



On Shape of Plane Elastic Curves

WASHINGTON MIO

Department of Mathematics, Florida State University, Tallahassee, FL 32306-4510

ANUJ SRIVASTAVA

Department of Statistics, Florida State University, Tallahassee, FL 32306-4330

SHANTANU JOSHI

Department of Electrical and Computer Engineering, Florida State University, Tallahassee, FL 32310-6046

Received September 19, 2005; Revised June 15, 2006; Accepted July 31, 2006

First online version published in September, 2006

Abstract. We study shapes of planar arcs and closed contours modeled on elastic curves obtained by bending, stretching or compressing line segments non-uniformly along their extensions. Shapes are represented as elements of a quotient space of curves obtained by identifying those that differ by shape-preserving transformations. The elastic properties of the curves are encoded in Riemannian metrics on these spaces. Geodesics in shape spaces are used to quantify shape divergence and to develop morphing techniques. The shape spaces and metrics constructed are novel and offer an environment for the study of shape statistics. Elasticity leads to shape correspondences and deformations that are more natural and intuitive than those obtained in several existing models. Applications of shape geodesics to the definition and calculation of mean shapes and to the development of shape clustering techniques are also investigated.

Keywords: planar shapes, shape geodesics, mean shape, shape analysis, clustering shapes

1. Introduction

Shapes and textures associated with a distribution of pixel values are two key elements in understanding the way an image is perceived. Many generative models attempt to capture the beautifully intricate, yet very natural, interaction between these somewhat independent components of images. Thus, the development of independent models of shapes and textures and the fusion of these to model images are natural problems in computer vision. This paper introduces a novel representation of plane shapes as curves obtained by stretching, compressing and bending elastic straight-line segments.

Using basic techniques from differential geometry, shape metrics and morphing techniques are developed to model energy-efficient deformations of shapes taking elasticity into account. As discussed in more detail below, the proposed shape model leads to shape correspondences and deformations that are more natural and intuitive than those obtained with many existing models making it better suited for many problems in computer vision.

The quantitative study of shapes dates back to the work of D'Arcy Thompson in the first half of the 20th century (Thompson, 1992). Research in the area has gained new impetus in recent years due to a large inflow

of new ideas from areas such as computer vision and medical imaging. Applications of algorithmic shape analysis include detection and recognition of objects in images and videos, algorithmic analysis of MRI scans, automated interpretation of imaged scenes, and morphometric studies of insects and fish. The seminal works of Bookstein (1986) and Kendall (1984) represent influential contributions to the modern theory of shapes with the introduction of methods and techniques derived from differential geometry. In their work, shapes are represented by collections of ordered landmark points; identifying sets of landmarks that differ by shape-preserving transformations, a quotient *shape space* is obtained whose geometry reflects properties of shapes. For example, geodesic distance is used to quantify shape divergence and as a basic tool for statistical shape analysis. Despite difficulties encountered in the selection of landmarks and the dependence of the resulting analysis of shapes on choices made, this approach has been used successfully in numerous applications. Of particular historical relevance is the fact that Kendall's shape manifolds have provided an environment for the development of a statistical theory of shapes via mean shapes and tangent-space probability models (Le and Kendall, 1993; Dryden and Mardia, 1998). The techniques extend to the analysis of data in more general Riemannian manifolds and have been used in various different contexts. Another approach, known as *active shape model*, uses principal component analysis on landmark representations to model shape variations in observed samples. Despite its simplicity and ease of use, the fact that it does not account for nonlinearity limits its scope (Cootes et al., 1995).

Grenander takes a different view and represents shapes using deformable templates (Grenander, 1993) with shape deformations modeled on the action of groups of diffeomorphisms. This approach has been widely explored (see e.g. Beg et al. (2005)), but typical computational costs tend to be somewhat high as compared to those associated with shapes represented as curves. This may represent a significant practical barrier in problems involving large collections of shapes.

In many applications, it is desirable that shape metrics and deformations respect certain shape correspondences. For example, in medical imaging, a correspondence between contours is often established to preserve some landmark points. Shape matching algorithms to produce correspondences that seek to optimally align elements such as velocity fields or curvature functions

of contours have been investigated in Cohen et al. (1992), Geiger et al. (1995), Tagare et al. (2002), and Sebastian et al. (2003). Such correspondences often require that curves be stretched or compressed non-uniformly along their extensions. Work toward a theory of shapes using elastic models has been carried out in Younes (1999, 1998), where some shape descriptors and metrics were derived. However, interpolation and statistical modeling problems were only partially addressed.

