

A Polynomial-Time Metric for Attributed Trees

Andrea Torsello¹, Džena Hidović², and Marcello Pelillo¹

¹ Dipartimento di Informatica, Università Ca' Foscari di Venezia
Via Torino 155, 30172 Venezia Mestre, Italy
{torsello,pelillo}@dsi.unive.it

² School of Computer Science, University of Birmingham
Edgbaston, Birmingham, B15 2TT, United Kingdom
D.Hidovic@cs.bham.ac.uk

Abstract. We address the problem of comparing attributed trees and propose a novel distance measure centered around the notion of a maximal similarity common subtree. The proposed measure is general and defined on trees endowed with either symbolic or continuous-valued attributes, and can be equally applied to ordered and unordered, rooted and unrooted trees. We prove that our measure satisfies the metric constraints and provide a polynomial-time algorithm to compute it. This is a remarkable and attractive property since the computation of traditional edit-distance-based metrics is NP-complete, except for ordered structures. We experimentally validate the usefulness of our metric on shape matching tasks, and compare it with edit-distance measures.

1 Introduction

Graph-based representations have long been used with considerable success in computer vision and pattern recognition in the abstraction and recognition of objects and scene structure. Concrete examples include the use of shock graphs to represent shape-skeletons [11,15], the use of trees to represent articulated objects [7] and the use of aspect graphs for 3D object representation [3]. The attractive feature of structural representations is that they concisely capture the relational arrangement of object primitives, in a manner which can be invariant to changes in object viewpoint. Using this framework we can transform a recognition problem into a relational matching problem. The problem of how to measure the similarity or distance of pictorial information using graph abstractions has been a widely researched topic of over twenty years.

The classic metric approach to graph comparison is edit-distance [4]. The idea behind this approach is that it is possible to identify a set of basic edit operations on nodes and edges of a structure, and to associate with these operations a cost. The edit-distance is found by searching for sequences of edit operations that will make the two graphs isomorphic with one-another, and the distance between the two graphs is then defined to be the minimum over all the costs of these sequences. By making the evaluation of structural modification explicit, edit-distance provides a very effective way of measuring the similarity

of relational structures. Moreover, the method has considerable potential for error tolerant object recognition and indexing problems. Unfortunately, the task of calculating edit-distance is an NP-hard problem [24], hence, goal-directed approximations are necessary to calculate it. The result is that the approximation almost invariably breaks the theoretical metric properties of the measure.

Recently, a new and more principled approach to the definition of distance measure has emerged. In [2], Bunke and Shearer introduce a distance measure on unattributed graphs based on the maximum common subgraph and prove that it is a metric. Wallis et al. [20] introduce a variant of this distance based on the size of the minimum common supergraph. Finally, Fernandez and Valiente [5] define a metric based on the difference in size between maximum common subgraph and minimum common supergraph. More recently, in [6] Hidović and Pelillo extend these metrics to the case of attributed graphs. Unfortunately all these metrics require the calculation of the maximum common subgraph, which is computationally equivalent to the calculation of edit-distance.

In many computer vision and pattern recognition applications, such as shape recognition [13,15,17], the graphs at hand have a peculiar structure: they are connected and acyclic, i.e., they are *trees*, either rooted or unrooted, ordered or unordered, and frequently they are endowed with symbolic and/or continuous-valued attributes. Most metrics on trees found in the literature are defined in terms of edit-distance [18,21]. Zhang and Shasha [23] have investigated a special case of edit-distance which involves trees with an order relation among sibling nodes in a rooted tree. This special case constrains the solution to maintain the order of the children of a node. They showed that this constrained tree-matching problem is solvable in polynomial time and gave an algorithm to solve it. Recently, Sebastian, Klein and Kimia [13] use a similar algorithm to compare shock trees. Unfortunately, in the general case the problem has been proven to be NP-complete both for rooted [24] and unrooted trees [25]. Recently, Valiente [19] introduced a bottom-up distance measure between trees that is an extension to trees of the graph metric introduced by Bunke and Shearer [2], proving that the measure can be calculated in polynomial time on trees, but falling short of proving that the measure is a metric. While this measure can be calculated efficiently both on ordered and unordered trees, it is limited to rooted and unattributed trees.

Motivated by the work described in [6], in this paper we propose a normalized distance measure for trees equipped with either symbolic or continuous-valued attributes. We prove that the proposed measure fulfills the properties of a metric, and provide a polynomial-time algorithm to compute it. At an abstract level, our approach involves the computation of a maximum similarity common subtree. This allows us to define equivalent variations of the metric on ordered and unordered, rooted and unrooted, and attributed and unattributed trees. Since edit-distance on ordered trees can be computed in polynomial time, in the paper we focus on the unordered case where our approach provides a clear computational advantage. To show the validity of the proposed measures, we present experiments on various shape matching tasks and compare our results with those obtained using edit-distance metrics.

