

Quantification and visualization of variation in anatomical trees

Nina Amenta¹ and Manasi Datar² and Asger Dirksen³
and Marleen de Bruijne^{4,11} and Aasa Feragen⁴ [‡] and Xiaoyin Ge⁵
and Jesper Holst Pedersen⁶ and Marylesa Howard⁷ and Megan Owen⁸ [‡]
and Jens Petersen⁴ and Jie Shi⁹ and Qiuping Xu¹⁰ *

June 2, 2014

Abstract

This paper presents two approaches to quantifying and visualizing variation in datasets of trees. The first approach localizes subtrees in which significant population differences are found through hypothesis testing and sparse classifiers on subtree features. The second approach visualizes the global metric structure of datasets through low-distortion embedding into hyperbolic planes in the style of multidimensional scaling. A case study is made on a dataset of airway trees in relation to Chronic Obstructive Pulmonary Disease.

1 Introduction

Tree-structured data appears in many medical imaging applications, e.g., airway trees [14], blood vessel trees [18], dendrites [32] and galactograms [26]. Typically, these anatomical trees vary both in tree topology and associated branch features such as branch length or shape, and as a result there is no straight-forward way to analyze the trees using standard Euclidean statistics. One way to integrate both tree topology and branch features in a single parametric framework is by modeling trees as residing in a non-linear, non-smooth *tree-space* [7, 14]. The non-linear, non-smooth nature of tree-space creates several problems for data analysis. First, statistics have to be redefined, as the standard statistical procedures such as finding an average or a principal component, or

¹University of California at Davis; ²Scientific Computing and Imaging Institute, University of Utah; ³Lungemedicinsk Afdeling, Gentofte Hospital, Denmark; ⁴Department of Computer Science, University of Copenhagen, Denmark; ⁵Department of Computer Science and Engineering, Ohio State University; ⁶Department of Cardiothoracic Surgery, Rigshospitalet, University of Copenhagen; ⁷National Security Technologies, LLC (A Department of Energy Contractor); ⁸ Department of Mathematics and Computer Science, Lehman College, City University of New York, USA; ⁹ School of Computing, Informatics, and Decision Systems Engineering, Arizona State University; ¹⁰Department of Mathematics, Florida State University; ¹¹Erasmus MC Rotterdam, The Netherlands; [‡]Corresponding authors: Megan Owen (megan.owen@lehman.cuny.edu) and Aasa Feragen (aasa@diku.dk).

performing classification, do not translate directly to the tree-space setting. Second, even if we define classification algorithms in tree-space, we do not know which parts of the anatomical tree are responsible for causing class differences, for example, because each tree-space point represents an entire tree structure. Third, due to the lack of statistical tools such as principal component analysis, it is hard to visualize how distributions of trees vary in tree-space. While recent work has resulted in basic statistical tools [5, 14, 27, 29], the two latter problems are still unsolved. In this paper we investigate two approaches to these two problems: First, we study the influence of local subtrees on the results of hypothesis testing and classification, and the identification of subtrees which are responsible for significant differences between two populations of trees. Second, we use *hyperbolic low-distortion embedding* to visualize the global metric structure of data living in tree-space. As a case study, we demonstrate the use of these techniques on a population of airway trees from a lung cancer screening study.

This paper presents results from the one-week collaboration workshop *Women in Shape: Modeling Boundaries of Objects in 2- and 3-Dimensions* held at the Institute of Pure and Applied Mathematics at UCLA, July 15-19 2013. At this workshop, most of the authors of this paper spent a week working together on two projects related to quantifying and visualizing variance in populations of trees, which are described in Sections 3 and 4, respectively.

1.1 Tree-space

A tree-space is any geometric space in which points represent trees. The tree-space used in this paper, described in [14, 15], is a generalization of the phylogenetic tree-space proposed by Billera et al. [7]. This tree-space, denoted \mathcal{T}_n , contains all rooted trees with n labeled leaves with vertices of degree at least 3, where the n leaf labels are given by a fixed set of cardinality n . In this paper, the root of a tree is not considered to be among the leaves. Furthermore, for any tree in this tree-space, each edge has a k -dimensional vector associated with it. An example of such an edge vector is a non-negative real number representing the edge length (i.e. $k = 1$); a second example is the vector of l 3-dimensional landmark points sampled along a branch centerline, giving $k = 3 \cdot l$. For each edge, the landmark points are translated so that the edge starts at the origin. We refer to the latter edge vector as the *shape* of the edge. The trees in such a tree-space can, for instance, be used to model airways in the lung. In this space, we will use edge shape with $l = 5$ to describe edges unless otherwise stated.

We now give a description of our tree-space, \mathcal{T}_n , for a fixed n , which is illustrated for $n = 4$ in Figure 1. All the trees in \mathcal{T}_n with the same *tree topology*, or branching order, form a lower-dimension Euclidean subspace in the tree-space. The dimension of this subspace is mk , where m is the number of edges in the tree topology and k is the dimension of the vector associated with each tree edge. Each tree edge in the tree topology is put into correspondence with k of the subspace's dimensions, and a particular tree with that topology can be written as a km -dimensional vector in that subspace, with the coordinates being the consecutive k -dimensional edge vectors. That is, if a tree has edges e_1, e_2, \dots, e_m , with corresponding edge vectors $\ell_1 = (\ell_1^1, \ell_2^1, \dots, \ell_5^1), \dots, \ell_m = (\ell_m^1, \ell_m^2, \dots, \ell_m^5)$, then that tree corresponds to the point (ℓ_1, \dots, ℓ_m) . All trees within the subspace must map their k -dimensional edge vectors to the km -dimensional vector in

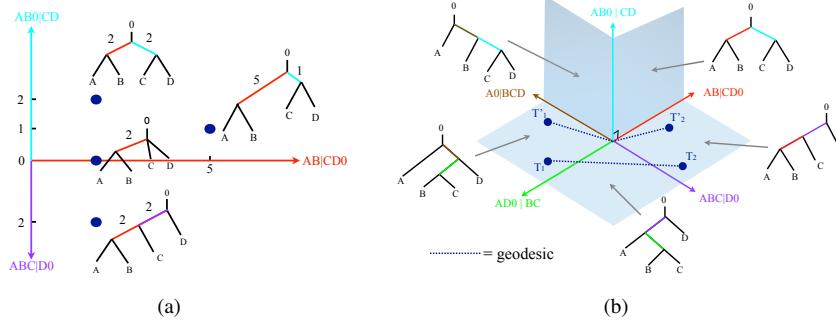


Figure 1: (a) Two adjacent quadrants in the tree space, \mathcal{T}_4 , of trees with 4 leaves and edge vectors of 1 dimension. Here, the edge vector, or length, is restricted to being non-negative and the pendant edges are ignored, so the Euclidean subspaces are represented as quadrants. A quadrant contains all trees with a given topology, and each tree with that topology is represented by the coordinates corresponding to its internal edge lengths. (b) Representation of 5 of the 15 quadrants in the tree space \mathcal{T}_4 , with edge vectors of 1 dimension.

the same order, but what this order is does not matter.

The Euclidean subspaces for each tree topology are glued together in the following way. Consider a tree containing the edges e_1, \dots, e_m , where each edge can be identified by the unique partition of the leaves it makes when it is removed from the tree (i.e. removing the edge forms a forest of two trees, each of whose leaves, including the root, forms one half of the partition). Let exactly one of the edges, say e_1 , have an all 0 edge vector. Then this tree lies on the boundary of the Euclidean subspace corresponding to its tree topology, and furthermore, it actually lies in a lower dimensional Euclidean subspace E corresponding to trees with only the edges e_2, \dots, e_m in their topologies. This lower dimensional subspace E is also on the boundary of two other Euclidean subspaces, and we identify all such common subspaces in all the Euclidean subspaces corresponding to tree topologies to form \mathcal{T}_n . See [15] and [7] for a more detailed description of the tree space.

