



A skeletal measure of 2D shape similarity

Andrea Torsello* and Edwin R. Hancock

Department of Computer Science, University of York, Heslington, York YO10 5DD, UK

Received 10 August 2001; accepted 25 March 2004

Available online 19 May 2004

Abstract

This paper presents a geometric measure that can be used to gauge the similarity of 2D shapes by comparing their skeletons. The measure is defined to be the rate of change of boundary length with distance along the skeleton. We demonstrate that this measure varies continuously when the shape undergoes deformations. Moreover, we show that ligatures are associated with low values of the shape-measure. The measure provides a natural way of overcoming a number of problems associated with the structural representation of skeletons. The first of these is that it allows us to distinguish between perceptually distinct shapes whose skeletons are ambiguous. Second, it allows us to distinguish between the main skeletal structure and its ligatures, which may be the result of local shape irregularities or noise. We illustrate how the new shape-measure can be used for the purposes of clustering shock-trees of the same shape class.

© 2004 Elsevier Inc. All rights reserved.

1. Introduction

The skeletal abstraction of 2D and 3D objects has proved to be an alluring yet highly elusive goal for over 30 years in shape analysis. The topic is not only important in image analysis, where it has stimulated a number of important developments including the medial axis transform and iterative morphological thinning operators, but is also an important field of investigation in differential geometry and biometrics where it has led to the study of the so-called Blum skeleton [8]. Because of this, the quest for reliable and efficient ways of computing skeletal shape

* Corresponding author.

E-mail addresses: atorsell@cs.york.ac.uk (A. Torsello), erh@cs.york.ac.uk (E.R. Hancock).

descriptors has been a topic of sustained activity. Recently, there has been a renewed research interest in the topic which has been aimed at deriving a richer description of the differential structure of the object boundary. This literature has focused on the so-called shock-structure of the reaction-diffusion equation for object boundaries [19,20].

Skeleton-based representations are just one of the possible abstractions of shape. Examples in the literature include boundary curves [6], regions [12], boundary statistics [7,18], and the medial axis [23,41]. Focusing on these methods in more detail, Basri et al. [6] opt for curve-based representations where object boundaries are matched by minimizing a curve functional that penalizes stretching and bending. Carson et al. [12], on the other hand, opt to represent a shape by clustering the set of interior pixels into a color–texture–position feature space. The clusters of pixels are then compared using a quadratic distance function. According to the shape-contexts of Belongie et al. [7] matching is effected using unordered boundary points labeled with a coarse histogram of the relative position of the remaining boundary points. Working at a more abstract structural level, Ioffe and Forsyth [18] learn an hierarchical distribution of image features represented as a mixture of trees. There are also examples in the literature of skeletal representations that are not based on the morphological skeleton. Among these are the shape axis representation of Liu and Geiger [23]. Here the skeleton is not defined using the symmetry axis, but as the midpoint between two corresponding boundary points on opposite sides of the shape. Another important skeleton-based representation is that used in the FORMS system [41]. In this work the medial axis is matched to a model skeleton using a branch and bound strategy.

1.1. Motivation

Broadly speaking the representation and recognition of 2D shapes based on the shock representation is a three-stage process. First, the skeleton must be computed from the available shape-boundary information [1,2,10,11,24,25,34]. The second issue is how to use the extracted skeleton to represent the differential structure of the original boundary [9,20,35]. The final step is the matching of the resulting shape representation [21,27,32,33,36,37]. Based on this three-step view, we provide a brief analysis of the related literature.

The idea of characterizing boundary shape using the differential singularities of the reaction equation was first introduced into the computer vision literature by Kimia et al. [20]. The idea is to evolve the boundary of an object to a canonical skeletal form using the reaction-diffusion equation. The skeleton represents the singularities in the curve evolution, where inward moving boundaries collide. The reaction component of the boundary motion corresponds to morphological erosion of the boundary, while the diffusion component introduces curvature dependent boundary smoothing. In practice, the skeleton can be computed in a number of ways [1,2,24]. Recently, Siddiqi, Tannenbaum and Zucker have shown how the eikonal equation which underpins the reaction-diffusion analysis can be solved using the Hamilton–Jacobi formalism of classical mechanics [11,34].

One of the criticisms that can be leveled at existing skeletonization methods is their sensitivity to small boundary deformations or ligatures. Although these can be reduced via curvature dependent smoothing, they may have a significant effect on the topology of the extracted skeleton.

Once the skeleton is on hand, the next step is to devise ways of using it to characterize the shape of the original boundary shape. Most of the approaches reported in the literature opt to use a structural characterization. For instance, Zucker, Siddiqi and others [36] have labeled points on the skeleton using so-called shock-labels. According to this taxonomy of local differential structure, there are different classes associated with behavior of the radius of the bitangent circle inscribed in the shape. The so-called shocks distinguish between the cases where the local bitangent circle has maximum radius, minimum radius, constant radius or a radius which is strictly increasing or decreasing. Kimia and Giblin opt for a simpler representation which is based just on the junctions and terminations of the skeleton [37].

Once the skeletal representation is on hand then shapes may be matched by comparing their skeletons. One approach in the literature adopts a structural approach to the matching problem. For instance, Pelillo et al. [27] use a sub-tree matching method. The shock-tree is attributed with the length, radii, velocities, and curvatures of the shocks. This method is potentially vulnerable to structural variations or errors due to local deformations, ligature instabilities or other boundary noise.

One of the criticisms of these structural-matching methods is that small boundary deformations may significantly distort the topology of the skeleton. To overcome the susceptibility of skeletal topology to noise and small deformation, Siddiqi and Zucker [33,35,36] label the shocks generated by the eikonal equation with their time of formation. The later the time of formation, and hence their proximity to the center of the shape, the higher the shock in the hierarchy. This temporal notion of relevance is based on the observation that the skeletal branches generated by noise and high-frequency features are always close to the border. Unfortunately, the converse does not hold. To give an example, a protrusion that ends on a vertex will always have the earliest time of creation, regardless of its relative relevance to the shape. For this reason the time of formation is not an effective measure of branch relevance in the presence of sharp boundary structure or high-curvature features.

Kimia and Klein, and their co-workers [21,29,37], have a potentially more robust method which matches by minimizing graph-edit distance. In particular, Sebastian et al. [29,30] have developed a variational method which can be used to measure the cost of boundary deformation, which they refer to as “edit distance” [29]. The cost of removing a branch of the skeleton is related to the associated boundary deformation. The distance measure based on this skeleton editing procedure has been successfully used to index and retrieve shapes from a large database [30]. However, the method is cumbersome since it requires alignment and explicit comparison of the boundary, and hence cannot be encoded on the skeleton alone.

Golland and Grimson [16] provide an interesting alternative: they minimize a boundary functional to find the optimal fit to a fixed model skeleton. This approach is very robust to boundary deformations, but is computationally very expensive. Therefore, it is not well suited to indexing large databases of shapes.

1.2. Contribution

We draw three observations from this review of the related literature. The first is that if a largely structural representation of the skeleton is used, then shapes which are perceptually different but which give rise to the same skeleton topology can be ambiguous with one-another. For this reason we would like to develop a representation which can be used to assess the differences in shape for objects which have topologically identical skeletons. Second, we would also like to be able to make comparisons between shapes that are perceptually close, but whose skeletons exhibit topological differences due to small but critical local shape deformations. Third, we aim to do this without making detailed boundary comparisons. In particular, we wish to construct a representation which dispenses with the boundary, but encodes information concerning its shape on the skeleton.

To meet these goals, our shape-measure must have three properties. First, it must be continuous over local regions in shape-space in which there are no topological transitions. If this were the case then it can be used to differentiate shapes with topologically identical skeletons. Second, it must vary smoothly across topological transitions. This is perhaps the most important property since it allows us to define distances across transitions in skeleton topology. In other words, we can traverse the skeleton without encountering singularities. Third, it must distinguish between the principal component of the skeleton and its ligatures [4,5]. This will allow us to suppress instabilities due to local shape deformations.

