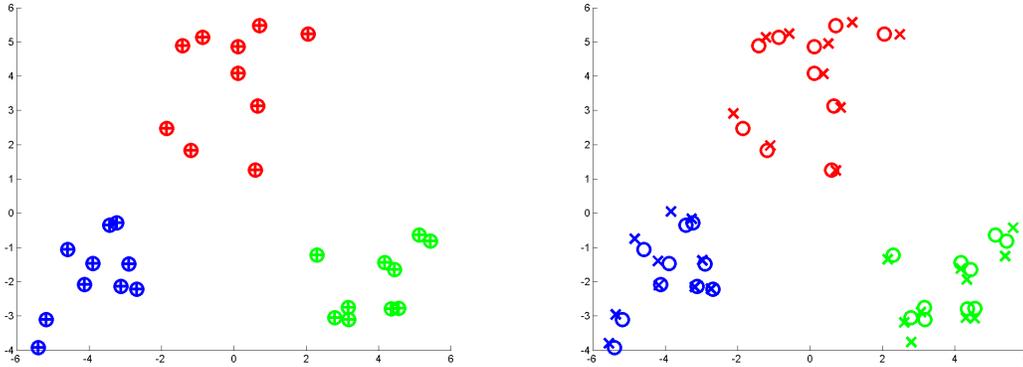


Figure 3.14: The above scatter plot matrix gives a visualization of 3-D MDS configuration on 3 clusters from front, side, and top views. Both distance MDS and classical MDS produce this same 3-D configuration. Three data clusters are colored red, green, and blue.

Euclidean distances, classical MDS is equivalent to PCA. In fact, the coordinates obtained in p dimensional configuration space for the N objects by classical MDS equal the component scores of the N objects on the first p principal components. However, distance MDS performs differently from both classical MDS and PCA even when the dissimilarities are given by Euclidean distances. In Figure 3.15(a), each point in 2-D classical MDS configuration of the above 3-cluster data set is represented by a “○” and the configuration of principal component 1 versus 2 scores are represented by “+”. The perfect overlap of these two configurations illustrates that classical MDS and PCA agree with each other on Euclidean data. In Figure 3.15(b), the same 3-cluster data set is represented using symbols “○” for the 2-D classical MDS configuration and using symbols “×” for the 2-D distance MDS configuration. The imperfect overlap of these two configurations indicates that classical MDS and distance MDS give different low rank approximation even in a Euclidean data space. Due to better algebraic properties, we will focus on classical MDS in tree space, and “MDS” will be used instead of “classical MDS” from now on for clarity.



(a) 2-D classical MDS configuration is plotted using “+” and 2-D PCA configuration is plotted using “O”. The symbols of “+” and “O” are perfectly overlapped, which indicates classical MDS and PCA are equivalent on Euclidean data.

(b) 2-D classical MDS configuration is plotted using “O” and 2-D distance MDS configuration is plotted using “x”. The symbols of “O” and “x” are not perfectly overlapped, which indicates classical MDS and distance MDS perform differently even on Euclidean data.

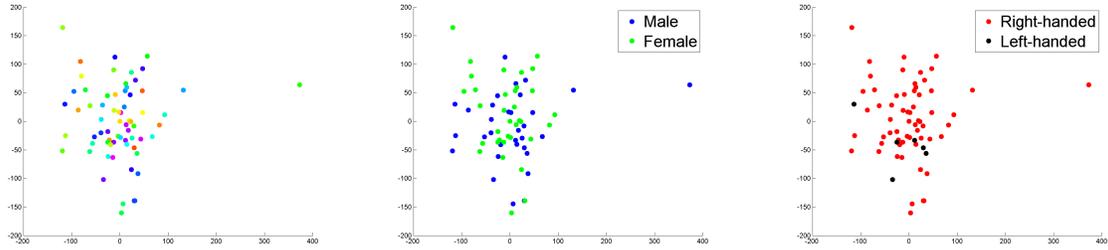
Figure 3.15: Comparison across classical MDS, distance MDS, and PCA on 3-cluster (colored red, green, and blue) simulated data set. Shows low rank approximations of Euclidean data are different between classical MDS and distance MDS.

3.2 MDS of Tree Data

Recall one of the original motivations in this thesis is to search for potential correlates of biological variables such as sex, age, and handedness from a set of human brain artery trees. To better detect the potential classes, MDS is applied to aid visualization by embedding trees into a 2-D or 3-D Euclidean space.

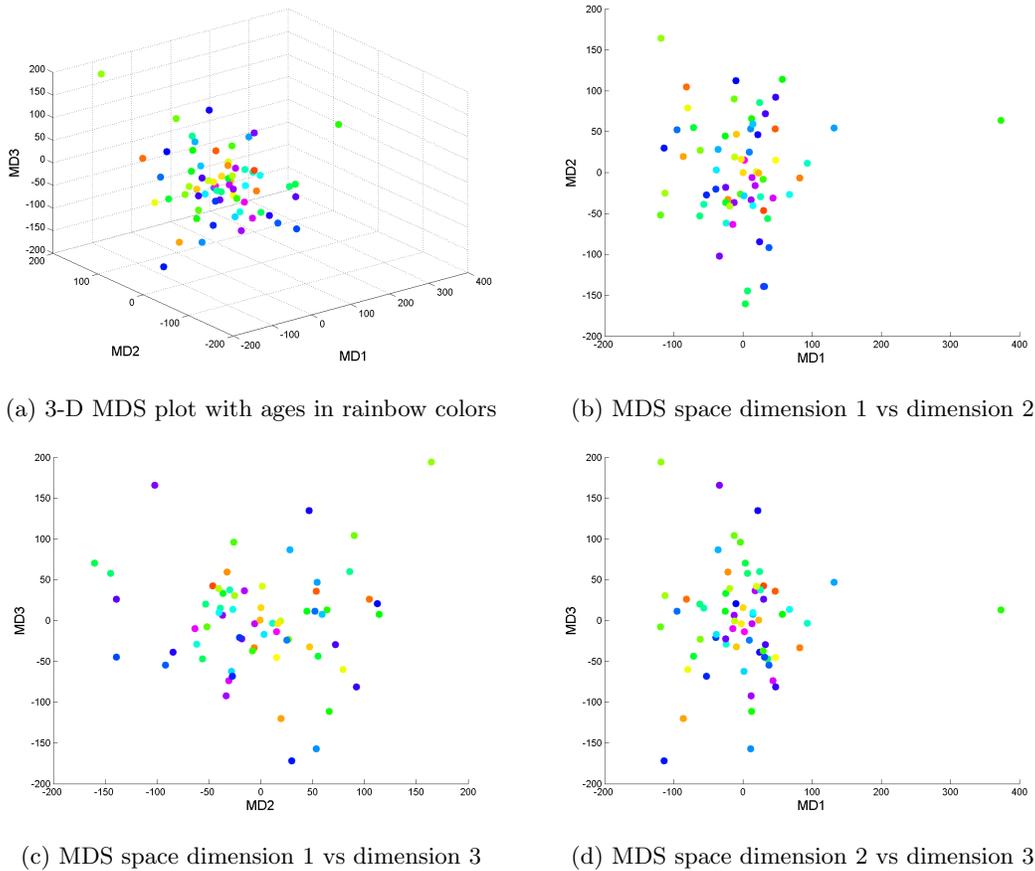
Recall MDS is based on the pairwise distance matrix of the data set. As in Section 1.3, distance between two trees is defined as the length of the geodesic connecting them, which can be computed by applying the linear time algorithm in [Owen and Provan, 2009]. Then MDS is performed on the geodesic distance matrix \tilde{D} with configuration space being 2-D or 3-D Euclidean space. The following Figure 3.16 shows the 2-D MDS plots of brain artery trees from 67 healthy people. All three sub-figures have the same scatter plot of 67 points, but each one is associated with a unique biological variable. In Figure 3.16(a) the points are colored with a rainbow color scheme: magenta for age 22, through blue, cyan, green, yellow, to red for age 72. This color scheme has been used in [Marron and Alonso, 2014] to show systematic changes in mortality across the years. Here we are hoping that the similar color scheme could help to show systematic change in MDS coordinates across the ages, but unfortunately the distribution of colors seems rather random. In Figure 3.16(b) males are denoted by blue dots and females are denoted by green dots. One might hope to see a distinction between two groups. However, males and females all seem to mix together without clear clusters. In Figure 3.16(c) we use red dots and black dots to denote right-handed people and left-handed people respectively. Although a faint potential separation of left-handedness from right-handedness is present, no

conclusion can be made due to the small number of left-handed people. While it may be challenging to approximate a 126-dimensional data set by using a 2-dimensional configuration, the 2-D plot can still be a helpful visualization tool.



(a) 2-D MDS plot with ages in rainbow colors: magenta for age 22, through blue, cyan, green, yellow, to red for age 72
 (b) Same 2-D MDS plot with males in blue dots and females in green dots
 (c) Same 2-D MDS plot with right-handedness in red dots and left-handedness in black dots

Figure 3.16: 2-D MDS plots associated with 3 biological variables: age, gender and handedness. No clear correlates of these biological variables can be drawn.



(a) 3-D MDS plot with ages in rainbow colors
 (b) MDS space dimension 1 vs dimension 2
 (c) MDS space dimension 1 vs dimension 3
 (d) MDS space dimension 2 vs dimension 3

Figure 3.17: 3-D MDS plots in rainbow color scheme across ages along with top, side, and front views in separate sub-figures. Very limited information gained from adding the third dimension.

The configuration space can be easily expanded into 3-D Euclidean space. The above Figure 3.17 shows a 3-D plot for the 67 brain artery trees in rainbow color scheme across age along with the orthographic views. Figure 3.17(a) shows a snapshot of the 3-D plot taken from a particular angle for the same data set in the above 2-D plot. Figure 3.17(b) shows the top view of Figure 3.17(a), which consists of the first and second dimensions in the MDS space and is also exactly the same plot as in Figure 3.16(a). Figure 3.17(c) and Figure 3.17(d) are the side view and front view of the 3-D plot in Figure 3.17(a). We are hoping to get better separation by introducing the third dimension in the MDS space, but the gain is quite limited.

The same visualization is shown for gender as well. Figure 3.18(a) shows a snapshot of the 3-D plot for the same data set. Figure 3.18(b), Figure 3.18(c), and Figure 3.18(d) show the top view, side view and front view of Figure 3.18(a) respectively. However, there is no apparent conclusion available about the separation between males and females.

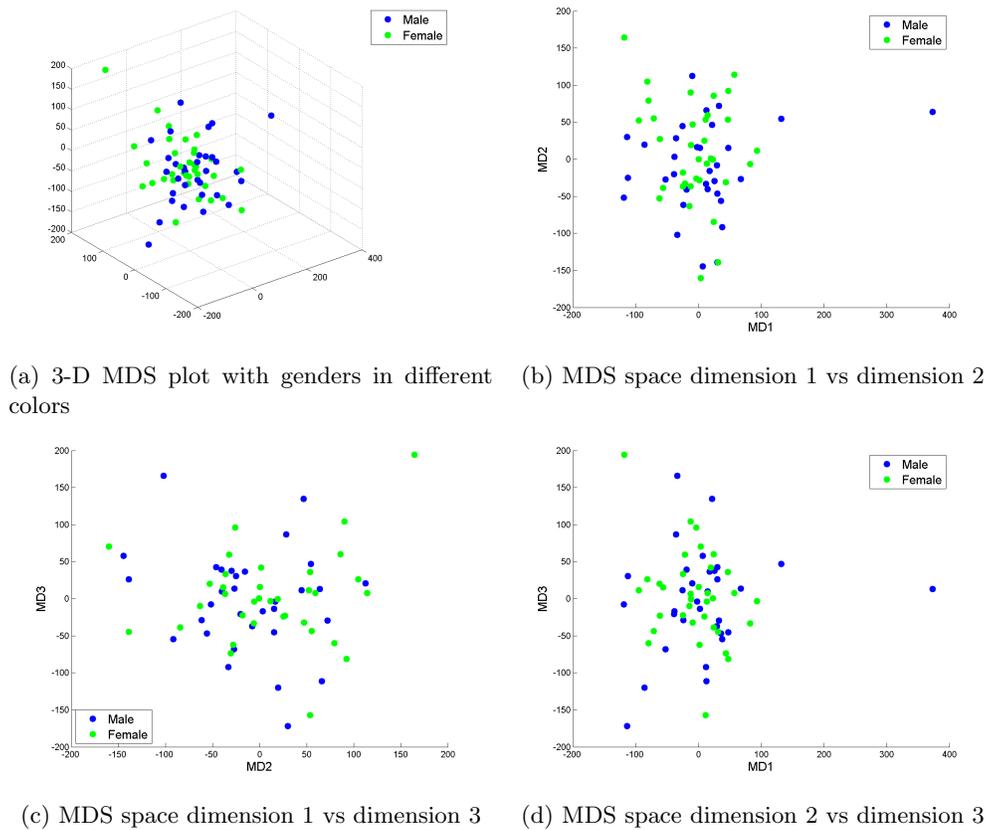


Figure 3.18: 3-D MDS plots with males in blue dots and females in green dots along with top, side, and front views in separate sub-figures. Very limited information gained from adding the third dimension.

Although no direct conclusion regarding classification is visually apparent from the 2-D or 3-D MDS plots, MDS is a useful method to represent tree data in a lower dimensional Euclidean space. In the next section, MDS will be applied to some specific tree space related problems.

3.3 MDS for Embedding Geodesics

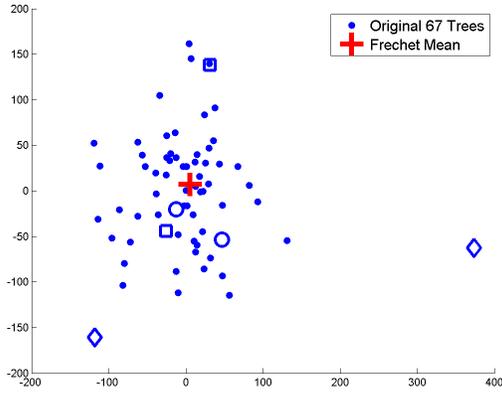
Since geodesics in tree space are thought of as analogs of straight lines in Euclidean space, it is interesting to see whether a geodesic turns out to be a straight line segment in the MDS configuration space. As noted in Section 2.4.2 geodesics between pairs of data trees are useful to consider. As three representatives, the minimum, median and maximum length geodesics are focused on. An equally spaced grid of points along a geodesic will be combined with the original data set, in a new MDS.

A preliminary task is to reconstruct the distance matrix \tilde{D} of the augmented data set (the union of the original data and the points on the geodesic). Suppose that there are n trees in the original brain artery data set with distance matrix D and k added grid points along the chosen geodesic. The distances between the added points and original trees are given by an $n \times k$ matrix B and the distances among the added points are denoted by a $k \times k$ matrix A . Then the augmented distance matrix \tilde{D} is given by

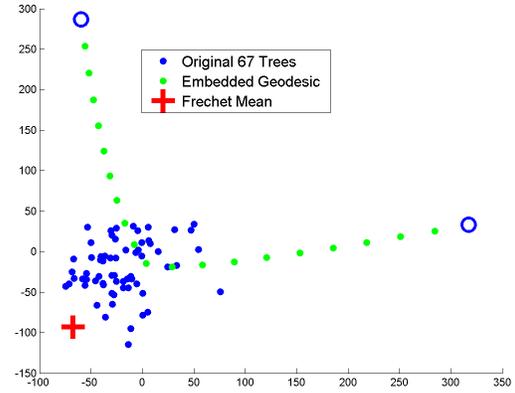
$$\tilde{D} = \begin{bmatrix} D & B \\ B^T & A \end{bmatrix}$$

MDS is applied to this augmented distance matrix \tilde{D} , and the plots in 2-D Euclidean configuration space are shown in Figure 3.19. Figure 3.19(a) shows the 2-D MDS plot based on only the 67 original trees shown as blue dots and their Fréchet mean as the red cross. Three pairs of trees marked as circles, squares, and diamonds, denote the end points of minimum, median, and maximum length geodesics, respectively. In Figure 3.19(b)(c)(d), 19 equally spaced grid points chosen from these three representative geodesics are embedded as green dots. It is surprising that the embedded geodesics completely dominate in these three MDS plots, and the original MDS configuration in (a) is seriously distorted by the embedded geodesics. This is because most data variation in the augmented data set is along the embedded geodesic and the first two MDS directions are strongly influenced by this path. It is also clear that all the embedded geodesics are heavily bent towards the cluster containing the rest of the data trees in MDS space. In addition, the Fréchet mean has been pushed away from the center position, and it is even outside the data cloud in (b) and (c). The location of the Fréchet mean is another indication of different levels of distortion to the original MDS configuration when embedding different geodesics.

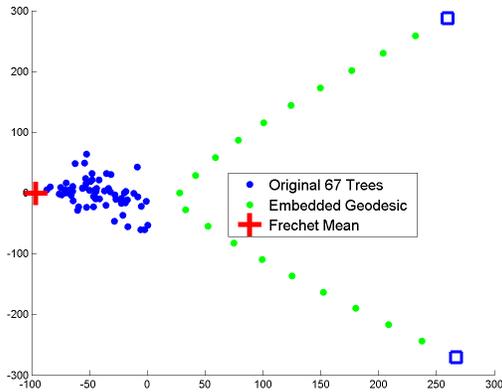
Recall the intuitive definition of $CAT(0)$ space from section 1.2.2: every triangle in a $CAT(0)$ space is skinnier than a triangle with exactly the same lengths of sides in Euclidean space. A more interesting question is to examine how a triangle in tree space looks in the MDS configuration space. Each triple of trees in the brain artery data set was taken and the angle sum of the triangle in tree space was calculated. To take some representative examples, triangles corresponding to the minimum, first quartile, median, third



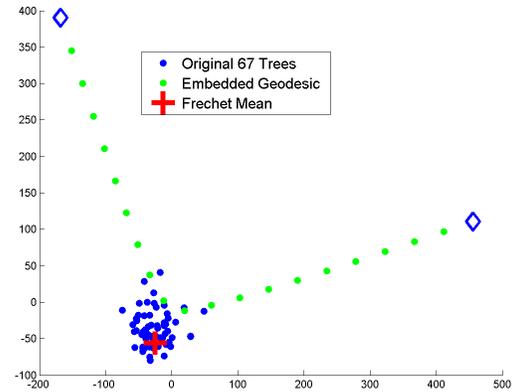
(a) 2-D MDS plot with Fréchet mean



(b) 2-D MDS plot with minimum length geodesic embedded



(c) 2-D MDS plot with median length geodesic embedded

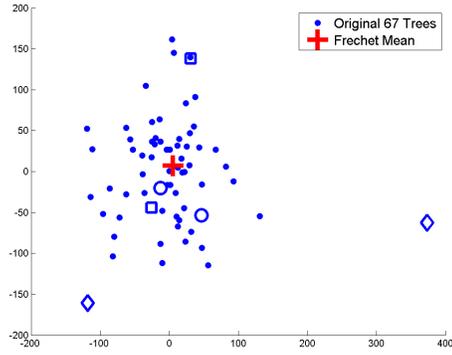


(d) 2-D MDS plot with maximum length geodesic embedded

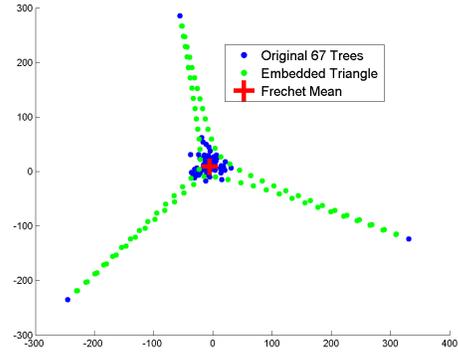
Figure 3.19: Examples of embedding geodesics in 2-D MDS plots

quartile and maximum angle sum were selected. To reconstruct the distance matrix \hat{D} , additional points were selected from each side of a triangle and then the augmented distance matrix \hat{D} was computed in the same way as \tilde{D} . The plots in Euclidean configuration space are displayed in Figure 3.20. Figure 3.20(a) again shows the 2-D MDS plot with the 67 original trees as blue dots and their Fréchet mean as the red cross. In Figure 3.20(b)(c)(d)(e)(f), 29 equally spaced grid points chosen from each side of these five representative geodesic triangles are embedded as green dots. It is clear that the angle sum of the embedded triangle in (a) is very small, and it gradually increases from there, until it is fairly large in (f). However, we notice that even the largest angle sum is obviously less than 180° because of the non-positive curvature in the tree space. Again we see that the embedded triangles dominate in these five MDS plots, and the original MDS configuration in (a) is seriously distorted. This can be explained by a similar reason to embedding a single geodesic: the two dimensional hyperplane with the largest data variation is heavily influenced by the

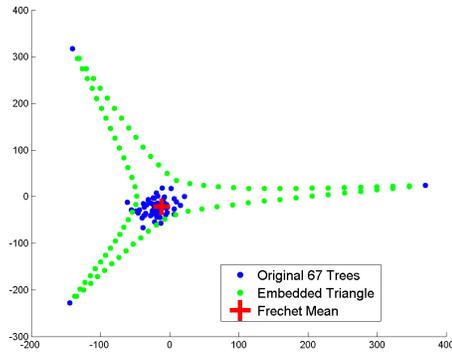
embedded triangle. In addition, the sides of each triangle are bent towards the data cluster, which is similar to the case of embedding a single geodesic. However, because the distorting effect of three geodesics are balanced with each other, the Fréchet mean is always in the middle of data cloud of the rest of the original trees.