The metric on \mathcal{T}_n is induced by the Euclidean metric on each of its constituent subspaces. Specifically, the distance between two trees with the same topology is the Euclidean distance between the two points representing those trees in the subspace for that tree topology. The distance between two trees with different topologies is the length of the shortest path joining their points in tree space. Such a path will consist of a sequence of line segments, each contained in exactly one of the subspaces, and thus the path length is just the sum of the Euclidean lengths of each segment.

Most importantly for this paper, tree space is a non-positively curved metric space [7], which implies that there is a unique shortest path within the space between any two trees, called the *geodesic*. The *geodesic distance* between two trees is the length of the geodesic between them, and it can be computed in polynomial time [30]. Certain statistics, such as means and first principal components, can also be computed on trees in this space [5, 14, 27, 29].

While sections of tree-space are identical to higher dimensional Euclidean spaces,

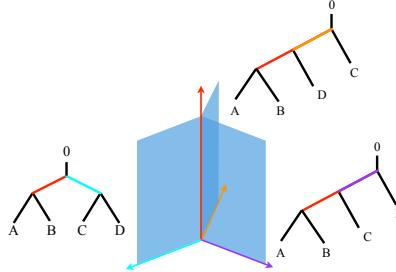


Figure 2: An occurrence of a half open book with 3 2-dimensional sheets in tree-space.

tree-space itself is not a manifold. In particular, it has several singularities, which have infinite negative curvature. One such singularity is at the origin, which corresponds to the tree which has all 0 edge lengths or vectors. A simpler model of this singularity in \mathcal{T}_4 is a *corner*: five Euclidean quadrants, glued together around a single origin, see Fig. 1(b). Singular points in tree-space also occur where the the higher dimensional Euclidean subspaces join together to form a space that locally resembles an *open book*. An open book is a set of Euclidean half planes, or *sheets*, which are identified along their boundary hyperplanes, which form the *spine*, see Fig. 2.

The Fréchet mean in tree-space

In Euclidean space, there are a number of equivalent definitions of the mean. Some of these definitions cannot be carried over to tree-space, while those that can be carried over are no longer equivalent. One of the definitions of mean in Euclidean space is the Fréchet mean, or barycenter, which minimizes the sum of square distances to the input set. If $\{T_1, T_2, \dots, T_r\}$ are a set of input trees, then their Fréchet mean in tree-space is the tree t which minimizes $\sum_{i=1}^r d(t, T_i)^2$, where d is the geodesic distance. The Fréchet mean was introduced for tree-space independently by [5] and [27], both of whom also gave an algorithm to approximate it based on a Law of Large Numbers holding in non-positively curved spaces [35].

2 Related work

This paper studies two problems related to understanding variance in datasets of trees: (i) detection of local subtree differences, and (ii) visualization of global population-level geometry.

Local significant differences

Many data types represent entities which can be decomposed into parts or regions. Examples are graph-structured data [13, 23], anatomical data which can be segmented into different organs [17, 20] or even single anatomical organs where additional spatial information is relevant; for instance, in the framework of shape analysis [9, 10] where

local analysis is made on correspondence points on biomedical shape surfaces. A typical problem when studying such data is *interpretability*: A classifier will often only predict a certain diagnosis or class, but in order to understand the cause of the result (and, e.g. in diagnostic settings, react on it) one also desires to know which parts of the collection caused a certain classification outcome.

While a large body of work has been done on classifying structured data, less is known about how to identify which parts of a structure are relevant for the classification problem. Most such work has been done in settings where there is a correspondence between the parts constituting the data object: In analysis of brain connectivity [19, 23], one usually has a matching between the nodes in the dataset, while in voxel-based morphometry [3] or shape analysis [10], registration is used to match different images to a template. A popular approach to such problems is *structured sparsity* [4, 22, 23], which detects discriminative substructures in data described by fixed-length Euclidean vectors with a known underlying structure relating the vector coordinates. However, anatomical trees usually cannot be described by fixed-length vectors without discarding parts of the tree. Thus, these methods are not directly applicable.

Low-distortion embeddings

The standard technique for visualising population structure in high-dimensional or non-Euclidean datasets is to extract the pairwise distances between data points, and then use *multidimensional scaling* (MDS), which attempts to embed the points into a lower dimensional Euclidean space such that the given distances between the points are preserved. This is expressed mathematically as minimizing the sum of the differences between original and embedded pairwise distances [8]. In a sequence of work [2, 21, 28] Amenta, St. John et al. investigate visualization of sets of phylogenetic, or evolutionary, trees using multidimensional scaling. In this work, inter-tree distances are given by the Robinson-Foulds distance [34], which only measures topological differences in the trees. More recently, Wilgenbusch et al. [41] compare several non-linear versions of MDS on phylogenetic trees, and find that a metric that places less weight on large distances gives more meaningful visualizations. Chakerian and Holmes [11] use MDS with the geodesic distance between trees [7]. A different approach is that of Sundberg et al. [36], who visualize phylogenetic trees by projecting them onto a hypersphere; this approach does not consider branch lengths, only tree topology.

All of these methods approach visualization through embedding into a Euclidean space in a low-distortion way. However, embedding spaces need not be restricted to only Euclidean spaces. For instance, low-distortion embedding of a general metric into a tree has been considered for various measures of distortion [1, 6]. Low-distortion embedding of general metrics into hyperbolic spaces has also been considered by Walter et al. [39, 40] and Cvetkovski and Crovella [12]. In this paper, we use hyperbolic MDS for more truthful visualizations of tree variation.

3 Quantification and visualization of local tree-shape differences

While previous work [14] developed methods for finding significant differences between populations of trees, this work did not address the question of where these changes came from. In this section we investigate different methods for detecting *where* in a tree significant differences appear. In Sec. 3.1 we perform hypothesis testing on nested subtrees in order to detect how significant changes take place in particular subtrees. In Sec. 3.2 we develop a structured sparsity framework which takes advantage of the tree-space geometry in order to handle the fact that subtrees have variable topological structure. In both of these sections, we obtain results on which subtrees induce significant differences. A disadvantage of the methods developed in Sections 3.1 and 3.2 is that they do not take correlation between different subtrees into account. In Sec. 3.3 we therefore develop a method that allows us to study how subtree differences correlate with each other.

Case study. We apply the developed methods to a case study of airway trees from subjects with and without Chronic Obstructive Pulmonary Disease (COPD). The 600 airway trees are from randomly selected subjects from the Danish Lung Cancer Screening Trial [31], of which 300 were diagnosed with COPD at scan time and 300 were symptom free. The hypothesis testing and classification experiments performed in this chapter all have the common goal of separating the class of COPD patients from the class of symptom free subjects.

The airway trees were extracted from low-dose (120 kV and 40 mAs) pulmonary CT scans. To extract the tree, the airway lumen surface was extracted from the images using the locally optimal path approach of [24] and then refined using the optimal surface approach of [33]. Afterwards centerlines were computed by front propagation within the refined lumen surface as described in [25]. The resulting centerlines were disconnected in bifurcation regions and so Dijkstra's algorithm was used to connect them along shortest paths within an inverted distance transform of the refined lumen surface. These centerlines were then represented by 6 equidistantly sampled landmark points. The airway trees were normalized by patient height as an affine scaling parameter.