Our approach to shape recognition and classification using a skeletal representation requires several components. In particular, we need to (a) extract the skeleton, (b) label the branches with some measure of shape-similarity, (c) calculate the global similarity of two shapes using an edit-distance where the shape-measure is used to determine edit cost, and, finally (d) use the overall similarity between shapes to learn shape categories. This paper focuses only on the computation of the shape-similarity measure. The remaining topics will be covered in separate publications. The experimental assessment of the shape-similarity measure requires as a prerequisite that we know the correct correspondences between skeletal branches are a prerequisite for the calculation of the similarity measure. In the complete shape-matching system, we would anticipate that the correspondences would be located using a tree-matching algorithm to minimize the edit-distance between structures. However, the description and analysis of such an algorithm is beyond the scope of this paper. For the purposes of making this paper self-contained, we perform most of our experiments with hand-picked correspondences.

We opt to use a shape-measure based on the rate of change of boundary length with distance along the skeleton. To compute the measure, we construct at each location on the skeleton the bitangent circle inscribed in the shape. This circle is centered on a skeletal point and is bitangent to the boundary at the two boundary points. Hence, each skeletal point is in correspondence with (at least) two points on the border. The rate of change of boundary length with distance along the skeleton is computed by taking neighboring points on the skeleton. The corresponding change in boundary length is computed by determining distance along the boundary between the corresponding

points of contact for the two bitangent circles. The boundary distances are averaged for the boundary segments either side of the skeleton.

This measurement has previously been used in the literature to express *relevance* of a branch when extracting or pruning the skeleton [24,25,31]. In [9], Blum and Nagel suggested that the border length to shock-length ratio could be used, together with other measures, to characterize the shape, but the reasons for its proposal were solely attributed to its ability to detect whether a skeletal section is a ligature. In practice they too used the measure only as a purely static measure of relevance, ignoring the properties of the measure when the shape undergoes deformation.

We show that the rate of change of boundary length with distance along the skeleton has a number of interesting properties. The consequence of these properties is that the descriptive content of the measure extends beyond simple feature saliency, and can be used to attribute the relational structure of the skeleton to achieve a richer description of shape. Furthermore, we demonstrate that there is an intimate relationship between the shape-measure and the divergence of the distance map. This is an important observation since the divergence plays a central role when the skeleton is computed using the Hamilton–Jacobi formalism to solve the eikonal equation.

Among the properties exhibited by this measure, we have that topological changes on the skeleton correspond to zero crossings. This means that ligatures are associated with a value of the measurement which is zero, and hence have neutral weight. Second, the measure does not change when the shape undergoes “bending.”

Hence, the contribution of the paper is as follows. Although the measure that we use has been known for some time, it has not been used for shape comparison. The novelty of our work resides in the use of the method to measure shape similarity. In particular, the method is simple. For instance, it does not require explicit boundary comparison. Moreover, it can be computed directly from the divergence analysis of the distance map. From a theoretical perspective, the contribution of this paper is to demonstrate the relationship of the measure to the divergence, and to illustrate a number of important properties that it possesses. Although this paper concerns shape-similarity, there is of course an underlying correspondence problem for the skeleton branches that must be solved if the measure is to be used for matching or recognition. When required we use the weighted tree-matching method outlined in [3] to compute correspondences.

The outline of the paper is as follows: Section 2 deals with skeleton extraction and introduces to the Hamilton–Jacobi framework. The shape-measure and its geometric properties are presented in Section 3. Section 4 builds on Sections 2 and 3 by showing how the measure extraction process can be integrated with the Hamilton–Jacobi approach to skeleton detection. Finally, experimental results are presented in Section 5 and Section 6 provides some conclusions and identifies directions for further investigation.

2. Skeleton detection

A great number of papers have been written on the subject of skeleton detection. The problem is a tricky one because it is based on the detection of singularities on the evolution of the eikonal equation on the boundary of the shape.

The eikonal equation is a partial differential equation that governs the motion of a wave front through a medium. In the case of a uniform medium the equation is

$$\frac{\partial}{\partial t} \vec{C}(t) = \alpha \vec{N}(t), \quad (1)$$

where $\vec{C}(t) : [0, s] \rightarrow \mathbb{R}^2$ is the equation of the front at time t , $\vec{N}(t) : [0, s] \rightarrow \mathbb{R}^2$ is the equation of the normal to the wave front in the direction of motion, and α is the propagation speed. As the wave front evolves, opposing segments of the wave front collide, generating a singularity. This singularity is called a shock and the set of all such shocks is the skeleton of the boundary defined by the original curve. This realization of the eikonal equation is also referred to as the reaction equation.

The geometric intuition underpinning the eikonal or reaction equation is fairly simple: each point on the shape-boundary moves at constant speed in a direction that is at each instant in time normal to the object boundary. When two such points collide a singularity arises. The skeleton is the set of such singularities. Alternatively, the skeleton can be thought of as the set of the centers of bitangent circles contained within the shape-boundary. It is easy to see the relationship between these two views. Since the boundary points move at a constant speed, they meet after having traveled the same distance. The distance traveled by the set of points that meet first is the radius of the bitangent circle and the original positions of these points are the points where the bitangent circle touches the shape-boundary.

2.1. The Hamilton–Jacobi framework

To detect the singularities in the eikonal equation we use the Hamilton–Jacobi approach presented by Siddiqi, Tannenbaum, and Zucker [11,34]. Here we review this approach.

We commence by defining a distance-map that assigns to each point on the interior of an object the closest distance D from the point to the boundary (i.e., the distance to the closest point on the object boundary). The gradient of this distance-map defines a field \vec{F} whose domain is the interior of the shape. The field is defined to be

$$\vec{F} = \nabla D, \quad (2)$$

where $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y})^T$ is the gradient operator. The trajectory followed by each boundary point under the eikonal equation can be described by the ordinary differential equation $\dot{\vec{x}} = \vec{F}(\vec{x})$, where \vec{x} is the coordinate vector of the point.

Siddiqi claims that this system is Hamiltonian at non-skeletal points. Hence at these non-skeletal points the field \vec{F} is conservative, or $\nabla \cdot \vec{F} = 0$ [26]. However, the total inward flux through the entire shape is non-zero. In fact, the flux is proportional to the length of the boundary.

The divergence theorem states that the integral of the divergence of a vector-field over an area is equal to the flux of the vector field over the enclosing boundary of that area. In our case,

$$\int_A \nabla \cdot \vec{F} d\sigma = \int_L \vec{F} \cdot \vec{n} dl = \Phi_A(\vec{F}), \quad (3)$$

where A is any area, \vec{F} is a field defined in A , $d\sigma$ is the area differential in A , dl is the length differential on the border L of A , and $\Phi_A(\vec{F})$ is the outward flux of F through the border L .

By virtue of the divergence theorem we have that, within the interior, there are points where the system is not conservative. The non-conservative points are those where the boundary trajectory is not well defined, i.e., where there are singularities in the evolution of the boundary. These points are the so-called shocks or skeleton of the shape- boundary. Shocks are thus characterized by locations where $\nabla \cdot \vec{F} < 0$.

Unfortunately, the assumption that the field \vec{F} is conservative does not take into account density effects due to a curved front in the boundary evolution. In such cases there still exists a conservative field $\vec{\zeta} = \rho\vec{F}$ whose direction is always parallel to \vec{F} . As a result, $\nabla \cdot \vec{F}$ can be negative at non-skeletal points corresponding to high-curvature fronts. In such cases, however, the skeletal points represent minima of the quantity $\nabla \cdot \vec{F}$ along the direction orthogonal to the skeleton. In this paper we will ignore the non-conservative effects at high-curvature fronts. However, in a recent paper we have performed a more detailed analysis which incorporates the effects of non-uniform density due to boundary curvature, and demonstrate that the resulting boundary evolution process leads to improvements in the properties of the located skeleton [40].

The extraction of skeletal points is reduced to the search for highly non-conservative points. Unfortunately, skeletal points are, also, ridges of the distance map D , that is $\vec{F} = \nabla D$ is not uniquely defined in those points, but have different values on opposite sides of the watershed. This means that the calculation of the derivatives of \vec{F} gives rise to numerical instabilities. To avoid this problem we can use the divergence theorem again. We approximate the divergence with the outward flux through a small area surrounding the point. That is $\nabla \cdot \vec{F}(\vec{x}) \cong \Phi_{\mathcal{U}}(\vec{F})(\vec{x})$, where \mathcal{U} is a *small* area containing x . Thus, calculating the flux through the immediate neighbors of each pixel we obtain a suitable approximation of $\nabla \cdot \vec{F}(\vec{x})$.