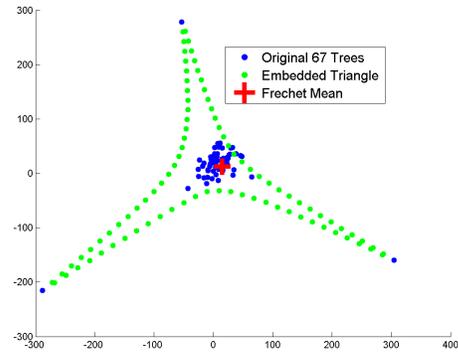
(a) 2-D MDS plot with Fréchet mean



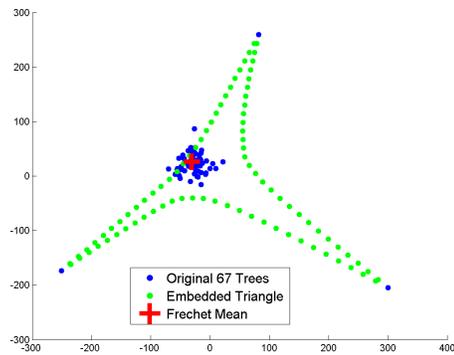
(b) 2-D MDS plot with the triangle with minimum angle sum embedded



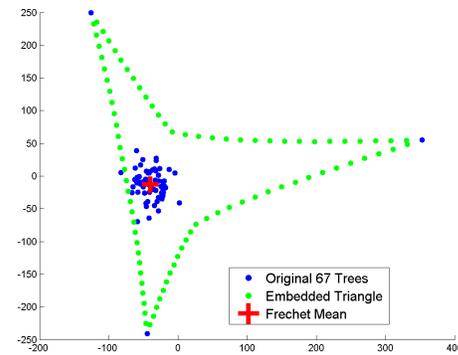
(c) 2-D MDS plot with the triangle with Q1 angle sum embedded



(d) 2-D MDS plot with the triangle with median angle sum embedded



(e) 2-D MDS plot with the triangle with Q3 angle sum embedded



(f) 2-D MDS plot with the triangle with maximum angle sum embedded

Figure 3.20: 2-D MDS plots with one triangle embedded

3.4 Out-of-Sample Embedding

In Figure 3.19, it is noticeable that the configuration of the original data points has been significantly altered after embedding geodesics. From the perspective of data visualization, it would be useful to embed geodesics while keeping the configuration of original data points fixed, i.e. we do not want the embedded geodesics to influence the original MDS configuration. One approach is to treat the embedded geodesics as *out-of-sample* objects without including them in the construction of the MDS configuration of the original data sample. Then, inserting the out-of-sample objects into the visualization is an *out-of-sample embedding* problem. Motivated by some semisupervised classification problems, [Trosset and Priebe, 2008] provided such an out-of-sample extension of classical MDS, which is formulated as an unconstrained nonlinear least-squares problem. The instant advantage of Trosset and Priebe’s approach is that the MDS configuration of the original sample is kept unchanged. A brief summary of Trosset and Priebe’s out-of-sample extension is given below.

That out-of-sample extension is based on the theory of classical MDS, hence it is convenient to start from there. Recalling notation from Section 3.1, let I denote the $n \times n$ identity matrix, let $e = (1, \dots, 1)^T \in \mathbb{R}^n$, and denote the centering matrix as $J = I - \frac{1}{n}ee^T$. By definition, a squared dissimilarity matrix $D^2 = [D_{i,j}^2]$ is a *Euclidean distance matrix* (EDM) if and only if there exist $x_1, \dots, x_n \in \mathbb{R}^p$ such that $D_{i,j}^2 = \|x_i - x_j\|^2$. The smallest such p is the *embedding dimension* of an EDM. Then the following theorem from classical geometry is very convenient for determining whether D^2 is an EDM.

Theorem 3.4.1. *A non-zero squared dissimilarity matrix D^2 is an EDM with embedding dimension p if and only if the symmetric matrix*

$$B = \tau(D^2) = -\frac{1}{2}JD^2J$$

is positive semidefinite and has rank p . Furthermore, if $D^2 = [D_{i,j}^2]$ is an EDM and

$$B = \tau(D^2) = \begin{bmatrix} x_1^T \\ \vdots \\ x_n^T \end{bmatrix} [x_1 \ \cdots \ x_n] = XX^T$$

then $D_{i,j}^2 = \|x_i - x_j\|^2$.

If D^2 is an EDM, then from the above theorem, τ maps squared Euclidean distances to Euclidean inner products. Since $Je = 0$, we have $Be = 0$. If D^2 is an EDM, and $X = [x_1, \dots, x_n]^T$, then $0 = e^TBe = e^TXX^Te = (X^Te)^T(X^Te)$, hence $X^Te = 0$. This means that τ maps squared Euclidean distances to Euclidean inner products of a configuration of points of which the centroid is the origin. We will see later

that the configuration of points made by classical MDS also have the origin as their centroid.

Now suppose that we are given a non-zero squared dissimilarity matrix D^2 which is not an EDM with embedding dimension $\leq p$, then $B = \tau(D^2)$ is not positive semidefinite with rank $\leq p$ any more, hence we can not factor B to obtain X . Classical MDS solves this problem by replacing B with \bar{B} , the closest (in the sense of Frobenius norm) positive semidefinite matrix with rank $\leq p$. Thus, classical MDS finds a configuration by approximating the inner products of complex components with those of real valued components in the least-squares sense. It turns out that a closed-form solution to the optimization problem can be obtained by discarding all the non-positive eigenvalues of $B = \tau(D^2)$ along with their eigenvectors, which gives the configuration as $X = Q_+ \Lambda_+^{1/2}$ in Section 3.1. This is a main advantage of classical MDS, however, we need to generalize this result to solve the out-of-sample problem.

The above Theorem 3.4.1 is a special case of a more general result. Given $w \in \mathbb{R}^p$, we say that the collection $x_1, \dots, x_n \in \mathbb{R}^p$ is w -centered if and only if $\sum_{j=1}^n w_j x_j = 0$. For w such that $e^T w \neq 0$, define a w -weighted centering matrix as

$$J_w = I - \frac{w e^T}{e^T w},$$

and then construct a w -weighted version of τ

$$\tau_w(D^2) = -\frac{1}{2} J_w^T D^2 J_w$$

Notice that τ_e is the τ in Theorem 3.4.1. Then we have the following more general formulation.

Theorem 3.4.2. *Let w be any vector not orthogonal to e , i.e. $e^T w \neq 0$, then the $n \times n$ squared dissimilarity matrix D^2 is an EDM with embedding dimension p if and only if there exists a w -centered spanning set of \mathbb{R}^p , $\{y_1, \dots, y_n\}$, for which*

$$\tau_w(D^2) = \begin{bmatrix} y_1^T \\ \vdots \\ y_n^T \end{bmatrix} [y_1 \ \cdots \ y_n].$$

In Theorem 3.4.2, the special case of $w = e$ was independently discovered by [Schoenberg, 1935] and [Young and Householder, 1938], and this special case was popularized in [Torgerson, 1952, 1958]. This general case of $e^T w \neq 0$ is due to [Gower, 1982, 1985].

For the out-of-sample problem, suppose there are n original objects and k out-of-sample objects. Let $D^2 = [D_{i,j}^2]$ denote the squared dissimilarities of the original n objects. Let the $n \times k$ matrix $\tilde{D}^2 = [\tilde{D}_{i,j}^2]$ denote the squared dissimilarities between the out-of-sample objects and the original objects. Let the $k \times k$ matrix $\hat{D}^2 = [\hat{D}_{i,j}^2]$ denote the squared dissimilarities of the out-of-sample objects. Let A^2 be the augmented

squared dissimilarity matrix of all $n + k$ objects, where

$$A^2 = \begin{bmatrix} D^2 & \tilde{D}^2 \\ (\tilde{D}^2)^T & \hat{D}^2 \end{bmatrix}.$$

By applying classical MDS to D^2 , we can get a configuration of n points x_1, \dots, x_n whose centroid is the origin of the embedding Euclidean space. However, by applying classical MDS to A^2 , we will get a configuration of $n + k$ points $x'_1, \dots, x'_n, y_1, \dots, y_k$ whose centroid is the origin of the embedding Euclidean space, so the centroid of x'_1, \dots, x'_n is not always the origin. Therefore, we can not solve this out-of-sample problem relative to D^2 by applying classical MDS to A^2 , because the center as well as the entire representation of the n original objects are often changed after embedding out-of-sample objects. Trosset and Priebe's out-of-sample extension circumvents this difficulty, as well as enables classification of new data beyond the training data, by preserving the original MDS configuration. Let $e = (1, \dots, 1)^T \in \mathbb{R}^n$, and let $w = (e^T, 0, \dots, 0)^T \in \mathbb{R}^{n+k}$. Using the notation of τ_w in Theorem 3.4.2, factoring $\tau_e(D^2)$ finds $X = [x_1, \dots, x_n]^T$, the $n \times p$ configuration matrix of the classical MDS embedding of the original n objects. Let $Y = [y_1, \dots, y_k]^T$ denote the $k \times p$ configuration matrix of the k out-of-sample objects to be solved for. Then the out-of-sample embedding is obtained by first computing the fallible (i.e. generally not positive semi-definite) inner product matrix

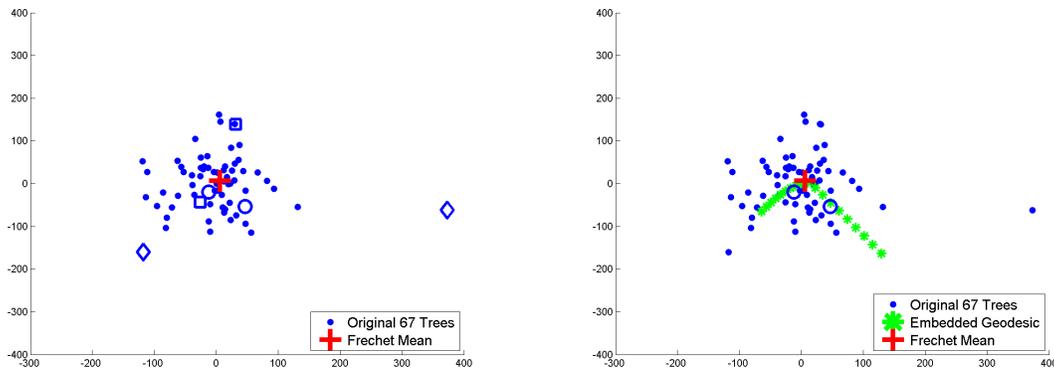
$$B = \tau_w(A^2) = \begin{bmatrix} \tau_e(D^2) & B_{xy} \\ B_{xy}^T & B_{yy} \end{bmatrix},$$

then solving for Y to minimize

$$\left\| B - \begin{bmatrix} X \\ Y \end{bmatrix} \begin{bmatrix} X^T & Y^T \end{bmatrix} \right\|^2 = 2 \left\| B_{xy} - XY^T \right\|^2 + \left\| B_{yy} - YY^T \right\|^2. \quad (3.6)$$

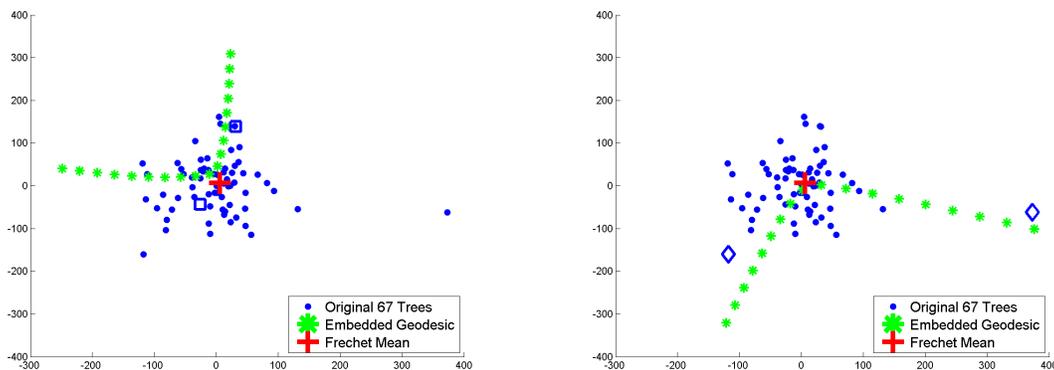
We then applied Trosset and Priebe's out-of-sample extension of classical MDS to our brain artery tree data. Recall in Figure 3.19, (a) shows the 2-D classical MDS configuration of 67 trees along with their Fréchet mean (original objects), and (b)(c)(d) illustrate these objects plus 19 equally spaced grid points chosen from the minimum, median, and maximum length geodesics respectively (out-of-sample objects). The goal now is to embed those grid points chosen from these three geodesics as out-of-sample objects into the configuration of 67 original trees along with their Fréchet mean. It is straightforward to adopt the above out-of-sample version of classical MDS except for the choice of the embedding dimension p . Since Trosset and Priebe were more interested in how to embed the out-of-sample objects into the original representation space in such a way that the original classifier can be applied, it was natural to choose p as the dimension

of the original representation space. However, our goal is to get a clear visualization of the embedded geodesic while keeping the MDS configuration of the original data points unchanged, so one might initially consider $p = 2$. Using the same notations as in loss function (3.6), let X be the 68×2 configuration matrix of the classical MDS embedding of the original objects, and let Y be the 19×2 configuration matrix of the out-of-sample objects. Let the 68×68 matrix D^2 denote the squared dissimilarities of the original objects, let the 87×87 matrix A^2 denote those of the augmented object set consisting of both the original and the out-of-sample objects. And let $e = (1, \dots, 1)^T \in \mathbb{R}^{68}$, let $w = (e^T, 0, \dots, 0)^T \in \mathbb{R}^{87}$. Then by applying classical MDS to D^2 , we can obtain X . To solve for Y , we need to minimize the loss function (3.6). Figure 3.21 shows the 2-D out-of-sample MDS representations of the embedded geodesics. In this



(a) 2-D MDS plot of the original 67 brain artery trees along with their Fréchet mean

(b) 2-D MDS plot with minimum length geodesic embedded as an out-of-sample object



(c) 2-D MDS plot with median length geodesic embedded as an out-of-sample object

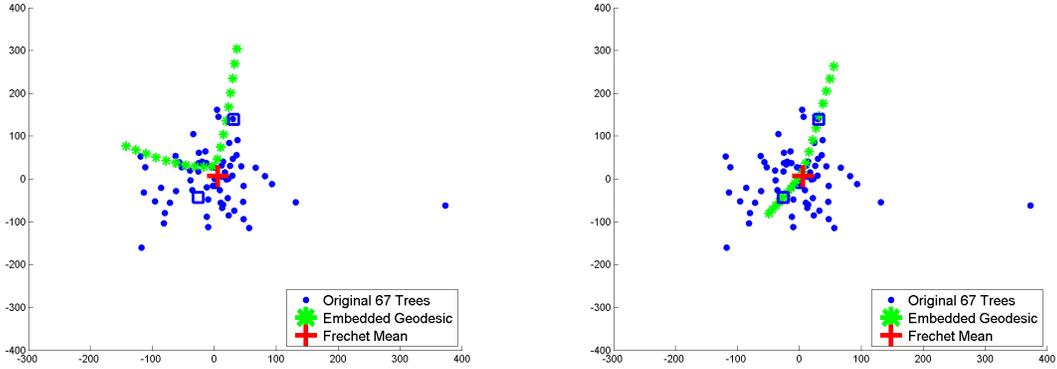
(d) 2-D MDS plot with maximum length geodesic embedded as an out-of-sample object

Figure 3.21: Examples of embedding geodesics as out-of-sample objects in 2-D MDS plots

figure, plot (a) shows the MDS representation of the 67 original brain artery trees (blue dots) along with their Fréchet mean (red cross). This plot is the same as Figure 3.19(a) except in different scale. Three different types of markers (circles, squares, and diamonds) represent the end points of three out-of-sample

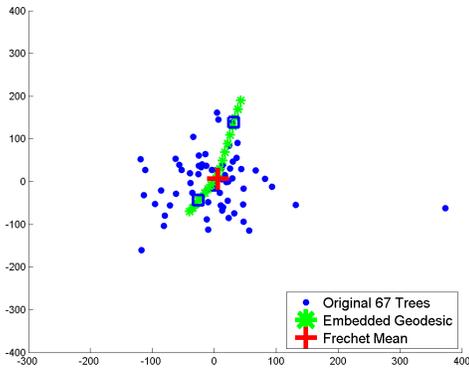
geodesics (minimum, median, and maximum length geodesic respectively). Plots (b), (c), and (d) illustrate the out-of-sample embedding of these geodesics into plot (a), and the embedded geodesics are displayed as green stars. Notice that the MDS configuration of the 67 original data points and their Fréchet mean stays unchanged in all three out-of-sample representations, which is the major advantage of Trosset and Priebe’s out-of-sample extension of classical MDS in this type of visualization. However, one aspect of the above out-of-sample embedding which does not make a lot of sense is that the embedded geodesics do not actually connect their end points, as clearly shown in Figure 3.21(b)(c)(d). One possible explanation is: when the out-of-sample extension tries to minimize the loss function (3.6), the coordinate matrix Y of the embedded grid points is naturally determined by the majority of the original data points rather than just the end points. This issue motivates our research topic of *high-dimensional out-of-sample (HDOS) embedding* for visualization purposes, i.e. we first solve the out-of-sample embedding problem with embedding dimension $p > 2$ and then choose the first two MDS coordinates to display in \mathbb{R}^2 . It is expected that the distortion in the first two MDS dimensions can be significantly reduced by allowing $p - 2$ extra variables in minimizing the loss function (3.6). For our brain artery data, we applied the HDOS technique to the embedding of the median length geodesic with the embedding dimensions $p = 3, 5, 10, 66$. Figure 3.22 shows the results of this HDOS embedding. We use the same marker types and colors as in Figure 3.21(c). First, as in Figure 3.21, since the embedded geodesic is treated as an out-of-sample object, the MDS configuration of the original trees along with their Fréchet mean stays the same for all 4 plots in the above figure. Second, the embedded geodesic systematically gets closer to connecting the two end points as d increases, especially when $p = 66$, the green stars connect with the two blue squares almost perfectly. This important observation confirmed our intuition that the distortion in the first two MDS dimensions can be significantly reduced by allowing the extra $p - 2$ dimensions when solving the out-of-sample embedding problem. However, this is not true in classical MDS: the first 2-dimensional MDS representation always stays the same no matter how large the embedding dimension is. This reveals an important property of the out-of-sample extension: although both classical MDS and the out-of-sample extension approximate an inner product matrix, in terms of visualization, the 2-dimensional out-of-sample representation significantly benefits from the above high-dimensional approach, but the classical MDS representation is not affected at all. This has not been pointed out explicitly in [Trosset and Priebe, 2008].

To quantify the benefit of HDOS in visualizing the out-of-sample embedding, we want to study the distortion of HDOS representation. Although the out-of-sample problem is solved by minimizing the distortion of the inner product matrix, we choose a more direct measure of distortion for the resulting HDOS representation. We define the above distortion to be the difference between the original augmented dissimilarity matrix and the distance matrix generated from the p -dimensional out-of-sample representation. Specifically, we will

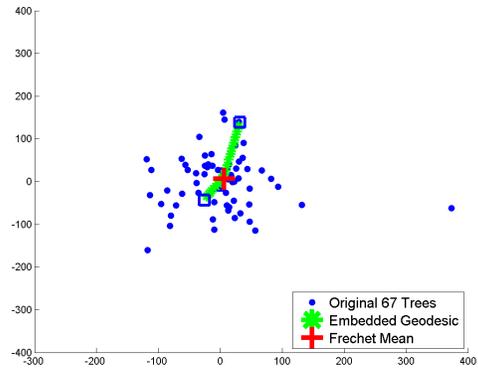


(a) HDOS embedding of median length geodesic with $p = 3$

(b) HDOS embedding of median length geodesic with $p = 5$



(c) HDOS embedding of median length geodesic with $p = 10$



(d) HDOS embedding of median length geodesic with $p = 66$

Figure 3.22: HDOS embedding of median length geodesic with embedding dimension $p = 3, 5, 10, 66$

compute Stress from equation (3.5) for $p = 2, 3, 5, 10, 66$, where $D_{i,j}$ is the original dissimilarity between objects i and j , and \mathbf{x}_i is the vector of MDS coordinates of object i in \mathbb{R}^2 . Table 3.1 shows the 2-dimensional Stress (i.e. only the first two MDS coordinates) of those five cases. We can see that the 2-dimensional Stress becomes larger when p increases. The reason is that the HDOS technique tries to minimize the loss function (3.6) in p -dimensional Euclidean space instead of just the 2-dimensional Euclidean space for visualization. More specifically, the contribution of the first two dimensions becomes less dominant as p increases, which leads to the ascending order of 2-dimensional Stresses shown in the table.

A reasonable alternative measure of out-of-sample distortion is the Stress of the p -dimensional out-of-sample representations (i.e. all p MDS coordinates). This time \mathbf{x}_i in equation (3.5) is the MDS coordinates of object i in \mathbb{R}^p . Table 3.2 shows the p -dimensional Stress of the same five cases. The value of the p -dimensional Stress decreases as p gets larger, because more flexibility is allowed in minimizing the loss function (3.6). Especially, the 66-dimensional Stress is only about 5% of the 2-dimensional Stress. This observation explains

why the HDOS representation produced better visualization as p increased.

Embedding dimension p	2	3	5	10	66
2-dimensional Stress/ 10^3	23.8	24.6	24.8	25.3	25.7

Table 3.1: Row 1 gives the values for the embedding dimension p , and Row 2 shows the Stress of the 2-dimensional out-of-sample representation for each p . The Stress gets slightly larger as p increases, which indicates that the Stress of 2-dimensional representation is not a good measure of HDOS distortion.

Embedding dimension p	2	3	5	10	66
p -dimensional Stress/ 10^3	23.8	22.3	19.8	15.7	1.32

Table 3.2: Row 1 gives the values for the embedding dimension p , and Row 2 shows the Stress of the p -dimensional out-of-sample representation. The Stress decreases as p increases, which is compatible with what we saw in Figure 3.22.