Airway trees are somewhat regular in the sense that some of the branches have anatomical names and can be found in most human lungs. The subtrees rooted at these branches feed different subdivisions of the lung at different hierarchical levels, as schematically illustrated in Fig. 3. The Trachea is the root branch that feeds both lungs. The left and right main bronchi (LMB, RMB) feed the left and right lungs. The left upper lobe and lower lobe branches (LUL and LLB) feed the left upper and lower lobes. The lower lobe splits into the branches L7-L10. The left upper lobe branches into two subsections; the first feeds the three segments L1, L2 and L3, and the branch feeding all of these is called L1+2+3. The second subsection feeds the segments L4 and L5, and their parent is called L4+5. The right lung is subdivided into the upper lobe, fed by the right upper lobe branch (RUL), and the middle- and lower lobes, both fed by the bronchus intermedius (BronchInt). The middle lobe consists of the segments R4-R5, fed by the parent R4+5, and the lower lobe is fed by the right lower lobe

branch (RLL), and splits into the segments R7-R10.

Due to variation in airway tree topology, these branches are not always all present. In our dataset, however, which has been automatically labeled using the algorithm presented in [15], the following branches are consistently present:

$$\text{Trachea, LMB, RMB, LUL, RUL, L1+2+3, LLB, BronchInt, and RLL.} \quad (1)$$

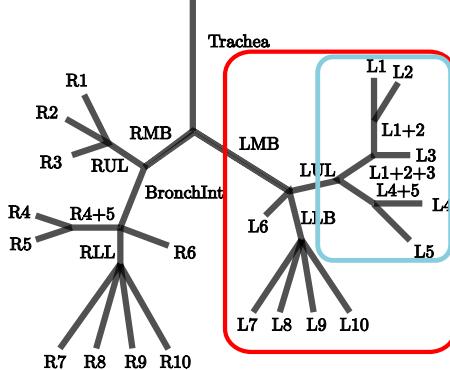


Figure 3: Airway tree (black) and sub-trees (LMB, red) and (LUL, blue).

LABEL	P-VALUE	P-VALUE
	mean	variance
full	0.0010	0.0060
RMB	0.0020	0.0939
RUL	0.2298	0.1668
BrInt	0.0050	0.1249
RLL	0.0300	0.0959
LMB	0.0859	0.0210
LUL	0.0320	0.0390
L123	0.0260	0.0410
LLB	0.5524	0.1588

Table 1: Case study: Group comparison showing results of permutation tests on subtrees of the full airway trees rooted at the branches listed in (1). The permutation test compares the populations of airway trees from COPD patients and symptom free subjects.

3.1 Permutation tests for subtree statistics

In this section we perform subtree hypothesis testing using the tree-shape permutation tests for equality of means and variances developed in [14] on the nested subtrees defined by the subtree root branches (1). These tests are standard permutation tests which, for samples G_1 and G_2 drawn from two different classes of trees (in our experiments: healthy subjects and COPD patients), use test statistics used for means and variances between classes defined as

$$t_m = d(\mu(G_1), \mu(G_2)) \text{ and } t_v = |v(G_1) - v(G_2)|,$$

respectively, where $\mu(G_i)$ is the Fréchet mean of the trees in the i^{th} class as defined on p. 4, and $v(G_i)$ denotes the variance of the i^{th} class $\frac{1}{(N-1)} \sum_{t \in G_i} d^2(t, \mu(G_i))$.

Under the null hypothesis, namely that there is no difference between the two classes, the samples G_1 and G_2 are drawn from the same distribution on \mathcal{T}_n , and randomly permuting the elements of G_1 and G_2 should not affect the value of the test statistic t_* .

Form the two-class data set $G = G_1 \cup G_2$ and consider partitions of G into subsets of size $N_1 = |G_1|$ and $N_2 = |G_2|$. Due to the size of G we cannot check all possible permutations, but instead compute the test statistics $(t_*)_m$ for means and variances

Table 2: Case study benchmark: Permutation tests for shape differences in individual branches between populations of airway trees from healthy individuals and COPD patients. P-values below a threshold of 0.05 are shown in bold.

LABEL	P-VALUE		LABEL	P-VALUE	
	mean	variance		mean	variance
RMB	0.163	0.621	LMB	0.020	0.786
RUL	0.134	0.416	LUL	0.297	0.118
R1	0.410	0.363	L1	0.163	0.391
R2	0.116	0.255	L2	0.324	0.017
R3	0.329	0.854	L3	0.968	0.800
BronchInt	0.001	0.764	L45	0.078	0.312
R4	0.134	0.190	L4	0.372	0.570
R5	0.027	0.175	L5	0.023	0.050
R6	0.992	0.135	L6	0.260	0.833
RLL	0.001	0.865	LLB	0.177	0.112
R7	0.058	0.325	L7	0.496	0.611
R8	0.014	0.207	L8	0.466	0.900
R9	0.037	0.127	L9	0.146	0.026
R10	0.308	0.652	L10	0.855	0.162
L123	0.393	0.361			

for the new subsets, $m = 1, \dots, M$, for M random partitions of G into sets of size N_1 and N_2 . Comparing the $(t_*)_m$ to the original statistic value t_* for the samples G_1 and G_2 , we obtain a p -value approximating the probability of observing t_* under the null hypothesis:

$$p = \frac{1 + \sum_{(t_*)_m \geq t_*, m \in \{1, \dots, M\}} 1}{M + 1},$$

where the additional 1 is added to avoid $p = 0$.

Permutation tests for the two statistics were performed on each subtree with $M = 1000$ permutations. The results are summarized in Table 1, and show significant differences in several subtrees. These subtrees are identified by their root branches, which are illustrated schematically in Fig. 3. In comparison, the same hypothesis test was made on the individual branches (1) along with the segment branches R1-R10, L1-L10, with results shown in Table 2. The tests on individually identified branches show relatively fewer significant differences between the two populations, emphasizing a need for considering the airway subtrees as entities rather than collections of independent branches.

This suggests that the permutation test described here can be applied to study local group differences between subtrees in a hierarchical manner.

3.2 Subtree classification

In the previous section we saw how hypothesis testing on subtrees allowed us to learn about which subtrees differed significantly between two populations of trees. While

Table 3: Mean \pm standard deviation of COPD classification accuracy using branch length (left) and branch shape (right).

Method	Accuracy length	Accuracy shape
LDA	0.56 ± 0.06	0.52 ± 0.06
QDA	0.55 ± 0.05	N/A
Mahalanobis	0.54 ± 0.05	N/A
k NN	0.53 ± 0.06	0.53 ± 0.06
SVM	0.56 ± 0.06	0.56 ± 0.06

significant differences are interesting in their own right, we are often particularly interested in finding *predictive* differences. In particular, we want to find subtrees such that restricting prediction to these subtrees results in good predictive performance, giving interpretable classifiers in the sense that we can detect *which* tree changes are predictive.

3.2.1 Classification on known branches

A straight-forward approach to tree classification and identification of discriminative substructures of trees is to use standard classification methods on vectors whose coordinates correspond to a fixed set of identified branches. In our case, these branches will be identified by their anatomical names, which create a natural matching between the branches of different trees. Classification of such vectors can return information about which branches are more discriminative, because classifiers such as the support vector machine (SVM), include coordinate weights that intuitively correspond to the relevance of the individual coordinate feature for the classification problem. This method is simple, but has the disadvantage that it can only use branches that are present in every single tree in the dataset. This method will form a baseline to which our proposed methods are compared.

