Geometries on spaces of treelike shapes

A. Feragen[a], F. Lauze[a], P. Lo[a], M. de Bruijne[a,b], M. Nielsen[a]

{aasa, francois, pechin, marleen, madsn}@diku.dk,
a) eScience Center, Dept. of Computer Science, University of Copenhagen, Denmark
b) Biomedical Imaging Group Rotterdam, Depts of Radiology & Medical Informatics, Erasmus MC – University Medical Center Rotterdam, The Netherlands

Abstract. In order to develop statistical methods for shapes with a tree-structure, we construct a shape space framework for treelike shapes and study metrics on the shape space. The shape space has singularities, which correspond to topological transitions in the represented trees. We study two closely related metrics, TED and QED. The QED is a quotient euclidean distance arising from the new shape space formulation, while TED is essentially the classical tree edit distance. Using Gromov's metric geometry we gain new insight into the geometries defined by TED and QED. In particular, we show that the new metric QED has nice geometric properties which facilitate statistical analysis, such as existence and local uniqueness of geodesics and averages. TED, on the other hand, has algorithmic advantages, while it does not share the geometric strongpoints of QED. We provide a theoretical framework as well as computational results such as matching of airway trees from pulmonary CT scans and geodesics between synthetic data trees illustrating the dynamic and geometric properties of the QED metric.

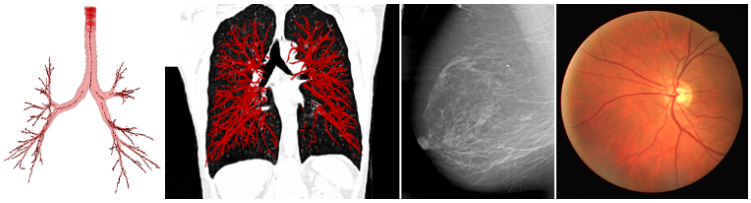

Fig. 1. Treelike structures found in airways and vessels in lungs [1], in breast cancer vascularization, and in retinal blood vessels [2].

1 Introduction

Trees are fundamental structures in nature, where they describe delivery systems for water, blood and other fluids, or different types of skeletal structures. Trees and graphs also track evolution, for instance in genetics [3]. In medical image analysis, treelike shapes appear when studying structures such as vessels [4, 5] or airways [6–8], see Fig. 1. In more general imaging problems, they appear in the form of shock graphs [9] describing general shapes.

An efficient model for tree-shape would have applications such as shape recognition and classification. In medical imaging, this could give new insight into

anatomical structure and biomarkers for diseases that affect shape – e.g. COPD affects the edge shapes and topological structures of airway trees [10].

We want to do statistical analysis on sets of treelike shapes, and in order to meaningfully define averages and modes of variation in analogy with the classical PCA analysis, we need at least local existence and uniqueness of geodesics and averages. Immediate questions are then: How do we model treelike structures and their variation? Can we encode global, tree-topological structure as well as local edgewise geometry in the geometry of a single shape space? Do geodesics exist, and are they unique? Do the dynamics of geodesics reflect natural shape variation?

Here we provide the theoretical foundations needed to answer these questions, accompanied by a first implementation. We define shape spaces of ordered and unordered treelike shapes, where transitions in internal tree-topological structure are found at shape space singularities in the form of self-intersections. This is a general idea, which can be adapted to other situations where shapes with varying topology are studied – e.g. graphs or $3D$ shapes described by medial structures.

The paper is organized as follows: In Section 1.1 we give a brief overview of related work. In Section 2 we define the treespace. Using Gromov's approach to metric geometry [11] we gain insight into the geometric properties of two different metrics; one which is essentially tree edit distance (TED) and one which is a quotient euclidean distance (QED). We pay particular respect to the properties of geodesics and averages, which are essential for statistical shape analysis. In Section 4 we discuss how to overcome the computational complexity of both metrics. In particular, the computational cost can be drastically reduced by interpreting order as an edge semi-labeling, which can often be obtained through geometric or anatomic restrictions. In Section 5 we discuss a simple QED implementation, and in Section 6 we illustrate the properties of QED through computed geodesics and matchings for synthetic planar trees and $3D$ pulmonary airway trees.

The paper contains mathematical results, but proofs are omitted due to length constraints.

1.1 Related work

Metrics on sets of trees have been studied by different research communities in the past 20 years. The best-known approach in the computer vision community is perhaps TED, as used for instance by Sebastian et al [9] for comparing shapes via their shock graph representations. TED performs well for tasks such as matching and registration [9]. Our goal is, however, to adapt the classical shape statistics to treelike shapes. The TED metric will nearly always have infinitely many geodesics between two given trees, and thus it is no longer sufficient, since it becomes hard to meaningfully define and compute an average shape or find modes of variation.

Another approach to defining a metric on treespace is that of Wang and Marron et al [12]. They define a metric on attributed trees as well as a "median-mean" tree, which is not unique, and a version of PCA which finds modes of variation in terms of so called *treelines*, encoding the maximum amount of structural and attributal variation. Wang and Marron analyze datasets consisting of

brain blood vessels, which are trees with few, long branches. The metric is not suitable for studying large trees with a lot of topological variation and noise such as airways, as it punishes structural changes much harder than shape variation with constant topological structure.