2.2. Locating the skeleton

The thinning of the points enclosed within the boundary to extract the skeleton is an iterative process which involves eliminating points with low inward flux. The steps in the thinning and localization of the skeleton are as follows:

- At each iteration of the thinning process we have a set of points that are candidates for elimination. We remove from this set the point with the lowest inward flux.
- Next we check whether the point is topologically simple, i.e., whether it can be eliminated without splitting the remaining point-set.
- If the point is not simple, then it must be part of the skeleton. Thus, we retain it.
- If the point is simple, then we check whether it is an endpoint. If the point is simple and not an endpoint, then we eliminate it from the image. If this were the case then we add to the candidate set the points in its 8-neighborhood that are still part of the thinned shape (i.e., points that were not previously eliminated).

- If a simple point is also an endpoint, then the decision of whether or not it will be eliminated is based on the inward flux value. If the flux value is below a certain threshold we eliminate the point in the manner described above. Otherwise we retain the point as part of the skeleton.

We initialize this iterative process by placing every boundary point in the candidate set. We iterate the process until we have no more candidates for removal. The residual points will all belong to the skeleton.

3. The shape-measure and its properties

When the skeleton is computed in this way, then the eikonal equation induces a map from a point in the skeleton to a set of points on the boundary of the shape. That is, there is a correspondence between a point on the skeleton and the set of points on the boundary whose trajectories intercept it under the motion induced by the eikonal equation. The cardinality of this set of corresponding points on the boundary can be used to classify the local topology of the skeleton in the following manner:

- the cardinality is greater than or equal to 3 for junctions.
- For endpoints the cardinality is a number from 1 to a continuum.
- For the general case of points on branches of the skeleton, the cardinality is exactly 2.

As a result of this final property, any segment of a skeleton branch s is in correspondence with two boundary segments l_1 and l_2 . This allows us to assign to a portion of the skeleton the portion of the boundary from which it arose (see Fig. 1). For each internal point in a skeleton branch, we can thus define the local ratio between the length of the generating boundary segment and the length of the generated skeleton segment. The rate of change of boundary length with skeleton length is defined to be

$$\frac{d}{ds} \|l\| = \frac{d}{ds} \|l_1\| + \frac{d}{ds} \|l_2\|, \quad (4)$$

where $\|l_1\|$ is the length of the segment l . This ratio is our measure of the relevance of a skeleton segment in the representation of the 2D shape-boundary.

Our proposal in this paper is to use this ratio as a measure of the local relevance of the skeleton to the boundary-shape description. In particular, we are interested in using the measure to identify ligatures [4]. Ligatures are skeleton segments that link the logically separate components of a shape (see Fig. 2). They are characterized by a high negative curvature on the generating boundary segment. The observation which motivates this proposal is that we can identify ligature by attaching to each infinitesimal segment of skeleton the length of the boundary that generated it. Under the eikonal equation, a boundary segment with high negative curvature produces a rarefaction front. This front will cause small segments to grow in length throughout their evolution, until they collide with another front and give rise to a so-called shock. This means that very short boundary segments generate very long skeleton branches. Consequently, when a skeleton branch is a ligature, then there is an asso-

ciated decrease in the boundary-length to shock-length ratio. As a result our proposed skeletal shape-measure “weights” ligature less than other points in the same skeleton branch.

To better understand the rate of decrease of the boundary length with skeletal length, we investigate its relationship to the local geometry of the bitangent circle inscribed within the object boundary. We have

$$\frac{d}{ds} \|l_1\| = \frac{\cos \theta}{1 - rk_1} \quad (5)$$

and, similarly,

$$\frac{d}{ds} \|l_2\| = \frac{\cos \theta}{1 - rk_2}, \quad (6)$$

where r is the radius of the bitangent circle and k_i is the curvature of the mapped segment on the boundary. The curvature is oriented inwards, that is, when the boundary bends towards the skeleton we have a positive curvature, while when the boundary bends away from the skeleton the curvature is negative. Finally, θ is the angle between the tangent to the skeleton and the tangent to the corresponding point on the boundary. These formulae show that the measure is inversely proportional to negative curvature and radius. That is, if we fix a negative curvature k_1 , the measure decreases as the skeleton gets further away from the border. Furthermore, the measure decreases faster when the curvature becomes more negative.

A second important property of the shape-measure is that its value varies smoothly across shape deformations, even when these deformations impose topological transitions to the skeleton. To demonstrate this property we make use of the taxonomy of topological transition of the skeleton compiled by Giblin and Kimia [15]. According to this taxonomy, a smooth deformation of the shape induces only two types of transition on the skeleton (plus their time reversals). The transitions are *branch contraction* and *branch splicing*. A deformation *contracts* a branch joining two junctions when it moves the junctions together. Conversely, it *splices* a branch when it reduces in size, smoothes out, or otherwise eliminates the protrusion or sub-part of a shape that generates the branch.

A deformation that contracts or splices a skeleton branch causes the global value of the shape-measure along the branch to go to zero as the deformation approaches the topological transition. This means that a decreasing length of boundary generates the branch, until the branch disappears altogether.

When a deformation causes a contraction transition, both the length of the skeleton branch and the length of the boundary segments that generate the branch go to zero. A more elusive case is that of splicing. Through a splicing deformation, a decreasing length of boundary maps to the skeleton branch. This is because either the skeleton length and its associated boundary length are both reduced, or because the deformation allows boundary points to be mapped to adjacent skeleton branches. For this reduction in the length of the generating boundary, we do not have a corresponding reduction of the length of the skeleton branch. In fact, in a splice operation the length of the skeleton branch is a lower bound imposed by the presence of

the ligature. This is the major cause of the perceived instability of the skeletal representation. Weighting each point on the boundary which gave rise to a particular skeleton branch allows us to eliminate the contributions from ligatures, thus smoothing the instability. Since a smooth shape deformation induces a smooth change in the boundary, the total shape-measure along the branch has to vary smoothly through any deformation.

Just like the radius of the bitangent circle, key shape elements such as necks and seeds are associated with local variations of the length ratio. For instance, a neck is a point of high rarefaction and, thus, a minimum of the shape-measure along the branch. A seed is a point where the front of the evolution of the eikonal equation concentrates, and so is characterized by a maximum of the ratio (see Figs. 1 and 2).

A third important property of the shape-measure is its invariance to bending of the shape. This invariance derives from the fact that, if we bend the shape, we lose from one side the same amount of boundary-length that we gain on the opposite side. This property was already identified in [8].




Fig. 1. Geometric quantities used in our analysis.




Fig. 2. Ligature points are generated by short boundary segments.

To prove the bending invariance, let k_s be the curvature on the skeleton, at point O . We can assume, without loss of generality, that at this point the skeleton is directed towards the border-segment dl_2 . Furthermore, let k_1 and k_2 be the inward curvatures on the corresponding boundary points, and let θ be the angle between the border tangents and the skeleton tangent. At the point O the tangent angle and the radius are linked by the relation $dr/ds = -\sin(\theta)$. We define the *radius curvature* k_r as

$$k_r = \frac{d\theta}{ds} = \frac{d^2r/d^2s}{\sqrt{1 - (dr/ds)^2}}.$$

This quantity represents the degree to which the boundary bends towards the skeleton. Positive values indicate that the boundary is convex with respect to the skeleton (i.e., bends towards the skeleton) and negative values that the boundary is concave with respect to the skeleton (i.e., bends away from the skeleton). Given the boundary segments dl_1 , we can perform a parallel transport of the segment along the radius, hence obtaining the segment dl_1^p that is parallel to dl_1 and that crosses the skeleton at point O . Similarly we can obtain the segment dl_2^p parallel to dl_2 . The length of these segments is $\|dl_1^p\| = \|dl_2^p\| = \cos(\theta) ds + O(ds^2)$, and their curvatures are k_1^p and k_2^p , respectively. Moving along the skeleton by a distance ds , the tangent to the skeleton rotates by an angle $d\alpha = k_s ds + O(ds^2)$, while the tangent at the corresponding border points dl_1^p rotates by an angle $d\beta = k_1^p dl_1^p + O(ds^2)$. As can be clearly seen in Fig. 3, we have $d\theta = d\beta - d\alpha$. Hence:

$$k_r ds = k_1^p \cos(\theta) ds - k_s ds + O(ds^2),$$

$$k_1^p = \frac{k_r + k_s}{\cos(\theta)}.$$




Fig. 3. Differential geometry of a skeletal branch.