Table 3.2 summarized the overall distortion between the original dissimilarity matrix and the Euclidean distance matrix produced from the out-of-sample MDS representations. Furthermore, it will be more helpful to examine the distortion for all 87×87 dissimilarities when embedding the median length geodesic, which will be displayed by using *heat maps*. A heat map is a graphical representation of data where the individual values contained in a matrix are represented as colors. Let D be the original 87×87 dissimilarity matrix, and let D_p^E be the Euclidean distance matrix constructed from the p -dimensional out-of-sample MDS representation of all 87 objects, then we construct a heat map for each distortion matrix $D_p^\Delta = D_p^E - D$, where $p = 2, 3, 5, 10, 66$. The color scheme is: *white* represents zero value; *red* represents negative value where $D > D_p^E$ and negative values with larger absolute value are colored darker red; *blue* is used in the corresponding way for positive value where $D < D_p^E$. Figure 3.23 demonstrates the color scheme and the heat map for D_2^Δ . Figure(a)

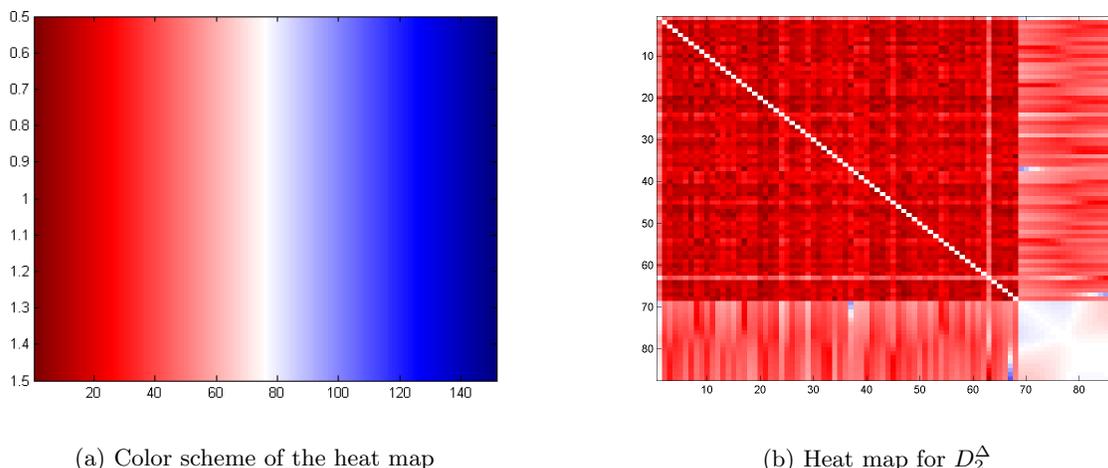


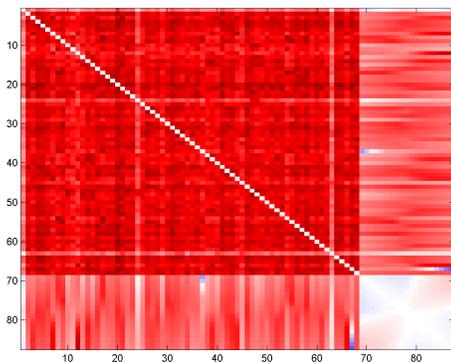
Figure 3.23: HDOS embedding of median length geodesic with embedding dimension $p = 3, 5, 10, 66$

displays the entire color range representing values from -600 (most negative) to 600 (most positive). Figure(b) demonstrates the heat map for the distortion matrix D_2^Δ . First, this heat map is symmetric and all diagonal elements are white (zero). Second, the two end points of the median length geodesic are labeled as 37 and 67 in the heat map. Third, the entire heat map can be divided into four regions and each region has its own character reflecting the corresponding portion of the distortion matrix:

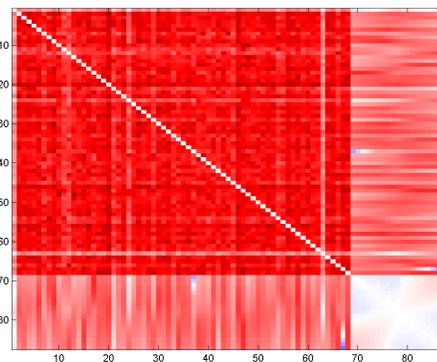
- Region 1: first row and first column represent the distortion of the dissimilarities between the Fréchet mean and all 87 objects. Region 1 is colored light red at the beginning, which indicates that 2-dimensional classical MDS tends to underestimate the dissimilarities between 67 original brain artery trees and their Fréchet mean. This is analogous to the PCA in Euclidean space, where if we project data onto the plane spanned by PC1 and PC2, the projected distances between the Fréchet mean and data points are smaller than the original distances. However, Region 1 is colored nearly white at the end, indicating that 2-dimensional out-of-sample MDS maintains the dissimilarities between the Fréchet mean and the out-of-sample objects quite well. This is because that the 2-dimensional out-of-sample MDS configuration is dominated by the embedded geodesic and the Fréchet mean happens to be the origin.
- Region 2: the square region consisting of rows 2 to 68 and columns 2 to 68 represents the distortion of the dissimilarities within 67 original brain artery trees. Region 2 except row and column labeled 63 is colored dark red, which indicates that 2-dimensional classical MDS tends to underestimate the dissimilarities among the original 67 brain artery trees. This underestimation can be explained by *Parseval identity*, and it is heavier than the underestimation in region 1 because the dissimilarity between two original trees is close to twice that between the Fréchet mean and one original tree. Row 63 (or column 63) corresponds to the the point on the far right in Figure 3.21(c), and the distortions of dissimilarities between this point and other original objects are much smaller than the distortions of those among other original objects. This can be explained by the fact that this point has a quite large norm compared to other original objects and the first MDS component direction is pulled towards this point.
- Region 3: the rectangle region containing rows 2 to 68 and columns 69 to 87 and its diagonal-symmetric part represent the distortion of the dissimilarities between 67 original brain artery trees and 19 embedded grid points from the median length geodesic. Region 3 other than the rows and columns corresponding to two end trees (labeled 37 and 67) are in light red, which indicates that 2-dimensional out-of-sample MDS also tends to underestimate the dissimilarities between the embedded objects and the original 67 brain artery trees. This is again closely related to Parseval identity. The rows and

columns corresponding to two end trees have some blues between each end tree and the embedded geodesic near the same end, which is consistent with Figure 3.21(c) since the embedded geodesic does not connect the two end points, thus creating large positive distortion.

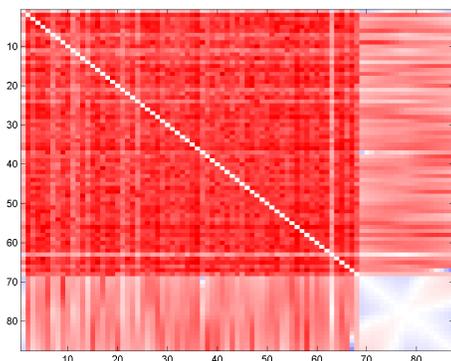
- Region 4: the square region at the lower right corner represents the distortion of the dissimilarities within 19 embedded points. Region 4 is colored either very light blue or very light red, which indicates that 2-dimensional out-of-sample MDS keeps the dissimilarities among the out-of-sample objects quite well. This is again due to the fact that the 2-dimensional out-of-sample MDS configuration is dominated by the embedded geodesic.



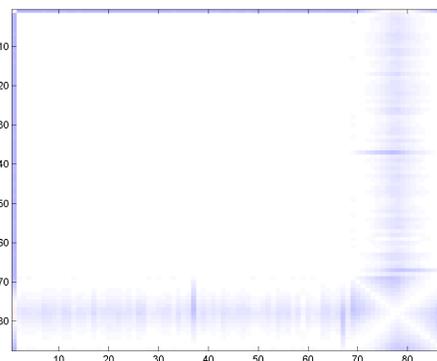
(a) HDOS embedding of median length geodesic with $p = 3$



(b) HDOS embedding of median length geodesic with $p = 5$



(c) HDOS embedding of median length geodesic with $p = 10$



(d) HDOS embedding of median length geodesic with $p = 66$

Figure 3.24: HDOS embedding of median length geodesic with embedding dimension $p = 3, 5, 10, 66$

To examine the effect of embedding dimension p on the distortion matrix, we will show heat maps for D_p^Δ , where $p = 3, 5, 10, 66$ in Figure 3.24. From the above plots, we can see that the red region systematically becomes lighter as p increases, which indicates that the distances become larger and closer to the tree space

distances in higher dimensional out-of-sample MDS space. This is consistent with Parseval identity. Especially when $p = 66$, red region completely disappears, meaning that the underestimation of the dissimilarities has been compensated by going into high enough embedding dimension. Note in plot (d), region 2 is white and this indicates that the distortion of dissimilarities within 67 brain artery trees is zero, further this means that the dissimilarity sub-matrix corresponding to region 2 is a Euclidean distance matrix. However, the dissimilarity sub-matrix corresponding to region 1 and 2, i.e. including the Fréchet mean, is not a Euclidean distance matrix any more.

As one approach of PCA in tree space, MDS is only based on the distance matrix of data. Since the actual locations of data in tree space are known, we will develop two line-based tree space PCA approaches in the rest of the dissertation. In the next section, we first define and characterize the notions of line in tree space.

CHAPTER 4: NOTIONS OF “LINE” IN TREE SPACE

Recall in Euclidean PCA, the principal component directions are defined as one dimensional lines in \mathbb{R}^n . However, there is no such Euclidean line in tree space since the extensions of a line is not well-defined outside the orthant in which it lies. In this chapter, notions of *tree-space lines* will be defined and characterized. We will first define what a general tree-space line is.

4.1 Tree-space Line

Definition 4.1.1. A connected set $L \subset \mathcal{T}_n$ is a tree-space line if the following three conditions are satisfied

- (C1) L is locally one dimensional, i.e. $L \cap \mathcal{O}$ is a one dimensional object for any orthant \mathcal{O} .
- (C2) L is closed under geodesics, i.e. $\Gamma(T^1, T^2) \subset L$ for each pair of trees $T^1, T^2 \in L$.
- (C3) L is maximal in the sense that there is no such set $L' \supset L$ satisfying properties **C1** and **C2**.

We use $\mathfrak{L}(n)$ to denote the set of all tree-space lines in \mathcal{T}_n .

This definition immediately gives the following property of a tree-space line.

Lemma 4.1.1. *A tree-space line $L \in \mathfrak{L}(n)$ has the property that for any two trees T^1, T^2 in L the only path between T^1 and T^2 in L is the (unique) geodesic $\Gamma(T^1, T^2)$ between them.*

Proof. From **C1** and **C2**, we know $L \cap \mathcal{O}$ must be either a line segment or a ray for any orthant \mathcal{O} that L intersects. Now suppose there is another path $P(T^1, T^2) \subset L$ between T^1 and T^2 other than $\Gamma(T^1, T^2)$. Since $P(T^1, T^2)$ is part of L , it must consist of a sequence of line segments. By the uniqueness of geodesic in tree space, $\Gamma(T^1, T^2) \neq P(T^1, T^2)$. Hence, we define $T = T^*$ to be the last point T along $P(T^1, T^2)$ such that $\Gamma(T^1, T) = P(T^1, T)$, as illustrated in Figure 4.25. Since every T on the first leg of $P(T^1, T^2)$ satisfies $\Gamma(T^1, T) = P(T^1, T)$, it must be that $T^* \neq T^1$. Now select $\tilde{T} \in P(T^*, T^2)$ close enough to T^* , so \tilde{T} is in the interior of the vial cell \mathcal{V} with respect to T^1 that contains T^* . Since $\Gamma(T^1, \tilde{T}) \neq P(T^1, \tilde{T})$, we must have $\Gamma(T^1, T^*) \not\subset \Gamma(T^1, \tilde{T})$. (Otherwise, $\Gamma(T^1, \tilde{T}) = P(T^1, \tilde{T})$.) Since $P(T^1, T^2) \cap \mathcal{V}$ must be a line segment, then $\Gamma(T^1, \tilde{T})$ can not go through T^* . Instead $\Gamma(T^1, \tilde{T})$ must go into the opposite direction along that line segment, passing through point $T' \in \Gamma(T^1, \tilde{T}) \cap \mathcal{V}$. Since T^* and \tilde{T} are in the same vial cell, $\Gamma(T^1, T^*)$ and $\Gamma(T^1, \tilde{T})$ must have the same support pair. Thus, $\Gamma(T^1, T^*)$ must also pass T' . However, we already assume $\Gamma(T^1, T^*) = P(T^1, T^*)$, which does not pass T' . Hence, it is a contradiction that there are two distinct geodesic between T^1 and T^* . Therefore, $\Gamma(T^1, T^2)$ must be the only path between T^1 and T^2 . \square

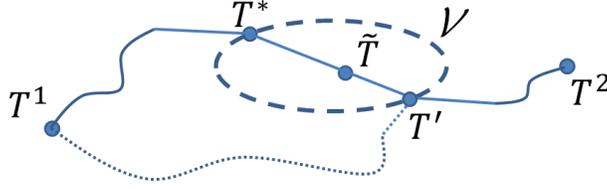


Figure 4.25: The solid path denotes $P(T^1, T^2)$ and $P(T^1, T^2) \neq \Gamma(T^1, T^2)$. The dashed ellipse represents the vial cell \mathcal{V} containing T^* , \tilde{T} and T' . The dotted path between T^1 and T' together with the solid segment between T' and \tilde{T} is $\Gamma(T^1, \tilde{T})$. We reach a contradiction that there are two distinct geodesics between T^1 and T^* .

In this chapter, $\mathfrak{L}(n)$ will be divided into two mutually exclusive types of tree-space lines, and both types will be completely characterized. Recall from Section 1.3, a geodesic is piecewise linear, and each linear portion of a geodesic is called a *leg*.

Definition 4.1.2. $L \in \mathfrak{L}(n)$ is a *type I tree-space line* if L contains a 3-leg geodesic, and $L \in \mathfrak{L}(n)$ is a *type II tree-space line* if L contains no 3-leg geodesic, i.e. all geodesics have one or two legs.

The concepts of ray and inner end point are useful when describing both types of tree-space lines, and the following definitions formally define a ray and inner end points in the tree space.

Definition 4.1.3. Assume $P, T \in \mathcal{T}_n$ are in the same orthant, and all edge lengths are non-decreasing from P to T along $\Gamma(P, T)$, then the *ray starting from P going through T* is denoted as $\vec{\mathbf{R}}(P, T) = \{P + \lambda \cdot \overrightarrow{(T - P)} \mid \lambda \geq 0\}$, where $\overrightarrow{(T - P)}$ denotes the usual Euclidean vector pointing from P to T .

Definition 4.1.4. Given a geodesic $\Gamma(T^1, T^2)$ with at least two legs, define the *inner end point* of $\Gamma(T^1, T^2)$ with respect to T^1 (or T^2) to be the opposite end of the leg adjacent to T^1 (or T^2).

In Figure 4.26, one example of each type of tree-space line is given in \mathcal{T}_4 . In Figure (a) the geodesic between $T^1 = \{|e_1| = 2, |e_2| = 1\}$ and $T^2 = \{|e_4| = 2, |e_5| = 1.5\}$ has inner end points $P = \{|e_1| = 0.4\}$ and $Q = \{|e_5| = 0.25\}$, and L_1 consists of $\Gamma(T^1, T^2)$ together with the rays $\vec{\mathbf{R}}(P, T^1)$ and $\vec{\mathbf{R}}(Q, T^2)$. L_2 in Figure (b) is comprised of three rays $\vec{\mathbf{R}}(S, T^3)$, $\vec{\mathbf{R}}(S, T^4)$, and $\vec{\mathbf{R}}(S, T^5)$, with $S = \{|e_6| = 1\}$, $T^3 = \{|e_6| = 1, |e_7| = 2\}$, $T^4 = \{|e_6| = 1, |e_8| = 2\}$, and $T^5 = \{|e_6| = 1, |e_9| = 2\}$. It is straightforward to see that L_1 contains a 3-leg geodesic $\Gamma(T^1, T^2)$ and L_2 contains no 3-leg geodesic. They will be shown to be tree-space lines in the following two sections. In Section 4.2, we will focus on constructing and characterizing type I tree-space lines. In Section 4.3, we will develop type II tree-space lines. From now on, “line” will be used to mean “tree-space line”.

4.2 Type I Lines

Since a geodesic in tree space satisfies **C1** and **C2**, then it could be a candidate for a line, as shown in Figure 4.26(a). However, the geodesic between two arbitrary trees is not maximal as defined in **C3** for two

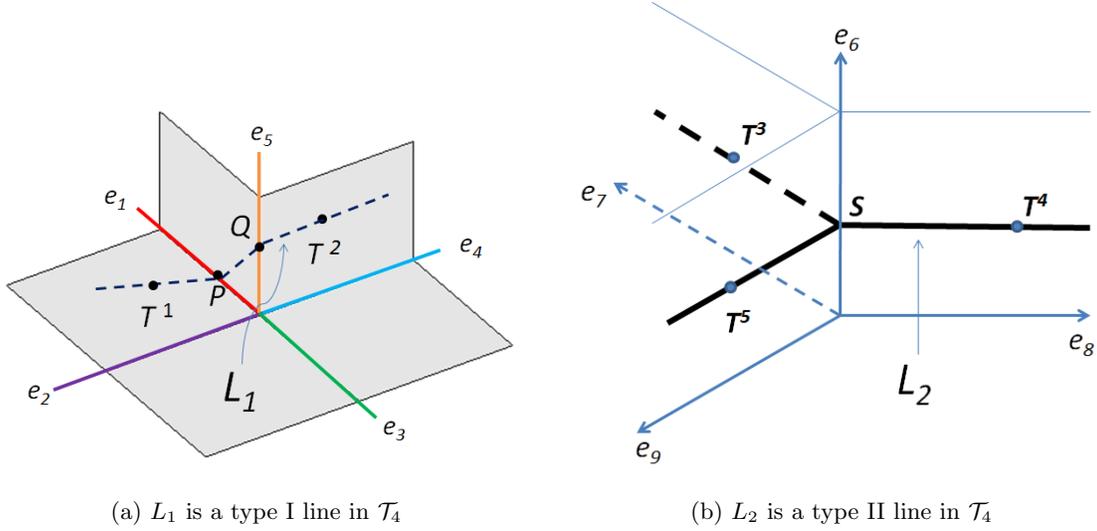


Figure 4.26: Examples of lines in \mathcal{T}_4

reasons. First, a line must be unbounded, thus a geodesic needs to be extended beyond its two end trees. Second, a line can bifurcate, that is, branch into multiple orthants when it hits an orthant boundary, as shown in Figure 4.26(b). In this section, we will see that some geodesics, which can be naturally extended into type I lines, avoid bifurcation both in the extensions and in the middle of a geodesic.

Definition 4.2.1. A geodesic $\Gamma(T^1, T^2)$ is *uniquely extendable* if each edge common to T^1 and T^2 has the same length in both trees, where common edges include edges with zero length as long as they are compatible with both trees. (Note that no 1-leg geodesic is uniquely extendable.)

Before characterizing a uniquely extendable geodesic, we would like to use a simple example again in \mathcal{T}_4 to present some intuition. Consider Figure 4.27 with the same three orthants shown in Figure 4.26(a). Thus $T^1 = \{|e_1| = 2, |e_2| = 1\}$ and $T^2 = \{|e_4| = 2, |e_5| = 1.5\}$ are the same as in Figure 4.26(a). T^6 is chosen to be $T^6 = \{|e_1| = 2, |e_5| = 0.3\}$. Then T^1 and T^2 have no common edges, and thus $\Gamma(T^1, T^2)$ is uniquely extendable. We will show later that $\Gamma(T^1, T^2)$ can be extended into a type I line. T^1 and T^6 share a common edge e_1 , but $|e_1| = 2$ for both trees, hence $\Gamma(T^1, T^6)$ is uniquely extendable. Finally T^2 and T^6 share a common edge e_5 , but $|e_5|_{T^2} \neq |e_5|_{T^3}$. Therefore, $\Gamma(T^2, T^6)$ is not uniquely extendable and the extension beyond T^6 hits the orthant boundary at T^7 having the single edge e_1 with $|e_1| = 3$. At this point, $\Gamma(T^2, T^6)$ has two possible extensions.

Lemma 4.2.1. A geodesic $\Gamma(T^1, T^2)$ is uniquely extendable if and only if neither $\vec{\mathbf{R}}(P, T^1)$ nor $\vec{\mathbf{R}}(Q, T^2)$ crosses any orthant boundary, where P and Q are the inner end points with respect to T^1 and T^2 .

Proof. “ \Rightarrow ” Suppose that $\Gamma(T^1, T^2)$ is uniquely extendable. We show that $\vec{\mathbf{R}}(Q, T^2)$ will not cross any orthant boundary, the case for $\vec{\mathbf{R}}(P, T^1)$ being symmetric. We need to prove that no edge in T^2 vanishes

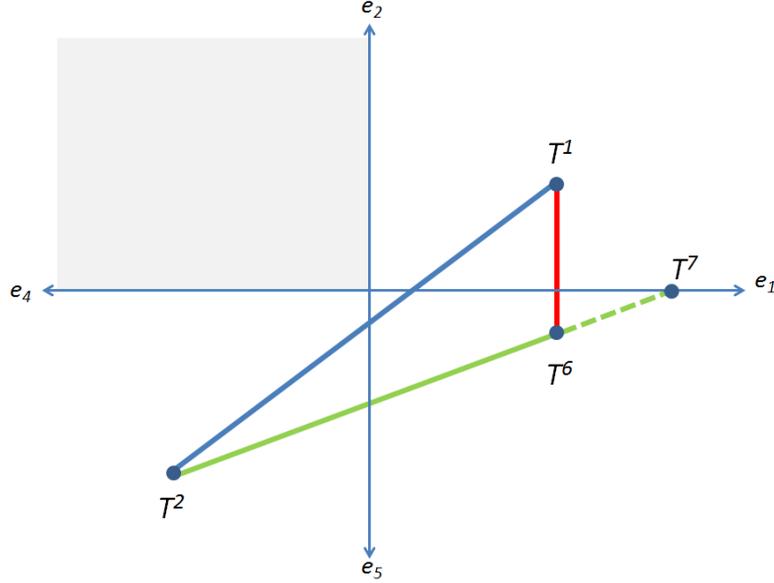


Figure 4.27: The above figure shows three geodesics in a portion of \mathcal{T}_4 : $\Gamma(T^1, T^2)$ is uniquely extendable and can be extended into a type I line; $\Gamma(T^1, T^6)$ is also uniquely extendable but can not be extended into a type I line; $\Gamma(T^2, T^6)$ is not even uniquely extendable.

along $\vec{\mathbf{R}}(Q, T^2)$. The common edges between T^1 and T^2 have the same edge lengths in both two end trees, and hence they will have the same constant edge lengths from Q to T^2 , and by the definition of a ray, they also have the same constant edge lengths in $\vec{\mathbf{R}}(Q, T^2)$. For the edges only in T^2 , each of them will increase from zero length at some point on the boundary of an orthant containing that edge, hence their lengths will only increase from Q to T^2 , and thus continue increasing along $\vec{\mathbf{R}}(Q, T^2)$. Therefore, $\vec{\mathbf{R}}(Q, T^2)$ will not cross any orthant boundary.

“ \Leftarrow ” Suppose neither $\vec{\mathbf{R}}(P, T^1)$ nor $\vec{\mathbf{R}}(Q, T^2)$ crosses any orthant boundary. Then all edge lengths in T^1 or T^2 will be nondecreasing along $\vec{\mathbf{R}}(P, T^1)$ and $\vec{\mathbf{R}}(Q, T^2)$ respectively. Suppose e is a common edge between T^1 and T^2 . If $|e|_{T^1} \neq |e|_{T^2}$, say $|e|_{T^1} > |e|_{T^2}$, then it must be that e is decreasing along $\vec{\mathbf{R}}(Q, T^2)$, which is a contradiction. We proved that each common edge between T^1 and T^2 have the same length in both, hence $\Gamma(T^1, T^2)$ is uniquely extendable. \square

As the name suggests, a uniquely extendable geodesic has unique extensions beyond its end points. In particular, we have

Definition 4.2.2. Let $\Gamma(T^1, T^2)$ be uniquely extendable, and let P and Q be the inner end points of the first and last leg of $\Gamma(T^1, T^2)$. Then the *extension* of $\Gamma(T^1, T^2)$ is denoted as

$$\Gamma_{Ext}(T^1, T^2) = \Gamma(T^1, T^2) \cup \vec{\mathbf{R}}(Q, T^2) \cup \vec{\mathbf{R}}(P, T^1)$$

Figure 4.28 gives examples of extensions of two uniquely extendable geodesics in Figure 4.27. The dashed lines represent the extensions of $\Gamma(T^1, T^2)$, and $\Gamma_{Ext}(T^1, T^2) = \Gamma(T^1, T^2) \cup \vec{\mathbf{R}}(Q, T^2) \cup \vec{\mathbf{R}}(P, T^1)$. The dotted lines represent the extensions of $\Gamma(T^1, T^3)$ and $\Gamma(T^1, T^2)$, and $\Gamma_{Ext}(T^1, T^3) = \Gamma(T^1, T^3) \cup \vec{\mathbf{R}}(S, T^3) \cup \vec{\mathbf{R}}(S, T^1)$.

