Elastic Shape Analysis of Boundaries of Planar Objects with Multiple Components and Arbitrary Topologies

Sebastian Kurtek[1], Hamid Laga[2,3], and Qian Xie[4]

[1] Department of Statistics, The Ohio State University
[2] Phenomics and Bioinformatics Research Centre, University of South Australia
[3] Australian Centre for Plant Functional Genomics, PtyLtd
[4] Department of Statistics, Florida State University

Abstract. We consider boundaries of planar objects as level set distance functions and present a Riemannian metric for their comparison and analysis. The metric is based on a parameterization-invariant framework for shape analysis of quadrilateral surfaces. Most previous Riemannian formulations of 2D shape analysis are restricted to curves that can be parameterized with a single parameter domain. However, 2D shapes may contain multiple connected components and many internal details that cannot be captured with such parameterizations. In this paper we propose to register planar curves of arbitrary topologies by utilizing the re-parameterization group of quadrilateral surfaces. The criterion used for computing this registration is a proper distance, which can be used to quantify differences between the level set functions and is especially useful in classification. We demonstrate this framework with multiple examples using toy curves, medical imaging data, subsets of the TOSCA data set, 2D hand-drawn sketches, and a 2D version of the SHREC07 data set. We demonstrate that our method outperforms the state-of-the-art in the classification of 2D sketches and performs well compared to other state-of-the-art methods on complex planar shapes.

1 Introduction

Shape is an important feature for characterizing objects in various fields of science. Analyzing objects based on their shapes and modeling the variability they exhibit within and across classes are fundamental problems in computer vision and pattern recognition. There has been an increasing interest in using Riemannian frameworks for shape analysis of objects because of the breadth of tools that they offer. First, they allow one to remove all of the shape preserving transformations from the representation space. Second, they allow for computing statistics (e.g. means, covariances, modes of variation) of shapes. However, most of this work is limited to shape analysis of the outer boundaries of objects, i.e. curves [1–3] and surfaces [4–6]. There are very few papers that study both shape boundaries and their interiors. Fuchs [7] considers such a case but their

(a) Subcortical structures in the brain (b) 2D sketch data set [8] (c) Arbitrary shapes from SHREC07 [9]

Fig. 1. Examples of planar shapes of arbitrary topology, which may contain multiple parts and complex internal features.

approach is not invariant to re-parameterization and also only considers shapes of the same topology.

In this paper, we propose a novel framework for analyzing shapes of planar objects of arbitrary topologies that can have multiple components as shown in Fig. 1. We propose to represent planar objects as level sets of their Euclidean distance functions. We consider each such function along with its smoothed gradient as an image $f : [0, 1]^2 \to \mathbb{R}^3$. We thus formulate the problem of analyzing curves as a problem in image analysis, and adapt a recently proposed Riemannian framework developed for statistical shape analysis of quadrilateral surfaces to the problem at hand. This framework is especially useful in registering multiple (non-intersecting) curves or curves with differing topologies. The registration step is very important in shape analysis as it allows one to generate meaningful comparisons of shape, where similar features are optimally matched across objects. Furthermore, the proposed registration criterion is a proper parameterization-invariant distance between images, and thus has nice mathematical properties. Such a framework, in principle, also allows one to perform subsequent statistical analysis of such objects such as computing their sample statistics. While at this stage we do not provide a setting for computing geodesic paths and proper statistics in this framework (future work), we display linear interpolations between registered shapes which can be used to asses the computed distances and registrations. We demonstrate this framework with multiple examples on toy shapes and real data from medical imaging and graphics. We also provide shape classification results on two very complex data sets: 2D sketches [8] and a 2D version of the SHREC07 data set [9]. We show that our framework performs well compared to state-of-the-art feature-based methods on the SHREC07 data set and significantly outperforms the state-of-the-art on the sketch data set. Note that feature-based methods are unable to perform subsequent statistical analysis such as computing shape means or covariances.

Related work. The concept of shape spaces with an associated metric was first proposed by Kendall [10] and then further developed by Dryden and Mardia [11]. In those frameworks, shapes were represented with a set of landmark points in the Euclidean space and compared using a Riemannian framework on the

representation space of such objects. The main drawback of this approach is that the landmarks have to be detected and labeled before one can analyze the shapes. Alternatively, one can look at curves as continuous objects and represent them as elements of infinite-dimensional Riemannian manifolds. Zahn and Roskies [12] performed Fourier analysis on angle functions of arc-length parameterized curves. However, restricting the analysis to arc-length parameterized curves can be very limiting in practice. It is generally accepted that a better approach is to search for optimal parameterizations using a proper distance. Younes [1] and Younes et al. [2] defined parameterization-invariant Riemannian metrics on shape spaces of planar curves. Along this line of research, Klassen et al. [13] introduced a family of elastic metrics that quantify the relative amounts of bending and stretching needed to deform planar, closed curves into each other. Joshi et al. [14] extended this framework to curves in $\mathbb{R}^n$. Later on, Srivastava et al. [3] proposed a special representation of planar curves, called the square-root velocity function (SRVF), which simplifies the computation of geodesics and geodesic distances between curves. The use of this representation chooses a specific instance of the elastic metric by fixing the parameters that control the bending and stretching energies.

The main limitation of these works is that they are based on a 1D parameterization of curves and thus only single boundaries of objects can be analyzed. There are few works that are able to analyze groups of curves simultaneously or curves of differing topologies. Fuchs et al. [7] handle the interior of the shapes instead of only the shape boundary. Thus, this framework is naturally defined for multiple components. However, it requires the topology to stay the same during the evolution, which can be limiting in real applications. In addition, this method is not invariant to re-parameterization. Kerr et al. [15] developed statistical models for multiple closed planar curves in a parameterization-invariant framework, but it cannot handle objects of differing topologies.