A rather different approach is that of Jain and Obermayer [13], who define metrics on graphs. Here, graphs are modeled as incidence matrices, and the space of graphs, which is a quotient by the group of vertex relabelings, is given a metric inherited from Euclidean space. Means are computed using Lipschitz analysis, giving fast computations; however the model is rigid when it comes to modeling topological changes, which is essential for our purposes.

We have previously [14] studied geodesics between small planar embedded trees in the same type of singular shape space as studied here. These geodesics are fundamental in the sense that they represent the possible structural changes found locally even in larger trees.

2 Geometry of treespace

Before defining a treespace and giving it a geometric structure, let us discuss which properties are desirable. In order to compute average trees and analyze variation in datasets, we require existence and uniqueness properties for geodesics. When geodesics exist, we want the topological structure of the intermediate trees to reflect the resemblance in structure of the trees being compared – in particular, *a geodesic passing through the trivial one-vertex tree should indicate that the trees being compared are maximally different*. Perhaps more important, we would like to compare trees where edge matching is inconsistent with tree topology as in Fig. 2(a); specifically, we would like to find geodesic deformations in which the tree topology changes when we have such edge matchings, for instance as in Fig. 2(b).

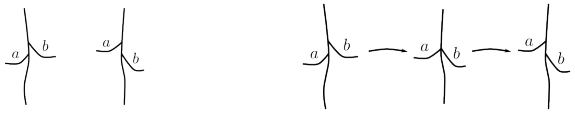

(a) Edge matching (b) Geodesic candidate

Fig. 2. A good metric must handle edge matchings which are inconsistent with tree topology.

2.1 Representation of trees

We represent any treelike (pre-)shape as a pair $(\mathscr{T}, x)$ consisting of a rooted, planar binary tree $\mathscr{T} = (V, E, r)$ with edge attributes describing edge shape. Here, $\mathscr{T}$ describes the tree topology, and the the attributes describe the edge geometry. The attributes are represented by a map $x\colon E \to A$ or, equivalently, by a point $x \in A^E$, where A is the attribute space. Tree-shapes which are *not* binary are represented by the maximal binary tree in a very natural way by allowing constant edges, represented by the zero scalar or vector attribute.

Endowing an internal edge with a zero attribute corresponds to collapsing that edge, see Fig. 4(a), and hence *an arbitrary attributed tree can be represented as an attributed binary tree.* The intuitive idea is illustrated in Fig. 3.

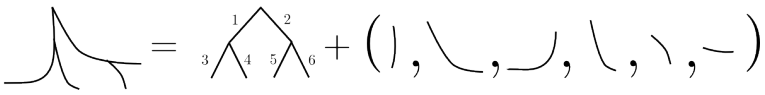

Fig. 3. Treelike shapes are encoded by an ordered binary tree and a set of attributes describing edge shape.

The attributes could be edge length, landmark points or edge parametrizations, for instance. In this work, we use trees whose attributes are open curves translated to start at the origin, described by a fixed number of landmark points. Thus, thorughout the paper, the attribute space is $(\mathbb{R}^d)^n$ where $d = 2$ or 3 and n is the number of landmark points per edge. Collapsed edges are represented by a sequence of origin points.

We need to represent trees of different sizes and structures in a unified way in order to compare them. As a first step, we describe all shapes using the *same* binary tree $\mathscr{T}$ to encode the topology of *all* the tree-shapes in the treespace – by choosing a sufficiently large binary tree we can represent all the trees in our dataset by filling out with collapsed edges. We call $\mathscr{T}$ *the maximal binary tree.*

When given a left-right order on every pair of children in a binary tree, there is an implicitly defined embedding of the tree in the plane. Hence we shall use the terms "planar tree" and "ordered tree" interchangingly, depending on whether we are interested in the planar properties or the order of the tree. We initially define our metric on the set of ordered binary trees; later we use it to compute distances between unordered trees by considering all possible orders. In Section 4 we discuss how to handle the computational challenges of the unordered case.

Given an ordered maximal binary tree $\mathscr{T}_n$ of depth n with edges E, any attributed tree T is represented by a point in $X = \prod_{e\in E}(\mathbb{R}^d)^n$. We call X the tree preshape space, since it parametrizes all possible shapes.

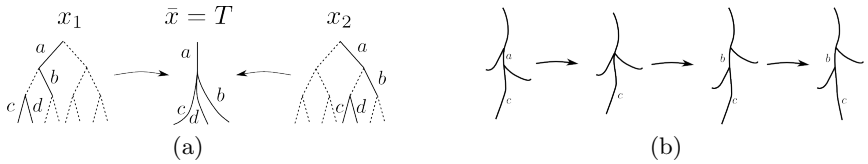

Fig. 4. (a) Higher-degree vertices are represented by collapsing internal edges. The dotted line indicates zero attribute, which gives a collapsed edge. We identify those tree preshapes whose collapsed structures are identical – here x_1 and x_2 represent the same tree T. (b) The TED moves: Remove an edge a, add an edge b, deform an edge c.