Although particular applications sometimes invoke only specific aspects of shape analysis, modern vision problems require a unifying framework for the statistical study of shapes and the development of computational strategies. This demand has spawned numerous studies of shapes in recent years. As pointed out earlier, the shape theory of Bookstein and Kendall meets many of the requirements, but the use of landmarks is a drawback. Additionally, while the simplicity of the shape representation is very attractive from a computational standpoint, it often leads to unsatisfactory interpolations, as illustrated below. Klassen et al. (2004) proposed an approach where shapes of curves are represented via angle (or curvature) functions associated with their arc-length parameterizations. The statistical approach to shapes of Dryden and Mardia (1998) and Le and Kendall (1993) was extended to this setting in Klassen et al. (2004). However, as deformations of curves respect the arc-length parameter, stretch elasticity is not incorporated to the model and resulting shape correspondences are sometimes far from optimal. Other recent studies of planar shapes include work by Mumford and Sharon based on conformal mappings (Sharon and Mumford, 2004), as well as several models involving representations of shapes as curves with various different metrics (Michor and Mumford, in press, ; Mennucci and Yezzi, 2004).

In this paper, we develop an algorithmic approach to shapes of planar contours modeled on elastic curves. Shape spaces will be constructed with geometric structures that will allow us to use geodesics to quantify shape divergence and model shape deformations. One of the goals is to retain some of the computational advantages associated with models such as the arclength-preserving, bending-only model of Klassen et al. (2004), while introducing elasticity to obtain improved, more natural shape correspondences. Although bearing some philosophical similarities to Klassen et al. (2004)—in the sense that shape spaces are constructed from spaces of parametric curves—elasticity

requires a completely different model and raises a series of new delicate issues to be dealt with both at the theoretical and computational levels. We present the heuristics of an infinite-dimensional model of elastic shapes and focus the discussion on finite-dimensional approximations that lead to a computational model.

A regular, parametric planar curve $\alpha: I \rightarrow \mathbb{R}^2$, where $I = [0, 1]$, will be represented by a pair (ϕ, θ) of functions that encode the velocity field $\alpha'(t)$, $0 \leq t \leq 1$, of the curve, as follows. The log-speed $\phi(t) = \log \|\alpha'(t)\|$ captures the rate at which the interval I is stretched or compressed at t to form α , and $\theta(t)$ measures the angle the velocity vector $\alpha'(t)$ makes with a horizontal axis, thus quantifying bending. We consider the infinite-dimensional manifold formed by all pairs (ϕ, θ) representing shapes and form a quotient shape space by identifying those that differ by shape-preserving transformations and by reparameterizations. This space will be equipped with a metric that captures the elastic properties of the curves. In this manner, we obtain an environment for the development of a statistical theory of elastic shapes. A fundamental ingredient is an algorithm to compute shape geodesics, which can be used to quantify shape divergence and to solve shape interpolation and extrapolation problems. With these elements in place, methods of Le and Kendall (1993), Dryden and Mardia (1998), and Srivastava et al. (2005) can be adapted to the study of shape statistics in the present context. Computational techniques for estimating shape geodesics will be developed using differential geometric techniques and dynamic programming.

A few words about the organization of the paper. In Section 2, we introduce the shape representation to be adopted in the paper and examine the effect of shape-preserving transformations and curve reparameterizations on the representation. The notion of pre-shape is introduced in Section 3. A pre-shape represents a class of curves that differ by shape-preserving transformations of the plane. Manifolds of pre-shapes of arcs and closed curves are constructed and equipped with a Riemannian metric that encodes the elasticity of the curves. An algorithm to calculate geodesics in pre-shape manifolds is developed in Section 4 and several examples are given. Dynamic programming shape matching is discussed in Section 5 and geodesics that preserve shape correspondences are used to quantify shape divergence and interpolate shapes. Shape spaces obtained as quotient spaces of pre-shape manifolds by

identifying pre-shapes that differ by parameterizations are introduced in Section 6, as well as shape metrics and geodesic morphing techniques. Applications of geodesics in pre-shape space to the definition and calculation of mean shapes, and to the development of shape clustering techniques are presented in Sections 7 and 8, respectively. This is followed by a few concluding remarks in Section 9.

2. Shape Representation

We begin with the construction of a Riemannian manifold of plane curves, which ultimately will allow us to model and analyze shapes of planar curves algorithmically. To guide the discussion, we first present a heuristic treatment of an infinite-dimensional model for smooth planar curves since a rigorous mathematical description would take us too far afield. Then, we discretize the model to obtain the algorithmic representation of shapes that will be adopted.