Another common representation of shape is using distance functions and their level sets [16], which can represent shapes of arbitrary topologies. For example, to register planar shapes, Paragios et al. [17] used Euclidean distance transforms while Munim and Farag [18] and Fahmi and Farag [19] used vector distance function-based representations. These three papers, however, assume that the global alignment between a pair of shapes can be solved by finding the optimal translation, rotation and anisotropic scaling between the shapes. Thus, these approaches are not applicable to shapes that undergo large articulated and elastic motions, such as the ones we consider in this paper. Although Yezzi and Saoto [20] used level sets, their approach is limited to planar shapes composed of a single closed curve. Also, while these methods are invariant to translation, rotation, and global scaling, none of them is invariant to re-parameterization.

There is a series of related work on Large Deformation Diffeomorphic Metric Mapping (LDDMM) [21] where planar objects of interest (curves, landmarks, etc.) are embedded in a 2D domain and the full domains are matched and compared. This type of approach is similar in flavor to the proposed work, but the matching and comparison are performed differently. The LDDMM approach searches for a geodesic between the two images in the group of diffeomorphisms

with an additional data matching term. In the proposed approach, the distance is computed on the space of level set functions modulo re-parameterization.

Another set of approaches represent complex shapes using point sets and compare them using the Hausdorff distance [22, 23] or the symmetric area difference [24], which do not require a 1D parameterization of the objects. While these methods have proven very useful in a number of tasks including partial object matching, they often do not allow for elastic deformations between the objects of interest. Several other papers used shape descriptors to compare 2D and 3D objects [25–30]. The Inner Distance-based Shape Context (IDSC) of Ling and Jacobs [25] is robust to articulated motion and aware of inner holes in the shapes. This descriptor, however, works on closed planar shapes, with a single connected component and with clearly defined interior regions. Our approach does not have this restriction and can even handle stroke-based drawings, such as human-drawn sketches, as demonstrated in the experimental results section.

Overview and contributions. We propose to represent planar objects as level sets of distance functions. To compare these functions, we adapt a recent framework for shape analysis of quadrilateral surfaces proposed in [31], which provides a recipe for generating parameterization-invariant comparisons between shapes of surfaces. We utilize the same metric but adapt it to our problem. A similar approach was taken by Xie et al. [32] in the context of image registration. In addition to parameterization, we remove variability due to other shape preserving transformations such as translation, scale and rotation. Our contributions can be summarized as follows:

1. We formulate the problem of performing computations on the space of planar objects of arbitrary topology as the problem of analyzing their associated level set distance functions. For this purpose, we utilize the square root function transformation of surfaces. To the best of our knowledge, this is the first time that a parameterization-invariant framework (where optimal registrations are computed) is being proposed for the analysis of planar objects of arbitrary topology using the distance function representation.
2. We demonstrate this framework with several examples using toy shapes and real data such as medical images, 2D sketches, and 2D projections of natural and manmade 3D shapes. We consider examples involving multiple simple planar closed curves and planar curves with different topologies.
3. We demonstrate the utility of the proposed distance in two shape classification studies and show that it performs well compared to the feature-based state-of-the-art methods for analyzing such data. In particular, we show that the classification performance of the proposed framework outperforms, by more than 13%, the state-of-the-art on the hand-drawn sketches of Eitz et al. [8], suggesting that the proposed framework is suitable for the analysis of complex shapes that do not have clearly-defined interior regions.

The rest of this paper is organized as follows. In Section 2, we provide details of the mathematical framework we will use to analyze shapes of complex planar contours with arbitrary topologies. In particular, we discuss how different

shape preserving transformations are removed from the representation space. In Section 3, we present comparison and classification results on different types of data. We conclude in Section 4 and outline directions for future work.

2 Mathematical Framework

Several past papers have considered a planar object as a parameterized curve $\beta : D \rightarrow \mathbb{R}^2$, where D is a certain domain for the parameterization (e.g. $D = [0, 1]$ for open curves or $D = \mathbb{S}^1$ for closed curves). Unfortunately, using such a representation does not allow one to analyze shapes with interior details or multiple components as those shown in Fig. 1. For this purpose, we propose to utilize the level set function representation (e.g. Euclidean distance transform), which we will refer to as ψ. That is, in our framework $\psi : [0, 1]^2 \rightarrow \mathbb{R}$, and the object β is defined as its zero-level (isocontour) set. In this paper, we propose to adapt a recently developed framework for statistical shape analysis of parameterized quadrilateral surfaces to the problem of registering and comparing level set representations of planar shapes with arbitrary topologies and multiple components. To do this, we build on the Riemannian framework proposed by Kurtek et al. [31]. A major advantage of this framework is that it searches for optimal parameterizations of the given objects using a parameterization-invariant metric. We take advantage of this useful property in the proposed framework.