On the opposite side of the skeleton, since k_s points towards dI_2 , we have $d\theta = d\beta - d\alpha$. Hence

$$k_r ds = k_2^p \cos(\theta) ds + k_s ds + k_s ds + O(ds^2),$$

$$k_2^p = \frac{k_r - k_s}{\cos(\theta)}.$$

Recalling that $\frac{1}{k_1} = \frac{1}{k_1^p} + r$ and $\|dI_1\|k_1 = \|dI_1^p\|k_1^p + O(ds^2)$, we have $\|dI_1\| = \|dI_1^p\|(1 + rk_1^p) + O(ds^2) = [\cos(\theta) + r(k_r + k_s)] + O(ds^2)$. Similarly, we have $\|dI_2\| = [\cos(\theta) + r(k_r - k_s)] + O(ds^2)$. Hence $\frac{d}{ds}\|I_1\| + \frac{d}{ds}\|I_2\| = 2(\cos(\theta) + rk_r) = 2\cos(\theta) + 2r\frac{d\theta}{ds}$ is independent of k_s since the factors in $\|dI_1\|$ and $\|dI_2\|$ that depend on k_s cancel out. In other words, if we bend the object sufficiently to cause a curvature k in the skeleton, the increase in boundary length on the one side is compensated by the decrease in boundary length on the opposite side.

4. Measure extraction

The extraction of the skeletal shape-measure is a natural by-product which comes for free when we use the Hamilton–Jacobi approach for skeleton extraction. This is a very important property of this shape-measure. Through the divergence theorem we can transport a quantity linked to a potentially distant border to a quantity local to the skeleton. Using this property, we can prove that the border length to shock-length ratio is proportional to the divergence of the gradient of the distance map.

The Hamilton–Jacobi approach ensures that the eikonal equation induces a system that is conservative everywhere except on the skeleton. As we have already mentioned, the conservative field is not $\vec{F} = \nabla D$ as indicated by Siddiqi, but the field $\vec{\zeta} = \rho \vec{F}$ which is parallel to it (where ρ is a scalar function that corrects the density effects due to a curved front). In this paper, we will ignore these curvature dependency effects by assuming that \vec{F} is conservative, or, equivalently, that $\nabla \cdot \vec{F} = 0$ everywhere except on the skeleton. When dealing with these curvature effects, arguments based on the assumption $\nabla \cdot \vec{F} = 0$ will also hold when we substitute $\vec{\zeta}$ for \vec{F} . More details of the curvature dependent analysis can be found in our recent paper [40].

To show how the shape-measure can be computed in the Hamilton–Jacobi setting, we consider a skeleton segment s and its ϵ -envelope. The ϵ -envelope of a function f is the set of points $\{(x, y) \mid |y - f(x)| < \epsilon\}$, that is the set of points inside a “tube” of radius ϵ around the value of the function. The ϵ -envelope around a skeletal segment s is, thus, an infinitesimally thin open area that includes every point in s .

The segment s maps to two segment borders l_1 and l_2 . The evolution of the points in these border segments defines two areas A_1^ϵ and A_2^ϵ enclosed within the ϵ -envelope of s , the segments of boundary l_1 and l_2 , and the trajectories b_1^1 and b_1^2 , and b_2^1 and b_2^2 of the endpoints of l_1 and l_2 . The geometry of these areas is illustrated in Fig. 4.




Fig. 4. The flux through the border is equal to the flux through ϵ .

Since $\nabla \cdot \vec{F} = 0$ everywhere in A_1^ϵ and A_2^ϵ , by virtue of the divergence theorem we can state that the flux from the two areas are both zero, i.e., $\Phi_{A_1^\epsilon}(\vec{F}) = 0$ and $\Phi_{A_2^\epsilon}(\vec{F}) = 0$. The trajectories of the endpoints of the border are, by construction, parallel to the field, so the normal is everywhere normal to the field and thus there is no flux through the segments $b_1^1, b_1^2, b_2^1,$ and b_2^2 . On the other hand, the field on the shape-boundary is always normal to the boundary. Hence, the flux through the border segments l_1 and l_2 is equal to the length $\|l_1\|$ and $\|l_2\|$ of the segments l_1 and l_2 , respectively.

Since $\Phi_{A_1^\epsilon}(\vec{F}) = 0$ and $\Phi_{A_2^\epsilon}(\vec{F}) = 0$ the flux that enters through the border segments l_1 and l_2 has to exit through the ϵ -envelope of s . That is, if ϵ_1 and ϵ_2 are the sides of A_1^ϵ and A_2^ϵ on the ϵ -envelope of s , we have $\Phi_{\epsilon_1}(\vec{F}) = \Phi_{l_1}(\vec{F})$ and $\Phi_{\epsilon_2}(\vec{F}) = \Phi_{l_2}(\vec{F})$. This, in turn, implies that the flux through the whole ϵ -envelope of s is $\Phi_\epsilon(\vec{F}) = \|l_1\| + \|l_2\|$.

Since

$$\lim_{\epsilon \rightarrow 0} \int_{\epsilon} \nabla \cdot \vec{F} \, d\epsilon = \int_s \nabla \cdot \vec{F} \, ds$$

and the value of the flux through the ϵ -envelope of s is independent of ϵ , we have

$$\int_s \nabla \cdot \vec{F} \, ds = \|l_1\| + \|l_2\|.$$

Taking the first derivative with respect to ds we have, for each non-singular point in the skeleton,

$$\nabla \cdot \vec{F} = \frac{d}{ds} \|l_1\| + \frac{d}{ds} \|l_2\| \tag{7}$$

Fig. 5 plots at each skeletal point the extracted value of the shape-measure of each of two sample shapes.

4.1. Computing the distance between skeletons

This result allows us to calculate a global shape-measure for each skeleton branch during the branch extraction process. For our matching experiments we have used mainly a simple graph representation where the nodes are junctions or endpoints, and the edges are branches of the skeleton. In experiments, where we needed to compare potentially very different shapes, we used the shock-graph representation: the nodes still represent skeletal segments, but branches are split at points which are extremal in the radius of the bitangent circle. When we have completed the thinning of the shape-boundary and we are left only with the skeleton, we select an endpoint and start summing the values of the length ratio for each skeleton point until we reach either a junction or an extremal point. This sum $\sum_{i \in s} \nabla \cdot \vec{F}(x_i)$ over every pixel x_i of our extracted skeleton branch is an approximation of

$$\int_s \nabla \cdot \vec{F} ds = \int_s \left(\frac{d}{ds} \|l_1\| + \frac{d}{ds} \|l_2\| \right) ds = \|l_1\| + \|l_2\|,$$

the length of the border that generates the skeleton branch.

At this point we have identified a branch and we have calculated the total value of the length-ratio along that branch, or, in other words, we have computed the total length of the border that generated the branch. We continue this process until we have spanned each branch in the entire skeleton. Thus, we obtain a weighted graph representation of the skeleton. In the case of a simple shape, i.e., a shape with no holes, the graph has no cycles and thus is an (unrooted) tree.

Given this representation we can cast the problem of computing distances between different shapes as that of the total difference in shape-measure between corresponding branches.




Fig. 5. Two sample shapes. The height and intensity of the skeleton at each point is proportional to the shape-measure.

5. Experimental results

In this section we experiment with the new skeletal similarity measure. The experimentation is divided into three parts. First, we assess the ability of the proposed measure to discriminate between deformed shapes that give rise to skeletons with the same topology. Second, we assess how smoothly the overall similarity measure varies through transitions in skeletal topology. Finally, we show how the similarity measure may be used to cluster similar shapes.

There is clearly an underlying correspondence problem that must be solved before the similarity between two skeletons can be computed. This arises because we need to know how to associate branches in the two skeletons being compared. To fully perform a shape recognition task we should recover these correspondences automatically. However, the aim of the work reported here was to analyze the properties of our length ratio measure independently of the correspondence process, and not to solve the full shape recognition problem. Thus for our first two set of experiments we have located the branch correspondences by hand. To assess the overall effectiveness of the augmented skeletal representation, in the third set of experiments the correspondences are not hand-picked, but are in fact automatically extracted using a minimum edit-distance approach.