Since we could only use branches that were present in every single dataset tree, we used the list of branches (1) along with the leaves {R1-R10, L1-L10}, which are guaranteed to be present. We performed classification with linear discriminant analysis (LDA), quadratic discriminant analysis (QDA), Mahalanobis distance, k -nearest neighbor (k NN) using $k = 5$, and support vector machine (SVM) using 10 repetitions of 10-cross validation. The corresponding classification accuracies are reported in Table 3. Note that the QDA and Mahalanobis distance are missing for the shape branch features; this is because these both require a positive definite covariance matrix, for which the data set was too small for the higher-dimensional shape vectors.

While the mean classification accuracy is above chance for all classifiers, none of them are significantly above chance. A common heuristic to find features which are important in classification is to study the magnitudes of the coordinates of the weight vector produced by the SVM algorithm, as shown in Table 4. Note that in addition to the classification accuracy being very low, the weights of high magnitude are scattered

around the airway tree, not adding much in terms of interpretation.

The poor performance of classifiers on the set of all branches could be explained by the fact that many branches are highly correlated. In such cases, the weights might not carry much information. The poor classification accuracy may be explained by the dimensionality of the data. This motivates our search for a more predictive and interpretable classification algorithm by including subtree information in the classifier.

Table 4: The mean and standard deviations of the SVM weight vectors on the COPD/healthy classification. The largest weight vectors as well as those falling within one standard deviation are highlighted.

Branch	SVM weight	Branch	SVM weight	Branch	SVM weight	Branch	SVM weight
RMB	2.0 ± 1.45	R5	2.8 ± 1.6	LMB	-2.6 ± 1.6	L4	-1.4 ± 1.2
RUL	4.3 ± 1.6	R6	-1.7 ± 1.6	LUL	1.7 ± 2.1	L5	2.3 ± 1.6
R1	-0.7 ± 1.2	L6	1.5 ± 1.6	L123	-0.8 ± 1.4	RLL	2.6 ± 1.6
R2	4.9 ± 1.2	R7	1.4 ± 1.5	L1	-1.6 ± 1.2	LLB	-3.6 ± 1.5
R3	-3.6 ± 1.4	R8	3.3 ± 2.2	L2	-3.2 ± 1.3	L7	-2.8 ± 1.3
BrInt	-5.0 ± 1.6	R9	7.3 ± 1.5	L3	1.4 ± 1.3	L8	2.2 ± 1.5
R4	-0.4 ± 1.6	R10	2.9 ± 1.6	L45	3.6 ± 1.8	L9	3.2 ± 1.5
						L10	5.2 ± 1.5

3.2.2 Structured sparse feature selection through regularized logistic regression on subtree similarity

Under a hypothesis that significant differences are found in local subtrees, we incorporate local subtree structure into classification through a sparse classifier taking subtree similarity as input. Logistic regression measures the relationship between a categorical dependent variable (class label) and one or more independent variables by using conditional probabilities as predicted values of the dependent variable. The L_1 regularized logistic regression, or the so-called sparse logistic regression [38], regularizes the classifier by forcing the weight vector of the classifier to have a small number of nonzero values. This results in implicit feature selection and robustness to noise, as well as interpretability through the selected subtree features. In addition to its solid theoretical foundation, this model is computationally efficient [16].

Consider a set of n training examples $T = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ from which a tree classifier $y = f(x)$ will be learned. The i^{th} tree is represented by a D -dimensional feature vector $x_i = [x_{i1}, x_{i2}, \dots, x_{id}]^T$ where $x_{ij} = d(S_j^i, \mu_j)$, where S_j^i is the j^{th} subtree of the i^{th} tree, μ_j is the Fréchet mean tree of all j^{th} subtrees in the training set, and d denotes geodesic distance between trees. The D subtrees are rooted at the branches listed in (1). The values $y_i \in \{0, 1\}$ indicate the class labels of the two groups, modeling the conditional probability distribution of the class label y given a feature vector x as:

$$p(y = 1|x, \beta) = \frac{1}{1 + \exp(-\beta^T x)},$$

where $\beta \in \mathbb{R}^D$ are the parameters of logistic model. The estimation of the parameters β is done by likelihood maximization, equivalent to minimizing the negative log-

Table 5: Results of the structured logistic classifier with lasso and elastic net regularization. Classification results are averaged over 10 randomized folds, and significant features are those where the distances to both class means were kept as features in the classifier in all folds.

Value of α	Classification result	Significant features
1 (lasso)	$65 \pm 2.7\%$	full
0.75	$64.3 \pm 2.1\%$	RMB, full
0.5	$62.5 \pm 2.4\%$	RMB, BronchInt, full
0.25	$62.5 \pm 2.3\%$	RMB, LMB, LLB, BronchInt, RLL, full

likelihood

$$\hat{\beta} = \arg \min \sum_{i=1}^n -\log p(y_i|x_i, \beta).$$

Applying a sparse regularizer we obtain feature selection, interpretability and reduced overfitting. This is done by adding a so-called *lasso* regularization term:

$$\hat{\beta} = \arg \min \sum_{i=1}^n -\log p(y_i|x_i, \beta) + \lambda \|\beta\|_1,$$

where λ is a parameter controlling the sparsity of β , in the sense that fewer nonzero coefficients of β remain as λ increases. The optimal λ is chosen to optimize classification accuracy by 5-fold cross validation.

The nested subtrees used will be correlated by definition. One way of handling this is by adding an l_2 norm regularization term as well, known as *elastic net regularization*. This leads to an objective function

$$\hat{\beta} = \arg \min \sum_{i=1}^n -\log p(y_i|x_i, \beta) + \lambda (\alpha \|\beta\|_1 + \frac{1-\alpha}{2} \|\beta\|_2^2).$$

The results of the sparse classifiers for different values of α are shown in Table 5. Note that the classification performance is significantly better than that of the standard classifiers on identified branches seen in Table 3. Moreover, note the discriminative subtrees selected by the classifier. The fact that the lasso regularizer results in only the full tree being selected is most likely a result of the correlation between subtrees. As an l_2 regularizer is also added, we obtain a tradeoff between sparsity and including correlated subtrees.

3.3 Subtree variance correlation testing

As the localized methods use features extracted from nested subtrees, we expect a high degree of correlation between overlapping subtrees. Most of the previously described methods do not take such correlations into account, and this may, in particular, be a problem for the interpretability through selected features. Moreover, it is interesting to know whether variation in non-overlapping subtrees is correlated. In this section we provide a method for testing the correlation between variance in the subtrees. We use the notation from Section 3.2.2.

To compare the variance between subtrees in the same airway tree, we use the distance from the j -th subtree S_j^i in tree i to the population mean μ_j of all j -th subtrees in some class as a measure of the amount of variation in that subtree S_j^i . We compute the correlation between these distances, represented by the random variable X and Y , for each pair of subtrees (j, k) , and measure whether deviation from the mean subtree μ_j is correlated with deviation from the mean subtree μ_k . To measure the correlation, we use Pearson's sample correlation coefficient,

$$r_{xy} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{(n-1)s_x s_y},$$

where \bar{x} and \bar{y} are the same means of the two distance variables X and Y , and s_x and s_y are the sample standard deviations of the X and Y . This is equivalent to the sample covariance divided by the sample standard deviations.

The results of applying this test to the airway data set is shown in Fig. 4. Many of the subtree pairs exhibiting correlation in their variance are nested, as expected. For example, variation is very correlated between the three subtrees RMB, Bronchint, and RLL, where RLL is a subtree of Bronchint, which is itself a subtree of RMB. Similarly, there is a high correlation in the variation between the nested pairs LMB and LUL, and LMB and LLB. However, not all nested subtrees are highly correlated. In particular, RUL is also a subtree of RMB, but variation in it is not very correlated with that in RMB – in fact, variance in RUL is not strongly correlated with any other subtree. The trees are not separated by class as there was no significant difference in the behavior of the two classes.