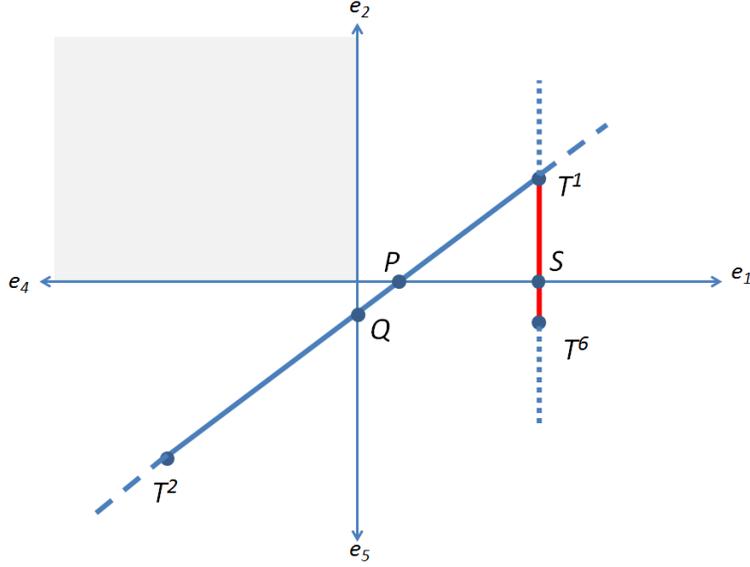


Figure 4.28: The above figure shows extensions of $\Gamma(T^1, T^2)$ and $\Gamma(T^2, T^3)$ which are uniquely extendable geodesics for the three trees in Figure 4.27.

Lemma 4.2.2. *If $\Gamma(T^1, T^2)$ is uniquely extendable, then $\Gamma_{Ext}(T^1, T^2)$ satisfies **C2**.*

Proof. To prove that $\Gamma_{Ext}(T^1, T^2)$ is closed under geodesics, let X and Y be two points on $\Gamma_{Ext}(T^1, T^2)$. Assume that $X \in \vec{\mathbf{R}}(P, T^1)$ and $Y \in \vec{\mathbf{R}}(Q, T^2)$, where P and Q are the inner end points of the first and last leg of $\Gamma(T^1, T^2)$, the other cases being similar. We need to show that $\Gamma(X, Y) = \Gamma(T^1, T^2) \cup \overline{(T^1, X)} \cup \overline{(T^2, Y)}$, where $\overline{(T^1, X)}$ denotes the line segment connecting T^1 and X , $\overline{(T^2, Y)}$ denotes the line segment connecting T^2 and Y .

From Theorem 1.3.1, $\Gamma(T^1, T^2)$ satisfies the following conditions:

- (a) $\Gamma(T^1, T^2)$ is contained entirely in the path space defined by the support $(\mathcal{A}, \mathcal{B})$.
- (b) **(P1)**: For each $i > j$, A_i and B_j are compatible.
- (c) **(P2)**: $\frac{\|A_1\|_{T^1}}{\|B_1\|_{T^2}} \leq \frac{\|A_2\|_{T^1}}{\|B_2\|_{T^2}} \leq \dots \leq \frac{\|A_k\|_{T^1}}{\|B_k\|_{T^2}}$.
- (d) **(P3)**: For each support pair (A_i, B_i) , there is no nontrivial partition $C_1 \cup C_2$ of A_i , and partition $D_1 \cup D_2$ of B_i , such that C_2 is compatible with D_1 and $\frac{\|C_1\|_{T^1}}{\|D_1\|_{T^2}} < \frac{\|C_2\|_{T^1}}{\|D_2\|_{T^2}}$.

(e) $\Gamma(T^1, T^2) = (\gamma(\lambda) : 0 \leq \lambda \leq 1)$ can be represented in \mathcal{T}_n with legs

$$\Gamma^i = \begin{cases} \left[\gamma(\lambda) : \frac{\lambda}{1-\lambda} \leq \frac{\|A_1\|_{T^1}}{\|B_1\|_{T^2}} \right], & i = 0 \\ \left[\gamma(\lambda) : \frac{\|A_i\|_{T^1}}{\|B_i\|_{T^2}} \leq \frac{\lambda}{1-\lambda} \leq \frac{\|A_{i+1}\|_{T^1}}{\|B_{i+1}\|_{T^2}} \right], & i = 1, \dots, k-1, \\ \left[\gamma(\lambda) : \frac{\lambda}{1-\lambda} \geq \frac{\|A_k\|_{T^1}}{\|B_k\|_{T^2}} \right], & i = k \end{cases}$$

(f) The length of $\Gamma(T^1, T^2)$ is

$$L(\Gamma(T^1, T^2)) = \left\| \left(\|A_1\|_{T^1} + \|B_1\|_{T^2}, \dots, \|A_k\|_{T^1} + \|B_k\|_{T^2} \right) \right\|$$

Now by the choice of the end trees X and Y , they will have the same edge sets as in T^1 and T^2 respectively. First, we want to show that the support $(\mathcal{A}, \mathcal{B})$ also satisfies **(P1)**, **(P2)**, and **(P3)**.

- It is straightforward that **(P1)** holds since the edge set remains the same between each end tree and its respective inner end tree.
- Let λ_i be the proportion of geodesic traversed from T^1 to T^2 when the edges in A_i contracted and edges in B_i start to grow. We use d_1 to denote the geodesic distance between X and T^1 , d_2 to denote the geodesic distance between Y and T^2 , L to denote $L(\Gamma(T^1, T^2))$. From the above (e), we have

$$\lambda_i = \frac{\|A_i\|_{T^1}}{\|A_i\|_{T^1} + \|B_i\|_{T^2}}$$

By the geometry of $\Gamma_{Ext}(T^1, T^2)$, we also have

$$\begin{aligned} \|A_i\|_X &= \left(1 + \frac{d_1}{\lambda_i L} \right) \|A_i\|_{T^1} = \|A_i\|_{T^1} + \frac{d_1}{L} (\|A_i\|_{T^1} + \|B_i\|_{T^2}) \\ \|B_i\|_Y &= \left(1 + \frac{d_2}{(1-\lambda_i)L} \right) \|B_i\|_{T^2} = \|B_i\|_{T^2} + \frac{d_2}{L} (\|A_i\|_{T^1} + \|B_i\|_{T^2}) \end{aligned}$$

Now we want to show that

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

After some algebra, we have the following result

$$\|A_i\|_X \cdot \|B_{i+1}\|_Y - \|A_{i+1}\|_X \cdot \|B_i\|_Y = \left(1 + \frac{d_1}{L} + \frac{d_2}{L} \right) (\|A_i\|_{T^1} \cdot \|B_{i+1}\|_{T^2} - \|A_{i+1}\|_{T^1} \cdot \|B_i\|_{T^2})$$

From part (c) above, we know $\|A_i\|_{T^1} \cdot \|B_{i+1}\|_{T^2} - \|A_{i+1}\|_{T^1} \cdot \|B_i\|_{T^2} \leq 0$, hence $\|A_i\|_X \cdot \|B_{i+1}\|_Y - \|A_{i+1}\| \cdot \|B_i\|_Y \leq 0$, i.e. $(\mathcal{A}, \mathcal{B})$ satisfies **(P2)**.

- Suppose $(\mathcal{A}, \mathcal{B})$ does not satisfy **(P3)** for X and Y , i.e. there exists a support pair (A_i, B_i) , and there is a nontrivial partition $C_1 \cup C_2$ of A_i , and a nontrivial partition $D_1 \cup D_2$ of B_i , such that C_2 is compatible with D_1 and $\frac{\|C_1\|_X}{\|D_1\|_Y} < \frac{\|C_2\|_X}{\|D_2\|_Y}$.

However, by the geometry of $\Gamma_{Ext}(T^1, T^2)$, we can get similar results for every single edge as in the proof of **(P2)**

$$\begin{aligned} \|e\|_X &= \left(1 + \frac{d_1}{\lambda_i L}\right) \|e\|_{T^1} && \text{for each } e \in A_i \\ \|e\|_Y &= \left(1 + \frac{d_2}{(1 - \lambda_i)L}\right) \|e\|_{T^2} && \text{for each } e \in B_i \end{aligned}$$

From this it is straightforward that

$$\begin{aligned} \|C_1\|_X &= \left(1 + \frac{d_1}{\lambda_i L}\right) \|C_1\|_{T^1} && \|C_2\|_X = \left(1 + \frac{d_1}{\lambda_i L}\right) \|C_2\|_{T^1} \\ \|D_1\|_Y &= \left(1 + \frac{d_2}{(1 - \lambda_i)L}\right) \|D_1\|_{T^2} && \|D_2\|_Y = \left(1 + \frac{d_2}{(1 - \lambda_i)L}\right) \|D_2\|_{T^2} \end{aligned}$$

As an immediate result of the above equations and the assumption that $\frac{\|C_1\|_X}{\|D_1\|_Y} < \frac{\|C_2\|_X}{\|D_2\|_Y}$, we know $\frac{\|C_1\|_{T^1}}{\|D_1\|_{T^2}} < \frac{\|C_2\|_{T^1}}{\|D_2\|_{T^2}}$ also holds. However, this is a contradiction to the fact that $\Gamma(T^1, T^2)$ satisfies **(P3)**. Therefore, $(\mathcal{A}, \mathcal{B})$ must satisfy **(P3)** for X and Y .

Now we know that support $(\mathcal{A}, \mathcal{B})$ satisfies **(P1)**, **(P2)**, and **(P3)**, hence determines a geodesic $\Gamma(X, Y)$. Next we need to show that $\Gamma(X, Y) = \Gamma(T^1, T^2) \cup \overline{(T^1, X)} \cup \overline{(T^2, Y)}$. From the above proof of **(P2)**, we have

$$\begin{aligned} \lambda_i &= \frac{\|A_i\|_{T^1}}{\|A_i\|_{T^1} + \|B_i\|_{T^2}} \\ \|A_i\|_X &= \left(1 + \frac{d_1}{\lambda_i L}\right) \|A_i\|_{T^1} \\ \|B_i\|_Y &= \left(1 + \frac{d_2}{(1 - \lambda_i)L}\right) \|B_i\|_{T^2} \end{aligned}$$

then after some algebra, we get

$$\|A_i\|_X + \|B_i\|_Y = \left(\frac{L + d_1 + d_2}{L}\right) (\|A_i\|_{T^1} + \|B_i\|_{T^2})$$

From the above part (f), we can compute the length of $\Gamma(X, Y)$ as

$$L(\Gamma(X, Y)) = \left\| \left(\|A_1\|_X + \|B_1\|_Y, \dots, \|A_k\|_X + \|B_k\|_Y \right) \right\| = \left(\frac{L + d_1 + d_2}{L} \right) \cdot L(\Gamma(T^1, T^2)) = L + d_1 + d_2$$

which is exactly the length of $\Gamma(T^1, T^2) \cup \overline{(T^1, X)} \cup \overline{(T^2, Y)}$. By the uniqueness of geodesic in the tree space, we have proven that $\Gamma(X, Y) = \Gamma(T^1, T^2) \cup \overline{(T^1, X)} \cup \overline{(T^2, Y)}$. Therefore, $\Gamma_{Ext}(T^1, T^2)$ is closed under geodesics. The same argument holds when $X \in \vec{\mathbf{R}}(P, T^1)$, $Y \in \Gamma(T^1, T^2)$ and $Y \in \vec{\mathbf{R}}(Q, T^2)$, $X \in \Gamma(T^1, T^2)$. The case $X, Y \in \Gamma(T^1, T^2)$ is trivial. \square

Now we are ready to introduce the type I line as a special case of extended geodesic.

Theorem 4.2.3. *If $\Gamma(T^1, T^2)$ is uniquely extendable, and has at least 3 legs, then $\Gamma_{Ext}(T^1, T^2)$ is a type I line.*

Proof. Let $\Gamma_{Ext}(T^1, T^2) = \Gamma(T^1, T^2) \cup \vec{\mathbf{R}}(Q, T^2) \cup \vec{\mathbf{R}}(P, T^1)$, where P and Q are the inner end points of the first and last leg of $\Gamma(T^1, T^2)$. From Definition 4.2.2, we know $\Gamma_{Ext}(T^1, T^2)$ satisfies **C1**. From Lemma 4.2.2, $\Gamma_{Ext}(T^1, T^2)$ also satisfies **C2**. Thus to prove that $\Gamma_{Ext}(T^1, T^2)$ is a line, we only need to prove that **C3** holds. Suppose not, that is, $\Gamma_{Ext}(T^1, T^2)$ is strictly contained in some line L which must then have a bifurcation point at some point S along $\Gamma_{Ext}(T^1, T^2)$. However, a bifurcation can not occur in the interior of any leg of $\Gamma(T^1, T^2)$, nor on $\vec{\mathbf{R}}(Q, T^2)$ or $\vec{\mathbf{R}}(P, T^1)$. Thus S must be at the intersection of two legs of $\Gamma(T^1, T^2)$.

Now let $(\mathcal{A}, \mathcal{B})$ be the support associated with $\Gamma(T^1, T^2)$, where $\mathcal{A} = (A_1, \dots, A_k)$ and $\mathcal{B} = (B_1, \dots, B_k)$, and let $\mathcal{O}_{i-1} = B_1 \cup \dots \cup B_{i-1} \cup A_i \cup \dots \cup A_k$ and $\mathcal{O}_i = B_1 \cup \dots \cup B_i \cup A_{i+1} \cup \dots \cup A_k$ be the two adjacent orthants which $\Gamma(T^1, T^2)$ traverses. By the above argument, we assume that S is where $\Gamma(T^1, T^2)$ leaves \mathcal{O}_{i-1} and enters \mathcal{O}_i . Let l be the second branch of the bifurcation, in orthant \mathcal{O}'_i obtained by replacing A_i by the set $B'_i \neq B_i$ but compatible with $\mathcal{O}_{i-1} \cap \mathcal{O}_i$. Select $\hat{T}^1 \in \Gamma(S, T^1)$, $\hat{T}^2 \in \Gamma(S, T^2)$, and $\hat{T}^l \in l$, all having equal distance ε away from S , where ε is small enough such that $\overline{(\hat{T}^1, S)} \cup \overline{(S, \hat{T}^2)}$, $\overline{(\hat{T}^1, S)} \cup \overline{(S, \hat{T}^l)}$, and $\overline{(\hat{T}^2, S)} \cup \overline{(S, \hat{T}^l)}$ are all 2-leg geodesics. See Figure 4.29 for an intuitive view.

Since $\Gamma(T^1, T^2)$ has at least three legs and \hat{T}^1, \hat{T}^2 are on two adjacent legs, then at least one of \hat{T}^1 and \hat{T}^2 must be on a leg of $\Gamma(T^1, T^2)$ which has neither T^1 nor T^2 as end points. Without loss of generality, assume \hat{T}^1 is on such a leg. Thus along this leg there must be at least one edge $e \in \mathcal{O}_{i-1} \cap \mathcal{O}_i$ which has $|e|_S \neq 0$ but goes to zero at the other end of the leg. It must follow that $|e|_S > |e|_{\hat{T}^1}$. Now by the choice of \hat{T}^1, \hat{T}^2 , and \hat{T}^l , e is a common edge for these three trees. Also \hat{T}^1, \hat{T}^2 , and \hat{T}^l are the same distance away from S . Along $\Gamma(\hat{T}^l, \hat{T}^1)$,

$$|e|_S = 0.5|e|_{\hat{T}^l} + 0.5|e|_{\hat{T}^1} \quad (4.7)$$

Likewise, along $\Gamma(\widehat{T}^1, \widehat{T}^2)$,

$$|e|_S = 0.5|e|_{\widehat{T}^1} + 0.5|e|_{\widehat{T}^2} \quad (4.8)$$

From equations (4.7) (4.8) and the fact that $|e|_S > |e|_{\widehat{T}^1}$, it follows that $|e|_{\widehat{T}^1} > |e|_S$ and $|e|_{\widehat{T}^2} > |e|_S$. But along $\Gamma(\widehat{T}^l, \widehat{T}^2)$,

$$|e|_S = 0.5|e|_{\widehat{T}^l} + 0.5|e|_{\widehat{T}^2}$$

which is a contradiction to the facts that $|e|_{\widehat{T}^l} > |e|_S$ and $|e|_{\widehat{T}^2} > |e|_S$. Therefore, $\Gamma_{Ext}(T^1, T^2)$ must satisfy **C3**, and hence is a type I line. \square

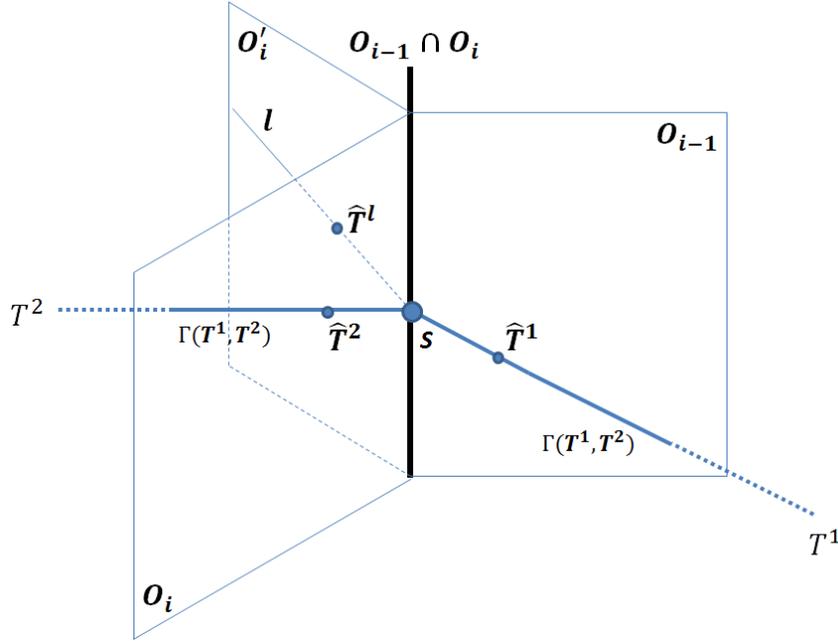


Figure 4.29: The above figure gives a detailed view of what happens when $\Gamma(T^1, T^2)$ crosses $O_{i-1} \cap O_i$ which is intuitively represented by the central vertical axis in bold.

The above theorem showed that a uniquely extendable geodesic with at least three legs can be extended into a type I line. To complete this section, we will also show that any type I line is an extension of some uniquely extendable geodesic.

Lemma 4.2.4. *If $L \in \mathfrak{L}(n)$ is a line and $\Gamma(T^1, T^2) \subset L$ is a geodesic with the maximum number of legs among all geodesics in L , then $\Gamma(T^1, T^2)$ is uniquely extendable.*

Proof. Suppose $\Gamma(T^1, T^2)$ is not uniquely extendable, then by Definition 4.2.1, there must be an edge e in both T^1 and T^2 with $\|e\|_{T^1} \neq \|e\|_{T^2}$. Without loss of generality, we assume $\|e\|_{T^1} > \|e\|_{T^2}$ and e is the first contracted edge if $\Gamma(T^1, T^2)$ is extended beyond T^2 . From the proof of Lemma 4.2.2 and the fact that L is a line, the above extension until e being contracted must be part of L . When e is contracted at \widetilde{T}^2 (see Figure

4.30), there will be a bifurcation and at least one branch is contained in L . Therefore we can find one line segment $\overline{(\tilde{T}^2, \hat{T}^2)}$ along that branch such that $\Gamma(T^1, \hat{T}^2) = \Gamma(T^1, \tilde{T}^2) \cup \Gamma(\tilde{T}^2, \hat{T}^2)$, which is a contradiction to the fact that $\Gamma(T^1, T^2)$ has the maximum number of legs among all geodesics in L . \square

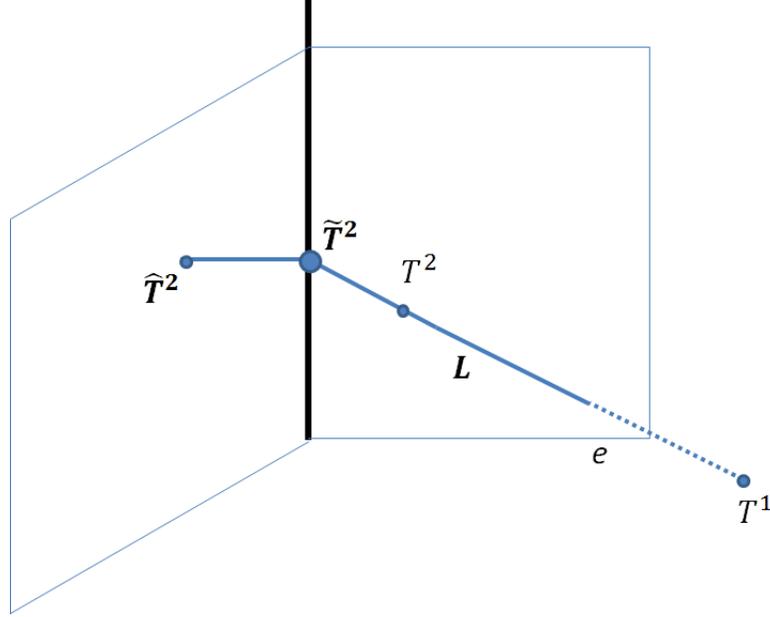


Figure 4.30: The above figure gives a detailed view of what happens when e is contracted.

Theorem 4.2.5. *If $L \in \mathfrak{L}(n)$ is a type I line, then $L = \Gamma_{Ext}(T^1, T^2)$ for some $T^1, T^2 \in L$ with $\Gamma(T^1, T^2)$ being uniquely extendable.*

Proof. Since L is a type I line, it must contain a 3-leg geodesic. Now select $\Gamma(T^1, T^2) \subset L$ with the maximum number of legs among all geodesics in L . From Lemma 4.2.4, we know that $\Gamma(T^1, T^2)$ must be uniquely extendable. Because L is a maximal set and closed under geodesic, $\Gamma_{Ext}(T^1, T^2)$ must be contained in L . However, Theorem 4.2.3 already showed that $\Gamma_{Ext}(T^1, T^2)$ is a type I line. Therefore, $L = \Gamma_{Ext}(T^1, T^2)$. \square

4.3 Type II Lines

In the above section, we characterized a type I line as an extension of a uniquely extendable geodesic with at least three legs. In this section, we characterize a type II line as the union of a set of rays which share the same starting point. Let L be a type II line, and let $G_2(L)$ denote the set of all 2-leg geodesics contained in L . For any $\Gamma(T^1, T^2) \in G_2(L)$, define its *center* to be the point where the two legs meet. The following lemma ensures that every member of $G_2(L)$ is uniquely extendable.