Tree edit distance is a generalization to trees of *String edit distance*. The edit distance is based on the existence of a set S of basic edit operations on a tree and a set C of costs, where $c_s \in C$ is the cost of performing the edit operation $s \in S$. The choice of the basic edit operations, as well as their cost, can be tailored to the problem, but common operations include leaf pruning, path merging, and, in the case of an attributed tree, change of attribute. Given two trees T_1 and T_2 , the set S of basic edit operations, and the cost of such operations $C = c_s, s \in S$, we call an *edit path* from T_1 to T_2 a sequence s_1, \dots, s_n of basic edit operations that transform T_1 into T_2 . The length of such path is $l = c_{s_1} + \dots + c_{s_n}$; the *minimum length edit path* from T_1 to T_2 is the path from T_1 to T_2 with minimum length. The length of the minimum length path is the tree edit distance.

With our measure assigned to each edge of the tree, i.e., branch of the skeleton, we define the cost of matching two edges as the difference of the total length ratio measure along the branches. The cost of eliminating an edge is equivalent to the cost of matching it to an edge with zero weight, i.e., one along which the total length ratio is zero.

5.1. Stability under deformation

As demonstrated earlier in the paper, we know that the length ratio measure should be stable to any local shape deformation, including those that exhibit an instability in shock length. This kind of behavior at local deformations is what has led to the idea that the skeleton is an unstable representation of shape.

To demonstrate the stability of the skeletal representation when augmented with the length ratio measurement, we have generated a sequence of images of a rectangle with a protrusion on one side (Fig. 6). The size of the protrusion is gradually reduced



Fig. 6. A “disappearing” protrusion which causes instability in shock-length, but not in our measure.

throughout the sequence, until it is completely eliminated in the final image. In Fig. 7 we plot the global value of the length ratio measure for the shock branch generated by the protrusion. It is clear that the value of the length ratio measure decreases monotonically and quite smoothly until it becomes zero when the protrusion disappears.

5.2. *Changes in skeleton topology*

In a second set of experiments we have aimed to assess the ability of the length ratio measure to distinguish between structurally similar shapes. To do this we selected two shapes that were perceptually different, but which possessed skeletons with a very similar topology. We, then, generated an image sequence in which the two shapes were morphed into one-another. Here the original shapes are the start and end frames of the sequence. At each frame in the sequence we calculated the distance between the start and end shapes.

We have repeated this experiment with two morphing sequences. The first sequence involved morphing a sand shark into a swordfish, while the second morphed a donkey into a hare (see Fig. 8).

To determine the difference between two shapes, we make use of hand-picked correspondences between skeletal branches. The distance between the complete skeletons is defined as the Euclidean distance between the normalized weights of matched edges (skeletal branches). In other words, the distance is

$$D(A, B) = \sqrt{\sum_i (e_i^A - e_i^B)^2}$$

where e_i^A and e_i^B are the normalized weights on the




Fig. 7. The measure of the skeleton segment generated by a protrusion.




Fig. 8. Morphing sequences and their corresponding skeletons. (A) Sand shark to swordfish sequence and (B) donkey to hare sequence.

corresponding edges indexed by i on the shapes denoted by A and B . The normalized weights are computed by dividing the raw weights by the sum of the weights of each tree.

We apply this normalized length ratio measure to ensure scale invariance. We note that two identical shapes scaled to different proportions would have different ratios due to the scale difference. However, the measure along equivalent branches of the two shapes would vary by a constant scale factor, namely the ratio of the lengths of the borders. Since the sum of the weights of the edges of a tree is equal to the total length of the border, dividing the weights in each branch by this quantity we have reduced the two measurements to the same scale. In this way the relevant quantity is not the absolute magnitude for a branch, but the magnitude ratio with other branches.

For each morphing sequence, in Fig. 9 we plot the distance between each frame in the sequence and the start and end frames. The monotonicity of the distance is evident throughout the sequences. This is a proof of capacity of our length ratio measure to disambiguate between shapes with topologically similar skeletons.




Fig. 9. Distances from first and last frame of the morphing sequences. (A) Distances in fish morphing sequence and (B) distances in donkey to hare morphing sequence.

To further assess the ability to discriminate between similar shapes, we selected a set of topologically similar shapes from a database of images of tools. As in the case of the previous experiments, the correspondences are hand-picked and the normalized Euclidean distance of the corresponding branch weights is used to measure the similarity of the skeletons. In the first column of Fig. 10 we show the selected shapes. To their right are the remaining shapes sorted by increasing normalized distance. Each shape is annotated by the value of the normalized distance.

It is clear that similar shapes are usually closest to one-another. However, there are problems due to a high sensitivity to occlusion. This can be seen in the high relative importance given to the articulation angle. This is due to the fact that, in the pliers images, articulation occludes part of the nose of pliers. While sensitivity to occlusion is, without a doubt, a drawback of the measure, we have to take into account that skeletal representation in general are highly sensitive to occlusion.

The reason that the monkey wrench is recorded as being more similar to the pliers than the second monkey wrench is due to sensitivity to articulation and the


















































						
	0.068	0.081	0.108	0.112	0.132	0.187
						
	0.053	0.054	0.069	0.080	0.087	0.112
						
	0.030	0.048	0.054	0.084	0.092	0.108
						
	0.021	0.030	0.054	0.081	0.086	0.110
						
	0.080	0.084	0.086	0.095	0.102	0.132
						
	0.088	0.092	0.095	0.110	0.128	0.186
						
	0.021	0.048	0.068	0.069	0.102	0.128

Fig. 10. Some tools and the normalized distance between them.

“closeness” of the head to the handles. Since the second monkey wrench is almost closed, the skeleton branches of the handles have a reduced overall weight. Thus, the process of normalizing the edge weights reduces the significance of the small yet salient head when it is compared to the remainder of the shape.

5.3. Clustering shock-graphs

The third set of experiments aims to establish the usefulness of the shape similarity measures as tool for clustering the shock-graphs associated with 2D shapes falling into distinct shape categories. We have compared the results obtained when the skeleton is both weighted with our measure, and when it is unweighted. In these experiments, we have used a database of shapes with very different topologies. Hence, there is no simple way to locate the correspondences between the skeletal branches. For this reason the correspondences were automatically extracted using a minimum edit distance graph-matching method [3]. Graph-edit distance allows the similarity between different graph structures to be measured. The idea underpinning the method 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 between two structures is found by searching for the sequence of edit operations that will make the two graphs isomorphic with one-another and which have minimum cost. In our application the available operations are node addition and removal. The cost of adding or removing a node with weight w is w , while the cost of matching two nodes with weights w and w' , respectively, is $|w - w'|$. To calculate the distance, we use the algorithm presented in [39].

Given the topological diversity of the shape skeletons, we have used a more powerful representation than the simple one used for the previous experiments. For this purpose we have a shock-graph representation in which we split skeletal branches at extrema of the radius of the bitangent circle.

The silhouettes used to generate the shock-graphs used in our experiments are shown in Fig. 11. There are 25 different shapes. These include brushes, tools, spectacles, various animals, and human hands. The figure is annotated with the pairwise similarity of the shapes. For the shapes indexed i and j , the similarity measure is defined to be

$$S_{i,j} = 1 - \frac{1}{2}d_{i,j},$$

where $d_{i,j}$ is the edit distance between shapes i and j .

For comparison purposes, Fig. 12 reports the similarity of the unweighted shock-trees. In this case the similarity between t_i and t_j is

$$S_{i,j} = \frac{1}{2} \frac{|t_i| + |t_j| - d_{i,j}}{2} \left(\frac{1}{|t_i|} + \frac{1}{|t_j|} \right),$$

where $|t|$ is the number of nodes in tree t and $d_{i,j}$ is the unattributed edit-distance between tree t_i and tree t_j .