2 Preliminaries

Let $G = (V, E)$ be a graph, where V is the set of nodes (or vertices) and E is the set of undirected edges. Two nodes $u, v \in V$ are said to be *adjacent* (denoted $u \sim v$) if they are connected by an edge. A *path* is any sequence of distinct nodes $u_0 u_1 \dots u_n$ such that for all $i = 1 \dots n$, $u_{i-1} \sim u_i$; in this case, the *length* of the path is n . If $u_n \sim u_0$ the path is called a *cycle*. A graph is said to be *connected* if any two nodes are joined by a path. Given a subset of nodes $C \subseteq V$, the *induced subgraph* $G[C]$ is the graph having C as its node set, and two nodes are adjacent in $G[C]$ if and only if they are adjacent in G . With the notation $|G|$ we shall refer to the cardinality of the node-set of graph G .

A connected graph with no cycles is called an unrooted tree. A rooted (or hierarchical) tree is a tree with a special node that can be identified as the root. In what follows, when using the word “tree” without qualification, we shall refer to both the rooted and unrooted cases. Given two nodes $u, v \in V$ in a rooted tree, u is said to be an *ancestor* of v (and similarly v is said to be a *descendent* of u) if the path from the root node to u is a subpath of the path from the root to v . Furthermore, if $u \sim v$, u is said to be the *parent* of v and v is said to be a *child* of u . Both ancestor and descendent relations are order relations in V .

Let $T_1 = (V_1, E_1)$ and $T_2 = (V_2, E_2)$ be two trees. Any bijection $\phi : H_1 \rightarrow H_2$, with $H_1 \subseteq V_1$ and $H_2 \subseteq V_2$, is called a *subtree isomorphism* if it preserves both the adjacency relationships between the nodes and the connectedness of the matched subgraphs. Formally, this means that, given $u, v \in H_1$, we have $u \sim v$ if and only if $\phi(u) \sim \phi(v)$ and, in addition, the induced subgraphs $T_1[H_1]$ and $T_2[H_2]$ are connected. Two trees or rooted trees T_1 and T_2 are *isomorphic*, and we write $T_1 \cong T_2$, if there exists an isomorphism between them that maps every node in T_1 to every node in T_2 . It is easy to verify that isomorphism is an equivalence relation. We shall use the notations $\text{Dom}(\phi)$ and $\text{Im}(\phi)$ to denote the domain and the image of ϕ , respectively.

Formally, an *attributed tree* is a triple $T = (V, E, \alpha)$, where (V, E) is the “underlying” tree and α is a function which assigns an attribute vector $\alpha(u)$ to each node $u \in V$. It is clear that in matching two attributed trees, our objective is to find an isomorphism which pairs nodes having “similar” attributes. To this end, let σ be any similarity measure on the attribute space, i.e., any (symmetric) function which assigns a positive number to any pair of attribute vectors. If $\phi : H_1 \rightarrow H_2$ is a subgraph isomorphism between two attributed trees $T_1 = (V_1, E_1, \alpha_1)$ and $T_2 = (V_2, E_2, \alpha_2)$, the overall similarity between the induced subtrees $T_1[H_1]$ and $T_2[H_2]$ can be defined as follows:

$$W_\sigma(\phi) = \sum_{u \in H_1} \sigma(u, \phi(u)) . \quad (1)$$

where, for simplicity, we define $\sigma(u, \phi(u)) \equiv \sigma(\alpha_1(u), \alpha_2(\phi(u)))$. The isomorphism ϕ is called a *maximum similarity subtree isomorphism* if $W_\sigma(\phi)$ is largest among all subtree isomorphisms between T_1 and T_2 . For the rest of the paper we will omit the subscript σ when the node-similarity used is clear from the context. Two isomorphic attributed trees $T_1 = (V_1, E_1, \alpha_1)$ and $T_2 = (V_2, E_2, \alpha_2)$,

with isomorphism ϕ , are said to be *attribute-isomorphic* if for all $u \in V_1$ we have $\alpha_1(u) = \alpha_2(\phi(u))$. In this case we shall write $T_1 \cong_a T_2$. Attribute-isomorphism is clearly an equivalence relation.

Note that the problem of determining a maximum similarity subtree isomorphism is a direct extension of the standard problem of finding a maximum (cardinality) common subtree, in fact the two problems are equivalent when the similarity σ is degenerate, i.e., $\sigma(u, v) = 1$.