4 Visualization of NPC information spaces via Low-Distortion Embedding into the Hyperbolic Disc

High-dimensional data is often embedded into lower dimensional spaces in order to improve the efficiency of computations, or, with a two- or three-dimensional embedding space, for visualization. Besides being high-dimensional and stratified, the tree space \mathcal{T}_n has negative curvature. While it is exactly this property that gives it unique geodesics, it also means that the number of trees within a given neighborhood can grow exponentially with the radius of the neighborhood. Our hypothesis was that embedding point sets in \mathcal{T}_n into lower-dimensional hyperbolic space, which also has negative curvature, would allow embeddings with lower distortion and/or lower total error. We explore the use of two different visualization techniques for general metric distance matrices, Multi-Dimensional Scaling and Isomap.

Multidimensional Scaling (MDS) is a classical approach that maps the original metric dataset to a target Euclidean space, usually of low dimension. It transforms the input metric distance matrix into a set of coordinate positions for the data points - in our case, each tree is a data point - such that the Euclidean distances between the coordinates approximate the input distances as well as possible. Using the new embedded coordinates, one can visualize dataset structure through the embedded dataset, where inter-point distances have been preserved as well as possible. Different definitions for

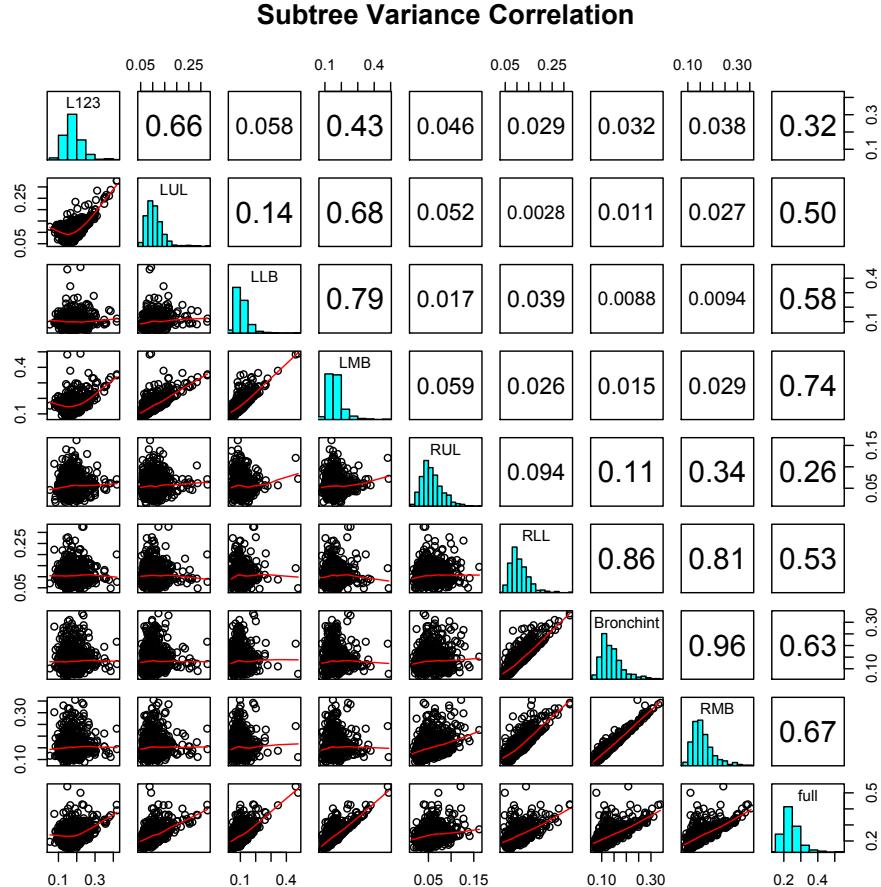


Figure 4: The plots in the lower triangle compare the distances between subtrees in the same tree to their corresponding mean subtree of healthy patients. The plots along the diagonal are the histograms of these distances when the subtree is fixed. The upper diagonal gives the correlation of the distances plotted in the corresponding plot in the lower triangle.

what it means to preserve the distances “as well as possible” produce different computational problems. Principle Components Analysis (PCA) can be seen as a version of MDS, for which the problem has a global solution, but other definitions of optimal distance preservation often lead to better visualizations. These versions are all non-linear, so both the optimization criterion and the method of optimization can lead to different results.

IsoMap [37] is a more recent method intended for points which lie on a lower-dimensional surface in the high-dimensional space. It begins by constructing a neighborhood graph connecting nearby points in the input space. Then, using this graph, it

approximates geodesic distances on the surface. Finally it applies MDS to the matrix of geodesic distances.

While the standard approach in both of these methods is to use a Euclidean target space for the embedding, in the past decades, hyperbolic multidimensional scaling has also been proposed. In a nutshell, the original Euclidean distance in the target space is replaced by the hyperbolic distance:

$$d(z_i, z_j) = 2 \tanh^{-1} \frac{|z_i - z_j|}{|1 - z_i \bar{z}_j|}, \quad (2)$$

where z_i and z_j denote two points in the target space. The modification in the distance metric makes the computation of gradients non-trivial. Our goal was to explore the question of whether hyperbolic space would be a more successful target space for the visualization of distributions of trees, since tree space and hyperbolic space are both non-positive curved.

Recently, Cvetkovski and Crovella [12] introduced a method MDS-PD (metric multidimensional scaling algorithm using the Poincaré disk model) which is based on a steepest descent method with hyperbolic line search. We adapted this software for our experiments with hyperbolic space, and we review the method here; more details can be found in [12]. Complex coordinates are used to present the points of the hyperbolic plane, making the Poincaré disk model a subset of the complex plane \mathbf{C} : $\mathbf{D} = \{z \in \mathbf{C} \mid |z| < 1\}$. The objective function to be minimized is the total embedding error

$$E = c \sum_{j=1}^n \sum_{k=j+1}^n c_{jk} (d_{jk} - \delta_{jk})^2.$$

where c and c_{jk} are constants, d_{jk} is the hyperbolic distance between points z_j and z_k (Equation 2), and δ_{jk} denotes the dissimilarity/distance between points z_j and z_k in the input dissimilarity/distance matrix. More specifically, we use the Sammon Stress Criterion, in which c and c_{jk} are fixed based on δ_{jk} as follows:

$$E = \frac{1}{\sum_{j=1}^n \sum_{k=j+1}^n \delta_{jk}} \sum_{j=1}^n \sum_{k=j+1}^n \frac{(d_{jk} - \delta_{jk})^2}{\delta_{jk}} \quad (3)$$

This criterion does not favor preserving large distances over small ones. The algorithm starts with a set of random points in the Poincaré disk. In each iteration, it moves each of the points along the gradient direction of the energy function shown in Equation 3 with a Möbius transform until one of the stopping tolerances is met or the maximum iteration number is reached.

4.1 Experiments on real and synthetic data

While much of tree-space looks locally like an Euclidean space, there are two local features which are decidedly not Euclidean: corners and open books. A corner is point concentration of negative curvature (see Fig. 1(b)), while an open book is a set of Euclidean half-space attached together along their axes, or "spine" (see Fig. 2). These two

Table 6: Multiplicative distortion of the embeddings.