2.2 The space of trees as a quotient – a singular treespace

We fix the maximal depth n binary tree $\mathscr{T}_n$ which encodes the connectivity of the trees. In the preshape space defined above, there will be several points corresponding to the same tree, as illustrated in Fig. 4(a). We go from preshapes to shapes by identifying those preshapes which define the same shape.

Consider two ordered tree-shapes where internal edges are collapsed. The orders of the original trees induce orders on the collapsed trees. We consider two ordered tree-shapes to be the same when their collapsed ordered topological structures are identical, and the edge attributes on corresponding non-collapsed edges are identical as well, as in Fig. 4(a). Thus, tree identifications come with an inherent bijection of subsets of E: If we identify $x, y \in X = (\mathbb{R}^N)^E$, denote $E_1 = \{e \in E | x_e \neq 0\}, E_2 = \{e \in E | y_e \neq 0\}$; the identification comes with an order preserving bijection $\varphi\colon E_1 \to E_2$ identifying those edges that correspond to the same edge in the collapsed tree-shape. Note that φ will also correspond to similar equivalences of pairs of trees with the same topology, but other attributes.

The bijection φ next induces a bijection $\Phi\colon V_1 \to V_2$ given by $\Phi\colon (x_e) \mapsto (x_{\varphi(e)})$. Here, $V_1 = \{x \in X | x_e = 0 \text{ if } e \notin E_1\}$ and $V_2 = \{x \in X | x_e = 0 \text{ if } e \notin E_2\}$ are subspaces of X where, except for at the axes, the topological tree structure is constant, and for $x \in V_1$, $\Phi(x) \in V_2$ describes the same shape as x.

We define a map Φ for each pair of identified tree-structures, and form an equivalence on X by setting $x \sim \Phi(x)$ for all x and Φ. For each $x \in X$ we denote by $\bar{x}$ the equivalence class $\{x' \in X | x' \sim x\}$. The quotient space $\bar{X} = (X/\sim) = \{\bar{x} | x \in X\}$ of equivalence classes $\bar{x}$ is the space of treelike shapes.

The geometric interpretation of the identification made in the quotient is that we are folding and gluing our space along the identified subspaces; i.e. when $x_1 \sim x_2$ we glue the two points x_1 and x_2 together. See the toy quotient space in Fig. 5 for an intuitive lower-dimensional illustration.

2.3 Metrics on treespace

Given a metric d on Euclidean space $X = \prod_{e \in E} \mathbb{R}^N$ we define the standard quotient pseudometric [15] $\bar{d}$ on the quotient space $\bar{X} = X/\sim$ by setting

$$\bar{d}(\bar{x}, \bar{y}) = \inf \left\{ \sum_{i=1}^{k} d(x_i, y_i) | x_1 \in \bar{x}, y_i \sim x_{i+1}, y_k \in \bar{y} \right\}. \tag{1}$$

This amounts to finding the optimal path from $\bar{x}$ to $\bar{y}$ passing through the identified subspaces, as shown in Fig. 5.

We define two metrics on X, which come from two different ways of combining the individual edge distances: The metrics d_1 and d_2 on $X = \prod_{e \in E} \mathbb{R}^N$ are induced by the norms $\|x-y\|_1 = \sum_{e \in E} \|x_e - y_e\|$ and $\|x-y\|_2 = \sqrt{\sum_{e \in E} \|x_e - y_e\|^2}$.

From now on, d and $\bar{d}$ will denote either the metrics d_1 and $\bar{d}_1$, or d_2 and $\bar{d}_2$.

We have the following:

Theorem 1. *The distance function $\bar{d}$ restricts to a metric on $\bar{X}$, which is a contractible, complete, proper geodesic space.* □

This means, in particular, that given any two trees, we can always find a geodesic between them in both metrics $\bar{d}_1$ and $\bar{d}_2$.

Note that $\bar{d}_1$ is essentially the classical tree edit distance (TED) metric, where the distance between two trees is the minimal sum of costs of edge edits needed to transform one tree into another, see Fig. 4(b). The $\bar{d}_2$ is a descent of the Euclidean metric, and geodesics in this metric are concatenations of straight lines in flat regions. We call it the QED metric, for quotient euclidean distance.

In Section 3 we describe the properties of these metrics with example geodesic deformations between simple trees.

Using methods from metric geometry [11] we obtain results on curvature and average trees. We consider two different types of average tree; namely the *centroid*, as defined and computed in [3], and the *circumcenter*, defined as follows: Given a set $S \subset \bar{X}$ with diameter D a circumcenter of S is defined as a point $\bar{x} \in \bar{X}$ such that $S \subset \bar{B}(\bar{x}, D/2) = \{\bar{z} \in \bar{X} | \bar{d}_2(\bar{x}, \bar{z}) \leq D/2\}$.