Now, given a set S , a function $d : S \times S \rightarrow \mathbb{R}$ is a *metric* on S if the following properties hold for any $x, y, z \in S$.

1. $d(x, x) \geq 0$ (non-negativity)
2. $d(x, y) = 0 \Leftrightarrow x = y$ (identity and uniqueness)
3. $d(x, y) = d(y, x)$ (symmetry)
4. $d(x, y) + d(y, z) \geq d(x, z)$ (triangular inequality).

Furthermore, if the function satisfies $d(x, y) \leq 1$ it is said to be a *normalized metric*.

If $d : S \times S \rightarrow \mathbb{R}_+$ is a normalized metric, then the similarity function derived from d , defined as $\sigma(x, y) = 1 - d(x, y)$ fulfills the identity, uniqueness and similarity properties. Furthermore, it fulfills the following variant of the triangular inequality: $\sigma(x, y) + \sigma(y, z) - \sigma(x, z) \leq 1$. In the rest of the paper, we shall assume that all similarity functions are indeed derived from normalized metrics.

It is straightforward to show that, with this assumption, we have

$$T_1 \cong_a T_2 \Leftrightarrow |T_1| = |T_2| = W(\phi) \tag{2}$$

where ϕ is a maximum similarity isomorphism between T_1 and T_2 .

3 Distance Metric

In this section, we define our measure for comparing attributed trees and prove that it fulfills the metric properties. First, we prove a lemma that turns out to be instrumental to prove our results, then, we introduce our measure and prove the metric properties.

Lemma 1. *Let T_1, T_2 and T_3 be three trees, and ϕ_{12}, ϕ_{23} , and ϕ_{13} be maximum similarity subtrees isomorphisms between T_1 and T_2 , T_2 and T_3 , and T_1 and T_3 , respectively. Then, we have: $|T_2| \geq W(\phi_{12}) + W(\phi_{23}) - W(\phi_{13})$.*

Proof. Let $V_2^1 = \text{Im}(\phi_{12}) \subseteq V_2$, $V_2^3 = \text{Dom}(\phi_{23}) \subseteq V_2$ be the sets of nodes in V_2 mapped by the isomorphisms ϕ_{12} and ϕ_{23} , respectively. Furthermore, let $\hat{V}_2 = V_2^1 \cap V_2^3$, be the set of vertices in V_2 that are mapped by both isomorphisms. It is clear that the subtrees $\hat{T}_1 = T_1[\phi_{12}^{-1}(\hat{V}_2)]$ and $\hat{T}_3 = T_3[\phi_{23}(\hat{V}_2)]$ are isomorphic to each-other, with isomorphism $\hat{\phi}_{13} = \phi_{12} \circ \phi_{23}$, where \circ denotes the standard function composition operator, restricted to the nodes of \hat{T}_1 . The similarity of this isomorphism is

$$W(\hat{\phi}_{13}) = \sum_{v \in \hat{V}_2} \sigma(\phi_{12}^{-1}(v), \phi_{23}(v)).$$

Since ϕ_{13} is a maximum similarity subtree isomorphism between T_1 and T_3 , we have $W(\phi_{13}) \geq W(\hat{\phi}_{13})$. Hence

$$\begin{aligned} W(\phi_{12}) + W(\phi_{23}) - W(\phi_{13}) &\leq W(\phi_{12}) + W(\phi_{23}) - W(\hat{\phi}_{13}) = \\ &\sum_{v \in V_2^1} \sigma(\phi_{12}^{-1}(v), v) + \sum_{v \in V_2^3} \sigma(v, \phi_{23}(v)) - \sum_{v \in \hat{V}_2} \sigma(\phi_{12}^{-1}(v), \phi_{23}(v)) = \\ &\sum_{v \in V_2^1 \setminus V_2^3} \sigma(\phi_{12}^{-1}(v), v) + \sum_{v \in V_2^3 \setminus V_2^1} \sigma(v, \phi_{23}(v)) + \\ &\sum_{v \in \hat{V}_2} [\sigma(\phi_{12}^{-1}(v), v) + \sigma(v, \phi_{23}(v)) - \sigma(\phi_{12}^{-1}(v), \phi_{23}(v))] \leq \\ &|V_2^1 \setminus V_2^3| + |V_2^3 \setminus V_2^1| + |V_2^1 \cap V_2^3| = |V_2^1 \cup V_2^3| \leq |T_2|, \end{aligned}$$

where the inequality follows from the triangular inequality for metric-derived similarities. \square

Let \mathcal{T} be the quotient set of trees modulo attribute-isomorphism, that is the set of trees on which two trees are considered the same if they are attribute-isomorphic.¹ For any $T_1, T_2 \in \mathcal{T}$ we define the following distance function

$$d(T_1, T_2) = 1 - \frac{W(\phi_{12})}{\max(|T_1|, |T_2|)}. \tag{3}$$

Theorem 1. *d is a normalized metric in \mathcal{T} .*

Proof.