	MDS	Isomap	HMDS	HIsomap
CORNER	3644.2	1224.9	64.2	119.7
3SHEETS_2D	71.4	98.9	76.3	44.0
3SHEETS_3D	44.0	189.54	54.1	68.2
5SHEETS_2D	551.9	567.4	87.8	76.1
5SHEETS_3D	2097.8	470.5	393.6	123.4
COPD_250	253.9	952.3	62.0	64.3

features, as well as the high dimension of the local Euclidean space, are the sources of error for the low-distortion embedding. We generate synthetic datasets that isolate the two features to determine how hyperbolic MDS (HMDS) and hyperbolic isomap (HIsomap) treat them. We compare the results both qualitatively and quantitatively with embeddings done with classical MDS and isomap. More specifically, the datasets are CORNER_2D, in which 50 points are generated in each of the 5 quadrants around the origin; 3SHEETS_2D, in which 50 points are generated in each of 3 2-dimensional sheets; 3SHEETS_3D, in which 50 points are generated in each of 3 3-dimensional sheets; 5SHEETS_2D, in which 50 points are generated in each of 5 2-dimensional sheets; 5SHEETS_3D, in which 50 points are in each of 5 3-dimensional sheets; and COPD, in which the lung airway trees of 125 healthy patients and 125 patients with COPD are randomly selected. Within each sheet or quadrant, the 50 points were generated by sampling from a symmetric normal distribution in the underlying Euclidean space that is centered at the origin.

The multiplicative distortion for each embedding approach is summarized in Table 6. The multiplicative distortion for a single distance between two points in the dataset is *original_distance/embedded_distance*. The distortion for the whole dataset is *max_distortion/min_distortion*, where *max_distortion* is the maximum distortion of any two points and *min_distortion* is the minimum distortion for any two points. HMDS and HIsomap perform the best for almost all of the datasets. The embedded visualizations and the histograms for each dataset are found in Figures 7 and 8.

Qualitatively, for CORNER_2D, while all methods were able to group the points in the same quadrant, only HMDS preserved the symmetry of the quadrants; we assume that this is because in the hyperbolic plane it is possible for five lines to meet at right angles, which is not possible in Euclidean space.

For the two-dimensional tree spaces 3SHEETS_2D and 5SHEETS_2D, all methods also grouped the points by their respective sheets. The two Euclidean methods overlaid all but two of the sheets, while the two hyperbolic methods kept the sheets distinct, particularly in 3SHEETS_2D, better representing the true geometry. Despite increasing the dimension only by one, for 3SHEETS_3D it was much harder for the methods to separate the distinct sheets. While MDS performed the best quantitatively, this was not the case qualitatively, where the two hyperbolic methods gave better sheet separation. All methods had trouble representing the 5SHEETS_3D dataset. Finally, no methods were able to visualize class separation in the COPD dataset.

The ideal histogram would place all of the distances in the column correspond-

ing to zero error. Although for the three-dimensional tree space 3SHEETS_3D and 5SHEETS_3D none of the visualizations provides much insight, the histograms show that again HMDS gives the most accurate reduction to two dimensions.

5 Discussion and conclusion

We have considered two different approaches for quantifying and visualizing variance in datasets of trees. In Sec. 3 the dataset trees were divided into nested subtrees, in order to quantify the contribution of different subtrees in distinguishing two populations of trees through either hypothesis testing or classification. These approaches were applied to populations of airway trees from COPD patients and healthy individuals, where the most discriminative subtrees were extracted for the different tasks. In Sec. 4 visualization of population structure for datasets of trees was studied through multidimensional scaling and isomap in a hyperbolic disc as opposed to in the Euclidean plane. The choice of a hyperbolic visualization space was motivated by the fact that tree-space itself has singular points which are hyperbolic, and it thus seems likely that a hyperbolic visualization space can give a more truthful rendering of the structure of the population of trees than a Euclidean space. We demonstrate a quantitative and visual improvement in dataset visualization on a set of synthetic datasets sampled from singular spaces representing the types of singularities found in tree-space, as well as on a set of airway trees.

These approaches supply a new set of tools, and give insight into new potential solutions, for analysis of tree-structured data. Future work includes development of structured sparsity methods using subtrees where the correlation between different subtrees is explicitly taken into account, as well as low-distortion embedding into more complex non-Euclidean visualization spaces whose geometry is similar to that of tree-space.

Acknowledgements

This work was supported by NSF IIS 111766; MO was supported through a Fields-Ontario Postdoctoral Fellowship; AF was supported by the Danish Council for Independent Research | Technology and Production. MH works for National Security Technologies, LLC, under Contract No. DE-AC52-06NA25946 with the U.S. Department of Energy/National Nuclear Security Administration, DOE/NV/25946--2016.

We thank the organizers, Kathryn Leonard and Luminita Vese, and the sponsors of the collaboration workshop Women in Shape(WiSh): Modeling Boundaries of Objects in 2- and 3- Dimensions that was held at the Institute for Pure and Applied Mathematics (IPAM) at UCLA July 15-19 2013, where the collaboration leading to this paper was established.

Table 7: The embedded datasets. For the CORNER, 3SHEETS_2D, 3SHEETS_3D, 5SHEETS_2D, and 5SHEETS_3D dataset embeddings, points have the same color if they are located in the same quadrant or sheet. For the COPD dataset embeddings, the class of healthy patients is colored in red, and the class of patients with COPD are colored in blue.