Theorem 2. *i) Endow $\bar{X}$ with the QED metric $\bar{d}_2$. A generic point $\bar{x} \in \bar{X}$ has a neighborhood $U \subset \bar{X}$ in which the curvature is non-positive. At non-generic points, the curvature of $(\bar{X}, \bar{d}_2)$ is unbounded.*

ii) Endow $\bar{X}$ with the QED metric $\bar{d}_2$. Given a generic point $\bar{x} \in \bar{X}$, there exists a radius $r_{\bar{x}}$ such that sets contained in the ball $B(\bar{x}, r_{\bar{x}})$ have unique centroids and circumcenters.

iii) Endow $\bar{X}$ with the TED metric $\bar{d}_1$. The curvature of $(\bar{X}, \bar{d}_1)$ is everywhere unbounded*, and $(\bar{X}, \bar{d}_1)$ does not have locally unique geodesics anywhere.* □

The practical meaning of Theorem 2 is that i) we can use techniques from metric geometry to look for QED averages, ii) for datasets whose scattering is not too large, there exist unique centroids and circumcenters for the QED metric, and iii) we cannot use the same techniques in order to prove existence or uniqueness of centerpoints for the TED metric; in fact, any geometric method which requires bounded curvature is going to fail for the TED metric. *This result motivates our study of the QED metric.*

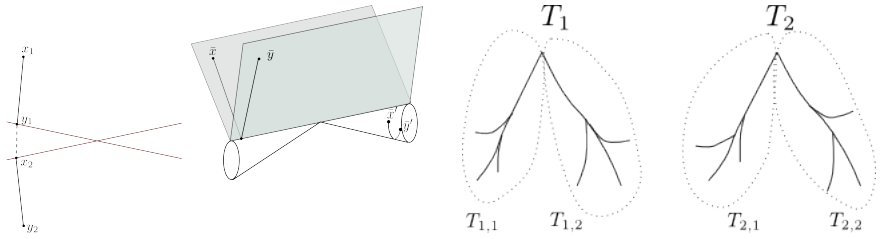

Fig. 5. (a) A 2-dimensional toy version of the folding of Euclidean space along two linear subspaces, along with geodesic paths from $\bar{x}$ to $\bar{y}$ and from $\bar{x}'$ to $\bar{y}'$. (b) Local-to-global properties of TED: $\bar{d}_1(T_1, T_2) = \bar{d}_1(T_{1,1}, T_{2,1}) + \bar{d}_1(T_{1,2}, T_{2,2})$.

2.4 From planar trees to spatial trees

The world is not flat, and for most applications it is necessary to study embedded trees in $\mathbb{R}^3$. As far as attributes in terms of edge embeddings and landmarks are concerned, this is not very different from the planar embedded trees – the main difference from the $\mathbb{R}^2$ case comes from the fact that trees in $\mathbb{R}^3$ have no canonical edge order. The left-right order on children of planar trees gives an implicit preference for edge matchings, and hence reduces the number of possible matches. When we no longer have this preference, we generally need to consider all orderings of the same tree and choose the one which minimizes the distance.

We define the space of spatial treelike shapes as the quotient $\bar{\bar{X}} = \bar{X}/G$, where G is the group of reorderings of the standard tree; G is a finite group. The metric $\bar{d}$ on $\bar{X}$ induces a quotient pseudometric $\bar{\bar{d}}$ on $\bar{\bar{X}}$. We can prove:

Theorem 3. *For $\bar{\bar{d}}$ induced by either $\bar{d}_1$ or $\bar{d}_2$, the function $\bar{\bar{d}}$ is a metric and the space $(\bar{\bar{X}}, \bar{\bar{d}})$ is a contractible, complete, proper geodesic space.* □

While considering all different possible orderings of the tree makes perfect sense from the geometric point of view, in reality this becomes an impossible task as the size of the trees grow beyond a few generations. In real applications we can, however, efficiently reduce complexity by taking tree- and treespace geometry into account. This is discussed in Section 4.

3 Comparison of the QED and TED metrics

In this section we list the main differences between the TED and QED metrics and compare their performance on the small trees studied in [14].

Geometry. As shown in Theorem 1 and Theorem 3 above, both $\bar{X}$ and $\bar{\bar{X}}$ are complete geodesic spaces. However, by Theorem 2, the QED metric gives locally non-positive curvature at generic points, while the TED metric gives unbounded curvature everywhere on $\bar{X}$. This means that we cannot imitate the classical statistical procedures on shape spaces using the TED metric.

Note also that the QED metric is the quotient metric induced from the Euclidean metric on the preshape space X, making it the obvious choice for a metric seen from the shape space point of view.

Computation. The TED metric has nice local-to-global properties, as illustrated in Fig. 5(b). If the trees T_1 and T_2 are decomposed into subtrees $T_{1,1}, T_{1,2}$ and $T_{2,1}, T_{2,2}$ as in Fig. 5(b) such that the geodesic from T_1 to T_2 restricts to geodesics between $T_{1,1}$ and $T_{2,1}$ as well as $T_{2,1}$ and $T_{2,2}$, then $d(T_1, T_2) = d(T_{1,1}, T_{2,1}) + d(T_{2,1}, T_{2,2})$. This property is used in many TED algorithms, and the same property does not hold for the QED metric.