Lemma 4.3.1. *Every $\Gamma(T^1, T^2) \in G_2(L)$ is uniquely extendable. Further, $\Gamma_{Ext}(T^1, T^2) = \overrightarrow{R}(C, T^1) \cup \overrightarrow{R}(C, T^2)$, where C is the center of $\Gamma(T^1, T^2)$.*

Proof. Since L does not contain any 3-leg geodesic, then $\Gamma(T^1, T^2) \in G_2(L)$ contains the maximum number of legs in L . From Lemma 4.2.4 in the previous section, we know $\Gamma(T^1, T^2)$ is uniquely extendable. As in Definition 4.2.2, $\Gamma_{Ext}(T^1, T^2) = \vec{R}(C, T^1) \cup \vec{R}(C, T^2)$ because $\Gamma(T^1, T^2)$ contains only two legs. \square

The concept of 2-leg uniquely extendable geodesic plays an essential role in characterizing type II lines, and it is a generalization of the following concept.

Definition 4.3.1. Two trees $T^1, T^2 \in \mathcal{T}_n$ are *antipodal* if the geodesic $\Gamma(T^1, T^2)$ is a cone path.

If $T^1, T^2 \in \mathcal{T}_n$ are antipodal, then they do not have common edges, hence $\Gamma(T^1, T^2)$ can be thought as a 2-leg uniquely extendable geodesic passing through the origin. However, a general 2-leg uniquely extendable geodesic does not necessarily go through the origin. The following lemma shows that we can consider 2-leg uniquely extendable geodesics as “shifted antipodal”, and its proof is obvious from the definition of unique extendability and antipodality.

Lemma 4.3.2. For $T^1, T^2 \in \mathcal{T}_n$, if $\Gamma(T^1, T^2)$ is a 2-leg uniquely extendable geodesic, then the trees obtained by contracting common edges in T^1 and T^2 are antipodal.

Define the *ray set* $R_2(L)$ for L to be the set of distinct rays $\vec{R}(C, T)$ as given in Lemma 4.3.1. The following lemma reveals that $R_2(L)$ has a simple and finite structure.

Lemma 4.3.3. $R_2(L)$ has a finite number of distinct rays $\vec{R}(C, T)$. Further, all centers C are the same.

Proof. Clearly there can be at most one ray of $R_2(L)$ in each orthant \mathcal{O} , since if there are two distinct rays in \mathcal{O} , then their convex hull must also be in \mathcal{O} , thus contradicting to the fact that L is locally 1-dimensional, thus we proved the claim. Now we know that the number of rays in $R_2(L)$ can not exceed the the number of orthants in \mathcal{T}_n , hence $R_2(L)$ has only a finite number r of rays.

Finally we need to prove that all centers are the same. For two distinct rays $\vec{R}(C, T)$ and $\vec{R}(C', T')$, we want to show that $C = C'$. By **C2**, $\Gamma(T, T')$ must be in L . By **C1**, $\overline{(T, C)}$ and $\overline{(T', C')}$ must both be parts of $\Gamma(T, T')$. Since L is a type II line, $\Gamma(T, T')$ must be a 2-leg geodesic, hence $C = C'$. Therefore, it must be true that all centers are the same. \square

Now we are ready to characterize the type II line as the union of a special set of rays in the following theorem.

Theorem 4.3.4. $L \in \mathfrak{L}(n)$ is a type II line if and only if L is of the form

$$L = \bigcup_{i=1}^r \vec{R}(C, T^i)$$

where C and $\{T^1, \dots, T^r\}$ are trees with the following properties:

1. for each pair (T^i, T^j) , $\Gamma(T^i, T^j) = \overline{(T^i, C)} \cup \overline{(C, T^j)}$,
2. each $\Gamma(T^i, T^j)$ is uniquely extendable,
3. $\{T^1, \dots, T^r\}$ is a maximal set of trees such that the above 2 conditions are satisfied.

Further in this case, L is uniquely determined by C and $\{T^1, \dots, T^r\}$.

Proof. “ \Leftarrow ” Suppose L has the form of $\cup_{i=1}^r \overrightarrow{R}(C, T^i)$ with C and $\{T^1, \dots, T^r\}$ satisfying the conditions stated in the theorem. From the proof of Lemma 4.3.3, there is at most one $\overrightarrow{R}(C, T^i)$ in each orthant, thus **C1** holds for L . By Lemma 4.2.2 and the fact that each $\Gamma(T^i, T^j)$ is uniquely extendable, we have that **C2** holds for L . To show that **C3** holds, suppose that L is strictly contained in some line L' . In particular, L' must have a bifurcation point along some geodesic in L . This bifurcation can not occur on the interior of any $\overrightarrow{R}(C, T^i)$, and thus it must occur at C . Let \widehat{T} be a tree on the additional branch, then by Lemma 4.1.1, $\Gamma(T^i, \widehat{T}) = \overline{(T^i, C)} \cup \overline{(C, \widehat{T})}$ for each $1 \leq i \leq r$. Because $\{T^1, \dots, T^r\}$ is already a maximal set satisfying conditions 1 and 2, there must be at least one T^i such that $\Gamma(T^i, \widehat{T})$ is not uniquely extendable. Hence, the extension of $\Gamma(T^i, \widehat{T})$ beyond \widehat{T} must hit an orthant boundary and create a third leg. From the previous section, we know if L' contains a 3-leg geodesic, then L' is a type I line. However, a type I line can not bifurcate, which is a contradiction to the fact that C is a bifurcation point of L' . Therefore, L must be maximal, i.e. **C3** holds for L . And all pairs (T^i, T^j) are 2-leg geodesics with the common center C , so there can not be a 3-leg geodesic. Therefore, L is a type II line.

“ \Rightarrow ” Since L is a type II line, L contains no 3-leg geodesic. Noticing that any 1-leg geodesic must be contained in some 2-leg geodesic, it is true that $L = \cup \Gamma(T', T'')$ for all $\Gamma(T', T'') \in G_2(L)$. Then from Lemma 4.3.3, we know that $L = \cup_{i=1}^r \overrightarrow{R}(C, T^i)$ for some r . And we also need to prove the 3 properties. First, by Lemma 4.1.1, we know $\Gamma(T^i, T^j) = \overline{(T^i, C)} \cup \overline{(C, T^j)}$ for each pair (T^i, T^j) . Second, by Lemma 4.3.1, each $\Gamma(T^i, T^j)$ is uniquely extendable. Third, the maximality of $\{T^1, \dots, T^r\}$ follows directly from condition **C3** of L . □

After completely characterizing lines in the tree space, we are interested constructive ways to use these lines as PCA objects. In the following two chapters, we will focus on the search algorithms for finding representative line objects of each type, and analyze our data sets by using those objects as principal component analogs.

CHAPTER 5: SAMPLE-LIMITED GEODESICS

In [Nye, 2011], geodesics were used as the first principal components in tree space to maximize the variance of data projections. The first principal component in [Nye, 2011] was constructed to be a general-positioned geodesic passing through the majority consensus tree by adding coordinates iteratively. To make it easier to use geodesics as principal components in tree space, using an idea from [Feragen et al., 2013], in this section we are going to focus on an even simpler set of geodesics, namely *sample-limited geodesics* which are geodesics between pairs of data trees. We will start by defining sample-limited geodesics in tree space.

Definition For T^1, T^2 in \mathcal{T}_n , the sample-limited geodesic associated with T^1 and T^2 is defined as the geodesic $\Gamma(T^1, T^2)$ between T^1 and T^2 .

To reduce dimension, a few sample-limited geodesics are needed in order to capture a good portion of total variation with respect to the Fréchet mean. Two aspects associated with sample-limited geodesic will be studied in this chapter. In section 5.1, the first sample-limited geodesic (SLG1) is developed and tested on the five data sets introduced in Chapter 2. In section 5.2, potential approaches to higher-order sample-limited geodesics are discussed.

5.1 The First Sample-limited Geodesic

Given a data set $\mathbf{T} = \{T^1, T^2, \dots, T^r\}$ in \mathcal{T}_n , for the geodesic $\Gamma(T^1, T^2)$ between any pair of trees T^1 and T^2 , the projection of any other tree $T \in \mathbf{T}$ onto $\Gamma(T^1, T^2)$ is defined as the tree T' along $\Gamma(T^1, T^2)$ which gives the shortest distance between T and T' . And the distance between T and T' , denoted as $d(T, T')$, is called *projection residual* of T onto $\Gamma(T^1, T^2)$. The first sample-limited geodesic (SLG1) is the sample-limited geodesic along which the projections of all trees in \mathbf{T} have the largest variation. The formal definition is given below.

Definition 5.1.1. For a given data set \mathbf{T} , the first sample-limited geodesic (SLG1) Γ^1 is defined as following:

$$\Gamma^1 = \underset{\Gamma(T^1, T^2): T^1, T^2 \in \mathbf{T}}{\operatorname{argmax}} \operatorname{Var}_{\Gamma(T^1, T^2)} \{P_{\Gamma(T^1, T^2)}(T) : T \in \mathbf{T}\}$$

where $P_{\Gamma(T^1, T^2)}(T)$ denotes the projection from T onto $\Gamma(T^1, T^2)$, and $\operatorname{Var}_{\Gamma(T^1, T^2)}$ indicates that the variance is calculated along geodesic $\Gamma(T^1, T^2)$.

However, given a geodesic segment $\Gamma(T^1, T^2)$ and a tree $T \notin \Gamma(T^1, T^2)$, it is not obvious to find the projection $P_{\Gamma(T^1, T^2)}(T)$. To find a good algorithm, it is beneficial to study the convexity of squared distance

$$d_T^2(\tilde{T}) = d^2(T, \tilde{T}) \text{ with } \tilde{T} \in \Gamma(T^1, T^2).$$

Definition 5.1.2. A subset $\mathbf{S} \subset \mathcal{T}_n$ is called convex if the whole geodesic $\Gamma(T^1, T^2)$ between any pair of trees $T_1, T_2 \in \mathbf{S}$ is contained in \mathbf{S} .

It is straightforward to show that any geodesic segment is a convex set in \mathcal{T}_n . Then the squared distance between a fixed tree and a variable tree on a given geodesic segment can be studied. The following lemma is critical to the development of the algorithm for finding the first sample-limited geodesic.

Lemma 5.1.1. *Given a geodesic segment $\Gamma(T^1, T^2)$ and a tree $T \notin \Gamma(T^1, T^2)$, the squared distance function $d_T^2(\tilde{T})$ is convex for all $\tilde{T} \in \Gamma(T^1, T^2)$.*

Proof. Take any two trees $T_a, T_b \in \Gamma(T^1, T^2)$ and let T_λ be a tree on $\Gamma(T^1, T^2)$ between T_a and T_b with distance to T_a being $\frac{\lambda}{1-\lambda}$ of length of $\Gamma(T^1, T^2)$. Need to show that

$$d_T^2(T_\lambda) \leq \lambda \cdot d_T^2(T_a) + (1 - \lambda) \cdot d_T^2(T_b) \quad \forall \lambda \in [0, 1]$$

Due to the non-positive curvature in \mathcal{T}_n , if the similar result can be proved in the comparison triangle of T_a, T_b , and T , then by the fact that triangle is skinnier in \mathcal{T}_n , the lemma will be completed. Now the rest needs to be proved is the following subproblem: in a 2 dimensional Euclidean space, fix two points $(a, 0)$ and $(b, 0)$, and choose a random point (c, d) , then

$$[\lambda \cdot a + (1 - \lambda) \cdot b - c]^2 + d^2 \leq \lambda \cdot [(c - a)^2 + d^2] + (1 - \lambda) \cdot [(c - b)^2 + d^2]$$

Let $f(\lambda) = LHS - RHS$, then by applying standard calculus, we can obtain $f'(\lambda) = (a - b)^2(2\lambda - 1)$, hence f reaches its maximum when $\lambda = 0$ or $\lambda = 1$. And $f(0) = f(1) = 0$, so $f(\lambda) \leq 0$ for all $0 \leq \lambda \leq 1$, which also proves the subproblem above.

Now by the non-positive curvature in \mathcal{T}_n , it is obvious that d_T^2 is convex on $\Gamma(T^1, T^2)$. □

There is an immediate result from the above lemma which will be stated as the following corollary.

Corollary 5.1.2. *Any local minimum of d_T^2 is also the global minimum on a given geodesic segment $\Gamma(T^1, T^2)$.*

Based on this corollary, it is not hard to prove the fact that the projection of a tree T onto any geodesic is unique. Given a geodesic segment $\Gamma(T^1, T^2)$ and a tree $T \notin \Gamma(T^1, T^2)$, by applying the uniqueness of projection, a sufficiently good heuristic solution T' for the projection of T can be obtained from the following algorithm.

Algorithm for Bisection Tree with Tolerance Level τ

Initialization: $\Delta = \infty$, $a = T_1$, $b = T_2$

while $\Delta > \tau$ **do**

Let \tilde{T} be the half-way bisection tree along the geodesic segment $\Gamma_{a,b}$

Calculate $d_T^2(a)$, $d_T^2(\tilde{T})$, $d_T^2(b)$

if $d_T^2(a) < d_T^2(\tilde{T})$ and $d_T^2(a) < d_T^2(b)$ **then**

$\Delta = |d_T^2(a) - d_T^2(\tilde{T})|$, $a = a$ and $b = \tilde{T}$

else if $d_T^2(b) < d_T^2(\tilde{T})$ and $d_T^2(b) < d_T^2(a)$ **then**

$\Delta = |d_T^2(b) - d_T^2(\tilde{T})|$, $b = b$ and $a = \tilde{T}$

else

Let \tilde{T}^- and \tilde{T}^+ be the trees with small perturbations toward a and b respectively

Calculate $d_T^2(\tilde{T}^-)$, $d_T^2(\tilde{T}^+)$

if $d_T^2(\tilde{T}^-) < d_T^2(\tilde{T}^+)$ **then**

$\Delta = |d_T^2(a) - d_T^2(\tilde{T}^-)|$, $a = a$ and $b = \tilde{T}^-$

else

$\Delta = |d_T^2(b) - d_T^2(\tilde{T}^+)|$, $b = b$ and $a = \tilde{T}^+$

end if

end if

end while

Let \tilde{T} be the half-way bisection tree along the geodesic segment $\Gamma(a, b)$

$T' = \tilde{T}$

By applying the above algorithm, the first sample-limited geodesic (SLG1) can be obtained for the same five data sets: WF2, WF10, WF40, brain artery data, and uniformly random data. One of the important criteria, for determining whether SLG1 is a good first-step approach to reduce the dimension in a data set, is the proportion of total variation captured. In this section the total variation of a data set is defined with respect to the Fréchet mean. For computing the Fréchet mean, the recently developed algorithm in [Skwerer, 2014] is used in this section. The following Table 5.3 shows the proportion of total variation captured by SLG1 in the five data sets.

Since the end points for each SLG1 are also data points, their projections onto SLG1 should be counted as part of the data projections. However, the projection variation contributed by end points is mainly

determined by the length of the SLG, which raises a concern that the long SLGs may dominate and prevent us from finding SLGs carrying more information. Therefore, the first column in Table 5.3 shows the proportion of variation captured by SLG1 if end points are not counted. On one hand, the decreasing ordering of proportions of total variation captured by SLG1 in the five data sets reflects the decreasing similarity ordering across data sets. On the other hand, the proportions of total variation captured by SLG1 in all five data sets are surprisingly low, with the largest being 1.9%. One major reason that SLG1 can only capture such a small amount of variation across five data sets is the intrinsic non-positive curvature of the tree space. More precisely, all the geodesics between pairs of data trees bend heavily towards the point cloud of other data trees as shown in Figure 3.19. In that figure, three representative SLGs are chosen in the brain artery data, namely minimum, median, and maximum length SLGs. The red cross represents the origin which is also the Fréchet mean, and the green dots represent the equally-spaced grid points along the SLG. It is clear that the bending behavior is consistent across all three representative SLGs, and this phenomenon is true for all five data sets.

	Proportion(without end points)	Proportion(with end points)	Fréchet sum of squares
WF2	1.9%	4.3%	4.9×10^2
WF10	1.2%	3.6%	1.8×10^4
WF40	0.6%	3.5%	9.1×10^4
Brain	0.2%	4.1%	8.8×10^4
Random	0.05%	4.6%	8.7×10^4

Table 5.3: Column 1 shows the performance of SLG1 for five data sets without counting end points. Column 2 shows the same statistics but counting end points. Column 3 lists the total variation with respect to the Fréchet mean in each data set. Overall, SLG1 explains a relatively small amount of the total variation in all five data sets.

The second column in Table 5.3 shows the proportion of variation captured by SLG1 obtained when end points are included. For WF2, the projections of end points contribute more than the total projection variation of all other data points; for WF10, this contribution is 2 times the total projection variation of all other data points; for the other three data sets, the projection variation contributed by the end points clearly dominates. And we notice that all proportions in column 2 are stable between 3% and 5%. As discussed in [Hall et al., 2005], data tend to lie deterministically at the vertices of a regular simplex when the sample size is fixed and the dimension goes to ∞ . Recall in Figure 2.13, the distribution of the pairwise geodesic lengths has small variance across the five data sets, which makes it reasonable to use the vertices of a Euclidean simplex to approximate our five tree space data sets. Assume there are n points in a high-dimensional Euclidean space, and their pairwise distances are all the same denoted as d . For the SLG connecting any

pair of points, all other points will project exactly onto the middle of that SLG. Under this model, all SLGs capture exactly the same amount of variation which can be calculated as:

$$\frac{1}{n-1} \left[\left(\frac{d}{2}\right)^2 \times 2 \right] = \frac{d^2}{2(n-1)}$$

And the total variation is

$$\frac{1}{n-1} \left[\left(\frac{d}{2}\right)^2 \times n \right] = \frac{nd^2}{4(n-1)}$$

hence the proportion of projection variation explained by each SLG is $2/n$. In the context of the five data sets, $n = 67$ and the proportion is about 3% which is surprisingly close to the proportions listed in column 2 of Table 5.3. This indicates that the SLGs across the five data sets act similarly to edges in a regular simplex in the sense of capturing data variation when including the two end point projections.

The third column lists the total variation in terms of mean square distances with respect to the Fréchet mean. We clearly see that among Wright-Fisher data this Fréchet sum of squares decreases when the width parameter increases, which can be explained by the following ANOVA calculation. Recall in Euclidean space, the total variation of a data set with respect to the origin can be decomposed into the sum of the total variation with respect to the mean and the effect of deviation of the mean from the origin. Precisely, given n points $\{x_1, \dots, x_n\} \in \mathbb{R}^k$, the following equality holds.

$$\sum_{i=1}^n |x_i|^2 = \sum_{i=1}^n |x_i - \bar{x}|^2 + n|\bar{x}|^2$$

Here we adopt some of the above notations from Euclidean space and use them in tree space. Specifically, $\sum |T^i|^2$, $\sum |T^i - \bar{T}|^2$, and $r|\bar{T}|^2$ represent the total variation of a set of trees with respect to the origin, the total variation of a set of trees with respect to the Fréchet mean, and the effect of deviation of the Fréchet mean from the origin, respectively. Table 5.4 lists the values of these three statistics across the five data sets. First, for three Wright-Fisher data sets, the values in Column 2 and Column 3 do not add up to the values

	$\sum T^i ^2$	$\sum T^i - \bar{T} ^2$	$r \bar{T} ^2$
WF2	5.15×10^6	0.03×10^6	5.12×10^6
WF10	5.36×10^6	1.17×10^6	4.18×10^6
WF40	6.97×10^6	6.12×10^6	0.79×10^6
Brain	5.88×10^6	5.88×10^6	0
Random	5.82×10^6	5.82×10^6	0

Table 5.4: Column 1, 2, and 3 show the total variation with respect to the origin, the total variation with respect to the Fréchet mean, and the effect of the Fréchet mean in each data set.

in Column 1 (for WF2, the equality is due to rounding), because of the non-positive curvature of the tree space. Second, for these three data sets, the norm of the Fréchet mean decreases as the width parameter increases. Recall in Figure 2.12, the spread of data increases dramatically from WF2 to WF40, and larger data spread is associated with stronger stickiness of the Fréchet mean. This is why the norm of the Fréchet mean decreases from WF2 to WF40. Third, we notice that the values of $r|\bar{T}|^2$ for the Brain Artery data and the Uniformly Random data are both zero, since their Fréchet means are the origin due to the stickiness property.

Other than the proportion of variation captured by SLG1, it is also interesting to look at the actual distribution of data projections along SLG1. Figure 5.31 below shows that all data projections concentrate on a relatively narrow segment of SLG1 across the five data sets. All SLGs are scaled to unit length, and the two end points are represented by 0 and 1. Each data projection onto SLG1 is assigned a score between 0 and 1 according to the relative position along SLG1. Each blue curve represents the kernel density estimate of the distribution of the projection scores for one data set with the bandwidth chosen by using the Sheather Jones plug-in method. The concentration of the projection scores increases among five data sets in the following order: WF2, WF10, WF40, brain artery data and uniformly random data, because the data sets become more and more spread out in this order. Also, the projection means for the brain artery data and the uniformly random data are very close to 0.5, but for WF2, WF10 and WF40, the means are clearly away from 0.5, which is possibly related to the non-origin Fréchet mean of the Wright-Fisher data.

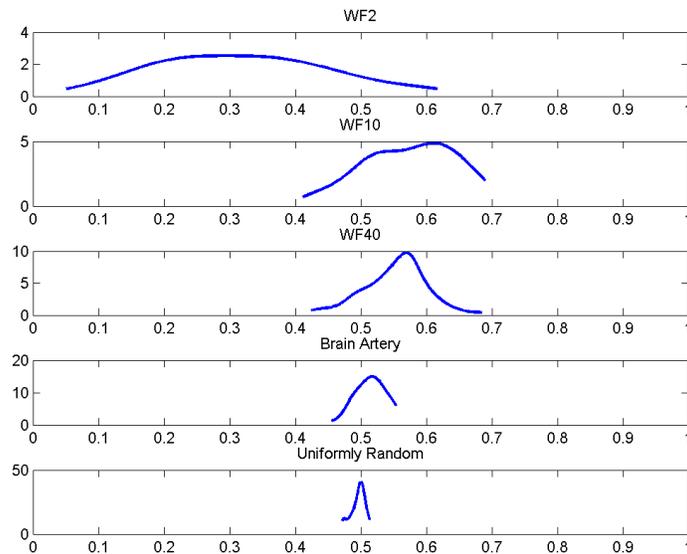


Figure 5.31: The above figures show the distributions of data projections along SLG1 across five data sets. The projections are only distributed within a relatively narrow range between the two end points of SLG1 for all five data sets.

5.2 Higher Order Sample-limited Geodesic

The SLG1 itself is not an effective data-summary tool unless this concept can be generalized to higher orders. It will be ideal that the portion of variation captured by SLGs of different orders do not overlap. Recall in Euclidean space PCA, different principal components are uncorrelated, or geometrically speaking, orthogonal with each other, hence they captured a non-overlapping portion of variation. However, in the tree space it is not easy to define the concept of orthogonality between geodesics so that SLGs will capture strictly non-overlapping portions of variation. Alternatively, it will be more practical to search for some heuristic higher order SLGs. From the discussion in the previous section, we know when the number of the leaves becomes large, the tree space data can be well approximated by the vertices of a regular simplex. Since all the edges of a regular simplex are orthogonal to each other, we can come up with the following simple heuristic:

1. Find the 1st SLG as the SLG with the largest projection variation.
2. Find the n^{th} SLG to be the SLG with the largest projection variation among all SLGs whose end points are not used in the $n - 1$ previously selected SLGs.
3. Repeat Step 2 until no new SLG can be found.