2.1. Continuous Model

Let I denote the unit interval $[0, 1]$ and $\alpha: I \rightarrow \mathbb{R}^2$ a smooth, regular parametric curve in the sense that $\alpha'(t) \neq 0, \forall t \in I$. One may think of the mapping α as a prescription for *stretching (or compressing)* and *bending* the interval I at varying rates to produce the curve. To quantify these two rather independent notions of elastic deformation, write the velocity vector as

$$\alpha'(t) = e^{\phi(t)} e^{j\theta(t)}, \quad (1)$$

where $\phi: I \rightarrow \mathbb{R}$ and $\theta: I \rightarrow \mathbb{R}$ are smooth, and $j = \sqrt{-1}$. Here, we are using the standard identification of \mathbb{R}^2 with the complex plane \mathbb{C} . The function ϕ can be interpreted as the *speed* of α expressed in logarithmic scale, and θ as a smooth measurement of the *angle* the velocity vector makes with a horizontal axis. Alternately, $\phi(t)$ may be viewed as a quantifier of the rate at which the interval I was stretched or compressed at t to form the curve α , and $\theta(t)$ as describing how the interval I was bent at t to produce the curve α . Note that $\phi(t) > 0$ indicates local stretching near t , and $\phi(t) < 0$ local compression. The arc-length element of α is $ds = e^{\phi(t)} dt$. Curves parameterized by arc length, i.e., traversed with constant

speed 1, are those with $\phi \equiv 0$. We shall represent α via the pair (ϕ, θ) and denote by \mathcal{H} the space of all such pairs. Note that the angle function associated with α is only defined up to the addition of integer multiples of 2π .

2.1.1. Shape-Preserving Transformations. Parametric plane curves that differ by the action of the group of transformations generated by (orientation preserving) rigid motions and homotheties of the plane are to be viewed as representing the same shape. Thus, we inspect the effect of these transformations on the representation (ϕ, θ) . Since the functions ϕ and θ encode properties of the velocity field of the curve α , the pair (ϕ, θ) is clearly invariant under translations of the curve. The effect of a rotation is to add a constant to θ keeping ϕ unchanged, and scaling the curve by a factor $k > 0$ changes ϕ to $\phi + \log k$ leaving θ unaltered. Depending on the application, one may wish to include orientation-reversing transformations such as reflections, as well.

2.1.2. Reparametrization. We now examine the action of reparameterizations on (ϕ, θ) . Reparameterizations of α that preserve the orientation of the curve and the property that it is regular are those obtained by composing α with an orientation-preserving diffeomorphism $\gamma: I \rightarrow I$ of the unit interval; the action of γ on α is to produce the curve $\beta(t) = \alpha(\gamma(t))$. Since

$$\begin{aligned} \beta'(t) &= \alpha'(\gamma(t)) \gamma'(t) = e^{\phi(\gamma(t))} e^{j\theta(\gamma(t))} \gamma'(t) \\ &= e^{\phi(\gamma(t)) + \log \gamma'(t)} e^{j\theta(\gamma(t))}, \end{aligned} \quad (2)$$

the curve β is represented by $(\phi \circ \gamma + \log \gamma', \theta \circ \gamma)$, where \circ denotes composition of maps; note that $\gamma'(t) > 0$ because γ is a diffeomorphism. Hence, reparameterizations define an action of the group \mathcal{D}_I of orientation-preserving diffeomorphisms of the interval I on \mathcal{H} by

$$(\phi, \theta) \cdot \gamma = (\phi \circ \gamma + \log \gamma', \theta \circ \gamma). \quad (3)$$

2.1.3. Riemannian Structure. In order to compare curves quantitatively, we assume that they are made of an elastic material and adopt a metric that measures

how difficult it is to reshape a curve into another taking elasticity into account. Infinitesimally, this can be done using a *Riemannian structure* on \mathcal{H} . Recall that a Riemannian metric on a manifold M consists of inner products $\langle \cdot, \cdot \rangle_x$ on the tangent spaces $T_x M$, $x \in M$, which vary smoothly along the manifold. Such structure allows us to define basic geometric quantities such as length of curves, not only infinitesimally, but also at larger scales via integration.

Since the tangent space to \mathcal{H} at any point is the space \mathcal{H} itself, for each (ϕ, θ) , we wish to define an inner product $\langle \cdot, \cdot \rangle_{(\phi, \theta)}$ on \mathcal{H} . We adopt the simplest Riemannian structure that will make the diffeomorphism group \mathcal{D}_I act as transformations that respect the Riemannian structure on \mathcal{H} , much like the way translations and rotations act on standard Euclidean spaces. Given $(\phi, \theta) \in \mathcal{H}$, let h_i and f_i , $i = 1, 2$, represent infinitesimal (first-order) deformations of ϕ and θ , resp., so that (h_1, f_1) and (h_2, f_2) are tangent vectors to \mathcal{H} at (ϕ, θ) . For $a, b > 0$, define

$$\begin{aligned} &\langle (h_1, f_1), (h_2, f_2) \rangle_{(\phi, \theta)} \\ &= a \int_0^1 h_1(t) h_2(t) e^{\phi(t)} dt + b \int_0^1 f_1(t) f_2(t) e^{\phi(t)} dt. \end{aligned} \quad (4)$$

This is a weighted sum of the standard \mathbb{L}^2 inner products of the h and f components with respect to the arc-length element $ds = e^{\phi(t)} dt$. A simple change-of-variables argument shows that reparameterizations indeed preserve the inner product. We sometimes omit the subscript (ϕ, θ) from the notation.