Performance. We have previously [14] studied fundamental geodesic deformations for the QED metric; that is, deformations between depth 3 trees which encode the local topological transformations found in generic tree-geodesics. These deformations are important for determining whether an internal tree-topological transition is needed to transform one (sub)tree into another – i.e. whether the correct edge registration is one which is inconsistent with the tree topology.

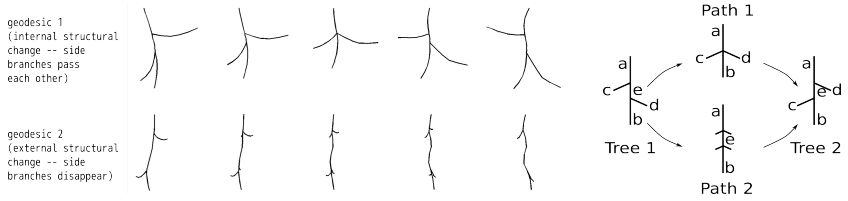

Fig. 6. (a) The fundamental geodesic deformations: geodesic 1 goes through an internal structural change, while geodesic 2 does not, as the disappearing branches approach the zero attribute. (b) Two options for structural change.

To compare the TED and QED metric on small, simple trees, consider the two tree-paths in Fig. 6(b), where the edges are endowed with non-negative scalar attributes a, b, c, d, e describing edge length. Path 1 indicates a matching of the identically attributed edges c and d, while Path 2 does not make the match. Now, the cost of Path 1 is $2e$ in both metrics, while the cost of Path 2 is $2\sqrt{c^2+d^2}$ in the QED metric and $2(c+d)$ in the TED metric. In particular, TED will chose to identify the c and d edges whenever $e^2 \leq c^2 + 2cd + d^2$, while QED makes the match whenever $e^2 \leq \frac{1}{2}(c^2+d^2)$. That is, TED will be more prone to internal structural changes than QED. This is also seen empirically in the comparison of TED and QED matching in Fig. 7. Note that although the TED is more prone to matching trees with different tree-topological structures, the matching is similar.

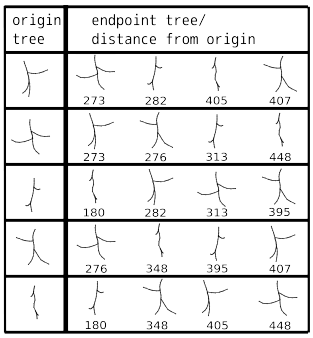

(a) Matching in the QED metric.

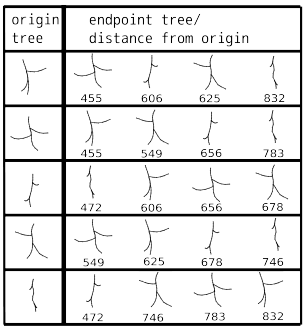

(b) Matching in the TED metric.

Fig. 7. Given a set of five data trees, we match each to the four others in both metrics.

4 Computation and complexity

Complexity is a problem with computing both TED and QED distances, in particular for $3D$ trees, which do not have a canonical planar order. Here we discuss how to use geometry and anatomy to find approximations of the metric whose complexity is significantly reduced.

Ordered trees: Reducing complexity using geometry. The definition of $\bar{d}$ in (1) opens for considering infinitely many possible paths. However, we can significantly limit the search for a geodesic by taking the geometry of treespace into account.

For instance, it is enough to consider paths going through each structural transition at most once:

Lemma 1. *For the metrics $\bar{d}_1$ and $\bar{d}_2$ on X, the shortest path between two points in $\bar{X}$ passes through each identified subspace at most once.* □

Another way of limiting the complexity of real computations is to find a generic shape structure. We show that locally in the geodesic, the only internal topological transitions found will be those described by the fundamental geodesic deformations in Section 3.

Theorem 4. *i) Treelike shapes that are truly binary (i.e. their internal edges are not collapsed) are generic in the space of all treelike shapes.*

ii) The tree-shape deformations whose local structure is described by the fundamental geodesic deformations in Fig. 6(a) are generic in the set of all tree deformations. Moreover, the set of pairs of trees whose geodesics are of this form, is generic in the space of all pairs of trees. □

Genericity basically means that when a random treelike shape is selected, the probability of that shape being truly binary is 1. This does not mean, however, that non-binary trees do not need to be considered. While non-binary trees do not appear as randomly selected trees, they *do* appear in paths between randomly selected pairs of trees, as in Fig. 6(a) above. This also means that we can, in fact, work on a less complicated semi-preshape-space where we, instead of identifying all possible representations of the same shape, restrict ourselves to making the identifications illustrated by Fig. 6(a).

This is similar to the generic form of shock graphs and generic transitions between shock graphs found by Giblin and Kimia [16]. The notions of genericity in the two settings are, however, different since in [16], genericity is defined with respect to the Whitney topology on the space of the corresponding shape boundary parametrizations.