Defining good higher order SLGs can be a challenging research topic. However, in Table 5.3, we already see that SLG1 is not very effective in terms of capturing data variation, due to the underlying geometrical properties of the five data sets considered in this dissertation. Hence, we do not spend more effort on SLGs, instead we will focus on a more promising principal component object coming from type II line, namely the principal ray set, in the next chapter.

CHAPTER 6: USING PRINCIPAL RAYS TO MODEL PCA IN TREE SPACE

In this chapter, we are going to use a special type II line to represent the first principal component in tree space. From Theorem 4.3.4, we know a type II line is a set of mutually antipodal rays with common center C . The special case we focus on in this chapter is when all the antipodal rays have the the origin O as their common center.

Let $\mathbf{T} = \{T^1, T^2, \dots, T^r\}$ be a data set representing a set of trees in \mathcal{T}_n . The *first principal ray set* (1st PR set) is the type II line centering at the origin O with the largest sum of squared projections from \mathbf{T} . Each ray in the 1st PR set is called a *principal ray*. In Section 6.1 we will first study projections of trees onto a ray with a fixed direction. In Section 6.2 we discuss how to find an optimal ray by maximizing the projection variance over direction-variable rays. Then in Section 6.3 we naturally expand our search to the 1st PR set. Finally, in Section 6.4 we explore the higher order *principal axis set* (PA set) in tree space, which is a special case of higher order PR set.

6.1 Fixed Rays

We first focus on the computation of projections onto a given ray in this section. Finding the projection of a tree onto a fixed ray is equivalent to finding the shortest geodesic between the tree and a point on the ray. Recall Theorem 1.3.1 in Section 1.3: a geodesic between two trees must satisfy (P1), (P2) and (P3). Given a tree $\tau \in \mathcal{T}_n$, the ray generated by τ is the set of all non-negative scalar multiples $\vec{\mathbf{r}} = \{\lambda \cdot \tau : \lambda \geq 0\}$, so all the trees on the ray will have exactly the same edge set. This leads to the following result regarding the topological structure of the geodesic between T and any point on $\vec{\mathbf{r}}$.

Lemma 6.1.1. *Given a tree T in \mathcal{T}_n , every geodesic between T and any point on the ray $\vec{\mathbf{r}} = \{\lambda \cdot \tau : \lambda \geq 0\}$ has the same support.*

Proof. Assuming we have already found a geodesic between T and τ , it means that there is a support $(\mathcal{A}, \mathcal{B})$ satisfying (P1), (P2) and (P3). Since all the trees on a ray have the same edge set, for any tree $\lambda \cdot \tau$ on the ray, we can come up with the same support $(\mathcal{A}, \mathcal{B})$ between $\lambda \cdot \tau$ and T which automatically satisfies (P1). Since all the edge lengths of $\lambda \cdot \tau$ are the same scalar multiple of those of τ , then all the inequalities in (P2) and (P3) are still satisfied. Therefore, the geodesics between T and all the trees on a ray will share the same support pairs. □

The above lemma is useful in determining the projection of any point onto a fixed ray, as the following

proposition shows.

Proposition 6.1.2. *Let τ be the tree with unit norm on the ray $\vec{\mathbf{r}} = \{\lambda \cdot \tau : \lambda \geq 0\}$ and let T be another tree in \mathcal{T}_n . Let \mathcal{C} be the common edge set of τ and T , and let $(\mathcal{A}, \mathcal{B}) = ((A_1, \dots, A_k), (B_1, \dots, B_k))$ be the support for the geodesic from τ to T satisfying (P1), (P2) and (P3). It follows that the projection of T onto $\vec{\mathbf{r}}$ is $\lambda^* \cdot \tau$, where*

$$\lambda^* = \begin{cases} \sum_{e \in \mathcal{C}} |e|_\tau |e|_T - \sum_{i=1}^k \|A_i\| \|B_i\|, & \text{if } \sum_{e \in \mathcal{C}} |e|_\tau |e|_T > \sum_{i=1}^k \|A_i\| \|B_i\|; \\ 0, & \text{otherwise.} \end{cases}$$

Proof. Let $T' = \lambda \cdot \tau$ be any tree on $\vec{\mathbf{r}}$ and let Γ be the geodesic between T and T' , then we are trying to find a λ^* which minimizes the $L(\Gamma)^2$, i.e.

$$\lambda^* = \operatorname{argmin}_\lambda L(\Gamma)^2 = \operatorname{argmin}_\lambda \left[\sum_{i=1}^k (\lambda \cdot \|A_i\| + \|B_i\|)^2 + \sum_{e \in \mathcal{C}} (\lambda \cdot |e|_\tau - |e|_T)^2 \right]$$

Taking first and second derivatives where,

$$\begin{aligned} \frac{d[L(\Gamma)^2]}{d\lambda} &= 2 \sum_{i=1}^k \|A_i\| (\lambda \cdot \|A_i\| + \|B_i\|) + 2 \sum_{e \in \mathcal{C}} |e|_\tau (\lambda \cdot |e|_\tau - |e|_T) \\ \frac{d^2[L(\Gamma)^2]}{d\lambda^2} &= 2 \sum_{i=1}^k \|A_i\|^2 + 2 \sum_{e \in \mathcal{C}} |e|_\tau^2 = 2 \end{aligned}$$

From the fact that the second order derivative is positive, we know that $L(\Gamma)^2$ is a convex function of λ , hence by setting the first derivative to zero, we can solve for λ^* . \square

The support $(\mathcal{A}, \mathcal{B})$ in the above proposition can be found in polynomial time by using the algorithm proposed in [Owen and Provan, 2009], and thus so can the projection of T onto $\vec{\mathbf{r}}$. The above proposition is also consistent with our intuition: if T has non-trivial projection onto a ray $\vec{\mathbf{r}}$, then T and any tree $\tau \in \vec{\mathbf{r}}$ must have some common edges, and the common edges should somehow dominate the edge sets of T and τ (in the sense that $\sum_{e \in \mathcal{C}} |e|_\tau |e|_T > \sum_{i=1}^k \|A_i\| \|B_i\|$).

6.2 Direction-variable Rays

In this section, we allow the direction of a ray to vary in \mathcal{T}_n . For a set of trees $\mathbf{T} = \{T^1, T^2, \dots, T^r\}$, we are going to search for the optimal ray in \mathcal{T}_n which maximizes the projection variance from \mathbf{T} . Since a direction-variable ray has different edge set in different orthants, we start by solving the problem constrained

to a specific orthant \mathcal{O} . Note that even when a ray is restricted in \mathcal{O} , the support of the geodesic between a fixed tree T and a tree τ on the ray will not always be the same. Recall in Section 1.5, we summarized the combinatorics of geodesics in tree space which is studied in [Miller et al., 2015, Section 3]. From their discussion, based on \mathbf{T} the orthant \mathcal{O} can be divided into a number of *multivistal cells*, which are convex polyhedral cones. A nice property of these multivistal cells is: for any tree τ in the multivistal cell, the support of the geodesic between T and τ stays the same. Hence, we want to further restrict our attention to a more specific problem of searching for the sub-optimal ray within a multivistal cell. Recalling Definition 1.5.4, each multivistal cell can be represented by $\mathcal{V}(\mathbf{T}; \mathcal{O}; \mathcal{A}^{\mathbf{T}}, \mathcal{B}^{\mathbf{T}}) = \bigcap_{l=1}^r \mathcal{V}(T^l, \mathcal{O}; \mathcal{A}^l, \mathcal{B}^l)$, where r is the number of trees in \mathbf{T} . Let \mathcal{C}^l be the common edge set between the ray and T^l and applying Proposition 6.1.2, we can formulate this most restricted problem into the following form:

$$\text{Maximize} \quad \sum_{l=1}^r \left[\left(\sum_{e \in \mathcal{C}^l} |e|_{\tau} |e|_{T^l} - \sum_{i=1}^{k_l} \|A_i^l\| \|B_i^l\| \right)^+ \right]^2 \quad (6.9)$$

$$\text{Subject to} \quad \tau \in \mathcal{V}(\mathbf{T}; \mathcal{O}; \mathcal{A}^{\mathbf{T}}, \mathcal{B}^{\mathbf{T}}) \quad (6.10)$$

$$\|\tau\| = 1 \quad (6.11)$$

Note that the above objective function only has a fixed form within each multivistal cell defined in equation (6.10). In other words, the form of the objective function will vary even within one single orthant, which results in little hope of getting a closed form solution. Therefore, we are going to focus on how to find a workable heuristic algorithm which gives approximate solutions for finding optimal rays in tree space. A natural approach is the *steepest descent* algorithm. However, the feasible region defined in equations (6.10) and (6.11) is not convex, which creates some difficulties for applying the steepest descent algorithm directly. Our remedy is that we make the feasible region convex by using a more complex objective function as in the following formulation.

$$\text{Maximize} \quad \sum_{l=1}^r \left[\frac{\left(\sum_{e \in \mathcal{C}^l} |e|_{\tau} |e|_{T^l} - \sum_{i=1}^{k_l} \|A_i^l\| \|B_i^l\| \right)^+}{\|\tau\|} \right]^2 \quad (6.12)$$

$$\text{Subject to} \quad \tau \in \mathcal{V}(\mathbf{T}; \mathcal{O}; \mathcal{A}^{\mathbf{T}}, \mathcal{B}^{\mathbf{T}}) \quad (6.13)$$

Here we relaxed the condition that τ must have unit norm by dividing the objective function by $\|\tau\|$. Although the above formulation is defined within a multivistal cell, in practice the supports along the geodesics between the moving ray and the data trees can be recalculated after each search step, even crossing a multivistal cell boundary. However, before we can solve the optimal ray problem by using the steepest descent algorithm, there are still two issues regarding its application in tree space. One is particularly related to the tree space

topology. When the steepest descent search hits an orthant boundary, we need to decide in which one of the adjacent orthants the search will continue. A *depth-first* search (DFS) strategy is used in our steepest descent algorithm. DFS searches along a particular sequence of orthants as far as possible until it hits a local optimal ray or enters a previously visited orthant. If it hits a local optimum which has the best objective value since the the search started, then we terminate the search. Otherwise, DFS backtracks to the orthant with the best objective value, and the search process continues from there. A special case may occur when DFS hits an orthant boundary: if the search within each of the adjacent orthants again hits the same boundary, then DFS will continue the search process within that orthant boundary. The other issue is common to the general steepest descent algorithm: the search process is often trapped in some local optimum. This issue can be avoided if the objective function is concave. However, by exploring the following example in \mathcal{T}_4 , it can be seen that the objective function (6.12) is generally not concave.

In \mathcal{T}_4 , 10 trees were randomly generated in terms of both topologies and edge lengths. We generated the topology for each of these 10 trees by randomly choosing one orthant from the five orthants in Figure 6.32 The edge lengths are integers randomly selected between 1 and 10. And the resulting data set is also displayed in the same picture.

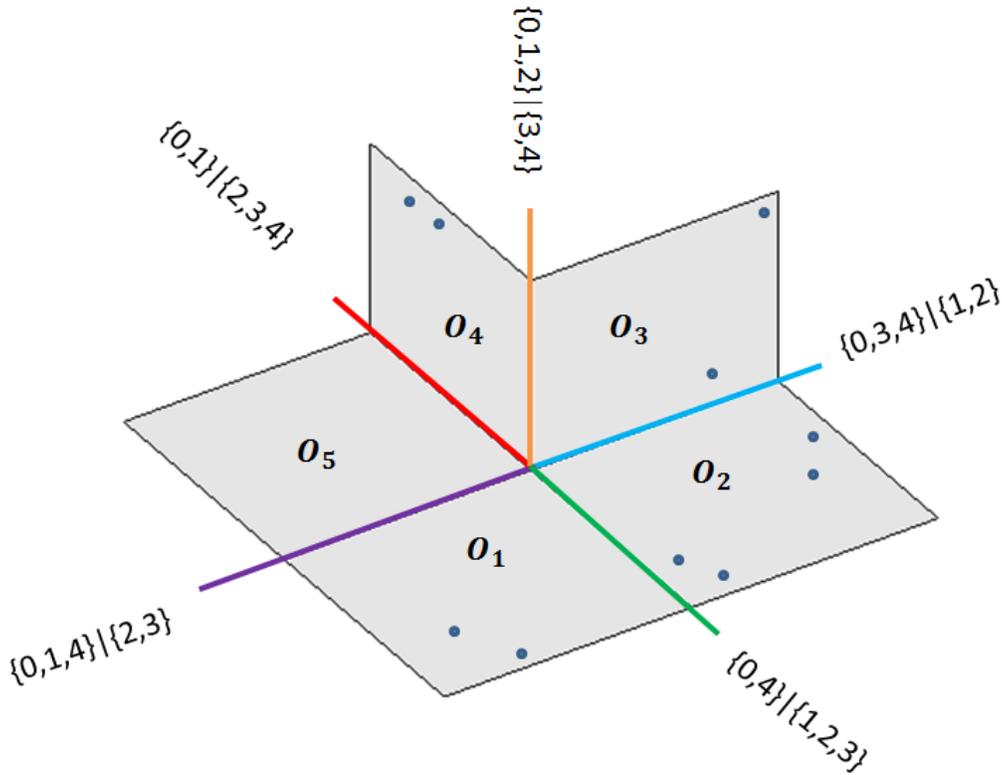


Figure 6.32: A randomly generated small data set in a portion of \mathcal{T}_4 : 10 trees are represented by blue dots, and all the edge lengths are integers between 1 and 10.

We tried our steepest descent algorithm in this small data set with the search being restricted to only the 5 orthants shown in the above figure. The starting rays were chosen to be the 45° ray in each of these 5 orthants. Two local optimal rays were detected: one is in \mathcal{O}_1 which is also the global optimum, and the other is in \mathcal{O}_4 . The optimal ray and the Fréchet mean are two measures summarizing a tree data set, and they both use the sum of squared distances as the objective. Recall that for data in n -dimensional Euclidean space \mathbb{R}^n , the best fit line (i.e. the first principal component) goes through the sample mean, which is also the Fréchet mean. Followed by this intuition in Euclidean space, we originally suspected that there might be some similar relationship between these two objects, such as the optimal ray going through the Fréchet mean or at least the two being in the same orthant. However, after actually calculating the Fréchet mean, it ends up in \mathcal{O}_2 , which is not even the same orthant as either of the two locally optimal rays. This observation also shows that the objective function (6.12) is not concave in general. Figure 6.33 illustrates the relative positions of the two locally optimal rays and the Fréchet mean in the same portion of \mathcal{T}_4 . The two locally optimal rays are represented by black dotted lines, and the Fréchet mean is plotted as a big filled circle in red.

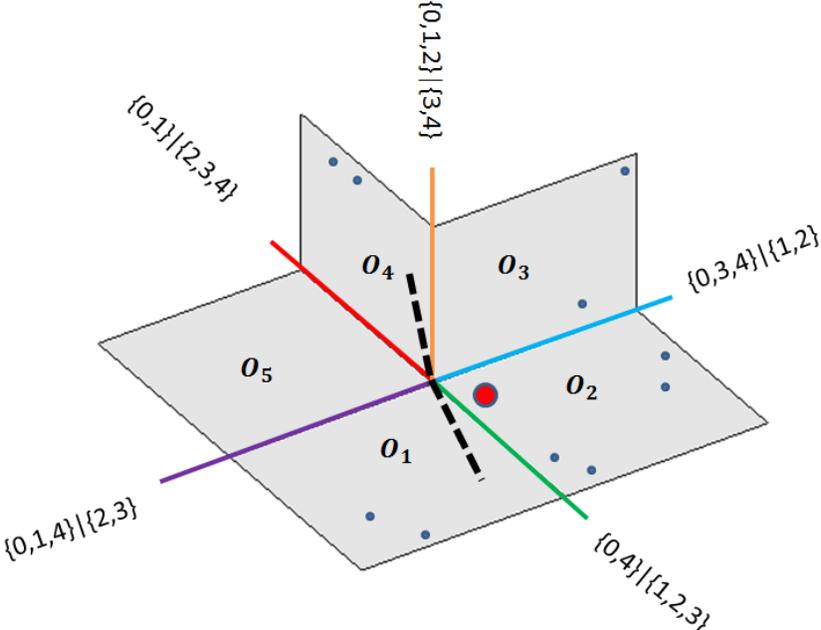


Figure 6.33: The black dotted lines represent locally optimal rays, and the red circle denotes the Fréchet mean. This plot shows that the locally optimal rays and the Fréchet mean do not locate in the same orthant.

One reason for first trying the steepest descent algorithm in a small portion of \mathcal{T}_4 is that this region is simple enough that we can actually see how the value of the objective function varies as a ray moves through this region. Let the ray start from the orthant boundary corresponding to the split $\{0, 1, 4|2, 3\}$ in Figure

6.33, and move counterclockwise. In Figure 6.34, the horizontal axis represents the angle from the start in units of radian/π passed by the ray, and the vertical axis represents the sum of squared projections from the 10 data trees onto the moving ray, which is the value of the objective function. The blue curve shows how the objective function behaves when the ray moves over the entire region. The vertical dashed lines denote the orthant boundaries, and the green vertical line indicates the Fréchet mean. One property of projection onto a ray is that a data point will project at the origin whenever the angle between the ray and the data point is larger than 90° . It is helpful to keep track of which data point projects positively as the ray moves through the region. This information is displayed by the red horizontal line segments near the top of this figure. The 10 data points are represented by those line segments with different heights. Each line segment covers a certain range of angles between the moving ray and the start, and the corresponding data point projects positively when the ray is in that range. When the global optimum is achieved in \mathcal{O}_1 , 5 data points project positively onto the ray. For the other local optimum in \mathcal{O}_4 , only 3 data points project positively. When the ray goes through the Fréchet mean, it gets 7 positive projections, but these are overall too small to attract the peak of the sum of squared projections curve. It is also noticeable that there is no positive projection when the ray is in a narrow region within \mathcal{O}_5 , i.e. the sum of squared projections is 0.

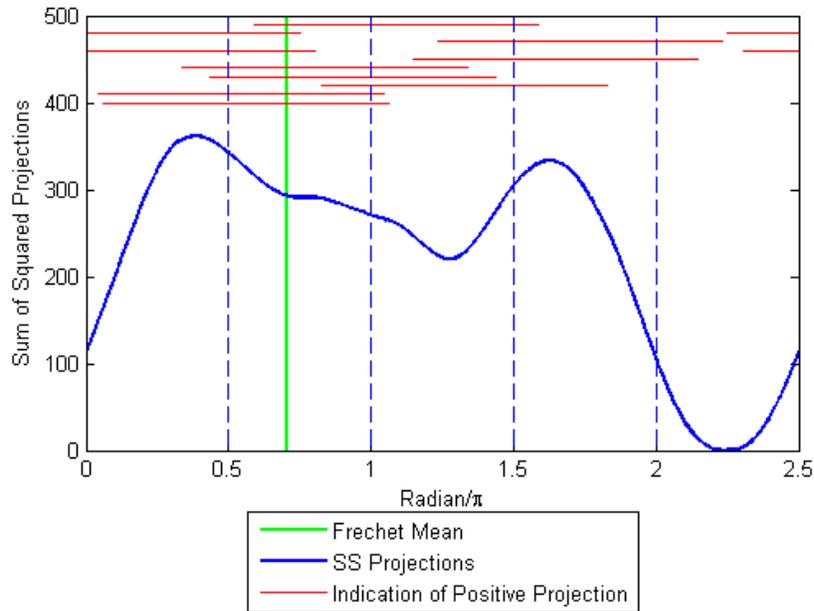


Figure 6.34: This figure shows the non-concavity of the objective function (6.12) as well as some other aspects of this toy example in \mathcal{T}_4 . The horizontal axis gives the angle travelled by the moving ray. The vertical axis represents the sum of squared projections from 10 data points onto the moving ray. The orthant boundaries are indicated by dashed vertical lines. The Fréchet mean is plotted as the green vertical line. The blue curve shows how the value of the objective function changes as the ray moves. The red horizontal line segments keep track of which data points project positively onto the ray.

Since the objective function (6.12) is not concave in general, it is very likely that the steepest descent algorithm can only find a local optimum. A common strategy for finding the global optimum or a good local optimum is to try different starting points. Recall that five data sets have been discussed in Chapter 2, and they are WF2, WF10, WF40, the brain artery data, and the uniformly random data. Trees contain 129 leaves across all five data sets. Ideally we want to try starting points from all different orthants, but the total number of orthants for the tree space containing each of those five data sets is $253!! \approx 10^{250}$, which is not manageable by using our available computing resources. One practical choice is to try starting points from each orthant which contains at least one data point. In particular, the starting point in each orthant is chosen to be the center ray which has the same angle between itself and all the orthant axes. Another aspect of applying the steepest descent algorithm is setting the criteria for terminating the search. One natural criterion is to stop the search when an acceptable local optimum is reached. In our particular version of the algorithm, an acceptable local optimum is a local optimum which has better objective value than any other rays along the search path. The number of orthants that needs to be considered, however, makes this unmanageable. Hence, we add the second criterion to terminate the search if 1000 distinct orthants have been visited. If this criterion is satisfied, the ray with the best objective value along the search path is returned. Now the numerical results are summarized below. Table 6.5 lists the best locally optimal ray among search paths with various starting points for each of those five data sets. This table summarizes the

	SS Projections	Total SS	Proportion	Angle
WF2	5.12×10^6	5.15×10^6	99.37%	0.01°
WF10	4.20×10^6	5.36×10^6	78.38%	0.55°
WF40	7.83×10^5	6.97×10^6	11.24%	21.44°
Brain	1.36×10^5	5.88×10^6	2.32%	—
Random	1.34×10^5	5.82×10^6	2.30%	—

Table 6.5: This table lists four summaries across five data sets and the proportion in the third column is used to measure the performance of the best locally optimal ray.

performance of the best locally optimal ray. The first column lists the sum of squared projections onto the best locally optimal ray across the five data sets. The second column gives the total sum of squared norms for all these data sets. The third column shows the proportion of variation captured by the best locally optimal ray, which is the quotient of the first column divided by the second column. This proportion is a key measurement of the performance of the best locally optimal ray. The fourth column gives the angle in degrees between the best locally optimal ray and the ray going through the Fréchet mean.