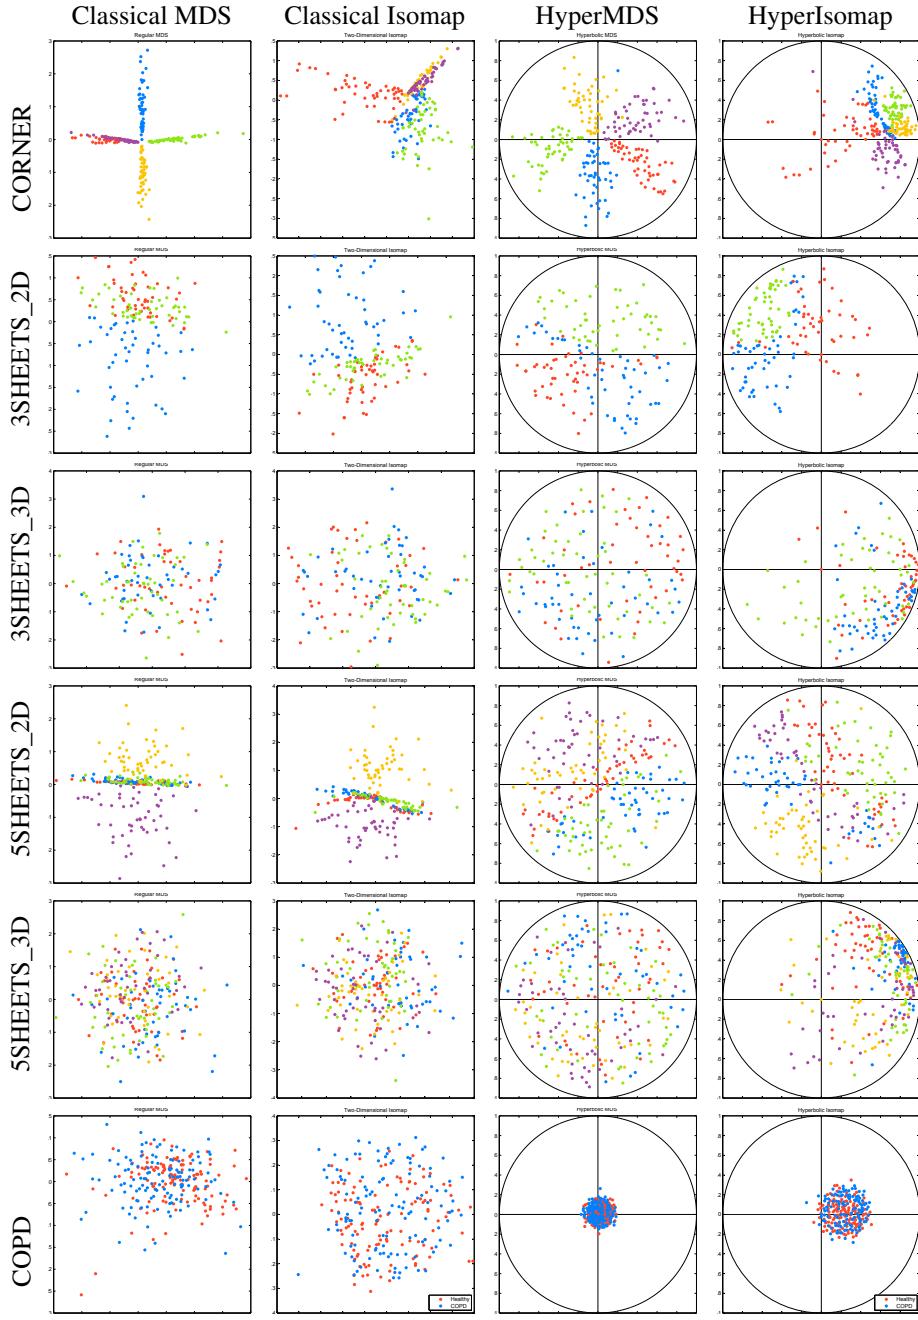
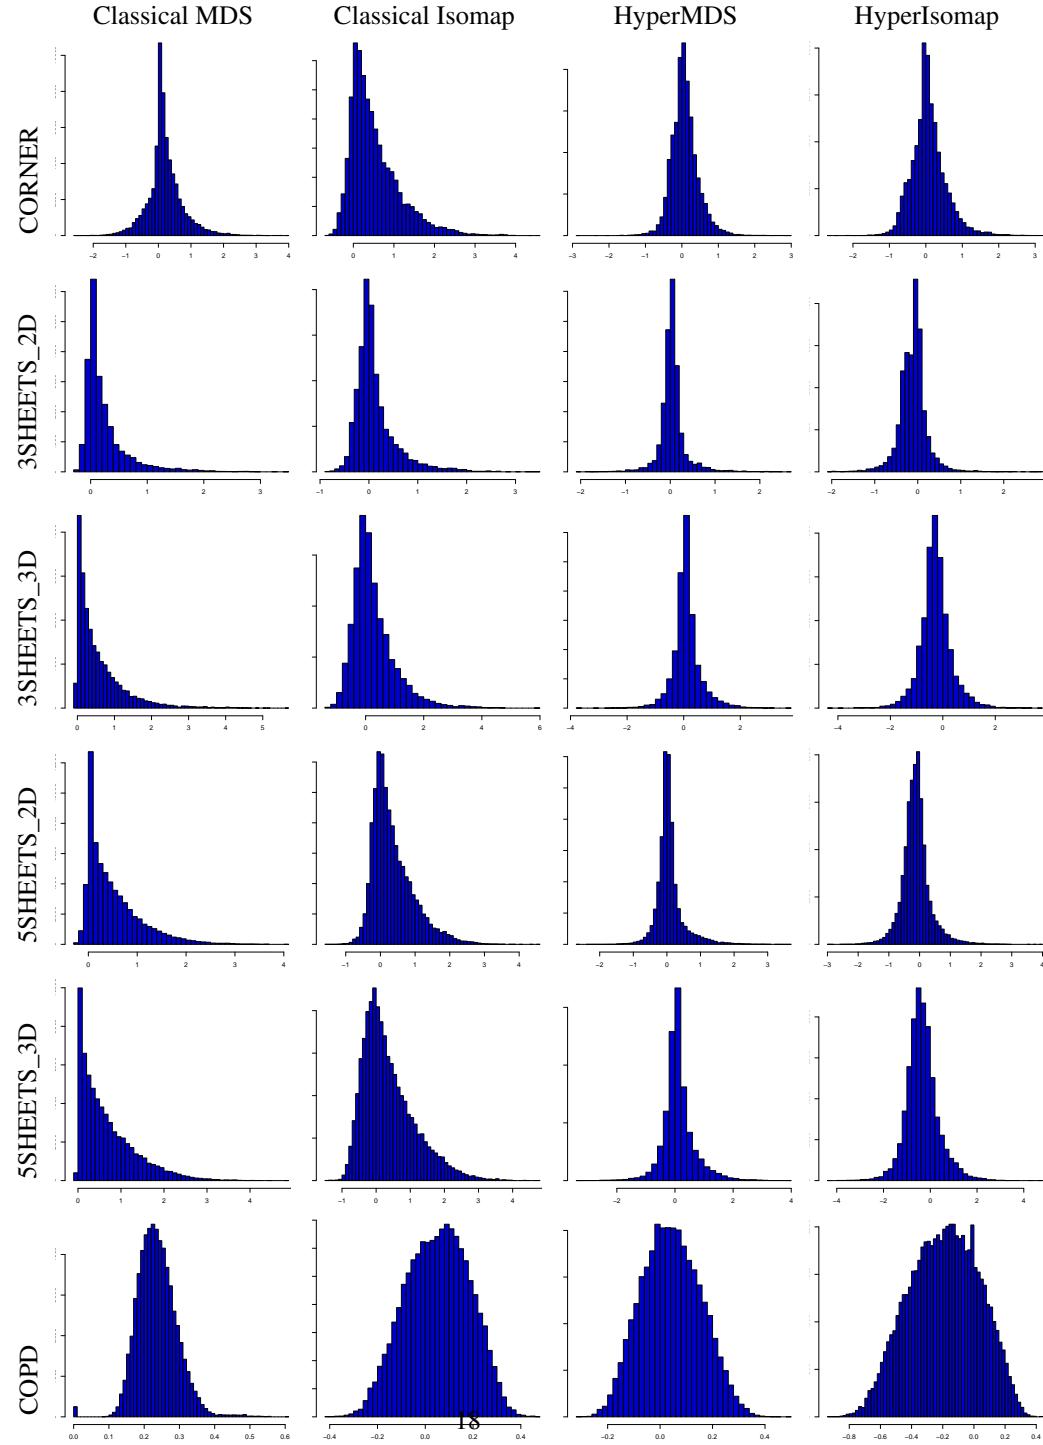


Table 8: The error histograms of the embedded datasets in Figure 7. Every pair of points is binned according to the error in the embedding, which is the difference between the original distance between the pair of points, and the distance between them in the embedding.