Although we do run into non-binary treelike shapes in real-life applications, for instance when studying airway trees, this can be interpreted as an artifact of resolution rather than as true higher-degree vertices. Indeed, the airway extraction algorithms record trifurcations when the relative distances are below certain treshold values.

Finally, it is often safe to assume that only very few structural changes will actually happen. For the airway trees studied in Section 6.2, we find empirically that it is enough to allow for one structural change in each depth 3 subtree.

Using these arguments, it becomes feasible to computationally handle semi-large structures. In many medical imaging applications, such as airways or blood vessels, we are not interested in studying tree-deformations between very large networks or trees, since beyond a certain point, the tree-structure stops following a predetermined pattern and becomes a stochastic variable.

Unordered trees: Reducing complexity using geometry- or anatomy-based semi-labeling schemes. It is well known that the general problem of computing TED-distances between unordered trees is NP-complete [17], and the QED metric is probably generally not less expensive to compute, as indicated also by Theorem 3. Here we discuss how to use geometry and anatomy to find approximations of the metric whose complexity is significantly reduced.

In particular, *trees appearing in applications are usually not completely unordered, but are often semi-labeled.* Semi-labelings can come from geometric or anatomical properties as in the pulmonary airway trees studied in Example 1 below, or may be obtained by a coarser registration method. (Semi-)labelings can also come from a TED distance computation or approximation, which is a reasonable way to detect approximate structural changes since the TED and QED give similar matchings, as seen in Fig. 7.

Example 1 (Semi-labeling of the upper airway tree). In Fig. 9(a), we see a "standard" airway tree with branch labels (or at least the first generations of it). Most airway trees, as the one shown to the left in Fig. 1, have similar, but not necessarily identical, topological structure to this one, and several branches, especially of low generation, have names and can be identified by experts.

The top three generations of the airway tree serve very clear purposes in terms of anatomy. The root edge is the trachea; the second generation edges are the left and right bronchi; and the third generation edges lead to the top and towards the middle and lower lung lobes on the left, and to the top and bottom lobes on the right. Thus we find a canonical semi-labeling of the airway tree, which we use to simplify the problem of computing airway distances in Section 6.2.

5 QED implementation

We explain a simple implementation of the QED metric. As shown in the next sections, our results are promising – the geodesics for synthetic data show the expected correspondences between edges, and in our experiments on 6 airway trees extracted from CT scans of 3 patients made at two different times, we recognize the correct patient in 5 out of 6 cases.

Alignment of trees. We translate all edges to start at 0. We do not factor out scale in the treespace, because in general, we consider scale an important feature of shape. Edge scale in particular is a critical property, as the dynamics of appearing and disappearing edges is directly tied to edge size.

Edge shape comparison. We represent each edge in a tree by a fixed number of landmark points – in our case 6 – evenly distributed along the edge, the first one at the origin (and hence neglected). The distance between two edge attributes $v_1, v_2 \in (\mathbb{R}^d)^5$ is defined as the Euclidean distance between them. Although simple, this distance measure does take the scale of edges into account to some degree since they are aligned.

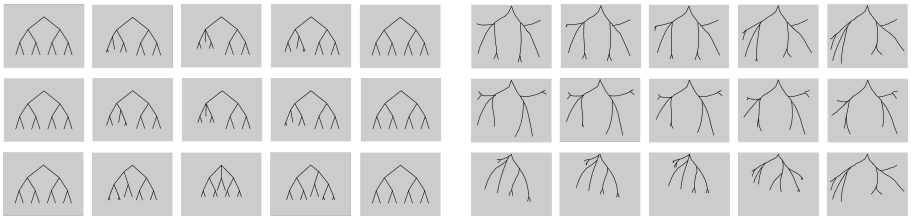

Fig. 8. Different options for structural changes in the left hand side of the depth 3 tree. Topological illustration on the left, corresponding tree-shape illustration on the right.

Implementation. In ordered depth 3 trees we make an implementation using Algorithm 1, where the number of structural changes taking place in the whole tree is limited to either 1 or 2, and in the final case we also rule out the case where the two changes happen along the same "stem". This leaves us with the options for structural changes illustrated in Fig. 8, applied to the left and right half tree. The complexity of Algorithm 1 is $O(n^k \cdot C(n,k))$, where n is the number of internal vertices, k is the number of structural changes and $C(n,k)$ is the maximal complexity of the optimization in line 8.