The elastic properties of the curves are built-in to the model via the parameters a and b , which can be interpreted as *tension* and *rigidity coefficients*, respectively. Large values of the ratio $\chi = a/b$ indicate that the material offers higher resistance to stretching and compression than to bending; the opposite holds for χ small.

To our knowledge, the only previously studied Riemannian metric on shape manifolds that account for both stretch and bending elasticity is due to Younes (1998). In the special case $a = 1$ and $b = 1$, the metric we adopt is similar to the one used by Younes. However, in his model, the penalty term on stretch elasticity is somewhat asymmetric; for example, stretching a curve near a point by a factor 2 is not penalized in the same way as compressing it by a factor 1/2. In our model, symmetry is achieved using a logarithmic scale.

2.2. Discrete Representation

Let α be a curve represented by the pair (ϕ, θ) . To obtain a discrete representation of the curve, we sample both ϕ and θ uniformly at n points $0 = t_1 < t_2 < \dots < t_{n-1} < t_n = 1$ on the interval I . Thus, the curve α is represented by a vector

$$(\phi, \theta) \in \mathbb{R}^n \times \mathbb{R}^n \cong \mathbb{R}^{2n}, \quad (5)$$

where $\phi = (\phi(t_1), \dots, \phi(t_n))$ and $\theta = (\theta(t_1), \dots, \theta(t_n))$. Throughout the paper, we abuse notation and use ϕ and θ for both continuous and discrete representations, but the meaning should be clear from the context. Based on the discussion in Section 2.1.3, the usual Euclidean metric on \mathbb{R}^{2n} is not the most appropriate for the present study of elastic shapes. A natural discretization of the inner product defined in (4) is given by

$$\begin{aligned} & \langle (h_1, f_1), (h_2, f_2) \rangle_{(\phi, \theta)} \\ &= a \sum_{i=1}^n h_i^1 h_i^2 e^{\phi_i} + b \sum_{i=1}^n f_i^1 f_i^2 e^{\phi_i}, \end{aligned} \quad (6)$$

where $h_j = (h_1^j, \dots, h_n^j) \in \mathbb{R}^n$ and $f_j = (f_1^j, \dots, f_n^j) \in \mathbb{R}^n$, for $j = 1, 2$. The space \mathbb{R}^{2n} equipped with this Riemannian structure will be denoted \mathcal{H}_n to distinguish it from \mathbb{R}^{2n} with the usual Euclidean metric.

In the discrete formulation, the diffeomorphisms used for reparameterizations will be replaced with a class of piecewise linear homeomorphism of the interval I to be described in more detail in Section 5.

3. Manifolds of Elastic Pre-Shapes

3.1. Open Pre-Shapes

The pair (ϕ, θ) associated with a parametric curve α is invariant under translations of α , but the representation is sensitive to rotations and scale. To make it scale invariant, we restrict (ϕ, θ) to represent curves of unit length. To get rotational invariance, we fix the average value of angle functions with respect to the arc-length element to be, say, π . In other words, we restrict shape representatives to pairs (ϕ, θ) satisfying the conditions

$$\int_0^1 e^{\phi(t)} dt = 1 \quad \text{and} \quad \int_0^1 \theta(t) e^{\phi(t)} dt = \pi, \quad (7)$$

whose discrete analogues are

$$\sum_{i=1}^n e^{\phi_i} = 1 \quad \text{and} \quad \sum_{i=1}^n \theta_i e^{\phi_i} = \pi. \quad (8)$$

Constraints (7) and (8) define subspaces $\mathcal{A} \subset \mathcal{H}$ and $\mathcal{A}_n \subset \mathcal{H}_n$, resp., which we refer to as *pre-shape spaces* of open planar elastic curves. We adopt this terminology because we have eliminated indeterminacies in the representation due to rigid motions and homotheties of the plane, but not those due to reparameterizations.