2.1 Translation, Scaling and Rotation Variability

The notion of shape is invariant to translation, scaling, rotation and re-parameterization. In this section, we provide details of how we remove some of these variabilities from the representation space. We assume that the data is originally given as a binary image, $I : [0, 1]^2 \rightarrow \{0, 1\}$. If the data is not given in this form, we begin by computing its binary image representation. For simplicity, we remove the translation and scaling variabilities at this stage using normalization. The area of an object present in a binary image can be computed using its zeroth moment, $A = \int_{[0,1]^2} I(s)ds$ where $s = (x, y)$ are the image coordinates. The centroid of a binary image can be computed using the first moments, $(\bar{x}, \bar{y}) = \frac{1}{A}(\int_{[0,1]^2} xI(s)ds, \int_{[0,1]^2} yI(s)ds)$. Thus, we first translate the object in the image such that its centroid has coordinates $(\bar{x}, \bar{y}) = (0.5, 0.5)$. Once centered, we normalize the scale of the object in the image by rescaling it to occupy a certain proportion of the area of the entire image. Note that this proportion is chosen based on the application of interest and affects the distance calculation. Thus, the computed distances are comparable within data sets but not across data sets. While, in principle, we could choose the same scale for all types of data (making all the distances comparable) this is not very practical when the objects of interest are very different across applications (see for example the 2D sketches vs. the medical imaging data). In both steps, we utilize nearest neighbor interpolation and a very high image resolution (1000×1000 pixels) to preserve all of the details of the given objects. These two steps ensure that our analysis is invariant to translation and rescaling of the objects present in the given images.

While we could also normalize the orientation of the objects at this stage using the second moment, this approach can be unstable in practice. That is, for a small perturbation of the object the rotational alignment can change drastically. Thus, we take a different approach where we exhaustively search for the optimal rotation in a pairwise manner using the level set function representation.

Given two normalized binary images, I_1 and I_2, we are interested in computing the rotational alignment of the objects present within. We proceed as follows. First, we compute their corresponding signed distance function representations on the unit disk ($\mathbb{D}$) domain: $\tilde{\psi}_1,\ \tilde{\psi}_2 : \mathbb{D} \to \mathbb{R}$. Second, we generate a set of area preserving diffeomorphisms by rotating the initial disk parameterization by a set of angles $\theta \in [0, 2\pi)$. Call this set $\mathcal{H}$. In our implementation, we utilize 360 equally spaced angles. Thus, the set $\mathcal{H}$ contains 360 initial grid alignments. Next, we exhaustively search for the optimal rotation that best aligns the two signed distance functions using $\hat{h} = \arg\min_{h \in \mathcal{H}} \|\tilde{\psi}_1 - \tilde{\psi}_2 \circ h\|$, which corresponds to an angle of rotation $\hat{\theta}$. Finally, we apply this rotation to the second binary image to result in $\hat{I}_2$.

2.2 Square Root Representation of Level Set Functions

In order to optimally register two shapes with arbitrary topologies or multiple components and compute the distance between them, we adapt the framework of Kurtek et al. [31], which was defined and used for statistical shape analysis of quadrilateral surfaces. Let ψ be the distance transform of a binary image I. We first define a new function $f = (\nabla\psi, \psi)^T : [0,1]^2 \to \mathbb{R}^3$, where $\nabla\psi$ is a smoothed gradient of the level set distance function ψ. With a slight abuse of notation we will refer to f as the level set function from now on. In our implementation, we smooth the gradient using a Gaussian filter. The gradient of the level set function provides important edge features of the objects of interest, which will be useful during the registration process. We let $\mathcal{F}$ represent the space of all such level set functions: $\mathcal{F} = \{f : [0,1]^2 \to \mathbb{R}^3 | f \text{ is differentiable almost everywhere}\}$. Let Γ be the set of all diffeomorphisms of $[0,1]^2$. Γ acts on $\mathcal{F}$ by composition: $(f, \gamma) \to f \circ \gamma$. One can define the standard $\mathbb{L}^2$ inner product on this space and utilize the resulting Riemannian structure for comparing level set functions. Unfortunately, the $\mathbb{L}^2$ Riemannian metric is not invariant to re-parameterizations (because $\|f_1 - f_2\| \neq \|f_1 \circ \gamma - f_2 \circ \gamma\|$) and thus cannot be used. Kurtek et al. [31] suggest an alternative approach based on a different representation termed the square root function. We present some details next.

Given a function f, its square root function (SRF) representation $q : [0,1]^2 \to \mathbb{R}^3$ is defined as

$$q(s) = \sqrt{|n(s)|} f(s), \tag{1}$$

where $|\cdot|$ denotes the Euclidean norm in $\mathbb{R}^3$, and $n(s) = \frac{\partial f}{\partial x}(s) \times \frac{\partial f}{\partial y}(s)$. The resulting space of SRFs is a subset of $\mathbb{L}^2([0,1]^2, \mathbb{R}^3)$, from now on simply referred to as $\mathbb{L}^2$. If a level set function f is re-parameterized to $f \circ \gamma$, its corresponding SRF changes to $(q, \gamma) = (q \circ \gamma)\sqrt{J_\gamma}$, where J_γ is the determinant of the Jacobian of γ. Given this new representation, it is easy to check that given two SRFs q_1

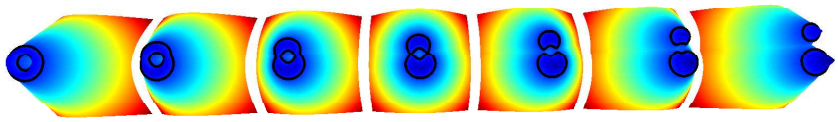

(a) Two simple shapes of arbitrary topology: $d([q_1],[q_2]) = 7.9863$.

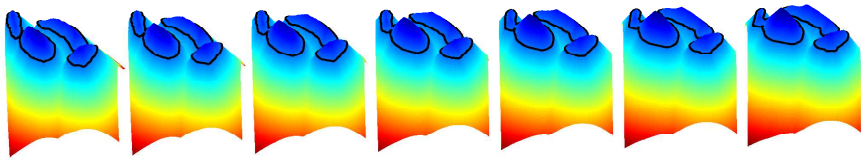

(b) Brain structures as multiple components; $d([q_1],[q_2]) = 6.1829$.