Algorithm 1 Computing geodesics between ordered, rooted trees with up to k structural changes

1: x, y planar rooted depth n binary trees
2: $\mathbf{S} = \{S\}$ set of ordered identified pairs $S = \{S_1, S_2\}$ of subspaces of X corresponding to internal topological changes, corresponding to a subspace S of $\bar{X}$, s.t. if $\{S_1, S_2\} \in \mathbf{S}$, then also $\{S_2, S_1\} \in \mathbf{S}$.
3: **for** $\tilde{S} = \{S^1, \ldots, S^s\} \subset \mathbf{S}$ with $|\tilde{S}| \leq k$ **do**
4: **for** $p^i \in S^i$ with representatives $p_1^i \in^i S_1$ and $p_2^i \in S_2^i$ **do**
5: $p = (p^1, p^2, \ldots, p^s)$
6: $f(p) = \min\{d_2(x, p_1) + \sum_{j=1}^{s-1} d_s(p_i^2, p_{i+1}^i) + d_2(p_s, y)\}$
7: **end for**
8: $d_{\tilde{S}} = \min\left\{f(p) | p = (p^1, \ldots, p^s), p^i \in S^i, \tilde{S} = \{S^1, \ldots, S^s\}\right\}$
9: $p_{\tilde{S}} = \{p_1^i, p_2^i\}_{i=1}^s = \operatorname{argmin} f(p)$
10: **end for**
11: $d = \min\{d_S | \tilde{S} \subset \mathbf{S}, |\tilde{S}| \leq k\}$
12: $p = \{p_1, p_2\}_{i=1}^s = \{p_{\tilde{S}} | d_{\tilde{S}} = d\}$
13: geodesic $= g = \{x \to p_1^1 \sim p_2^1 \to p_1^2 \sim p_2^2 \to \ldots \to p_1^s \sim p_2^s \to y\}$
14: **return** d, g

6 Experimental results

The QED metric is new, whereas the matching properties of the TED metric are well known [9]. In this section we present experimental results on real and synthetic data which illustrate the geometric properties of the QED metric. The experiments on airway trees in Section 6.2 show, in particular, that it is feasible to compute the metric distances on real $3D$ data trees.

6.1 Synthetic planar trees of depth 3

We have uploaded movies illustrating geodesics between planar depth 3 trees, as well as a matching table for a set of planar depth 3 trees, to the webpage http://image.diku.dk/aasa/ACCVsupplementary/geodesicdata.html. Note the geometrically inutitive behaviour of the geodesic deformations.

6.2 Results in 3D: Pulmonary airway trees

We also compute geodesics between subtrees of pulmonary airway trees. The airway trees were first segmented from low dose screening computed tomography (CT) scans of the chest using a voxel classification based airway tree segmentation algorithm by Lo et al. [18]. The centerlines were then extracted from the segmented airway trees using a modified fast marching algorithm based on [19], which was originally used in [18] for measuring the length of individual branches. Since the method also gives connectivity of parent and children branches, a tree structure is obtained directly. Leaves with a volume less than 11 mm^3 were assumed to be noise and pruned away, and the centerlines were sampled with 5 landmark points on each edge.

Six airways extracted from CT scans taken from three different patients at two different times were analyzed, restricting to the first six generations of the airway tree. The first three generations were identified and labeled as in Ex. 1, leaving us with four depth 3 subtrees representing the 4^{th} to 6^{th} generations of a set of $3D$ pulmonary airway trees. Algorithm 1 was run with both one and two structural changes, and no subtree geodesics (out of 24) had more than one structural change. Thus, we find it perfectly acceptable to restrict our search to paths with few structural changes.

In the depth 3 subtrees we expect the deformation from one lung to the other to include topological changes, and we compare them using the unordered QED metric implementation on the subtrees which are cut off after the first 3 generations. As a result we obtain a deformation from the first depth 6 airway tree to the second, whose restriction to the subtrees is a geodesic. In this experiment we study four airways originating from CT scans of two different patients acquired at two different times, and we obtain the geodesic distances and corresponding matching shown in Fig. 9(b).

The five subtree distances (top three generations and four subtrees) were added up to give the distance measure in Fig. 9(b), and we see that for 5 out of 6 images, the metric recognizes the second image from the same patient.

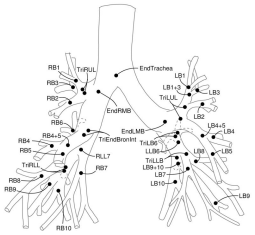

(a)

	P(1,1)	P(1,2)	P(2,1)	P(2,2)	P(3,1)	P(3,2)
P(1,1)	0	309.09	437.58	452.62	375.40	378.19
P(1,2)	309.09	0	435.11	402.71	400.41	349.41
P(2,1)	437.58	435.11	0	400.91	448.45	392.69
P(2,2)	452.62	402.71	400.91	0	456.69	411.24
P(3,1)	375.40	400.41	448.45	456.69	0	324.43
P(3,2)	378.19	349.41	392.69	411.24	324.43	0

(b)

Fig. 9. (a) Standard airway tree with branch labels. Drawing from [6] based on [20]. (c) Table of sum of QED distances for the five airway subtrees for 6 airway trees retrieved from CT scans of three patients taken at two different times, denoted $\mathrm{P}(a, i)$ where a denotes patient and i denotes time.

7 Conclusions and future work

Starting from a purely geometric point of view, we define a shape space for treelike shapes. The intuitive geometric framework allows us to simultaneously take both global tree-topological structure and local edgewise geometry into account. We define two metrics on this shape space, QED and TED, which give the shape space a geodesic structure (Theorems 1 and 3).