1.000	1.000	1.000	0.774	1.000	0.786	0.889	0.889	0.706	0.643	0.850	0.536	0.536	0.565	0.818	0.792	0.889	0.850	0.804	0.719	0.635	0.640	0.706	0.694	0.684					
1.000	1.000	1.000	0.774	1.000	0.786	0.889	0.889	0.706	0.643	0.729	0.536	0.536	0.565	0.818	0.792	0.889	0.850	0.804	0.719	0.635	0.640	0.706	0.694	0.684					
1.000	1.000	1.000	0.774	1.000	0.786	0.889	0.889	0.706	0.643	0.729	0.536	0.536	0.565	0.701	0.792	0.889	0.850	0.804	0.719	0.635	0.640	0.706	0.694	0.684					
0.774	0.774	0.774	1.000	0.774	0.833	0.833	0.833	0.676	0.714	0.800	0.714	0.714	0.500	0.773	0.625	0.694	0.667	0.875	0.688	0.615	0.620	0.676	0.667	0.658					
1.000	1.000	1.000	0.774	1.000	0.786	0.889	0.889	0.706	0.643	0.850	0.536	0.536	0.565	0.818	0.792	0.889	0.850	0.804	0.719	0.635	0.640	0.706	0.694	0.684					
0.786	0.786	0.786	0.833	0.786	1.000	0.722	0.722	0.618	0.643	0.700	0.643	0.643	0.667	0.682	0.667	0.722	0.700	0.750	0.625	0.577	0.580	0.618	0.611	0.605					
0.889	0.889	0.889	0.833	0.889	0.722	1.000	1.000	0.765	0.730	0.950	0.548	0.548	0.486	0.808	0.681	0.778	0.739	0.826	0.781	0.673	0.680	0.765	0.750	0.737					
0.889	0.889	0.889	0.833	0.889	0.722	1.000	1.000	0.765	0.730	0.950	0.548	0.548	0.486	0.808	0.681	0.778	0.739	0.826	0.781	0.673	0.680	0.765	0.750	0.737					
0.706	0.706	0.706	0.676	0.706	0.618	0.765	0.765	1.000	0.782	0.715	0.456	0.456	0.498	0.674	0.640	0.680	0.715	0.643	0.667	0.730	0.791	0.765	0.743	0.724					
0.643	0.643	0.643	0.714	0.643	0.643	0.730	0.730	0.782	1.000	0.857	0.500	0.500	0.619	0.568	0.542	0.639	0.600	0.688	0.938	0.714	0.724	0.586	0.635	0.744					
0.850	0.729	0.729	0.800	0.850	0.700	0.950	0.950	0.715	0.857	1.000	0.600	0.600	0.458	0.764	0.642	0.739	0.700	0.788	0.812	0.692	0.700	0.794	0.778	0.763					
0.536	0.536	0.536	0.714	0.536	0.643	0.548	0.548	0.456	0.500	0.600	1.000	1.000	0.619	0.487	0.387	0.456	0.429	0.589	0.469	0.495	0.669	0.651	0.635	0.620					
0.536	0.536	0.536	0.714	0.536	0.643	0.548	0.548	0.456	0.500	0.600	1.000	1.000	0.619	0.487	0.387	0.456	0.429	0.589	0.469	0.495	0.669	0.651	0.635	0.620					
0.565	0.565	0.565	0.500	0.565	0.667	0.486	0.486	0.498	0.619	0.458	0.619	0.619	1.000	0.523	0.500	0.583	0.550	0.521	0.438	0.609	0.678	0.640	0.625	0.612					
0.818	0.818	0.701	0.773	0.818	0.682	0.808	0.808	0.674	0.568	0.764	0.487	0.487	0.523	1.000	0.784	0.909	0.764	0.864	0.767	0.647	0.655	0.824	0.806	0.789					
0.792	0.792	0.792	0.625	0.792	0.667	0.681	0.681	0.640	0.542	0.642	0.387	0.387	0.500	0.784	1.000	0.875	0.917	0.729	0.656	0.548	0.617	0.711	0.694	0.680					
0.889	0.889	0.889	0.694	0.889	0.722	0.778	0.778	0.680	0.639	0.739	0.456	0.456	0.583	0.909	0.875	1.000	0.844	0.826	0.694	0.598	0.604	0.765	0.750	0.737					
0.850	0.850	0.850	0.667	0.850	0.700	0.739	0.739	0.715	0.600	0.700	0.429	0.429	0.550	0.764	0.917	0.844	1.000	0.788	0.731	0.623	0.700	0.715	0.622	0.687					
0.804	0.804	0.804	0.875	0.804	0.750	0.826	0.826	0.643	0.688	0.788	0.589	0.589	0.521	0.864	0.729	0.826	0.788	1.000	0.656	0.654	0.660	0.735	0.722	0.711					
0.719	0.719	0.719	0.688	0.719	0.625	0.781	0.781	0.667	0.938	0.812	0.469	0.469	0.438	0.767	0.656	0.694	0.731	0.656	1.000	0.757	0.666	0.667	0.708	0.633					
0.635	0.635	0.635	0.615	0.635	0.577	0.673	0.673	0.730	0.714	0.692	0.549	0.495	0.609	0.647	0.548	0.598	0.623	0.654	0.757	1.000	0.706	0.730	0.752	0.729					
0.640	0.640	0.640	0.620	0.640	0.580	0.680	0.680	0.791	0.724	0.700	0.669	0.669	0.678	0.655	0.617	0.604	0.700	0.660	0.666	0.706	1.000	0.840	0.860	0.880					
0.706	0.706	0.706	0.676	0.706	0.618	0.765	0.765	0.765	0.586	0.794	0.651	0.651	0.640	0.824	0.711	0.765	0.715	0.735	0.667	0.730	0.840	1.000	0.972	0.836					
0.694	0.694	0.694	0.667	0.694	0.611	0.750	0.750	0.743	0.635	0.778	0.635	0.635	0.625	0.806	0.694	0.750	0.622	0.722	0.708	0.752	0.860	0.972	1.000	0.811					
0.684	0.684	0.684	0.658	0.684	0.605	0.737	0.737	0.724	0.744	0.763	0.620	0.620	0.612	0.789	0.680	0.737	0.687	0.711	0.633	0.729	0.880	0.836	0.811	1.000					

Fig. 12. Pairwise similarities between shapes for the unweighted shock-trees.

study. From the third row down errors begin to emerge. For instance, a monkey wrench (object 6) matches to a hammer (object 11), and a horse (object 22) matches to a hand (object 25). Although there are six such errors in the third row (objects 6, 10, 11, 14, 16, and 22), several of those are associated with small differences in similarity. This is the case with object 6, where a monkey wrench is matched to a hammer. In both objects the dominant feature is the long handle. Additionally, for four

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	
0.981	0.981	0.973	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.981	0.896	0.976	0.848	0.999	0.998	1.000	1.000	0.933	0.998	1.000	0.992	0.975	0.981	1.000
0.844	0.844	0.823	0.809	0.809	0.813	0.857	0.857	0.796	0.612	0.813	0.853	0.853	0.788	0.825	0.788	0.825	0.844	0.844	0.707	0.693	0.765	0.895	0.895	0.863	
0.823	0.748	0.736	0.790	0.755	0.790	0.714	0.796	0.727	0.606	0.651	0.756	0.747	0.784	0.755	0.747	0.784	0.744	0.748	0.613	0.586	0.718	0.863	0.817	0.817	
0.656	0.744	0.612	0.651	0.750	0.755	0.684	0.750	0.714	0.592	0.530	0.744	0.706	0.756	0.733	0.733	0.747	0.707	0.727	0.610	0.584	0.711	0.741	0.718	0.765	
0.552	0.736	0.594	0.648	0.727	0.608	0.666	0.720	0.648	0.584	0.506	0.741	0.692	0.755	0.727	0.678	0.700	0.700	0.688	0.610	0.577	0.693	0.711	0.687	0.747	
0.646	0.685	0.584	0.643	0.666	0.572	0.656	0.670	0.578	0.554	0.505	0.687	0.670	0.664	0.692	0.671	0.688	0.678	0.671	0.577	0.576	0.663	0.706	0.670	0.747	