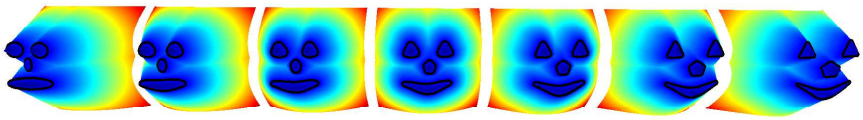

(c) Synthetic human face. $d([q_1],[q_2]) = 6.645$.

Fig. 2. Linear deformation path between the most left and the most right level set functions. The zero-level set is highlighted in black. (The figure is best viewed in color.)

and q_2 we have $\|q_1 - q_2\| = \|(q_1 \circ \gamma)\sqrt{J_\gamma} - (q_2 \circ \gamma)\sqrt{J_\gamma}\|$. This property ensures that our framework is invariant to re-paramterizations of the level set functions. Thus, we will utilize the SRF representation and the natural $\mathbb{L}^2$ metric to register and compute distances between level set functions, which represent the complex planar objects of interest.

2.3 Registration and Distance Calculation

We have already removed the translation, scale, and rotation variabilities from the representation space by normalizing the given binary images and aligning them rotationally in a pairwise manner. Thus, we are left with removing the parameterization variability of the level set functions. This step can also be thought of as the registration process, where similar structures are matched together across the given objects. To do this, we will utilize the notion of equivalence classes. That is, we will define two level set functions, f_1 and f_2, as equivalent if they are within a re-parameterization of each other. This provides us with the following definition of an equivalence class: $[q] = \{(q \circ \gamma)\sqrt{J_\gamma} | q \in \mathbb{L}^2,\ \gamma \in \Gamma\}$. Thus, we can define a parameterization-invariant distance between level set functions by minimizing over the equivalence classes:

$$d([q_1],[q_2]) = \inf_{\gamma \in \Gamma} \|q_1 - (q_2, \gamma)\|. \tag{2}$$

$d([q_1],[q_2])$ defines the parameterization-invariant (extrinsic) geodesic distance between SRF representations of level set functions. We use it as a measure of

similarity between planar objects that have arbitrary topologies and multiple components.

In addition to computing the distance, we would also like to display the path of deformation between the zero-levels of the registered level set functions. While the geodesic path under the pullback metric on $\mathcal{F}$ is a natural choice to do this, it is computationally expensive to compute [5]. Instead, we approximate these deformations using linear interpolation paths. As can be seen in the experimental results section, this does not seem to have any adverse effects. We note that the quality of the metric and registration is very closely related to how natural the deformations between the given objects are.

The computation of the distance in Equation 2 requires solving an optimization problem over the re-parameterization group Γ. Kurtek et al. [31] outline a gradient descent approach to do this and we summarize it here for convenience. Begin by constructing a geodesic (straight line) between q_1 and the current element of $[q_2]$, call it r. If the geodesic is perpendicular to the equivalence class $[q_2]$, then r is the optimally registered (and closest) element of $[q_2]$ to q_1. If not, we update r in the direction of the projection of the geodesic while staying within $[q_2]$. This is accomplished in three steps. First, define an orthonormal basis for the tangent space $T_{\gamma_{id}}(\Gamma)$ using products of adapted Fourier bases. One cannot use the Fourier basis as given due to the fact that this vector field must be tangential to the boundary of $[0,1]^2$. Then, use the differential of the mapping $\phi(\gamma) = (q \circ \gamma)\sqrt{J_\gamma}$ to find an update vector field b on $T_{\gamma_{id}}(\Gamma)$. Compute an infinitesimal update to the parameterization of r using $\gamma_{new} = \gamma_{id} + \epsilon b,\ \epsilon > 0$ and small, and compute the corresponding element of $[q_2]$ (a new version of r) using the mapping ϕ. Repeat these steps until the geodesic is perpendicular to $[q_2]$. This procedure reduces the distance at each iteration. It is computationally efficient but is not guaranteed to converge to the global solution. However, in practice we found that it produces natural correspondences and a measure of dissimilarity that outperforms the state-of-the-art.

Fig. 2 shows three examples of deformation paths between synthetic shapes of arbitrary topology. For each example, we show the deformation field between the level set functions (plotted as surfaces) and between their corresponding zero-level set (highlighted in black), which corresponds to the boundaries of the 2D shapes of interest.

3 Experimental results

We demonstrate the performance of the proposed framework using two types of results. First, we show several examples of computing deformation paths between 2D shapes that have fixed or varying topology and that may contain multiple parts (Section 3.1). The visual quality of the deformation paths is an indication of the quality of the computed correspondences. Next, we report experimental results on the classification and retrieval of 2D shapes (Section 3.2). We use three different data sets for evaluation, see Fig. 1:

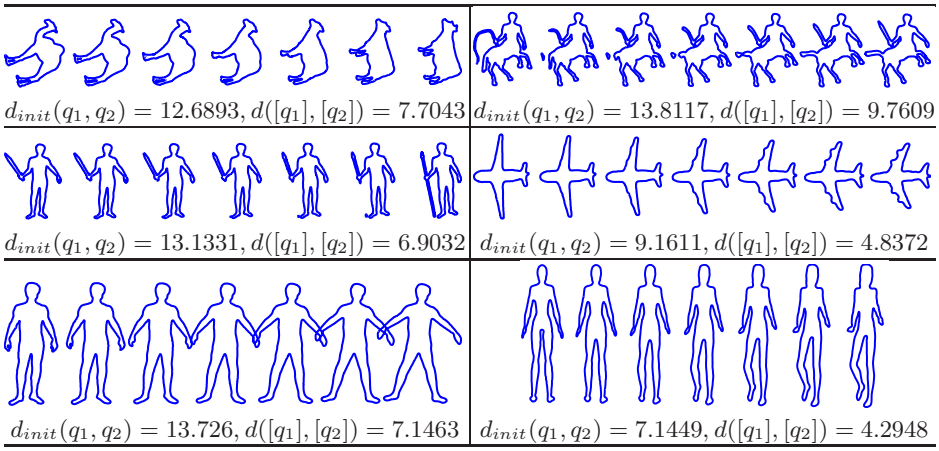

Fig. 3. Deformation paths between shapes of the the same topology, with no internal details.