A slightly more formal description of the pre-shape spaces will help elucidate some of their geometric properties. Define a map $\mathcal{F} = (\mathcal{F}_1, \mathcal{F}_2): \mathcal{H} \rightarrow \mathbb{R}^2$ by

$$\begin{aligned} \mathcal{F}_1(\phi, \theta) &= \int_0^1 e^{\phi(t)} dt \quad \text{and} \\ \mathcal{F}_2(\phi, \theta) &= \int_0^1 \theta(t) e^{\phi(t)} dt, \end{aligned} \quad (9)$$

and let $F = (F_1, F_2): \mathcal{H}_n \rightarrow \mathbb{R}^2$ given by

$$F_1(\phi, \theta) = \sum_{i=1}^n e^{\phi_i} \quad \text{and} \quad F_2(\phi, \theta) = \sum_{i=1}^n \theta_i e^{\phi_i} \quad (10)$$

be its discrete counterpart. The pre-shape spaces just introduced may be viewed as the level sets

$$\mathcal{A} = \mathcal{F}^{-1}(1, \pi) \quad \text{and} \quad \mathcal{A}_n = F^{-1}(1, \pi),$$

respectively. An inspection of (11) shows that these two constraints on (ϕ, θ) are independent, so that \mathcal{A}_n is a submanifold of \mathcal{H}_n . In particular, at any point $(\phi, \theta) \in \mathcal{A}_n$, the normal (or, perpendicular) space to \mathcal{A}_n in \mathcal{H}_n is 2-dimensional and spanned by the gradient vectors

$$\begin{aligned} \nabla F_1(\phi, \theta) &= (1/a, 0) \quad \text{and} \\ \nabla F_2(\phi, \theta) &= (\theta/a, 1/b). \end{aligned} \quad (11)$$

Here, 1 denotes the constant n -tuple $(1, \dots, 1)$, and similarly 0.

This description of the pre-shape space \mathcal{A}_n has several useful practical implications. The explicit calculation of gradient vectors in (11) yields a simple algorithmic procedure for projecting vectors $(h, f) \in \mathcal{H}_n$

orthogonally onto the tangent space $T_{(\phi,\theta)}\mathcal{A}_n$ by subtracting normal components. This will be needed, e.g., in numerical calculations of geodesics in \mathcal{A}_n .

Algorithm 3.1.1. Orthogonal Projection of $(h, f) \in \mathcal{H}_n$ onto $T_{(\phi,\theta)}\mathcal{A}_n$.

- (i) Apply Gram-Schmidt to $\{\nabla F_1(\phi, \theta), \nabla F_2(\phi, \theta)\}$, with respect to the inner product $\langle \cdot, \cdot \rangle_{(\phi,\theta)}$, to obtain an orthonormal basis $\{e_1(\phi, \theta), e_2(\phi, \theta)\}$ of $T_{(\phi,\theta)}\mathcal{A}_n$.
- (ii) The orthogonal projection of (h, f) is given by

$$\begin{aligned} \Pi_{(\phi,\theta)}(h, f) &= (h, f) - \sum_{i=1}^2 \langle (h, f), e_i(\phi, \theta) \rangle_{(\phi,\theta)} e_i(\phi, \theta). \end{aligned}$$

Our calculations of geodesics in pre-shape manifolds will also require a mechanism to project points in \mathcal{H}_n onto \mathcal{A}_n . This is because, during numerical integrations of the differential equation that governs geodesics in \mathcal{A}_n , points on \mathcal{A}_n typically evolve to points slightly off the pre-shape manifold. The projection will then place them back on \mathcal{A}_n . For $(\phi, \theta) \in \mathcal{H}_n$, the residual vector

$$r(\phi, \theta) = (1, \pi) - F(\phi, \theta) \quad (12)$$

quantifies how far off (ϕ, θ) is from \mathcal{A}_n and it is zero if and only if $(\phi, \theta) \in \mathcal{A}_n$. To project (ϕ, θ) onto \mathcal{A}_n , we use Newton's method to search for a (nearby) zero of r , initializing the search with (ϕ, θ) , as explained next.

The interesting infinitesimal variations of F occur along directions perpendicular to its level sets. Thus, given (ϕ, θ) , we calculate the differential dF restricted to the normal space $N(\phi, \theta)$ at (ϕ, θ) . The Jacobian of the mapping F restricted to $N(\phi, \theta)$ can be expressed in the basis $\{\nabla F_i(\phi, \theta), 1 \leq i \leq 2\}$ of $N(\phi, \theta)$ and the standard basis of \mathbb{R}^2 as the 2×2 symmetric matrix $J(\phi, \theta)$ whose (i, j) entry is

$$J_{ij}(\phi, \theta) = \langle \nabla F_i(\phi, \theta), \nabla F_j(\phi, \theta) \rangle_{(\phi,\theta)}. \quad (13)$$

Let $\epsilon > 0$ be a small number.

Algorithm 3.1.2. Projection of $(\phi, \theta) \in \mathcal{H}_n$ onto \mathcal{A}_n .

1. Compute $F(\phi, \theta)$ and the residual vector according to Eq. (12).
2. If $\|r(\phi, \theta)\| < \epsilon$, stop. Else, continue.
3. Use (11) to calculate $\nabla F_i(\phi, \theta)$, $1 \leq i \leq 2$, and the Jacobian matrix $J(\phi, \theta)$ whose entries are given by (13).
4. Solve the linear equation $J(\phi, \theta)x^T = r^T(\phi, \theta)$, where $x = (x_1, x_2)$ and T denotes transposition.
5. Update $(\phi, \theta) = (\phi, \theta) + \sum_{j=1}^2 x_j \nabla F_j(\phi, \theta)$. Return to Step 1.