Fig. 13. Top six matches for each shape for the weighted shock-trees.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
1.000	1.000	1.000	0.875	1.000	0.833	1.000	1.000	0.791	0.938	0.950	1.000	1.000	0.678	0.909	0.917	0.909	0.917	0.875	0.938	0.757	0.880	0.972	0.972	0.880
1.000	1.000	1.000	0.833	1.000	0.786	0.950	0.950	0.782	0.857	0.950	0.714	0.714	0.667	0.864	0.875	0.889	0.850	0.864	0.813	0.752	0.860	0.840	0.86	0.836
1.000	1.000	1.000	0.833	1.000	0.786	0.889	0.889	0.765	0.782	0.857	0.669	0.669	0.640	0.824	0.792	0.889	0.850	0.826	0.781	0.730	0.840	0.836	0.811	0.811
0.889	0.889	0.889	0.833	0.889	0.786	0.889	0.889	0.765	0.744	0.850	0.651	0.651	0.625	0.818	0.792	0.889	0.850	0.826	0.781	0.730	0.791	0.824	0.806	0.789
0.889	0.889	0.889	0.800	0.889	0.786	0.889	0.889	0.765	0.730	0.850	0.643	0.643	0.619	0.818	0.792	0.889	0.850	0.826	0.767	0.729	0.724	0.794	0.778	0.763

Fig. 14. Top six matches for each shape for the unweighted shock-trees.

of the objects the correct matches appear in the fourth (object 6, 16), fifth (object 14), or sixth (object 22) position. It is only the two hammers that pose a real problem. This is due to the fact that the handle, the main feature on both objects, shows variation in its differential properties. Specifically, object 10 bulges on the grip, creating a type one shock that splits the handle, whereas the handle of object 11 generates a single shock segment. The problem could be solved by allowing the edit distance calculation to merge segments, but this is beyond the scope of this paper.

Fig. 14 displays the top matches obtained using unattributed shock-trees. Here too the top match is again always a perfect fit. However, the performance degrades

more quickly as we descend the rows of the table. In fact, the first error emerge in the second row of the figure.

To visualize the pairwise relationships between the different shapes, we have performed multidimensional scaling on the set of pairwise similarities. Multidimensional scaling is a well-known statistical technique for visualizing data which exists in the form of pairwise similarities rather than ordinal values. Stated simply, the method involves embedding the objects associated with the pairwise distances in a low-dimensional space. This is done by performing principal components analysis on the matrix of pairwise similarities, and projecting the original objects into the resulting eigenspace. The objects are visualized by displaying their positions in the space spanned by the leading eigenvectors. The method has been widely exploited for data analysis in the psychology literature. A comprehensive review can be found in the recent book of Cox and Cox [14]. Details of the procedure can be found in Appendix A.

The projections of the edit distances onto the 2D space spanned by the two leading eigenvectors are shown in Figs. 15 (when the skeleton is weight with the edit distance) and 16 (when it is not). When the skeleton is weighted with the measure (Fig. 15), then the MDS projection reveals some class structure emerging. However, the full shape-structure is not captured by the two leading eigenvectors. For instance, the hands, the fishes, the tools, and the brushes all appear close to each other. However, there is no clear delineation of the shape-classes. When the skeleton is not weighted using the measure (Fig. 16) then the grouping of the shapes is even poorer, with only the spectacles forming a well-separated group.

Encouraged by these results, we have performed a detailed pairwise clustering of the pattern of similarities. Here we use the method recently described by Robles-




Fig. 15. First and second principal components of the edit distances of the shapes for the weighted shock-trees.



Fig. 16. First and second principal components of the edit distances of the shapes for the unweighted shock-trees.

Kelly and Hancock [28]. Details of the clustering algorithm are outside the scope of this paper. However, the method uses an iterative log-likelihood algorithm to identify the pairwise clusters via matrix factorization. The initial and final matrices of pairwise distance are shown in Fig. 17 for the measure-weighted skeleton and Fig. 18 for the unweighted skeleton. In the case of the weighted skeleton the initial pairwise similarity matrix shows a strong separation of the shape-groups, which is further re-enforced by the iterative clustering method. Based on the block structure of the final matrix of pairwise distances, we identify eight clusters. In the order of importance, they are




Fig. 17. (A) Initial similarity matrix for the weighted tree edit distances; (B) final similarity matrix for the weighted tree edit distances.




Fig. 18. (A) Initial similarity matrix for the unweighted tree edit distances; (B) final similarity matrix for the unweighted tree edit distances.

-
-
-
-
-
-
-
-


In other words, the hands, tools, spectacles, and animals form clusters. However, there are shapes which leak between these clusters. The problems encountered above are due to the fact that certain shapes straddle the true shape-classes and cause cluster-merging. When a pruned set of 16 shapes is used, then the following set of clusters emerges:

-
-
-
-
-
-
-




This is a much better set of clusters, that reflect the true shape-classes in the data.

We have repeated these clustering experiments with the unweighted skeletons. Here the initial pairwise similarity matrix contains less structure than in the weighted case, and iteration of the clustering algorithm results in a noisier set of final cluster membership indicators (Fig. 12). In particular, the clusters extracted from

unweighted shock-trees do not appear to correlate well with the shape classes in the database. The ordered clusters are listed below:

- 
- 
- 
- 
- 

Clearly there is a considerable merging of and leakage between clusters. As illustrated below, even reducing the size of the database does not improve the classification:

- 
- 
- 
- 
- 
- 

6. Conclusions

In this paper we presented a shape-measure defined on the skeleton. This quantity has been used in the literature as a branch relevance measure during skeleton extraction and pruning. We show that the measure has greater informational utility, and can be used to augment the purely structural information residing in a skeleton in order to perform shape indexation and matching tasks. We show that the shape-measure has a number of interesting properties that allow it to distinguish between structurally similar shapes. In particular, the measure (a) changes smoothly through topological transitions of the skeleton, (b) is able to distinguish between ligature and non-ligature points and to weight them accordingly, and (c) it exhibits invariance under “bending.” What makes the use of this measure particularly appealing is the fact that it can be calculated with no added effort when the skeleton is computed using the Hamilton–Jacobi method of Siddiqi et al. [34].

We have performed some experiments which verify the usefulness of the proposed shape-measure. These showed that the distance increases monotonically as shapes undergo smooth deformation. Moreover, the shape-measure can be used to cluster perceptually similar shapes. We acknowledge that the method has been evaluated on a relatively small set of shapes, and that large-scale experiments are needed. Our future plans involve investigating whether the measure can be used to identify small but perceptually salient features. Moreover, we are currently studying the effects of curvature in the computation of Hamilton–Jacobi skeleton and how curvature impacts on the computation of the measure.

Appendix A. Multidimensional scaling

Multidimensional scaling (MDS) [13] is a procedure which allows data specified in terms of a matrix of pairwise distances to be embedded in a Euclidean space. The classical multidimensional scaling method was proposed by Torgenson [38] and Gower [17]. Shepard and Kruskal [22] developed a different scaling technique called ordinal scaling. Here we intend to use the method to embed shock-trees in a low-dimensional space.

Suppose that $d_{i1,i2}$ is the edit-distance between the shock-trees indexed $i1$ and $i2$. The first step of MDS is to calculate a matrix T whose element with row r and column c is given by

$$T_{rc} = -\frac{1}{2}[d_{rc}^2 - \hat{d}_r^2 - \hat{d}_c^2 + \hat{d}_{..}^2], \quad (\text{A.1})$$

where

$$\hat{d}_r = \frac{1}{N} \sum_{c=1}^N d_{rc} \quad (\text{A.2})$$

is the average dissimilarity value over the r th row, \hat{d}_c is the similarly defined average value over the c th column and

$$\hat{d}_{..} = \frac{1}{N^2} \sum_{r=1}^N \sum_{c=1}^N d_{r,c} \quad (\text{A.3})$$

is the average similarity value over all rows and columns of the similarity matrix T .