Medical Imaging Data Set (Fig. 1(a)). We have manually extracted a set of four subcortical structures (putamen, hippocampus, thalamus and caudate) from 2D slices of 10 structural MRI images. Fig. 1(a) displays an example with each of the subcortical structures labeled. There are two types of variation in this data. The first is topological because in some images the structures are represented by three closed curves while in others with four. The second type is in the shape and the relative locations, rotations and scales of the structures.

2D Human sketches (Fig. 1(b)). Eitz et al. [8] provide a data set of $20,000$ objects sketched by humans. To demonstrate the performance of our approach, we use a subset of 100 sketches evenly distributed over 10 shape categories. Similar to the medical data set, the sketch data exhibit large topological variation as well as variations in the locations, scales and rotations of the different components of the objects. More importantly, most of the images in this data set are composed of strokes, with no clear definition of the interior and exterior regions of the sketched shapes.

SHREC07 watertight data set (Fig. 1(c)). This data set contains 400 three-dimensional objects evenly distributed into 20 classes [9]. For each of the 3D models, we generate a thumbnail image by rendering its frontal view into a binary image of size 450×600. Since the 3D data set contains complex shapes with arbitrary topology and pose, the resulting 2D images are of arbitrary topology. We use a subset of 100 images (the first five images of each class) in our analysis.

3.1 Shape matching and comparison

First, we focus on shapes that have the same topology as shown in Figures 3, 4 and 5. In the first case, we compare shapes that can be represented by their

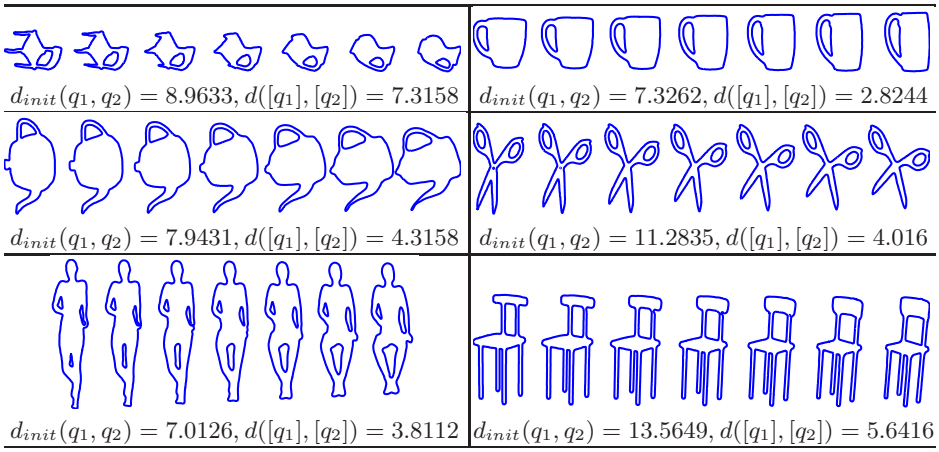

Fig. 4. Deformation paths between shapes of the same topology with internal details.

outer boundaries only. Fig. 3 shows several of such examples. Each row of the figure is a deformation path between the most-left and the most-right shapes. Note that the quality of the deformation is highly dependent on the quality of the computed correspondences between the level set functions; any error in the correspondences will result in distorted intermediate shapes. For each example, we also show the computed initial distance $d_{init} = \|q_1 - q_2\|$ (after rotational alignment, prior to optimization over Γ) and the computed parameterization-invariant distance d (after optimization over Γ) between the level set functions. In all of these examples, we used the same scale for each data set making the distances comparable.

Fig. 4 shows examples of deformation paths between shapes with the same topology but of high genus, i.e. with internal details. Observe that the deformations are very natural. It is important to note that the topology is preserved along the deformation path.

Fig. 5 shows deformation paths between shapes of the same topology but that contain several components. In each of these examples, we compare two synthetic faces sketched by hand. Each face contains four components (corresponding to different face parts). Observe that our approach is able to match these shapes correctly and generate natural deformations. In all of the three examples, we observe a decrease in the distance between shapes of approximately 40% due to the optimization over Γ.