3.2. Closed Pre-Shapes

We now consider a similar manifold \mathcal{C}_n of closed pre-shapes where, in addition to (7), curves satisfy a *closure condition*. A curve α is closed if and only if $\int_0^1 \alpha'(s) ds = 0$. If α is represented by the pair (ϕ, θ) , the closure condition can be expressed as $\int_0^1 e^{\phi(t)} e^{j\theta(t)} dt = 0$, or equivalently,

$$\int_0^1 \cos \theta(t) e^{\phi(t)} dt = 0 \quad \text{and} \quad \int_0^1 \sin \theta(t) e^{\phi(t)} dt = 0.$$

This leads us to consider the mapping $\mathcal{G}: \mathcal{H} \rightarrow \mathbb{R}^4$ defined by

$$\begin{aligned} \mathcal{G}_1(\phi, \theta) &= \int_0^1 e^{\phi(t)} dt; \\ \mathcal{G}_2(\phi, \theta) &= \int_0^1 \theta(t) e^{\phi(t)} dt; \\ \mathcal{G}_3(\phi, \theta) &= \int_0^1 \cos \theta(t) e^{\phi(t)} dt; \\ \mathcal{G}_4(\phi, \theta) &= \int_0^1 \sin \theta(t) e^{\phi(t)} dt. \end{aligned} \quad (14)$$

As before, $\mathcal{G}_1(\phi, \theta) = 1$ and $\mathcal{G}_2(\phi, \theta) = \pi$ ensure that curves have length 1 and angle functions have average π with respect to the arc-length parameter, respectively. $\mathcal{G}_3(\phi, \theta) = \mathcal{G}_4(\phi, \theta) = 0$ guarantee that the closure condition is satisfied. The finite-dimensional analogue of \mathcal{G} is a mapping $G: \mathcal{H}_n \rightarrow \mathbb{R}^4$ obtained by discretizing (14). The pre-shape spaces of closed plane curves are defined as

$$\mathcal{C} = G^{-1}(1, \pi, 0, 0) \quad \text{and} \quad \mathcal{C}_n = G^{-1}(1, \pi, 0, 0).$$

The space \mathcal{C}_n is a submanifold of \mathcal{H}_n whose normal space is 4-dimensional at any point. We compute the derivative of G explicitly. The derivatives of G_1 and G_2 were computed in (11). In the continuous formulation,

$$\begin{aligned} d\mathcal{G}_3(h, f) &= \int_0^1 h(t) \cos \theta(t) e^{\phi(t)} dt \\ &\quad - \int_0^1 f(t) \sin \theta(t) e^{\phi(t)} dt \\ &= \left\langle (h, f), \left(\frac{\cos \theta}{a}, -\frac{\sin \theta}{b} \right) \right\rangle_{(\phi, \theta)}; \end{aligned} \quad (15)$$

$$\begin{aligned} d\mathcal{G}_4(h, f) &= \int_0^1 h(t) \sin \theta(t) e^{\phi(t)} dt \\ &\quad + \int_0^1 f(t) \cos \theta(t) e^{\phi(t)} dt \\ &= \left\langle (h, f), \left(\frac{\sin \theta}{a}, \frac{\cos \theta}{b} \right) \right\rangle_{(\phi, \theta)}. \end{aligned}$$

The corresponding calculation for G yields:

$$\begin{aligned} \nabla G_1(\phi, \theta) &= (1/a, 0); & \nabla G_2(\phi, \theta) &= (\theta/a, 1/b); \\ \nabla G_3(\phi, \theta) &= (\cos \theta/a, -\sin \theta/b); & & (16) \\ \nabla G_4(\phi, \theta) &= (\sin \theta/a, \cos \theta/b). \end{aligned}$$

As in the previous case, two projection algorithms can be derived from this calculation. The extension of Algorithm 3.1.1. to closed shapes is straightforward, the only difference being that Gram-Schmidt is applied to a set of four vectors instead of two. We only present the details of the calculation of the Jacobian matrix of G needed for the projection of \mathcal{H}_n onto \mathcal{C}_n . Other than that, Algorithm 3.1.2. can be reproduced almost word by word.