1. $d(T_1, T_2) \geq 0$

We have $0 \leq W(\phi_{12}) \leq \max(|T_1|, |T_2|)$. Hence, $0 \leq d(T_1, T_2) = 1 - \frac{W(\phi_{12})}{\max(|T_1|, |T_2|)} \leq 1$.

2. $d(T_1, T_2) = 0 \iff T_1 \cong_a T_2$

Let us consider the direction of implication \Leftarrow (identity). From (2), we have $T_1 \cong_a T_2 \Rightarrow |T_1| = |T_2| = W(\phi_{12})$. Hence $d(T_1, T_2) = \frac{\max(|T_1|, |T_2|) - W(\phi_{12})}{\max(|T_1|, |T_2|)} = 0$

For the reverse implication (uniqueness), we have $d(T_1, T_2) = 0 \Rightarrow W(\phi_{12}) = \max(|T_1|, |T_2|)$. Since $W(\phi_{12}) \leq \min(|T_1|, |T_2|) \leq \max(|T_1|, |T_2|)$, we have $W(\phi_{12}) = \min(|T_1|, |T_2|) = \max(|T_1|, |T_2|)$. Hence, (2) yields $T_1 \cong_a T_2$.

3. $d(T_1, T_2) = d_1(T_2, T_1)$

This follows directly from the symmetry of the maximum similarity graph and of the function \max .

4. $d(T_1, T_2) + d(T_2, T_3) \geq d(T_1, T_3)$

¹ The quotient set formalizes the intuitive idea that two attributed trees are indistinguishable when they are attribute-isomorphic. Furthermore, it is needed in order to fulfill the uniqueness property of a metric.

The triangular inequality can be simplified to the inequality

$$\begin{aligned} & \max(|T_1|, |T_2|) \max(|T_2|, |T_3|) \max(|T_1|, |T_3|) \geq \\ & W(\phi_{12}) \max(|T_2|, |T_3|) \max(|T_1|, |T_3|) + W(\phi_{23}) \max(|T_1|, |T_2|) \max(|T_1|, |T_3|) - \\ & W(\phi_{13}) \max(|T_1|, |T_2|) \max(|T_2|, |T_3|) \quad (4) \end{aligned}$$

To prove this we need to separately analyze each of the six possible cases

1. $|T_1| \geq |T_2| \geq |T_3|$ 2. $|T_1| \geq |T_3| \geq |T_2|$ 3. $|T_2| \geq |T_1| \geq |T_3|$
4. $|T_2| \geq |T_3| \geq |T_1|$ 5. $|T_3| \geq |T_1| \geq |T_2|$ 6. $|T_3| \geq |T_2| \geq |T_1|$.

However, the roles of T_1 and T_3 in our proofs are symmetric, hence we can use this symmetry to reduce the analysis to three cases: (a) $|T_2| \geq |T_1| \geq |T_3|$, (b) $|T_1| \geq |T_2| \geq |T_3|$, and (c) $|T_1| \geq |T_3| \geq |T_2|$.

- a) $|T_2| \geq |T_1| \geq |T_3|$

The triangular inequality reduces to $|T_1||T_2| \geq W(\phi_{12})|T_1| + W(\phi_{23})|T_1| - W(\phi_{13})|T_2|$.

$$\begin{aligned} |T_1||T_2| & \geq |T_1|(W(\phi_{12}) + W(\phi_{23}) - W(\phi_{13})) \geq \\ & W(\phi_{12})|T_1| + W(\phi_{23})|T_1| - W(\phi_{13})|T_2| \end{aligned}$$

- b) $|T_1| \geq |T_2| \geq |T_3|$

Equation (4) reduces to $|T_1||T_2| \geq W(\phi_{12})|T_2| + W(\phi_{23})|T_1| - W(\phi_{13})|T_2|$.

$$\begin{aligned} |T_1||T_2| & = |T_2|(|T_1| - |T_2|) + |T_2|^2 \geq W(\phi_{23})(|T_1| - |T_2|) + |T_2|^2 \geq \\ & W(\phi_{23})(|T_1| - |T_2|) + |T_2|(W(\phi_{12}) + W(\phi_{23}) - W(\phi_{13})) = \\ & W(\phi_{12})|T_2| + W(\phi_{23})|T_1| - W(\phi_{13})|T_2| \end{aligned}$$