In Fig. 6, we compare planar objects that have different topologies. Figures 6-(a) to (c) show a comparison between various brain structures. These shapes are naturally similar, but they exhibit two major types of variation. The first is topological because the structures are represented by two, three or four closed curves. The second type is in shape and the relative locations, rotations and scales of the structures. When we perform the optimization over Γ, the distance between the left-most and the right-most shapes reduces from 12.52 to 6.18, 13.33

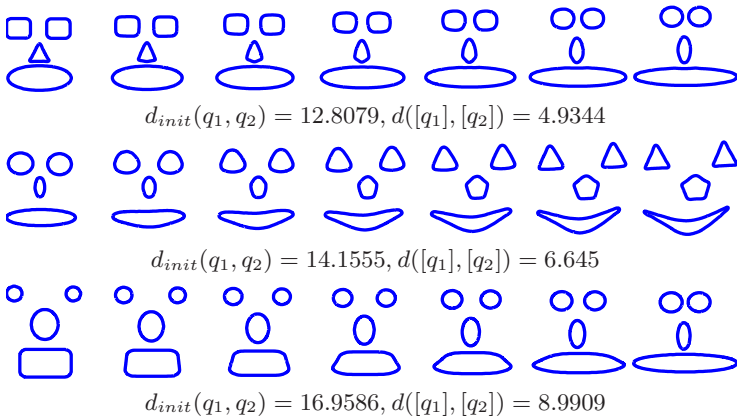

Fig. 5. Deformation paths between shapes of the same topology with multiple components.

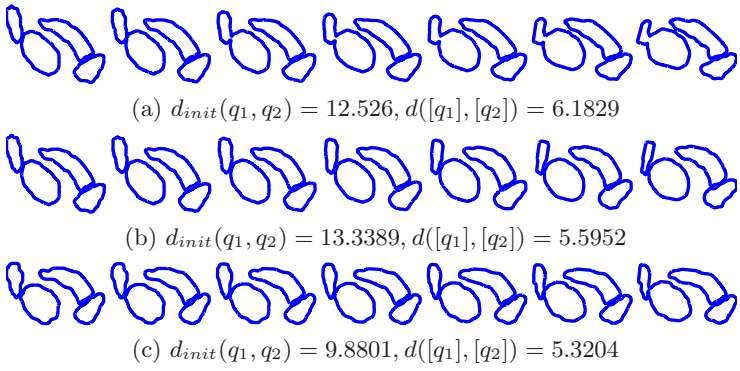

Fig. 6. Deformation paths between brain structures with multiple components and small topological changes.

to 5.59, and 9.88 to 5.32 for the cases (a), (b) and (c) respectively. In particular, observe how the topological change is carried in the intermediate shapes along the deformation paths. More examples of complex topological variations are shown in Fig. 7.

Finally, we compare in Fig. 8 the quality of the deformation paths that are generated without optimization over Γ vs. the deformation paths obtained using the full elastic metric, i.e. with optimization over Γ. We can clearly see that our elastic metric provides natural deformations and thus it finds correct correspondences. More examples are shown in the attached supplementary materials.

These examples show that our approach is able to handle shapes with arbitrary topology and with multiple structures. It is also able to compare shapes that have different topologies, which is a significant deviation from previous work [3, 4]. Furthermore, the presented examples clearly demonstrate that the

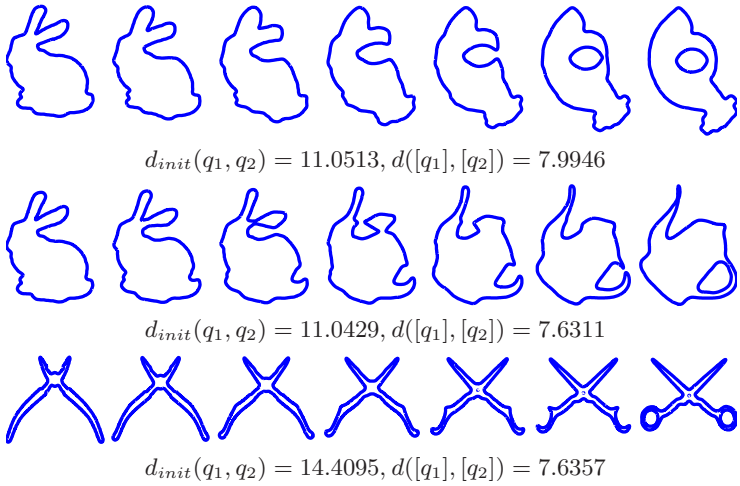

Fig. 7. Deformation paths between shapes of different topologies.

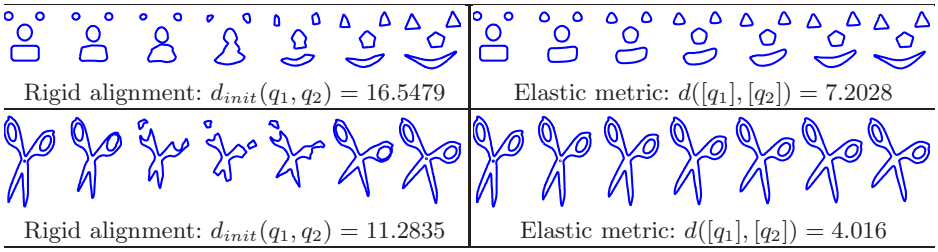

Fig. 8. Comparison between deformation paths obtained with rigid alignment and with the elastic metric proposed in this paper. More examples are included in the supplementary materials.

generated comparisons are natural and thus the computed distance is a good measure of differences between these shapes.

3.2 Classification performance

To quantitatively evaluate the performance of the proposed metric, we consider the classification of hand-written sketches [8] and 2D projections of the 3D models in SHREC07 [9] data set described above. For each set, we computed the pairwise distances using the approach proposed in this paper and compared its classification performance against six algorithms in the literature, namely: (1) the rigid alignment defined as the Euclidean distance between rigidly-aligned pairs of shapes (unlike the proposed elastic metric, this metric does not search for the optimal re-parameterization of the level set functions), (2) The Hard and Soft Histogram of Oriented Gradients (HOG-hard and HOG-soft) of Eitz et al. [8], (3) the Modified Hausdorff distance [33], which is equivalent to the

standard Hausdorff distance computed after normalizing the shapes for translation, scale and rotation, (4) the Inner Distance-based Shape Context (IDSC) [25], which is a very popular descriptor used in the analysis of planar shapes undergoing articulated motion, (5) the D2 shape distribution [34], and (6) the Gaussian Euclidean Transform (GEDT) of the shape boundaries [34]. We use the leave-one-out 1,3,5-nearest neighbor (LOO 1,3,5-NN) classifiers in all examples, where the class is determined by a majority vote. Below, we discuss the classification performance on each data set.