The normal space $N(\phi, \theta)$ to the level set of G at (ϕ, θ) is spanned by the vectors $\nabla G_i(\phi, \theta)$, $1 \leq i \leq 4$. The Jacobian of the mapping G restricted to $N(\phi, \theta)$ can be expressed in this basis of $N(\phi, \theta)$ and the standard basis of \mathbb{R}^4 as the 4×4 symmetric matrix J whose (i, j) -entry is

$$J_{ij}(\phi, \theta) = \langle \nabla G_i(\phi, \theta), \nabla G_j(\phi, \theta) \rangle. \quad (17)$$

For implementation purposes, we write J more

explicitly. Let

$$\begin{aligned} A(\phi, \theta) &= \sum_{i=1}^n \theta_i \cos \theta_i e^{\phi_i}; \\ B(\phi, \theta) &= \sum_{i=1}^n \theta_i \sin \theta_i e^{\phi_i}; \\ C(\phi, \theta) &= \sum_{i=1}^n \cos^2(\theta_i) e^{\phi_i}; \\ D(\phi, \theta) &= \sum_{i=1}^n \sin(2\theta_i) e^{\phi_i}; \\ E(\phi, \theta) &= \sum_{i=1}^n \theta_i^2 e^{\phi_i}. \end{aligned}$$

Then,

$$J = \begin{bmatrix} G_1/a & G_2/a & G_3/a & G_4/a \\ G_2/a & G_1/b + E/a & A/a - G_4/b & B/a + G_3/b \\ G_3/a & A/a - G_4/b & G_1/b + C \left(\frac{1}{a} - \frac{1}{b} \right) & D \left(\frac{1}{2a} - \frac{1}{2b} \right) \\ G_4/a & B/a + G_3/b & D \left(\frac{1}{2a} - \frac{1}{2b} \right) & G_1/a + C \left(\frac{1}{b} - \frac{1}{a} \right) \end{bmatrix}. \quad (18)$$

4. Geodesics in Pre-Shape Spaces

Let $p_0 = (\phi_0, \theta_0)$ and $p_1 = (\phi_1, \theta_1)$ represent pre-shapes in \mathcal{A}_n or \mathcal{C}_n . The *geodesic distance*, $d(p_0, p_1)$, between p_0 and p_1 is defined as

$$d(p_0, p_1) = \inf_{\gamma} \ell(\gamma), \quad (19)$$

where γ ranges over all piecewise smooth paths in \mathcal{A}_n or \mathcal{C}_n from p_0 to p_1 , and $\ell(\gamma)$ is the length of γ defined as

$$\ell(\gamma) = \int_I \|\gamma'(t)\|_{\gamma(t)} dt = \int_I \langle \gamma'(t), \gamma'(t) \rangle_{\gamma(t)}^{1/2} dt.$$

It is well known that, locally, the geodesic distance is realized by a geodesic. For finite-dimensional, complete Riemannian manifolds, the Hopf-Rinow Theorem asserts that the same holds in general (see e.g. do Carmo (1994)). In the remainder of the paper, we will work under the assumption that distances in \mathcal{A}_n and \mathcal{C}_n are realized by geodesics. Thus, our next goal is to develop an algorithm to calculate geodesics in pre-shape spaces with prescribed initial and terminal points. Our general strategy is similar to that adopted in Klassen et al. (2004), however, the arguments are much more elaborate and the details differ significantly

due to the facts that the inner product on the ambient space \mathcal{H}_n varies from point to point and reparameterizations involve an action of the diffeomorphism group \mathcal{D}_1 .

We give a detailed description of the construction of geodesics in the pre-shape space \mathcal{A}_n ; the corresponding construction for closed pre-shapes only requires minor changes. We begin with the simpler problem of constructing geodesics in \mathcal{A}_n with prescribed first-order initial conditions; that is, initial position $(\phi_0, \theta_0) \in \mathcal{A}_n$ and initial velocity $(h, f) \in T_{(\phi_0, \theta_0)}\mathcal{A}_n$.

4.1. Geodesics with Given Initial Conditions

Suppose $p_0 = (\phi_0, \theta_0) \in \mathcal{A}_n$ and $(h, f) \in T_{(\phi_0, \theta_0)}\mathcal{A}_n$ are given. Infinitesimally, we approximate the geodesic emanating from (ϕ_0, θ_0) with initial velocity (h, f) with a small straight line segment in $\mathcal{H}_n = \mathbb{R}^{2n}$. The terminal point of this segment typically falls slightly off of \mathcal{A}_n , so we use Algorithm 3.1.2. to project it back onto \mathcal{A}_n to obtain (ϕ, θ) . To iterate the construction, we need the velocity vector at the new point, which is the parallel transport of (h, f) along the infinitesimal arc just constructed. As a numerical approximation to the parallel transport, we first update the velocity vector as

$$(\tilde{h}, \tilde{f}) = (h, f) - \delta \left(\frac{h^2}{2} - \frac{b}{2a} f^2, hf \right), \quad (20)$$

where multiplication of vectors is performed coordinate by coordinate. This update rule is derived from the system of differential Eq. (31) satisfied by geodesics in \mathcal{H}_n (see Appendix 9). Using Algorithm 3.1.1, project (\tilde{h}, \tilde{f}) onto $T_{(\phi, \theta)}\mathcal{A}_n$ and scale the vector to have the same norm as the velocity vector at the original point to obtain the new (h, f) . Scaling is used because geodesics are parametric curves with no intrinsic acceleration; in particular, they are traversed with constant speed.