- c) $|T_1| \geq |T_3| \geq |T_2|$

We need to prove $|T_1||T_3| \geq W(\phi_{12})|T_3| + W(\phi_{23})|T_1| - |T_3|W(\phi_{13})$.

$$\begin{aligned} |T_1||T_3| & \geq |T_1||T_2| - |T_2||T_3| + |T_2||T_3| \geq W(\phi_{23})(|T_1| - |T_3|) + |T_3||T_2| \geq \\ & W(\phi_{23})(|T_1| - |T_3|) + |T_3|(W(\phi_{12}) + W(\phi_{23}) - W(\phi_{13})) = \\ & W(\phi_{12})|T_3| + W(\phi_{23})|T_1| - |T_3|W(\phi_{13}). \quad \square \end{aligned}$$

4 Extracting the Maximum Similarity Common Subtree

In this section we give a polynomial-time algorithm for finding a maximum similarity subtree. The algorithm is based on the subtree identification algorithm presented by Matula [9], extending it in two ways. First, it generalizes it to deal with attributed trees and, second, it extends it to solve the more general problem of extracting the maximum (similarity) subtree and not merely to verify whether one tree is a subtree of the other. We give an algorithm to find the maximum similarity common subtree problem for rooted trees, and then we show how the same algorithm can be used for the unrooted tree case.

Let $T_1 = (V_1, E_1)$ and $T_2 = (V_2, E_2)$ be two rooted trees, and let $u \in V_1$ and $w \in V_2$. We say that a subtree isomorphism between T_1 and T_2 is *anchored* at nodes u and w , if the subtrees of T_1 and T_2 induced by the isomorphism are rooted at u and w , respectively. In this case, we shall write $\phi^{(u,w)}$ to refer to any isomorphism anchored at u and w . Clearly, if ϕ is a maximum similarity subtree isomorphism, we have

$$W(\phi) = \max_{(u,w) \in V_1 \times V_2} \max_{\phi^{(u,w)}} W(\phi^{(u,w)}).$$

To determine the maximum similarity subtree isomorphism anchored at nodes u and w we adopt a divide-and-conquer approach. Let u_1, \dots, u_n be the children of node u in T_1 , and w_1, \dots, w_m the children of node w in T_2 . Without loss of generality, we can assume $n \leq m$. Moreover, let us assume that we know, for each $i = 1, \dots, n$ and $j = 1, \dots, m$, a maximum similarity subtree isomorphism $\widehat{\phi}^{(u_i, w_j)}$ anchored at u_i and w_j . Let W_{ij} be the similarity of $\widehat{\phi}^{(u_i, w_j)}$, then the computation of a maximum similarity subtree isomorphism anchored at u and w can be reduced to an assignment problem on the children of u and w , i.e.,

$$W(\phi^{(u,w)}) = \sigma(u, w) + \max_{\pi \in \Sigma_n^m} \sum_{i=1}^n W_{i\pi(i)}, \tag{5}$$

where Σ_n^m is the space of all possible assignments between a set of cardinality n and one of cardinality m . As a consequence, if π is the optimal assignment, the function $\phi^{(u,w)}$ defined as:

$$\phi^{(u,w)}(x) = \begin{cases} w & \text{if } x = u \\ \widehat{\phi}^{(u_i, w_{\pi(i)})}(x) & \text{if } x \in \text{Dom}(\widehat{\phi}^{(u_i, w_{\pi(i)})}) \end{cases} \tag{6}$$

turns out to be a maximum similarity subtree isomorphism anchored at u and w .

Figure 1 shows the resulting algorithm for determining a maximum similarity subtree isomorphism of two rooted attributed trees. Since in the rest of the paper we only need the maximum similarity induced by an isomorphism, and not the isomorphism itself, for simplicity the main procedure **Similarity** accepts as input a pair of attributed rooted trees and returns only the similarity value. It makes use of a recursive procedure **AnchoredSimilarity** that accepts as input two vertices, one from T_1 and the other from T_2 and returns the similarity of the maximum isomorphism anchored at the input vertices, according to (5). To this end, it needs a procedure for solving an assignment (or, equivalently, a bipartite matching) problem, of which the algorithms literature abound (see., e.g., [1]). The calculation of the maximum similarity common subtree of two trees with N and M nodes respectively, is reduced to at most NM weighted assignments problems of dimension at most b , where b is the maximum branching factor of the two trees. The computational complexity of our algorithm heavily depends on the actual implementation of the assignment procedure. A popular way of solving it, and the one we actually employed, is the so-called Hungarian algorithm, which has complexity $O(n^2m)$, n and m being the number of children of u and v as used

```

Similarity( $T_1, T_2$ )
  maxsim=0
  for each node  $u$  in  $T_1$ 
    sim=AnchoredSimilarity( $u, \text{root}(T_2)$ )
    if sim > maxsim
      maxsim=sim
  for each node  $w$  in  $T_2$ 
    sim=AnchoredSimilarity( $\text{root}(T_1), w$ )
    if sim > maxsim
      maxsim=sim
  return maxsim