Recall from Section 2.4.1, the pairwise angles of WF2 are all less than 10° , which means that all data

trees stay within a relatively narrow cone. This agrees with the fact that 99.37% of the total variation is captured by the best locally optimal ray. Trees in WF10 also stay relatively close to each other and their pairwise angles are mostly between 30° and 50° , on the other hand, the best locally optimal ray can still capture 78.38% of the total variation. For these two data sets, their Fréchet means are far away from the origin, and the best locally optimal rays almost go through the Fréchet means, which is shown by the angles in Column 4. WF40 becomes much more scattered and most of its pairwise angles are between 80° and 120° , and as expected the corresponding proportion drops to only 11.24%. The Fréchet mean of WF40 gets pretty close to the origin, and the best locally optimal ray is less influenced by the Fréchet mean, with the angle being more than 20° . For the brain artery data and the uniformly random data, their pairwise angles are so close to 180° that the best locally optimal ray captures less than 3% of the total variation. The angles in Column 4 are not available for these two data sets because for each the Fréchet mean is the origin [Skwerer et al., 2014a,b]. The performance of the best locally optimal rays for these five data sets is consistent with the previous angle-based data summaries.

In order to understand more details about the structure of the five data sets, it is helpful to look at the returned objective values from all different starting points. Again we choose the proportion of variation captured by the ray as the comparative measure. This proportion is calculated for each returned ray and together there are 67 proportions (one for each ray) for each data set. For WF2, all returned optimal rays capture between 99.33% and 99.37% of the total variation, which reflects the small variation of the data set. For WF10, those proportions are still between 76.5% and 78.5%, which is expected as these data are more spread out. When the trees become really scattered as in WF40, the proportions are only roughly between 1% and 11%. Furthermore, when trees are antipodal or nearly antipodal to each other in the brain artery data and the uniformly random data, all returned rays only capture between 1% and 2.5% of the total data variation.

Among those five data sets, we are most interested in the brain artery data. However, a closer look at the returned objective values for the brain artery data set reveals the fact that the stickiness (as mentioned in Section 2.5) of the tree space plays a critical role in searching for the best locally optimal ray. For the 67 returned locally optimal rays, their proportions are sorted in ascending order and they are plotted against this ordering as red dots in Figure 6.35. We can also divide the squared norm of each tree by the total sum of squared norms in the brain artery data set to obtain another 67 proportions, which are then sorted and plotted as blue circles. These two plots are overlaid in Figure 6.35, and we can see that the centers of blue circles perfectly align with the red dots. It is not a surprise that the two plots are the same. Recall from Figure 2.12(a), the pairwise angles in the brain artery data set are all greater than 120° . This property simplifies the search significantly: after starting off from the starting point in an orthant, the moving ray

will quickly move towards the data point in the same orthant and eventually go through it, since other data points only project zero onto this ray. Hence, for the brain artery data containing 67 trees, the 67 rays returned by the steepest descent algorithm are basically the rays going through those 67 trees. In other words, the PCA type analysis reveals nothing that is not visible in the raw data. Parallel to the result for HDLSS data in [Hall et al., 2005], it is not surprising that the brain artery data tend to be very scattered.

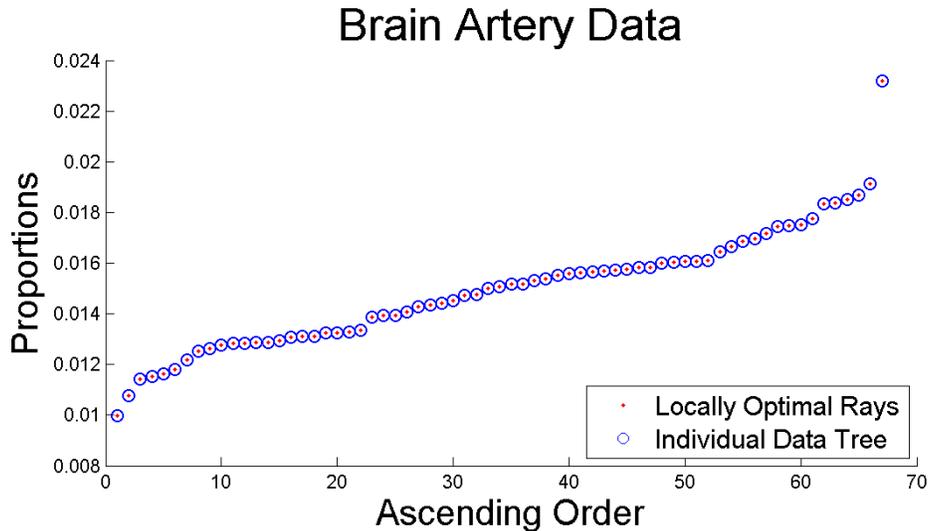


Figure 6.35: Red dots represent the sorted proportions of the total variation in the brain artery data captured by 67 locally optimal rays which are returned by the steepest descent algorithm. Blue circles represent the sorted proportions obtained in the way that the squared norm of each tree is divided by the total squared norm of the brain artery data. The two plots are the same, which indicates that locally optimal rays go through trees in the data set.

To reduce the effect of high dimension, it is interesting to study some data sets in tree space with a relatively smaller number of leaves. Recall in Section 2.5, we introduced a series of landmark-reduced data sets inherited from an earlier version of the brain artery data with 85 trees. In those data sets, the number of leaves are reduced in the sense that the total sum of squared norms is maximized. Then we choose some representative data sets with number of leaves between 5 and 40 with increment 5. Again when a data set has only 5 leaves, it must be in \mathcal{T}_4 . The advantage of \mathcal{T}_4 is that we can represent the whole space by the *Peterson graph* shown in Figure 6.36 and hence see things intuitively. In the Peterson graph, each line segment represents an orthant in \mathcal{T}_4 . For the landmark-reduced data set with only 5 leaves, all trees are distributed in 5 out of 15 orthants, which are represented by 5 dashed line segments in the Peterson graph. Then we run the steepest descent algorithm with starting point at the center ray in each of all 15 orthants. We found that there are only 3 different objective values returned depending on the starting orthant. If the search starts from orthants plotted as red dotted lines, then the algorithm will be terminated immediately since no data point projects positively onto the starting ray. If the search starts from orthants plotted

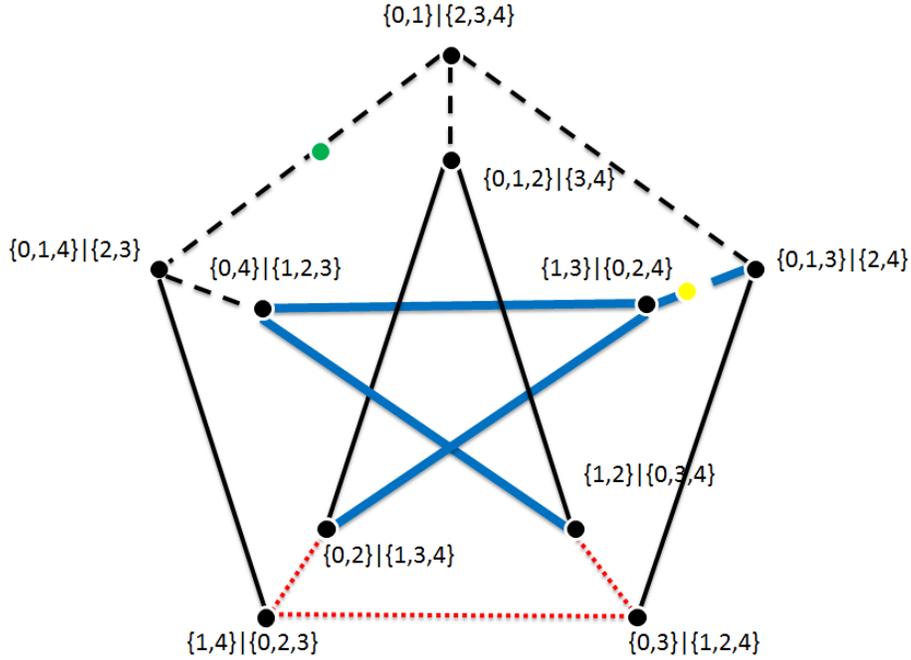


Figure 6.36: Peterson graph illustrating how the steepest descent search algorithm works for the landmark-reduced brain artery data set with only 5 leaves. All 85 data trees are contained in 5 orthants represented by dashed line segments. The search algorithm will end at locally optimal rays in different orthants depending on where the search process starts.

as thick blue lines (both solid and dashed), then the same locally optimal ray will be reached, colored in yellow in the figure, which captures only 3.37% of the total variation. If the search starts from all other orthants colored in black, then the same locally optimal ray will be reached, colored in green in the figure, which captures 85.86% of the total variation. This latter ray is the global optimum. This proportion seems sufficiently good for a single ray, however, this proportion is expected to decrease when the number of leaves gets larger. The bar graph in Figure 6.37 shows the proportion of the total variation captured by the best locally optimal ray across all 8 representative landmark-reduced brain data sets. Overall, the proportion decreases as the number of leaves increases. Although the increments of leaves between consecutive data sets are consistent, the amount of decrease in proportion is not steady. When the number of leaves increases from 5 to 10, the proportion decreases dramatically. When it keeps increasing to 15, we can see another big drop in the proportion. However, when the number of leaves is between 15 and 30, the proportion stays relatively stable with small fluctuations. And we notice another drop when the number of leaves jumps from 30 to 35 and 40. Also if we compare Figure 6.37 with Table 6.5, in terms of the proportion of variation explained by the best ray, the landmark-reduced data set with 5 leaves is somewhere between WF2 and WF10; the data set with 10 leaves is between WF10 and WF40; the data sets with 15, 20, 25, and 30 leaves are similar to WF40; the data sets with 35 and 40 leaves are close to the full brain artery data.

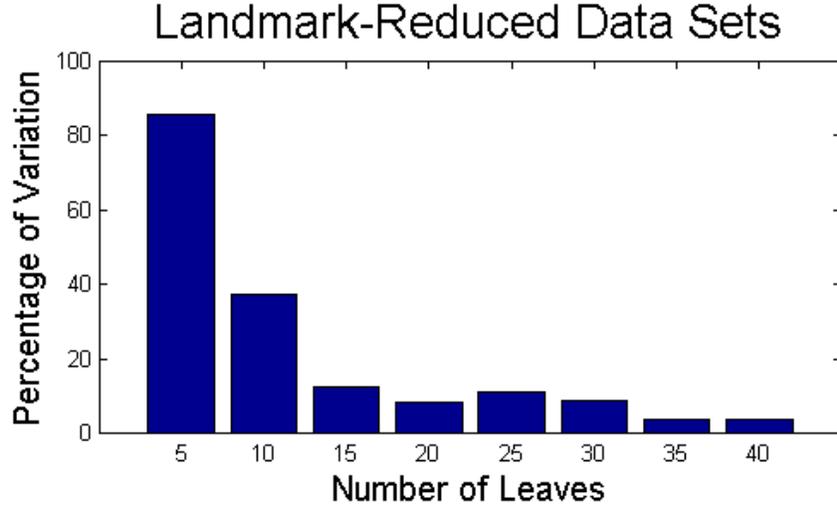


Figure 6.37: Bar graph showing the comparison of proportions of total variation captured by the best locally optimal ray across 8 landmark-reduced brain data sets.

6.3 The First Principal Ray Sets

Up to this point we have concentrated on finding a good *single* ray. However, it is not very effective to use one single ray as a principal component when there is only a small proportion of variation captured, which is caused by the fact that a large portion of data points may not project positively onto the ray. In this section we are going to define the analog of the first principal component as a set of rays, called the *first principal ray set* (1st PR set). Of course, we want the set of rays to have the property that makes them useful for PCA type analysis. A major concern about using a set of rays as a principal component is the uniqueness of projections. If a whole set of rays is considered as one principal component, then the projection of any data tree onto this set of rays should be unique. To address this issue, we need to introduce the concept of antipodality between rays.

Definition 6.3.1. In \mathcal{T}_n , two rays \vec{r}_1 and \vec{r}_2 are *antipodal* if any two trees $T^1 \in \vec{r}_1$ and $T^2 \in \vec{r}_2$ are antipodal.

Now we can define a mutually antipodal ray set (MAR set) as a maximal subset of rays which are pairwise antipodal to each other. One nice property about antipodality is that the projection of any data tree onto a MAR set is always unique. This property is stated formally in the following lemma.

Lemma 6.3.1. *Given an MAR set \mathfrak{R} , any tree T projects positively onto at most one ray of \mathfrak{R} .*

Proof. To prove the lemma, we will use the fact that the angle sum of any triangle in $CAT(0)$ space is no larger than 180° . Suppose that T has positive projections P and Q onto two rays \vec{r}_1 and \vec{r}_2 of \mathfrak{R} , respectively, and let O denote the origin (see Figure 6.38), then $\angle TPO = 90^\circ$ and $\angle TQO = 90^\circ$. By the above fact, we

have $\angle TOP + \angle OTP \leq 90^\circ$ and $\angle TOQ + \angle OTQ \leq 90^\circ$, hence $\angle TOP < 90^\circ$ and $\angle TOQ < 90^\circ$. Therefore, $\angle POQ \leq \angle TOP + \angle TOQ < 180^\circ$, and this contradicts the fact that \vec{r}_1 and \vec{r}_2 are antipodal. \square

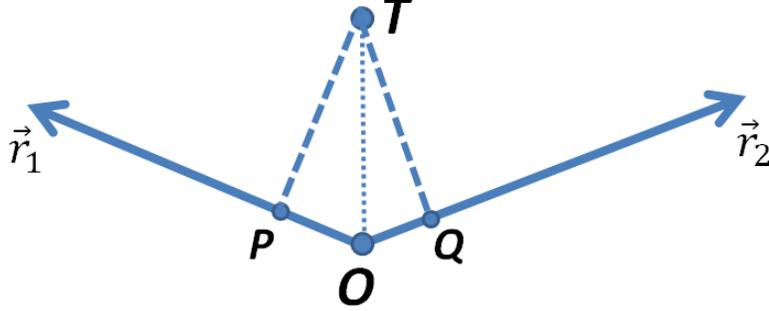


Figure 6.38: This plot shows that there will be a contradiction if tree T projects positively onto two antipodal rays \vec{r}_1 and \vec{r}_2 .

Due to this nice property of MAR sets, it is natural to define the first principal ray set as the best MAR set in terms of capturing data variation.

Definition 6.3.2. Given a set of trees $\mathbf{T} = \{T^1, T^2, \dots, T^r\}$ in \mathcal{T}_n , the *first principal ray set* (1st PR set) is defined as the MAR set which has the largest sum of squared projections from \mathbf{T} .

Now we are ready to modify the optimization formulation defined by (6.12) and (6.13) in order to search for the 1st PR set. Suppose we look for the 1st PR set containing at most m rays, there will be m variable trees, and to enforce the antipodality, a large penalty will be added into the objective function for any pair of rays being not antipodal. The modified formulation is as follows

$$\text{Maximize} \quad \sum_{j=1}^m \sum_{l=1}^r \left[\frac{\left(\sum_{e \in \mathcal{C}^{(j,l)}} |e|_{\tau^j} |e|_{T^{(j,l)}} - \sum_{i=1}^{k^{(j,l)}} \|A_i^{(j,l)}\| \|B_i^{(j,l)}\| \right)^+}{\|\tau^j\|} \right]^2 - P(\tau^1, \dots, \tau^m) \quad (6.14)$$

$$\text{Subject to} \quad \tau^j \in \mathcal{V}_j(\mathbf{T}; \mathcal{O}; \mathcal{A}_j^{\mathbf{T}}, \mathcal{B}_j^{\mathbf{T}}), \quad \text{for all } 1 \leq j \leq m \quad (6.15)$$

where

$$P(\tau^1, \dots, \tau^m) = H \sum_{1 \leq i < j \leq m} \left[\|\tau^i\| + \|\tau^j\| - L(\tau^i, \tau^j) \right]$$

is the penalty function, and H is a large positive constant to prevent any pair of rays from being not antipodal. However, if the number of rays m is large, it can be computationally expensive to apply the steepest descent algorithm, especially when multiple rays approach the orthant boundaries at the same time.

For this reason, we propose two simple heuristics.

The first heuristic is a trial-and-error based greedy algorithm, which is implemented as follows:

1. Obtain a set of n locally optimal rays Ω by solving the single ray problem defined in (6.12) and (6.13) for n various starting points.
2. Sort the rays in Ω in descending order according to the proportion of data variation captured, and label the rays from 1 to n .
3. Get i^{th} candidate MAR set ω_i by starting from ray i in sorted Ω , and keep adding rays with label larger into ω_i than i as long as ω_i is a MAR set.
4. Choose the set with the largest total proportion of data variation among $\omega_1, \dots, \omega_n$ to be the approximate 1st PR set.

This heuristic is easy to implement because it completely avoids solving the penalized optimization problem defined in (6.14) and (6.15). However, one clear drawback is that the accuracy of the heuristic highly depends on the set Ω obtained in step 1.

The second heuristic is a sequentially greedy algorithm, containing:

1. Start with a ray by solving the single ray problem defined in (6.12) and (6.13), label the ray as R^1 .
2. Suppose we already have a set of rays R^1, \dots, R^k , to search for the next ray R^{k+1} , we solve the optimization problem defined in (6.14) and (6.15) with the penalty function being

$$P(\tau^{k+1}) = H \sum_{1 \leq i \leq k} \left[\|\tau^i\| + \|\tau^{k+1}\| - L(\tau^i, \tau^{k+1}) \right], \quad \text{where } \tau^i \in R^i.$$

3. Repeat Step 2 until no more rays can be added to increase the total variation captured.

Although this heuristic solves the penalized optimization problem, it only deals with one variable tree at a time, hence it is easier to implement than the original formulation.

6.4 Higher Order Principal Axis Set

In the last section, we defined the 1st PR set and discussed how to search for a heuristic solution efficiently. As for Euclidean PCA, it is natural to study the higher order PR sets. Recall in Euclidean space, all PCs are orthogonal to each other such that the data variations captured by the PCs do not overlap. In case of PR sets, each of them usually contains multiple rays, so the orthogonality between two PR sets can be defined as the pairwise orthogonality between each ray in one set and each ray in the other set. From the last section, it is seen that significant computational effort is needed to find a heuristic solution for just the 1st PR set. In this section, for simplicity, we will illustrate the idea of higher order PR sets by restricting our candidate ray set to include only orthant axes. In the following part of this section, we will focus on

how to define higher order *principal axis sets* (PA sets). Similar to PR sets, all the axes within a PA set are required to be pairwise antipodal, which leads to the following definition of mutually antipodal axis sets.

Definition 6.4.1. A *mutually antipodal axis set* (MAA set) is an MAR set where each ray is an axis in tree space.

Working with orthant axes has at least four computational advantages. First, there is a finite list of candidate axes, so computational effort to search for PA sets will be much smaller. Second, it is easy to characterize MAA sets as the following lemma shows.

Lemma 6.4.1. *Two axes e_1 and e_2 are antipodal if and only if the corresponding splits σ_1 and σ_2 are incompatible.*

Proof. “ \Rightarrow ” Suppose σ_1 and σ_2 are compatible, then σ_1 and σ_2 can be in the same orthant, which indicates that the angle between e_1 and e_2 is at most 90° . This is a contradiction to the definition of antipodality.

“ \Leftarrow ” Suppose e_1 and e_2 are not antipodal, then the cone path will not be the geodesic. But since the edge sets for both e_1 and e_2 only contain a single edge, at least one support pair will include both σ_1 and σ_2 , hence they are compatible. This is a contradiction to the assumption that σ_1 and σ_2 are incompatible. \square

Thirdly, if two distinct axes are not antipodal, then they share the same orthant, and are therefore orthogonal in the Euclidean sense. Thus we can say that two disjoint MAA sets have the property that every element in one is always either orthogonal or antipodal to the elements in the other. Finally, the projection of a tree onto an axis is simple to obtain. Now we are ready to define PA sets in terms of MAA sets.

Definition 6.4.2. Given a set of trees $\mathbf{T} = \{T^1, T^2, \dots, T^r\}$ in \mathcal{T}_n , the 1^{st} *principal axis set* (1^{st} PA set) is defined as the MAA set which has the largest sum of squared projections from \mathbf{T} among all MAA sets. The i^{th} *principal axis set* (i^{th} PA set) is defined as the MAA set which has the i^{th} largest sum of squared projections from \mathbf{T} among all MAA sets which do not have common axes with the $i - 1$ previously selected PA sets.

Since no tree projects positively onto an axis outside the orthant containing itself, we can simply put together all axes which appear in at least one tree in the data set as the candidate axis pool. To select PA sets in terms of MAA sets, there are two approaches. One is the traditional “forward approach” (using the familiar terminology from variable selection in multiple regression analysis), which is how Euclidean PCA is typically explained. The forward approach chooses the MAA set with the largest sum of squared projections from candidate axes as the 1^{st} PA set, and deletes all the axes included in this set from the candidate pool.

Then it chooses the MAA set from the updated axis pool with the largest sum of squared projections as the 2^{nd} PA set, and deletes all the associated axes from the axis pool. We can iteratively continue this process until the candidate axis pool is empty. The other approach constructs PA sets in the opposite “backward direction” (also using the regression terminology). This idea was studied empirically in [Marron et al., 2010; Jung et al., 2010] and was later investigated more theoretically in [Damon and Marron, 2014] by using a nested sequences of relations. The backward PCA approach was also used in [Jung et al., 2012; Pizer et al., 2013] to construct a nested sequence of subspaces decreasing in dimension. Instead of searching for the 1^{st} PA set, the backward approach starts with the last PA set as the MAA set with the smallest sum of squared projections, and updates the candidate pool. Similar to the forward approach, it will iteratively continue the process until the axis pool is empty. We will analyze both forward and backward approaches in the rest of this section.

Unfortunately, the sub-problem for finding the i^{th} PA set is not an easy task. It has been proven that this sub-problem is equivalent to the problem of finding a maximum weight clique in a weighted graph, which is NP-complete. Nevertheless, there are good, though not polynomial-time, algorithms for this type of problem, (see [Bomze et al., 1999] for surveys of these algorithms). As for finding the 1^{st} PR set, it is more practical to apply some heuristic algorithm, and specifically we will choose some greedy approaches. We propose four greedy algorithms to select a MAA set from a candidate axis pool: two using a forward approach and the other two using a backward approach. Since the objective of a PA set is to capture as much data variation as possible, the first two greedy algorithms use the sum of squared data projections as the criterion to select axes. The forward approach chooses axes with the largest sum of squared projections:

1. Initialize the MAA set as an empty set.
2. Choose the axis with the largest sum of squared projections as the first axis in the MAA set, and update the candidate axis pool by removing the chosen axis.
3. Suppose there are already k axes in the MAA set, the $(k + 1)^{st}$ axis is chosen to be the one being antipodal to the k previously selected axes and also having the largest sum of squared projections. Update the candidate axis pool accordingly.
4. Repeat step 3 until no antipodal axis can be added.

The corresponding backward approach chooses axes with the smallest sum of squared projections, hence we can change the “largest” to “smallest” in the above procedure to obtain the second greedy algorithm.