References

- [1] R. Agarwala, V. Bafna, M. Farach, M. Paterson, and M. Thorup. On the approximability of numerical taxonomy (fitting distances by tree metrics). *SIAM J. Comput.*, 28(3):1073–1085, 1999.
- [2] N. Amenta and J. Klingner. Case study: Visualizing sets of evolutionary trees. In *IEEE Symposium on Information Visualization, 2002. INFOVIS 2002.*, pages 71–74. IEEE, 2002.
- [3] J. Ashburner and K.J. Friston. Voxel-based morphometry—the methods. *Neuroimage*, 11:805–821, 2000.
- [4] C.-A. Azencott, D. Grimm, M. Sugiyama, Y. Kawahara, and K.M. Borgwardt. Efficient network-guided multi-locus association mapping with graph cuts. *Bioinformatics*, 29(13):171–179, 2013.
- [5] M. Bacak. A novel algorithm for computing the Fréchet mean in Hadamard spaces. *Preprint*, <http://arxiv.org/abs/1210.2145>, 2012.
- [6] M. Badoiu, P. Indyk, and A. Sidiropoulos. Approximation algorithms for embedding general metrics into trees. In *SODA*, pages 512–521, 2007.
- [7] L. J. Billera, S. P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.*, 27(4):733–767, 2001.
- [8] I. Borg and P. Groenen. *Modern multidimensional scaling*. Springer-Verlag, Heidelberg, 1997.
- [9] J.R. Busch, P.A. Ferrari, G.A. Flesia, R. Fraiman, and S. Grynberg. Testing statistical hypothesis on Random Trees, 2006.
- [10] J. Cates, P.T. Fletcher, M. Styner, M. Shenton, and R. Whitaker. Shape modeling and analysis with entropy-based particle systems. In Nico Karssemeijer and Boudewijn Lelieveldt, editors, *Information Processing in Medical Imaging*, volume 4584 of *Lecture Notes in Computer Science*, pages 333–345. 2007.
- [11] J. Chakerian and S. Holmes. Computational tools for evaluating phylogenetic and hierarchical clustering trees. *Journal of Computational and Graphical Statistics*, 21(3):581–599, 2012.
- [12] A. Cvetkovski and M. Crovella. Multidimensional scaling in the poincaré disk. *CoRR*, abs/1105.5332, 2011.
- [13] A. Feragen, N Kasenburg, Petersen J., M. de Bruijne, and K. Borgwardt. Scalable kernels for graphs with continuous attributes. In *NIPS*, 2013.
- [14] A. Feragen, M. Owen, J. Petersen, M.M.W. Wille, L.H. Thomsen, A. Dirksen, and M. de Bruijne. Tree-space statistics and approximations for large-scale analysis of anatomical trees. In *IPMI*, pages 74–85, 2013.

- [15] A. Feragen, J. Petersen, M. Owen, P. Lo, L.H. Thomsen, M.M.W. Wille, A. Dirksen, and M. de Bruijne. A hierarchical scheme for geodesic anatomical labeling of airway trees. In *MICCAI (3)*, pages 147–155, 2012.
- [16] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010.
- [17] E. Gasparovic. The blum medial linking structure for multi-region analysis, 2012.
- [18] A.J. Geers, I. Larabide, A.G. Radaelli, H. Bogunovic, M. Kim, H.A.F. Gratama van Andel, C.B. Majoie, E. VanBavel, and A.F. Frangi. Patient-specific computational hemodynamics of intracranial aneurysms from 3D rotational angiography and CT angiography: An in vivo reproducibility study. *American Journal of Neuroradiology*, 32(3):581–586, 2011.
- [19] Yasser Ghanbari and et al. Connectivity subnetwork learning for pathology and developmental variations. In *MICCAI (1)*, pages 90–97, 2013.
- [20] K. Gorczowski, M. Styner, J-Y. Jeong, J.S. Marron, J. Piven, H. Cody Hazlett, S.M. Pizer, and G. Gerig. Multi-object analysis of volume, pose, and shape using statistical discrimination. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(4):652–661, 2010.
- [21] D.M. Hillis, T.A. Heath, and K. St John. Analysis and visualization of tree space. *Systematic Biology*, 54(3):471–482, 2005.
- [22] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, pages 417–424, 2009.
- [23] R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, E. Eger, F. Bach, and B. Thirion. Multiscale mining of fmri data with hierarchical structured sparsity. *SIAM J. Imaging Sciences*, 5(3):835–856, 2012.
- [24] Pechin Lo, Jon Sporring, Jesper Johannes Holst Pedersen, and Marleen Bruijne. Airway tree extraction with locally optimal paths. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2009*, volume 5762 of *Lecture Notes in Computer Science*, pages 51–58. 2009.
- [25] Pechin Lo, Bram van Ginneken, Joseph M. Reinhardt, T. Yavarna, P. A. de Jong, B. Irving, C. Fetita, M. Ortner, R. Pinho, J. Sijbers, M. Feuerstein, A. Fabijanska, C. Bauer, R. Beichel, C. S. Mendoza, R. Wiemker, J. Lee, A. P. Reeves, S. Born, O. Weinheimer, E. M. van Rikxoort, J. Tschirren, K. Mori, B. Odry, D. P. Naidich, I. Hartmann, E. A. Hoffman, M. Prokop, J. H. Pedersen, and M. de Bruijne. Extraction of airways from CT (EXACT'09). *IEEE Transactions on Medical Imaging*, 31(11):2093–2107, 2012.

- [26] V. Megalooikonomou, M. Barnathan, D. Kontos, P. Bakic, and A. Maidment. A representation and classification scheme for tree-like structures in medical images: Analyzing the branching pattern of ductal trees in x-ray galactograms. *TMI*, 28(4):487–493, 2009.
- [27] E. Miller, M. Owen, and J. S. Provan. Averaging metric phylogenetic trees. *Preprint*, <http://arxiv.org/abs/1211.7046>, 2012.
- [28] I. Montealegre and K. St John. Visualizing restricted landscapes of phylogenetic trees. In *Proc. of the European Conference for Computational Biology (ECCB 03)*. Citeseer, 2002.
- [29] T. M. W. Nye. Principal components analysis in the space of phylogenetic trees. *Ann. Statist.*, 39(5):2716–2739, 2011.
- [30] M. Owen and J.S. Provan. A fast algorithm for computing geodesic distances in tree space. *ACM/IEEE Trans. Comp. Biol. Bioinf.*, 8:2–13, 2011.
- [31] J. Pedersen, H. Ashraf, A. Dirksen, K. Bach, H. Hansen, P. Toennesen, H. Thorsen, J. Brodersen, B. Skov, M. Døssing, J. Mortensen, K. Richter, P. Clementsen, and N. Seersholm. The Danish randomized lung cancer CT screening trial - overall design and results of the prevalence round. *J Thorac Oncol*, 4(5):608–614, May 2009.
- [32] C. Perez-Cruz, M. Simon, B. Czéh, G. Flügge, and E. Fuchs. Hemispheric differences in basilar dendrites and spines of pyramidal neurons in the rat prelimbic cortex: activity- and stress-induced changes. *Eur J Neurosci*, 2009.
- [33] J. Petersen, M. Nielsen, P. Lo, L. H. Nordenmark, J. H. Pedersen, M. M. W. Wille, A. Dirksen, and M. de Bruijne. Optimal surface segmentation using flow lines to quantify airway abnormalities in chronic obstructive pulmonary disease. *Medical Image Analysis*, 2014 (In Press).
- [34] D.F. Robinson and L.R. Foulds. Comparison of phylogenetic trees. *Math. Biosciences*, 53:131–147, 1981.
- [35] K.-T. Sturm. Probability measures on metric spaces of nonpositive curvature. volume 338 of *Contemp. Math.*, pages 357–390. 2003.
- [36] K. Sundberg, M. Clement, and Q. Snell. On the use of cartographic projections in visualizing phylogenetic tree space. *Algorithms for Molecular Biology*, 5(26), 2010. <http://www.almob.org/content/5/1/26>.
- [37] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [38] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1996.

- [39] J. Walter. H-MDS: a new approach for interactive visualization with multidimensional scaling in the hyperbolic space. *Information systems*, 29(4):273–292, 2004.
- [40] J. Walter and H. Ritter. On interactive visualization of high-dimensional data using the hyperbolic plane. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 123–132. ACM, 2002.
- [41] J.C. Wilgenbusch, W. Huang, and K.A. Gallivan. The evaluation of dimensionality reduction methods to characterize phylogenetic tree-space, 2010. Poster at Evolution 2010. Available at <http://www.math.fsu.edu/whuang2/>.