2D sketch data set [8]. We use a subset of 100 sketches evenly distributed over 10 shape categories. Figure 1(b) shows a few examples from this data set. For these 100 objects, we first compute their Euclidean distance transforms and use the algorithm described in Section 2 to compute the pairwise distances. The resulting LOO 1-NN classification rate was 85%, outperforming the best state-of-the-art descriptor by approximately 13%, see Table 1. Note that HOG-hard achieved a 65% LOO 1-NN classification rate while HOG-soft achieved a 72.0% classification rate. These two methods use supervised learning to build the codebook. Our approach, which obtained 85.0% LOO 1-NN performance, is completely unsupervised. The IDSC descriptor, which is extensively used for the analysis of planar shapes undergoing articulated motion, achieved 25% LOO 1-NN classification rate, which is significantly below the proposed approach. Finally, note that the proposed elastic metric significantly outperforms all of the other methods when more neighbors are considered, and deteriorates much more slowly.

The SHREC07 data set. We performed a similar evaluation on the SHREC07 data set and report the performance in Table 2. Again, our metric outperforms the state-of-the-art by more than 5% on LOO 1-NN classification, except the IDSC, which performed 3% better.

Comparison with the Inner Distance-based Shape Context (IDSC) [25]. IDSC is a very popular descriptor that has been used for the registration of planar shapes in the presence of articulated motion. When used with shapes that have well defined interior regions, such as the SHREC07 data set, it slightly outperforms the metric proposed in this paper, see Table 2. IDSC, however, fails when used in the analysis of the 2D sketch data set as shown in Table 1. 2D sketches are composed of strokes with large topological variations and without clearly defined interior regions. Our approach outperformed all of the state-of-the-art methods on this data set. Note also that unlike the IDSC, our approach is invariant to all shape-preserving transformations including re-parameterization.

Computation time. The total computation time for comparing two objects is approximately 51s. This can be improved by optimally finding rotations without doing an exhaustive search as currently implemented.

	Hausdorff	D2	GEDT	HOG-hard	HOG-soft	Rigid align.	IDSC	**Proposed**
LOO-1	65.0%	49.0%	68.0%	65.0%	72.0%	70.0%	25.0%	**85.0%**
LOO-3	57.0%	46.0%	54.0%	56.0%	64.0%	73.0%	17.0%	**76.0%**
LOO-5	49.0%	38.0%	44.0%	48.0%	54.0%	67.0%	5.0%	**73.0%**

Table 1. LOO 1,3,5-NN classification on the 2D human sketch data set [8].

	Hausdorff	D2	GEDT	HOG-hard	Rigid align.	IDSC	**Proposed**
LOO-1	74.0%	40.0%	75.0%	61.0%	68.0%	83.0	**80.0%**
LOO-3	57.0%	24.0%	64.0%	37.0%	47.0%	79.0	**71.0%**
LOO-5	40.0%	12.0%	45.0%	20.0%	31.0%	69.0	**56.0%**

Table 2. LOO 1,3,5-NN classification on the SHREC07 data set.

4 Conclusions

We have proposed a novel framework for simultaneous registration and comparison of planar objects with multiple components and differing topologies. To accomplish this, we use the distance function representation and a parameterization-invariant framework for elastic shape analysis of surfaces. We validated our framework on different types of examples, which included objects of the same topology, high genus objects, objects with multiple components, and objects of different topologies. The resulting natural deformations show the strength of our method and the benefit of optimizing over the re-parameterization group when generating such comparisons. We also used the distance function for two classification experiments. The results show that we outperform the current state-of-the-art methods in the classification of 2D sketches as well as arbitrary planar shapes. There are many directions for future work. First, we would like to compute geodesics between distance functions using the metric proposed in [5]. This will enable us to generate means and covariances of shapes with different topologies. Second, we plan to extend this framework to the analysis of 3D objects that have arbitrary topology and multiple components.

Acknowledgement. This work is partially funded by the Australian Research Council (ARC) and the South Australian Government, Department of Further Education, Employment, Science and Technology.