We subject the matrix T to an eigenvector analysis to obtain a matrix of embedding coordinates X . If the rank of T is k , $k \leq N$, then we will have k non-zero eigenvalues. We arrange these k non-zero eigenvalues in descending order, i.e., $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0$. The corresponding ordered eigenvectors are denoted by \vec{e}_i where λ_i is the i th eigenvalue. The embedding coordinate system for the shock-trees is

$$X = [\vec{f}_1, \vec{f}_2, \dots, \vec{f}_k], \quad (\text{A.4})$$

where $\vec{f}_i = \sqrt{\lambda_i} \vec{e}_i$ are the scaled eigenvectors. For the shock-tree indexed i , the embedded vector of coordinates is

$$\vec{x}_i = (X_{i,1}, X_{i,2}, X_{i,3})^T. \quad (\text{A.5})$$

References

- [1] C. Arcelli, G. Sanniti di Baja, A width-independent fast thinning algorithm, IEEE Trans. PAMI 7 (4) (1985) 463–474.
- [2] C. Arcelli, G. Sanniti di Baja, Ridge points in Euclidean distance maps, Pattern Recogn. Lett. 13 (1992) 237–243.
- [3] A. Torsello, E.R. Hancock, Computing approximate tree edit distance using relaxation labelling, Pattern Recogn. Lett. 24 (2003) 1089–1097.
- [4] J. August, K. Siddiqi, S. Zucker, Ligature instabilities in the perceptual organization of shape, Comput. Vision Image Understand. 76 (3) (1999) 231–243.

- [5] J. August, A. Tannenbaum, S.W. Zucker, On the evolution of the skeleton, in: Seventh Internat. Conf. on Computer Vision, IEEE, IEEE Computer Society, New York, Silver Spring, MD, September 1999, pp. 315–322.
- [6] R. Basri, L. Costa, D. Geiger, D. Jacobs, Determining the similarity of deformable shapes, *Vision Res.* 38 (1998) 2365–2385.
- [7] S. Belongie, J. Malik, J. Puzicha, Shape matching and object recognition using shape contexts, *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (4) (2002).
- [8] H. Blum, Biological shape and visual science (part I), *J. Theor. Biol.* 38 (1973) 205–287.
- [9] H. Blum, R.N. Nagel, Shape description using weighted symmetric axis features, *Pattern Recogn.* 10 (1978) 167–180.
- [10] G. Borgefors, G. Ramella, G. Sanniti di Baja, Multi-scale skeletons via permanence ranking, in: *Advances in Visual Form Analysis*, World Scientific, Singapore, 1997, pp. 31–42.
- [11] S. Bouix, K. Siddiqi, Divergence-based medial surfaces, in: *Computer Vision ECCV 2000*, vol. 1, LNCS, Springer, 2000, pp. 603–618, LNCS, 1842.
- [12] C. Carson, S. Belongie, H. Greenspan, J. Malik, Blobworld: color- and texture-based image segmentation using EM and its application to image querying and classification, *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (8) (2002) 1026–1038.
- [13] C. Chatfield, A.J. Collins, *Introduction to Multivariate Analysis*, Chapman & Hall, London, 1980.
- [14] T.F. Cox, M.A.A. Cox, *Multidimensional Scaling*, Chapman & Hall, London, 1994.
- [15] P.J. Giblin, B.B. Kimia, On the local form and transitions of symmetry sets, medial axes, and shocks, in: Seventh Internat. Conf. on Computer Vision, IEEE, IEEE Computer Society, New York, Silver Spring, MD, September 1999, pp. 385–391.
- [16] P. Golland, E.L. Grimson, Fixed topology skeletons, in: *Conf. on Computer Vision and Pattern Recognition*, vol. 1, June 2000, pp. 10–17.
- [17] J.C. Gower, Some distance properties of latent root and vector methods used in multivariate analysis, *Biometrika* 23 (1964) 325–328.
- [18] S. Ioffe, D.A. Forsyth, Mixtures of trees for object recognition, in: *IEEE Comput. Soc. Conf. on Computer Vision and Pattern Recognition*, vol. II, 2001, pp. 180–185.
- [19] B.B. Kimia, K. Siddiqi, Geometric heat equation and nonlinear diffusion of shapes and images, *Comput. Vision Image Understand.* 64 (3) (1996) 305–322.
- [20] B.B. Kimia, A.R. Tannenbaum, S.W. Zucker, Shapes, shocks, and deformations I, *Int. J. Comput. Vision* 15 (1995) 189–224.
- [21] P. Klein, S. Tirthapura, D. Sharvit, B.B. Kimia, A tree-edit-distance algorithm for comparing simple, closed shapes, in: *ACM-SIAM Symp. on Discrete Algorithms*, 1999.
- [22] J.B. Kruskal, R.N. Shepard, Nonmetric methods for scaling and for factor analysis, *American Psychologist* 29 (1964) 557–558.
- [23] T. Liu, D. Geiger, Approximate tree matching and shape similarity, *Int. Conf. Comput. Vision* (1999) 456–462.
- [24] R.L. Ogniewicz, A multiscale mat from voronoi diagrams: the skeleton-space and its application to shape description and decomposition, in: *Aspects of Visual Form Processing*, 2nd Internat. Workshop on Visual Form, World Scientific, Singapore, 1994, pp. 430–439.
- [25] R.L. Ogniewicz, O. Kübler, Hierarchic voronoi skeletons, *Pattern Recogn.* 28 (3) (1995) 343–359.
- [26] S.J. Osher, J.A. Sethian, Fronts propagating with curvature dependent speed: algorithms based on Hamilton–Jacobi formulations, *J. Comput. Phys.* 79 (1988) 12–49.
- [27] M. Pelillo, K. Siddiqi, S.W. Zucker, Matching hierarchical structures using association graphs, *IEEE Trans. Pattern Anal. Mach. Intell.* 21 (11) (1999) 1105–1120.
- [28] A. Robles-Kelly, E.R. Hancock, A maximum likelihood framework for iterative eigendecomposition, in: *Internat. Conf. on Computer Vision*, IEEE Computer Society Press, Silver Spring, MD, 2001, pp. 654–661.
- [29] T.S. Sebastian, P.N. Klein, B.B. Kimia, Recognition of shapes by editing shock graphs, in: *Internat. Conf. on Computer Vision*, vol. I, 2001, pp. 755–762.
- [30] T.S. Sebastian, P.N. Klein, B.B. Kimia, Shock-based indexing into large shape databases, in: *Eur. Conf. on Computer Vision*, vol. III, 2002, pp. 731–746.

- [31] D. Shaked, A.M. Bruckstein, Pruning medial axes, *Comput. Vision Image Understand.* 69 (2) (1998) 156–169.
- [32] D. Sharvit, J. Chan, H. Tek, B.B. Kimia, Symmetry-based indexing of image database, *J. Vis. Commun. Image Representation* 9 (4) (1998) 366–380.
- [33] A. Shokoufandeh, S.J. Dickinson, K. Siddiqi, S.W. Zucker, Indexing using a spectral encoding of topological structure, *Conf. Comput. Vision Pattern Recogn.* (1999).
- [34] K. Siddiqi, S. Bouix, A. Tannenbaum, S.W. Zucker, The Hamilton–Jacobi skeleton, in: *Seventh Internat. Conf. on Computer Vision, IEEE, IEEE Computer Society, New York, Silver Spring, MD, September 1999*, pp. 828–834.
- [35] K. Siddiqi, B.B. Kimia, A shock grammar for recognition, in: *Computer Vision and Pattern Recognition, IEEE Computer Society Press, Silver Spring, MD, 1996*, pp. 507–513.
- [36] K. Siddiqi, A. Shokoufandeh, S.J. Dickinson, S.W. Zucker, Shock graphs and shape matching, *Int. J. Comput. Vision* 35 (1) (1999) 13–32.
- [37] S. Tirthapura, D. Sharvit, P. Klein, B.B. Kimia, Indexing based on edit-distance matching of shape graphs, in: *SPIE Internat. Symp. on Voice, Video, and Data Communications, 1998*, pp. 25–36.
- [38] W.S. Torgerson, Multidimensional scaling. I: theory and method, *Psychometrika* 17 (1952) 401–419.
- [39] A. Torsello, E.R. Hancock, Efficiently computing weighted tree edit distance using relaxation labeling, in: *Energy Minimization Methods in Computer Vision and Pattern Recognition, 2001*, pp. 438–453.
- [40] A. Torsello, E.R. Hancock, Curvature correction of the Hamilton–Jacobi skeleton, in: *IEEE Comput. Soc. Conf. on Computer Vision and Pattern Recognition, vol. I, 2003*, pp. 828–834.
- [41] S.C. Zhu, A.L. Yuille, FORMS: a flexible object recognition and modelling system, *Int. J. Comput. Vision* 20 (3) (1996) 187–212.