```

```

AnchoredSimilarity( $u, w$ )
   $C_u = \text{children}(u)$ 
   $C_w = \text{children}(w)$ 
  for each  $u_i$  in  $C_u$ 
    for each  $w_j$  in  $C_w$ 
       $w_{ij} = \text{AnchoredSimilarity}(u_i, w_j)$ 
  return  $\sigma(u, w) + \text{Assign}(\{w_{ij}\})$ 

```

Fig. 1. A polynomial-time algorithm for computing the similarity between two trees.

in (5), with $n \leq m$. It is simple to show that, using the Hungarian algorithm, our algorithm has overall complexity of $O(bNM)$. Of course, the algorithm can be sped up by using more sophisticated assignment procedures [1].

Finally, if we have two unrooted trees $T_1 = (V_1, E_1)$ and $T_2 = (V_2, E_2)$, we can still pick two nodes $r_1 \in V_2$ and $r_2 \in V_2$, and consider the trees $T_1^{r_1} = (V_1, E_1)$ and $T_2^{r_2} = (V_2, E_2)$ rooted at r_1 and r_2 , respectively. Note that if ϕ is an isomorphism between $T_1^{r_1}$ and $T_2^{r_2}$ with similarity W , then it is an isomorphism between T_1 and T_2 with the same similarity. This yields a straightforward $O(bN^3M)$ algorithm for unrooted trees, which consists of iteratively calling $\text{Similarity}(T_1^u, T_2^w)$ for all $u \in V_1$ and $w \in V_2$, and taking the maximum. However, we do not actually need to try all possible pairs of roots since by simply fixing the root in one tree and let the other vary among all possible vertices in the other tree, the algorithm is still guaranteed to achieve the maximum similarity. This yields an $O(bN^2M)$ algorithm for unrooted trees.

5 Experimental Results

We evaluated the new metric on three different tree-based shape representations. The first is the shock tree representation used by Pelillo, Siddiqi and Zucker in [11], which is based on the differential structure of the boundary of a 2D shape. It is obtained by extracting the skeleton of the shape, determined as the set of singularities (shocks) arising from the inward evolution of the shape boundary, and then examining the differential behavior of the radius of the bitangent circle




Fig. 2. Distance matrices from the first experiment. Left: Our metric. Right: Edit-distance.

to the object boundary, as the skeleton is traversed. This yields a classification of local differential structure into four different classes [15]. The so-called shock-classes, distinguish between the cases where the local bitangent circle has maximum, minimum, constant, or monotonic radius. The labeled shock-groups are then abstracted using a rooted tree where two vertices are adjacent if the corresponding shock-groups are adjacent in the skeleton, and the distance from the root is related to the distance from the shape barycenter. Here, we used the same attributes and node-distances employed in [11]. Each shock was attributed with its coordinates, distance from the border, and propagation velocity and direction. The distance between two nodes, was defined as a convex combination of the (normalized) Euclidean distances of length, distance to the border, propagation speed, and curvature.

We compared our distance metric with edit-distance. To approximate the edit-distance we used the relaxation labeling algorithm presented in [17] with the following costs: we defined the cost of matching node u to node w to be equal to the distance between their attributes, while the cost of removing any node to be equal to 1. Note that, with these costs, edit-distance is not normalized.

Our shape database contained 29 shapes from 8 different classes. Figure 2 shows the distance matrices obtained using our metric and edit-distance. Here, lighter colors represent lower distances while darker colors represent higher distances. As can be seen, the same block structure emerges in both matrices. Essentially, the most significant difference among the two metrics is the dark bands clearly visible in the edit-distance matrix.

In order to assess the ability of the distances to preserve class structure, we performed pairwise clustering. In particular, we used two pairwise clustering algorithms: Shi and Malik’s Normalized Cut [14], and Pavan and Pelillo’s Dominant Sets [10]. Figure 3 shows the clusters obtained with both algorithms, displayed in order of extraction. While the performance of the clustering algorithms, on this shape recognition task, varied significantly, the dependency on the choice of the distance measure was less pronounced. Nonetheless, some differences can be observed. In particular, we notice how Normalized Cut exhibits