QED is the geometrically natural metric, which turns out to have excellent geometric properties which are essential for statistical shape analysis. In particular, the QED metric has local uniqueness of geodesics and local existence and uniqueness for two versions of average shape, namely the circumcentre and the centroid (Theorem 2). TED does not share these properties, but has better computational properties.

Both metrics are generally NP hard to compute for $3D$ trees. We explain how semi-labeling schemes and geometry can be used to overcome the complexity problems, and illustrate this by computing QED distances between trees extracted from pulmonary airway trees as well as synthetic planar data trees.

Our future research will be centered around two points: Development of non-linear statistical methods for the singular tree-shape spaces, and finding fast algorithms for an approximate QED metric. The latter will allow us to actually compute averages and modes of variation for large, real $3D$ data trees.

8 Acknowledgements

This work is partly funded by the Lundbeck foundation, the Danish Council for Strategic Research (NABIIT projects 09 − 065145 and 09 − 061346) and the Netherlands Organisation for Scientific Research (NWO). We thank Dr. J.H. Pedersen from the Danish Lung Cancer Screening Trial (DLCST) for providing the CT scans.

We thank Marco Loog, Søren Hauberg and Melanie Ganz for helpful discussions during the preparation of the article.

References

1. Pedersen, J., Ashraf, H., Dirksen, A., Bach, K., Hansen, H., Toennesen, P., Thorsen, H., Brodersen, J., Skov, B., Døssing, M., Mortensen, J., Richter, K., Clementsen, P., Seersholm, N.: The Danish randomized lung cancer CT screening trial - overall design and results of the prevalence round. J Thorac Oncol **4** (2009) 608–614
2. Staal, J., Abramoff, M., Niemeijer, M., Viergever, M., van Ginneken, B.: Ridge based vessel segmentation in color images of the retina. TMI **23** (2004) 501–509
3. Billera, L.J., Holmes, S.P., Vogtmann, K.: Geometry of the space of phylogenetic trees. Adv. in Appl. Math. **27** (2001) 733–767
4. Charnoz, A., Agnus, V., Malandain, G., Soler, L., Tajine, M.: Tree matching applied to vascular system. GbRPR, LNCS **3434** (2005) 183–192
5. Chalopin, C., Finet, G., Magnin, I.E.: Modeling the 3D coronary tree for labeling purposes. MIA **5** (2001) 301–315
6. Tschirren, J., McLennan, G., Palágyi, K., Hoffman, E.A., Sonka, M.: Matching and anatomical labeling of human airway tree. TMI **24** (2005) 1540–1547
7. Kaftan, J., Kiraly, A., Naidich, D., Novak, C.: A novel multi-purpose tree and path matching algorithm with application to airway trees. In: SPIE. Volume 6143 I. (2006)
8. Metzen, J.H., Kröger, T., Schenk, A., Zidowitz, S., Peitgen, H.O., Jiang, X.: Matching of tree structures for registration of medical images. GbRPR, 5th IAPR Int. WS. LNCS **4538** (2007) 224–231
9. Sebastian, T.B., Klein, P., Kimia, B.: Recognition of shapes by editing their shock graphs. TPAMI **26** (2004) 550–571
10. Nakano, Y., Muro, S., Sakai, H., Hirai, T., Chin, K., Tsukino, M., Nishimura, K., Itoh, H., Paré, P.D., Hogg, J.C., Mishima, M.: Computed tomographic measurements of airway dimensions and emphysema in smokers. Correlation with lung function. Am. J. Resp. Crit. Care Med **162** (2000) 1102–1108
11. Gromov, M.: Hyperbolic groups. In: Essays in group theory. Volume 8 of Math. Sci. Res. Inst. Publ. (1987) 75–263
12. Wang, H., Marron, J.S.: Object oriented data analysis: sets of trees. Ann. Statist. **35** (2007) 1849–1873
13. Jain, B.J., Obermayer, K.: Structure spaces. J. Mach. Learn. Res. **10** (2009) 2667–2714
14. Feragen, A., Lauze, F., Nielsen, M.: Fundamental geodesic deformations in spaces of treelike shapes. In: Proc ICPR. (2010)
15. Bridson, M.R., Haefliger, A.: Metric spaces of non-positive curvature. Springer-Verlag (1999)
16. Giblin, P.J., Kimia, B.B.: On the local form and transitions of symmetry sets, medial axes, and shocks. IJCV **54** (2003) 143–156
17. Zhang, K., Statman, R., Shasha, D.: On the editing distance between unordered labeled trees. Inf. Process. Lett. **42** (1992) 133–139
18. Lo, P., Sporring, J., Ashraf, H., Pedersen, J.J., de Bruijne, M.: Vessel-guided airway tree segmentation: A voxel classification approach. Medical Image Analysis **14** (2010) 527–538
19. Schlathölter, T., Lorenz, C., Carlsen, I.C., Renisch, S., Deschamps, T.: Simultaneous segmentation and tree reconstruction of the airways for virtual bronchoscopy. Volume 4684., SPIE (2002) 103–113
20. Boyden, E.A.: Segmental anatomy of the lungs. McGraw-Hill. (1955)