References

1. Younes, L.: Computable elastic distance between shapes. SIAM Journal of Applied Mathematics **58** (1998) 565–586
2. Younes, L., Michor, P.W., Shah, J., Mumford, D., Lincei, R.: A metric on shape space with explicit geodesics. Matematica E Applicazioni **19** (2008) 25–57
3. Srivastava, A., Klassen, E., Joshi, S.H., Jermyn, I.H.: Shape analysis of elastic curves in Euclidean spaces. IEEE Trans. Pattern Analysis Machine Intelligence **33** (2011) 1415–1428
4. Jermyn, I., Kurtek, S., Klassen, E., Srivastava, A.: Elastic shape matching of parameterized surfaces using square root normal fields. In: European Conference on Computer Vision. (2012) 804–817
5. Kurtek, S., Klassen, E., Gore, J.C., Ding, Z., Srivastava, A.: Elastic geodesic paths in shape space of parameterized surfaces. IEEE Trans. Pattern Analysis and Machine Intelligence **34** (2012) 1717–1730
6. Kurtek, S., Srivastava, A., Klassen, E., Laga, H.: Landmark-guided elastic shape analysis of spherically-parameterized surfaces. Computer Graphics Forum (Proceedings of Eurographics 2013) **32** (2013) 429–438
7. Fuchs, M., Jüttler, B., Scherzer, O., Yang, H.: Shape metrics based on elastic deformations. Journal of Mathematical Imaging and Vision **35** (2009) 86–102
8. Eitz, M., Hays, J., Alexa, M.: How do humans sketch objects? ACM Transactions on Graphics (Proceedings SIGGRAPH) **31** (2012) 44:1–44:10
9. Giorgi, D., Biasotti, S., Paraboschi, L.: Shape retrieval contest 2007: Watertight models track. SHREC competition **8** (2007)
10. Kendall, D.G.: Shape manifolds, Procrustean metrics, and complex projective spaces. Bulletin of the London Mathematical Society **16** (1984) 81–121
11. Dryden, I.L., Mardia, K.V.: Statistical Shape Analysis. John Wiley & Son (1998)
12. Zahn, C.T., Roskies, R.Z.: Fourier descriptors for plane closed curves. IEEE Trans. Computers **21** (1972) 269–281
13. Klassen, E., Srivastava, A., Mio, W., Joshi, S.H.: Analysis of planar shapes using geodesic paths on shape spaces. IEEE Trans. Pattern Analysis and Machine Intelligence **26** (2004) 372–383
14. Joshi, S., Klassen, E., Srivastava, A., Jermyn, I.: Removing shape-preserving transformations in square-root elastic (SRE) framework for shape analysis of curves. In: EMMCVPR. (2007) 387–398
15. Kerr, G., Kurtek, S., Srivastava, A.: A joint model for boundaries of multiple anatomical parts. In: SPIE Medical Imaging. (2011)
16. Osher, S., Fedkiw, R.: Level Set Methods and Dynamic Implicit Surfaces. Springer Verlag (2003)
17. Paragios, N., Rousson, M., Ramesh, V.: Non-rigid registration using distance functions. Computer Vision and Image Understanding **89** (2003) 142–165
18. Munim, H.A.E., Farag, A.: Shape representation and registration using vector distance functions. In: IEEE Conf. on Computer Vision and Pattern Recognition. (2007)
19. Fahmi, R., Farag, A.A.: A global-to-local 2d shape registration in implicit spaces using level sets. In: IEEE International Conference on Image Processing. Volume 6., IEEE (2007) VI–237
20. Yezzi, A.J., Soatto, S.: Deformotion: Deforming motion, shape average and the joint registration and approximation of structures in images. International Journal of Computer Vision **53** (2003) 153–167

21. Glaunes, J., Qiu, A., Miller, M.I., Younes, L.: Large deformation diffeomorphic metric curve mapping. International Journal of Computer Vision **80** (2008) 317–336
22. Huttenlocher, D.P., Klanderman, G.A., Kl, G.A., Rucklidge, W.J.: Comparing images using the hausdorff distance. IEEE Trans. Pattern Analysis and Machine Intelligence **15** (1993) 850–863
23. Davis, E.: Continuous shape transformation and metrics on regions. Fundamenta Informaticae **46** (2001) 31–54
24. Berkels, B., Linkmann, G., Rumpf, M.: Shape median based on symmetric area differences. In Deussen, O., Keim, D.A., Saupe, D., eds.: Vision, Modeling and Visualization. (2008) 399–407
25. Ling, H., Jacobs, D.W.: Shape classification using the inner-distance. Pattern Analysis and Machine Intelligence, IEEE Transactions on **29** (2007) 286–299
26. Tabia, H., Laga, H., Picard, D., Gosselin, P.H., et al.: Covariance descriptors for 3d shape matching and retrieval. IEEE Conference on Computer Vision and Pattern Recognition (2014)
27. Wang, X., Feng, B., Bai, X., Liu, W., Jan Latecki, L.: Bag of contour fragments for robust shape classification. Pattern Recognition **47** (2014) 2116–2125
28. Laga, H., Takahashi, H., Nakajima, M.: Spherical wavelet descriptors for content-based 3d model retrieval. In: Shape Modeling and Applications, 2006. SMI 2006. IEEE International Conference on, IEEE (2006) 15–15
29. Laga, H., Nakajima, M.: Supervised learning of salient 2D views of 3D models. The Journal of The Society for Art and Science **7** (2008) 124–131
30. Laga, H.: Semantics-driven approach for automatic selection of best views of 3d shapes. In: Proceedings of the 3rd Eurographics Workshop on 3D Object Retrieval, Eurographics Association (2010) 15–22
31. Kurtek, S., Klassen, E., Ding, Z., Srivastava, A.: A novel Riemannian framework for shape analysis of 3D objects. In: IEEE Computer Vision Pattern Recognition. (2010) 1625–1632
32. Xie, Q., Kurtek, S., Christensen, G.E., Ding, Z., Klassen, E., Srivastava, A.: A novel framework for metric-based image registration. In Dawant, B.M., Christensen, G.E., Fitzpatrick, J.M., Rueckert, D., eds.: Workshop on Biomedical Image Registration. Volume 7359 of Lecture Notes in Computer Science. (2012) 276–285
33. Dubuisson, M.P., Jain, A.K.: A modified hausdorff distance for object matching. In: Proceedings of the 12th IAPR International Conference on Pattern Recognition. Volume 1., IEEE (1994) 566–568
34. Laga, H., Kurtek, S., Srivastava, A., Golzarian, M., Miklavcic, S.J.: A riemannian elastic metric for shape-based plant leaf classification. In: DICTA. (2012) 1–7