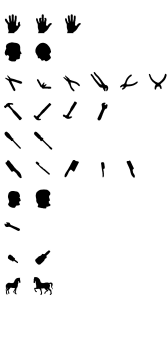
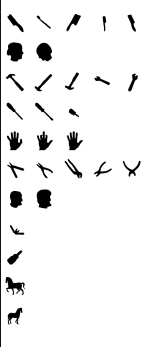

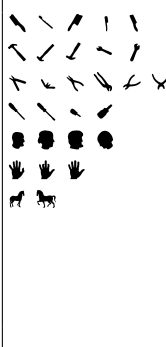
Normalized Cut		Dominant Sets	
Our metric	Edit-distance	Our metric	Edit-distance
			

Fig. 3. Clusters obtained with Normalized Cut and Dominat Sets in the first experiment.

a well-known tendency to over-segment the data. The clusters obtained with the Dominant Sets approach are much better, with our metric providing results almost identical to edit-distance.

As for the running times, on a Pentium 4 2.5GHz PC, the maximum similarity algorithm presented in Section 4, took around 8 seconds to compute our metric, while the relaxation labeling algorithm computed edit-distance in over 30 minutes.

Our second set of experiments used a larger database of shapes abstracted again in terms of shock-trees. Here, however, we used a different set of attributes recently analyzed in [16], i.e., the proportion of the shape boundary generating the corresponding shock-group. The database consisted of 150 shapes divided into 10 classes of 15 shapes each, and presented a higher structural noise than the previous one. Here the node distance and node-matching cost for edit-distance was defined as the absolute difference between the attributes, while the node removal cost was the value of the attribute itself. With this edit costs edit-distance is a normalized metric.

Figure 4 shows the distance matrices obtained using our metric and edit-distance. Note that, as before, both matrices exhibit the same block structure. We applied the same clustering algorithms used in the previous series of experiments. In order to assess the quality of the groupings, we used two well-known cluster-validation measures [8]. The first is the standard misclassification rate. We assigned to each cluster the class that has most members in the cluster. The members of the cluster that belong to a different class are considered misclassified. The misclassification rate is the percentage of misclassified shapes over the total number of shapes. To avoid the bias towards higher segmentation that this measure exhibits, we also used a second validation measure, i.e., the Rand index. We count the number of pairs of shapes that belong to the same class and that are clustered together and the number of pairs of shapes belonging to different




Fig. 4. Distance matrices from the second experiment. Left: Our metric. Right: Edit-distance

classes that are in different clusters. The sum of these two figures divided by the total number of pairs gives us the Rand index. Here, the higher the value, the better the classification.

Table 1 summarizes the results obtained using Normalized Cut and Dominant Sets. Here the two metrics generate clusters with comparable validation measures regardless of the clustering algorithm used.

Table 1. Validation measures of clusters obtained in the second experiment.

	Misclassification rate		Rand index	
	Normalized Cut	Dominant Sets	Normalized Cut	Dominant Sets
Our metric	23.3%	21.3%	90.3%	90.8%
Edit-distance	22.7%	24.0%	90.4%	90.8%

The last set of experiments was performed on a tree representation of Northern Lights [12]. As in the previous experiments, the representation used is derived from the morphological skeleton, but the choice of structural representation was different from the one adopted for shock-graphs, and the extracted trees tend to be larger. The database consisted of 1440 shapes. Using our metric we were able to extract the full distance matrix within a few hours, but it was unfeasible to compute edit-distance on the entire database. For this reason, in order to be able to compare the results with edit-distance, we also performed experiments using a smaller database consisting of 50 shapes. The calculation of edit-distance, even on this reduced database, took a full weekend.

In this case, we did not have the ground truth for the class memberships, so we needed a different cluster-validation measure. We opted for a standard measure that favors compact and well-separated clusters: the Davies-Bouldin index [8]. Let e_i be the average distance between elements in class i , and d_{ij} the average distance between elements in cluster i and elements in cluster j . The




Fig. 5. Distance matrices from the second experiment. Left: Our metric. Right: Edit-distance.

Davies-Bouldin index is

$$DB = \frac{1}{c} \sum_{i=1}^c \max_j R_{ij} \tag{7}$$

where c is the number of clusters and $R_{ij} = \frac{e_i + e_j}{d_{ij}}$ is the cluster separation measure. Clearly, lower values correspond to better separated and more compact clusters.

Table 2 provides the values of the Davies-Bouldin index on the clusters extracted using Normalized Cut and the Dominant Sets algorithm. As was the case with the previous experiments, both metrics produced comparable results.

Table 2. Davies-Bouldin index of clusters obtained in the third experiment.