The above two greedy algorithms use the sum of squared data projections as the axis selection criterion. Two more greedy algorithms can be constructed by using the number of positive data projections as their

criterion. Similarly, the forward and backward algorithms can be simply obtained by replacing “sum of squared projections” by “number of positive data projections” in Step 2 and Step 3 of the corresponding two algorithms discussed previously. After repeatedly applying one of the above four algorithms, we will group all the candidate axes into MAA sets. Then we sort these MAA sets according to the descending order of data variation captured by each MAA set, and define the i^{th} -order PA set as the i^{th} MAA set in the sorted list.

We applied the above four heuristic algorithms to the same five data sets: WF2, WF10, WF40, brain artery data, and uniformly random data. First, we study the 1st PA set, and the following two tables illustrate how the resulting 1st PA set performs by comparing two types of information for 4 different greedy approaches across the 5 data sets. The first basis of comparison is the proportion of data variation captured by the 1st PA set shown in Table 6.6. The second basis is the number of antipodal axes in the 1st PA set shown in Table 6.7. For both tables, the column labels give the names of the 5 data sets and the row labels specify which greedy approach is used: “SS” stands for “Sum of Squared projections”, and “Nr” stands for “Number of positive projections”.

Proportion of Variation	WF2	WF10	WF40	Brain	Random
Forward SS	7.76%	7.47%	3.76%	1.83%	5.27%
Backward SS	7.76%	7.47%	3.61%	1.22%	0.97%
Forward Nr	7.76%	7.47%	3.76%	1.74%	1.00%
Backward Nr	7.76%	7.47%	3.61%	1.13%	1.47%

Table 6.6: Comparison of the proportions of data variation captured by the 1st PA set using 4 greedy approaches across 5 data sets. Suggests that the forward approaches are better than the backward approaches in terms of capturing data variation.

Number of Axes	WF2	WF10	WF40	Brain	Random
Forward SS	1	1	6	15	67
Backward SS	1	1	1	3	33
Forward Nr	1	1	6	13	46
Backward Nr	1	1	1	2	67

Table 6.7: Comparison of the numbers of antipodal axes in the 1st PA set using 4 greedy approaches across 5 data sets explains what we have seen in Table 6.6.

In Table 6.6, for WF2 and WF10, the 1st PA sets can capture between 7% and 8% of the total variation and all four greedy approaches give the same percent. This is because the data trees stay in a relatively narrow cone, which can be seen from the distribution of pairwise angles in Figure 2.12. The 1st PA set

explains between 3% and 4% of the total variation for WF40, and this percent drops to less than 2% for the Brain data. The reason is that these two data sets have more spread. Also for these two data sets, the two forward approaches are slightly better than the two backward ones, since the forward approaches have a better chance to collect the axes with larger projection variation. For the Uniformly Random data, “Forward SS” overwhelms the other 3 approaches by capturing more than 5% of the data variation, compared to less than 1.5% for others. This again can be explained by the pairwise angles summarized in Figure 2.12. We see a large proportion of angles are 180° or near 180° , which means a lot of antipodality occurs among data points. By aggregating antipodality, “Forward SS” includes more axes with larger projection variation than the other 3 approaches. In general, if we compare Table 6.6 with Table 6.5, the 1st PA set does not capture as much data variation as the best single ray for the five data sets, since axes are more restricted than general rays. The exception is the “Forward SS” approach in Uniformly Random data, and the reason is still under investigation.

In Table 6.7, for WF2 and WF10, the data distribute in a relatively small region, so that their 1st PA sets contain only one axis and all four approaches give the same axis. This also explains why the percents are all the same for these two data sets in Table 6.6. For more spread WF40, the two forward approaches give the same set of 6 axes and the two backward approaches give the same single axis set, which is why there are two pairs of equal percents for WF40 in Table 6.6. Note that the single axis found by the backward approaches is included in the axis set found by the forward approaches. Similarly for the Brain data, the forward approaches include more axes, because they search for the 1st PA set at very beginning. Since the Uniformly Random data set has the largest data spread, the antipodality is more frequently found among axes, the 1st PA set contains many axes for all 4 approaches.

Next we want to compare the performance of all the PA sets obtained from the four greedy approaches. For WF2 and WF10, the PA sets coming from the four approaches are very much the same, hence we will focus on the other 3 data sets. Since all the PA sets are already sorted in descending order of the proportion of variation captured, a curve is created by plotting these proportions against the descending ordering for each greedy approach. The curves corresponding to “Forward SS”, “Backward SS”, “Forward Nr”, and “Backward Nr” are colored red, green, blue, and black respectively in Figure 6.39. In plot (a), it is clear that for the Brain Artery data “Forward SS” gives the best result, “Forward Nr” is the second best, then followed by “Backward SS” and “Backward Nr”. In plot (b), for the Uniformly Random data, “Forward SS” overwhelms the other three approaches, and it is hard to rank those three approaches. Plot (c) shows the same comparison for WF40, and it seems that all four approaches produce quite similar results. However, in plot (d), a zoom-in version of (c) reveals that two forward approaches are doing slightly better than two backward approaches. Based on the above comparison, “Forward SS” seems to be the best choice among

the four greedy approaches.

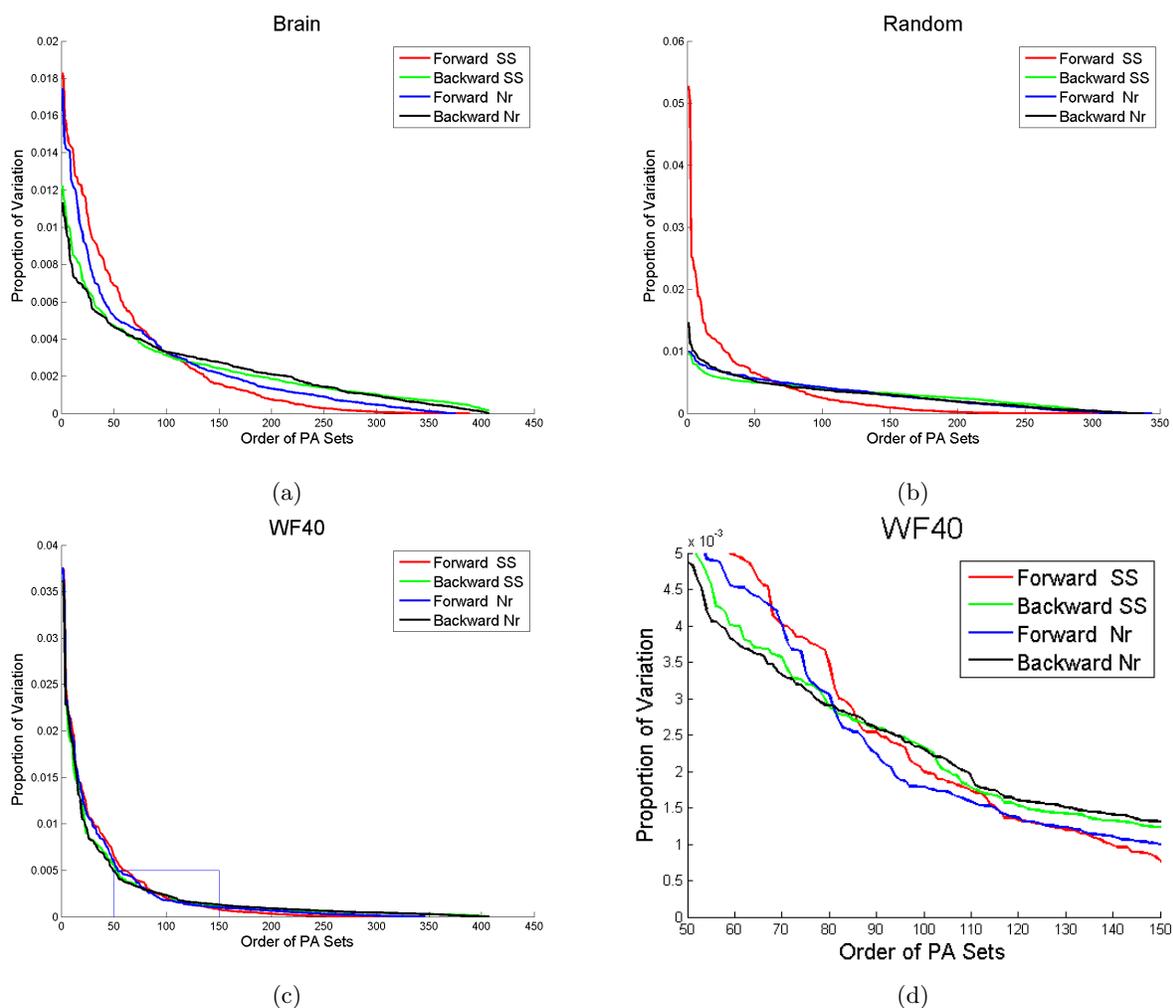


Figure 6.39: Shows the comparison of the four greedy approaches for three data sets: (a)Brain Artery data, (b)Uniformly Random data, (c)WF40, and (d)a zoom-in version of WF40. Suggests that the “Forward SS” is the best among four greedy approaches.

Another criterion for comparing the performance of the above four greedy approaches is the number of PA sets needed to cover at least 80% of the total variation. Using this criterion, this number from all four approaches is calculated for five data sets in Table 6.8. For WF2 and WF10, all four approaches find the same group of PA sets because these PA sets are all single-axis sets. For WF40, the Brain data and the Uniformly Random data, the “Forward SS” approach is better than the other approaches in terms of using fewer PA sets. Also, along each row, we see that the number of PA sets generally gets larger as the spread of the data set increases. This reverses for the Brain data and the Uniformly Random data, since PA sets for the Uniformly Random data contain many more axes than those for the Brain data.

From the discussion in this section, we see that PA sets are computationally more convenient to find

than PR sets. However, they are not as effective as the more general PR sets in terms of capturing data variation. Therefore, an interesting future research direction is to refine the search algorithm for general PR sets in order to better summarize tree structured data.

Number of PA Sets	WF2	WF10	WF40	Brain	Random
Forward SS	34	44	68	105	75
Backward SS	34	44	107	205	181
Forward Nr	34	44	77	146	159
Backward Nr	34	44	113	200	165

Table 6.8: Comparison of the numbers of PA sets needed to capture at least 80% of the data variation by using the 4 greedy approaches across the 5 data sets. Reveals that the “Forward SS” use fewer principal components than the other 3 approaches.

REFERENCES

- Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Upper Saddle River, NJ, 1993. 1.2.1
- A.D. Alexandrov. A theorem on triangles in a metric space and some of its applications. *Trudy Mat. Inst. Steklov*, 38:5–23, 1951. 2.4.1
- B. Aydin, G. Pataki, H. Wang, E. Bullitt, and J.S. Marron. A principal component analysis for trees. *The Annals of Applied Statistics*, 3:1597–1615, 2009a. 1.8
- B. Aydin, G. Pataki, H. Wang, A. Ladha, E. Bullitt, and J.S. Marron. Visualizing the structure of large trees. *Electron. J. Statist*, 5:405–420, 2009b. 2.1.1
- M. Bačák. A novel algorithm for computing the Fréchet mean in Hadamard spaces, 2012. URL <http://arxiv.org/pdf/1210.2145v1.pdf>. 1.4
- Dennis Barden, Huiling Le, and Megan Owen. Central limit theorems for Fréchet means in the space of phylogenetic trees. *Electronic Journal of Probability*, 18:1–25, 2013. 2.5
- J.-P. BARTHÉLÉMY. The median procedure for n -trees. *J. Classification*, 3:329–334, 1986. 1.8
- B. Basrak. Limit theorems for the inductive mean on metric trees. *Journal of Applied Probability*, 47:1136–1149, 2010. 2.5
- Mokhtar S. Bazaraa, John J. Jarvis, and Hanif D. Sherali. *Linear Programming and Network Flows*. John Wiley & Sons, Hoboken, N.J., 4th edition, 2010. 1.2.1
- L. Billera, S. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Advances in applied mathematics*, 27:733–767, 2001. 1.1, 1.2, 1.2.2, 1.3
- Immanuel M. Bomze, Marco Budinich, Panos M. Pardalos, and Marcello Pelillo. The maximum clique problem. *Handbook of Combinatorial Optimization*, pages 1–74, 1999. 6.4
- I. Borg and P. Groenen. *Modern Multidimensional Scaling : Theory and Applications*. Springer, New York, 2nd edition, 2005. 1.7, 1.9, 3.1, 3.1
- Andreas Buja, Deborah F. Swayne, Michael L. Littman, Nathaniel Dean, Heike Hofmann, and Lisha Chen. Data visualization with multidimensional scaling. *Journal of Computational and Graphical Statistics*, 17:444–472, 2008. 3.1
- E Bullitt, D Zeng, G Gerig, S Aylward, S Joshi, JK Smith, W Lin, and MG Ewend. Vessel tortuosity and brain tumor malignancy: A blinded study. *Academic Radiology*, 12:1232–1240, 2005. 2.1
- E. Bullitt, J. Smith, and W. Lin. Designed database of MR brain images of healthy volunteers, 2008. URL <http://www.insight-journal.org/midas/community/view/21>. 2.1
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, 3rd edition, 2009. 1.2.1
- T.F. Cox and M.A.A. Cox. *Multidimensional Scaling*. Chapman & Hall/CRC, Boca Raton, F.L., 2nd edition, 2001. 1.7, 1.9, 3.1, 3.1, 3.1
- James Damon and J.S. Marron. Backwards principal component analysis and principal nested relations. *Journal of Mathematical Imaging and Vision*, 50(1):107–114, 2014. 6.4
- C. Eckart and G. Young. Approximation of one matrix by another of lower rank. *Psychometrika*, 1:211–218, 1936. 3.1

- Aasa Feragen, Megan Owen, Jens Petersen, Mathilde M. W. Wille, Laura H. Thomsen, Asger Dirksen, and Marleen de Bruijne. Tree-space statistics and approximations for large-scale analysis of anatomical trees. In *Information Processing in Medical Imaging*, volume 7917 of *Lecture Notes in Computer Science*, pages 74–85. Springer Berlin Heidelberg, 2013. 4.3
- N.I. Fisher. *Statistical Analysis of Circular Data*. Cambridge University Press, 1995. 1.7
- N.I. Fisher, T. Lewis, and B.J. Embleton. *Statistical Analysis of Spherical Data*. Cambridge University Press, 1993. 1.7
- R.A. Fisher. *The Genetical Theory of Natural Selection*. Clarendon Press, 1st edition, 1930. 2.3
- P.T. Fletcher, C. Lu, S.M. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23:995–1005, 2004. 1.7
- S. Geneser, J. Hinkle, R. Kirby, B. Wang, B. Salter, and S. Joshi. Quantifying variability in radiation dose due to respiratory-induced tumor motion. *Medical Image Analysis*, 15:640–649, 2011. 1.1
- J.C. Gower. Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika*, 53:325–338, 1966. 3.1, 3.1
- J.C. Gower. Euclidean distance geometry. *Mathematical Scientist*, 7:1–14, 1982. 3.4
- J.C. Gower. Properties of Euclidean and non-Euclidean distance matrices. *Linear Algebra and Its Applications*, 67:81–97, 1985. 3.4
- M. Gromov. Hyperbolic groups. In *Essays in Group Theory*, pages 75–263. Springer-Verlag, New York, 1987. 1.3
- Peter Hall, J.S. Marron, and Amnon Neeman. Geometric representation of high dimension, low sample size data. *Journal of the Royal Statistical Society: Series B*, 67:427–444, 2005. 5.1, 6.2
- J. Hein, M. H. Schierup, and C. Wiuf. *Gene Genealogies, Variation and Evolution – A Primer in Coalescent Theory*. Oxford University Press, 2005. 2.3.1
- Thomas Hotz, Stephan Huckemann, Huiling Le, J. S. Marron, Jonathan C. Mattingly, Ezra Miller, James Nolen, Megan Owen, Vic Patrangenaru, and Sean Skwerer. Sticky central limit theorems on open books. *The Annals of Applied Probability*, 23:2238–2258, 2013. 2.5
- S. Huckmann, T. Hotz, and A. Munk. Intrinsic shape analysis: Geodesic PCA for Riemannian manifolds modulo isometric lie group actions. *Statistica Sinica*, 5:1–58, 2010. 1.7
- I. Jolliffe. *Principal Component Analysis*. Wiley Online Library, 2005. 1.6, 1.9
- S. Jung, X. Liu, J.S. Marron, and S. Pizer. Generalized PCA via the backward stepwise approach in image analysis. In *Brain, Body and Machine*, volume 83 of *Advances in Intelligent and Soft Computing*, pages 111–123. Springer, 2010. 6.4
- S. Jung, M. Foskey, and J.S. Marron. Principal arc analysis on direct product manifolds. *The Annals of Applied Statistics*, 5:578–603, 2011. 1.7
- S. Jung, I.L. Dryden, and J.S. Marron. Analysis of principal nested spheres. *Biometrika*, 99:551–568, 2012. 1.7, 6.4
- J.B. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:1–27, 1964a. 3.1
- J.B. Kruskal. Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29:115–129, 1964b. 3.1
- J. Lingoes, E.E. Roskam, and I. Borg. *Geometric Representations of Relational Data*. Mathesis Press, Ann Arbor, M.I., 2nd edition, 1979. 1.7, 1.9

- K.V. Mardia. *Statistics of Directional Data*. Academic Press, New York, 1972. 1.7
- K.V. Mardia, J.T. Kent, and J.M. Bibby. *Multivariate Analysis*. Academic Press, New York, 1979. 1.6
- J. S. Marron and Andrés M. Alonso. Overview of object oriented data analysis. *Biometrical Journal*, 56: 732–753, 2014. 1.1, 1.7, 1.8, 3.2
- J.S. Marron, S. Jung, and I.L. Dryden. Speculation on the generality of the backward stepwise view of PCA. In *Proceedings of the international conference on multimedia information retrieval*, pages 227–230, New York, 2010. ACM. 6.4
- E. Miller, M. Owen, and J.S. Provan. Polyhedral computational geometry for averaging metric phylogenetic trees. *Advances in Applied Mathematics*, 68:51–91, 2015. 1.4, 1.5, 1.5.1, 1.5.2, 1.5.1, 1.5.2, 1.5.3, 1.5, 1.5.4, 6.2
- T.M.W. Nye. Principal components analysis in the space of phylogenetic trees. *The Annals of Statistics*, 39: 2716–2739, 2011. 1.8, 4.3
- I. Oguz, J. Cates, T. Fletcher, R. Whitaker, D. Cool, S. Aylward, and M. Styner. Cortical correspondence using entropy-based particle systems and local features. *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on Biomedical Imaging*, pages 1637–1640, 2008. 1.1, 2.1.2
- M. Owen. Computing geodesic distances in tree space. *SIAM Journal on Discrete Mathematics*, 25:1506–1529, 2011. 1.3
- M. Owen and J.S. Provan. A fast algorithm for computing geodesic distances in tree space. *Computational Biology and Bioinformatics*, 8:2–13, 2009. 1.1, 1.3, 1.3, 3.2, 6.1
- S.M. Pizer, S. Jung, D. Goswami, X. Zhao, R. Chaudhuri, J.N. Damon, S. Huckemann, and J.S. Marron. Nested sphere statistics of skeletal models. In *Innovations for Shape Analysis*, Mathematics and Visualization, pages 93–115. Springer Berlin Heidelberg, 2012. 1.7
- S.M. Pizer, S. Jung, D. Goswami, J. Vicory, X. Zhao, R. Chaudhuri, J.N. Damon, S. Huckemann, and J.S. Marron. Nested sphere statistics of skeletal models. In *Innovations for Shape Analysis*, Mathematics and Visualization, pages 93–115. Springer, 2013. 6.4
- J. Ramsay and B. Silverman. *Applied Functional Data Analysis*. Springer, 2002. 1.1
- J. Ramsay and B. Silverman. *Functional Data Analysis*. Springer, 2005. 1.1
- I.J. Schoenberg. Remarks to Maurice Fréchet’s article Sur la définition axiomatique d’une classe d’espaces distanciés vectoriellement applicable sur l’espace de Hilbert. *Annals of Mathematics*, 38:724–732, 1935. 3.4
- E. Schröder. Vier combinatorsche probleme. *Z. Math. Phys.*, 15:361–376, 1870. 1.2.2
- D. Shen, H. Shen, S. Bhamidi, Y.M. Maldonado, Y. Kim, and J.S. Marron. Functional data analysis of tree data objects. *Journal of Computational and Graphical Statistics*, 2013. 1.8
- R.N. Shepard. The analysis of proximities: multidimensional scaling with an unknown distance function, I and II. *Psychometrika*, 27:125–139 & 219–246, 1962. 3.1
- N. Singh, T. Fletcher, Preston J., L. Ha, R. King, J. Marron, M. Wiener, and S. Joshi. Multivariate statistical analysis of deformation momenta relating anatomical shape to neuropsychological measure. *Medical Image Computing and Computer-Assisted Intervention*, 6363:529–537, 2010. 1.1
- S. Skwerer. *Tree oriented data analysis*. PhD thesis, The University of North Carolina at Chapel Hill, 2014. 1.1, 2.5, 5.1

- Sean Skwerer, Elizabeth Bullitt, Stephan Huckemann, Ezra Miller, Ipek Oguz, Megan Owen, Vic Patrangenaru, Scott Provan, and J. S. Marron. Tree-oriented analysis of brain artery structure. *Journal of Mathematical Imaging and Vision*, 50:126–143, 2014a. 1.1, 6.2
- Sean Skwerer, J. S. Marron, and Scott Provan. Optimization method for Fréchet means in BHV treespace, 2014b. URL <http://arxiv.org/pdf/1411.2923>. 6.2
- K. Sturm. Probability measures on metric space of nonpositive curvature in , 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*, 338:357–390, 2003. 1.4
- W.S. Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17:401–419, 1952. 3.1, 3.1, 3.4
- W.S. Torgerson. *Theory and Methods of Scaling*. John Wiley & Sons, New York, 1958. 3.1, 3.1, 3.4
- Michael W. Trosset and Carey E. Priebe. The out-of-sample problem for classical multidimensional scaling. *Computational Statistics & Data Analysis*, 52:4635–4642, 2008. 3.4, 3.4
- M.P Wand and M.C. Jones. *Kernel Smoothing*. Chapman & Hall/CRC, London;New York, 1995. 2.2
- H. Wang and J.S. Marron. Object oriented data analysis: Sets of trees. *The Annals of Statistics*, 35:1849–1873, 2007. 1.1, 1.1, 1.8
- I.J. Wilson, M.E. Weale, and D.J. Balding. Inferences from DNA data: population histories, evolutionary processes and forensic match probabilities. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 166:155–188, 2003. 2.3.1
- S. Wright. Evolution in mendelian populations. *Genetics*, 16:97–159, 1931. 2.3
- F.W. Young and R.M. Hamer. *Multidimensional Scaling: History, Theory and Application*. L. Erlbaum Associates, Hillsdale, N.J., 1987. 1.7, 1.9
- G. Young and A.S. Householder. Discussion of a set of points in terms of their mutual distances. *Psychometrika*, 3:19–22, 1938. 3.1, 3.4