	Normalized Cut	Dominant Sets
Our metric	0.0486	0.0723
Edit-distance	0.0232	0.0635

6 Conclusions

In this paper we have presented a novel distance measure for attributed trees based on the notion of a maximum similarity subtree isomorphism, and provided a polynomial-time algorithm to calculate it. We have proven that this measure satisfies the metric properties and have experimentally validated its usefulness by comparing it with edit-distance on three different shape recognition tasks. Our experimental results show that, in terms of quality, the proposed metric compares well with edit-distance, its computation being, however, orders of magnitude faster.

References

1. R. K. Ahuja, T. L. Magnanti, J. B. Orlin. *Network Flows*. Prentice-Hall, Upper Saddle River, NJ, 1993.
2. H. Bunke and K. Shearer. A graph distance metric based on the maximal common subgraph. *Pattern Recognition Letters*, 19:255–259, 1998.
3. S. J. Dickinson, A. P. Pentland, and A. Rosenfeld. 3-D shape recovery using distributed aspect matching. *PAMI*, 14(2):174–198, 1992.
4. M. A. Eshera and K.-S. Fu. An image understanding system using attributed symbolic representation and inexact graph-matching. *PAMI*, 8:604–618, 1986.
5. M. L. Fernandez and G. Valiente. A graph distance metric combining maximum common subgraph and minimum common supergraph. *Pattern Recognition Letters*, 22:753–758, 2001.
6. D. Hidović and M. Pelillo. Metrics for attributed graphs based on the maximal similarity common subgraph. *Int. J. Pattern Recognition Artif. Intell.*, 2004 (in press).
7. S. Ioffe and D. A. Forsyth. Human tracking with mixtures of trees. in *Proc. ICCV*, Vol. I, pp. 690–695, 2001.
8. A. K. Jain and R. C. Dubes. *Algorithms for Clustering Data*. Prentice Hall, Englewood Cliffs, NJ, 1988.
9. D. W. Matula. An algorithm for subtree identification, *SIAM Review*, 10:273–274, 1968.
10. M. Pavan and M. Pelillo. A new graph-theoretic approach to clustering and segmentation. In *Proc. CVPR*, Vol. I, pp. 145–152, 2003.
11. M. Pelillo, K. Sidiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *PAMI*, 21(11):1105–1120, 1999.
12. M. Peura. Attribute trees in image analysis: Heuristic matching and learning techniques. In *Proc. Int. Conf. Image Anal. Processing*, pp. 1160–1165, 1999.
13. T. B. Sebastian, P. N. Klein, and B. B. Kimia. Recognition of shapes by editing their shock graphs. *PAMI*, to appear, 2004.
14. J. Shi and J. Malik, Normalized cuts and image segmentation, *PAMI*, 22(8):888–905, 2000.
15. K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. *Int. J. Computer Vision*, 35(1):13–32, 1999.
16. A. Torsello and E. R. Hancock. A skeletal measure of 2D shape similarity. In C. Arcelli, L. P. Cordella, G. Sanniti di Baja (Eds.), *Visual Form 2001*. Springer-Verlag, Heidelberg (LNCS 2059), pages 260–271, 2001.
17. A. Torsello and E. R. Hancock. Efficiently computing weighted tree edit-distance using relaxation labeling. In M. Figueiredo, J. Zerubia, A.K. Jain (Eds.), *Energy Minimization Methods in Computer Vision and Pattern Recognition*. Springer-Verlag, Heidelberg (LNCS 2134), pp. 438–453, 2001.
18. W. H. Tsai and K.-S. Fu. Error-correcting isomorphism of attributed relational graphs for pattern analysis. *IEEE Trans Syst. Man Cybern.*, 9:757–768, 1979.
19. G. Valiente. An efficient bottom-up distance between trees, in *Proc. Int. Symp. String Processing Information Retrieval*, pp. 212–219, 2001.
20. W. D. Wallis, P. Shoubridge, M. Kraetz, and D. Ray. Graph distances using graph union. *Pattern Recognition Letters*, 22:701–704, 2001.
21. J. T.-L. Wang and K. Zhang. Finding similar consensus between trees: An algorithm and a distance hierarchy. *Pattern Recognition*, 34:127–137, 2001.
22. K. Zhang. A constrained edit-distance between unordered labeled trees. *Algorithmica*, 15:205–222, 1996.

23. K. Zhang and D. Shasha. Simple fast algorithms for the editing distance between trees and related problems. *SIAM J. Comput.*, 18:1245–1262, 1989.
24. K. Zhang, R. Statman, and D. Shasha. On the editing distance between unordered labeled trees. *Inform. Process. Letters*, 42:133–139, 1992.
25. K. Zhang, J. T. L. Wang, and D. Shasha. On the editing distance between undirected acyclic graphs. *Int. J. Found. Computer Sci.*, 7(1):43–57, 1996.