We denote the geodesic starting at $(\phi_0, \theta_0) \in \mathcal{A}_n$ with initial velocity $(h, f) \in T_{(\phi_0, \theta_0)}\mathcal{A}_n$ by $\Psi(\phi_0, \theta_0, h, f; t)$, where t represents the time parameter. The time-1 map $\exp_{p_0} : T_{p_0}\mathcal{A}_n \rightarrow \mathcal{A}_n$ given by

$$\exp_{p_0}(h, f) = \Psi(\phi, \theta, h, f; 1)$$

is known as the *exponential map* at $p_0 = (\phi_0, \theta_0)$.

4.2. Geodesics with Given Boundary Conditions

Constructing a geodesic in \mathcal{A}_n from $p_0 = (\phi_0, \theta_0)$ to $p_1 = (\phi_1, \theta_1)$ is equivalent to finding a direction $(h, f) \in T_{p_0}\mathcal{A}_n$ such that the geodesic starting from p_0 with initial velocity (h, f) reaches p_1 in unit time; that is,

$$\exp_{p_0}(h, f) = p_1. \quad (21)$$

To find such direction, consider the functional $E : T_{p_0}\mathcal{A}_n \rightarrow \mathbb{R}$ defined by

$$E(h, f) = \|\exp(h, f) - p_1\|_{p_0}^2. \quad (22)$$

Notice that $E(h, f) = 0$ if and only if $\exp_{p_0}(h, f) = p_1$. Thus, the desired geodesic is determined by the direction (h, f) that annihilates the functional E . The search for the zeros of E on the inner product space $T_{p_0}\mathcal{A}_n$ can be approached with Newton's method. We outline the estimation of the gradient of E needed in the implementation of Newton's method.

The normal space to \mathcal{A}_n in \mathcal{H}_n at p_0 is spanned by $\nabla F_1(p_0)$ and $\nabla F_2(p_0)$ given by Eq. (11). Define a linear map $L : \mathcal{H}_n \simeq \mathbb{R}^{2n} \rightarrow \mathbb{R}^2$ by

$$L(h, f) = (\langle (h, f), \nabla F_1(p_0) \rangle_{p_0}, \langle (h, f), \nabla F_2(p_0) \rangle_{p_0}).$$

Then, $T_{p_0}\mathcal{A}_n = \ker L$. Find an orthonormal basis $\{(h_1, f_1), \dots, (h_{2n-2}, f_{2n-2})\}$ of $\ker L$ using standard methods. Then, any vector $(h, f) \in T_{p_0}\mathcal{A}_n$ can be written uniquely as

$$(h, f) = x_1(h_1, f_1) + \dots + x_{2n-2}(h_{2n-2}, f_{2n-2}).$$

Hence, (h, f) can be represented by the $(2n - 2)$ -tuple $x = (x_1, \dots, x_{2n-2}) \in \mathbb{R}^{2n-2}$ and E can be treated as a function of x . The partial derivatives of E can be estimated as

$$\frac{\partial E}{\partial x_i} \sim \frac{E(h + \epsilon h_i, f + \epsilon f_i) - E(h, f)}{\epsilon}, \quad (23)$$

for $1 \leq i \leq 2n - 2$. This gives a numerical estimation of the gradient of E needed in the implementation of Newton's method.

As pointed out earlier, the calculation of geodesics in \mathcal{C}_n can be approached with identical methods. The only difference is that \mathcal{C}_n is a $(2n - 4)$ - dimensional

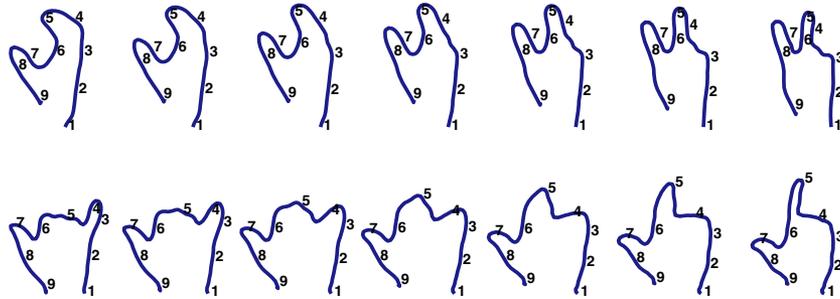


Figure 1. Examples of geodesics in the pre-shape space \mathcal{A}_n .

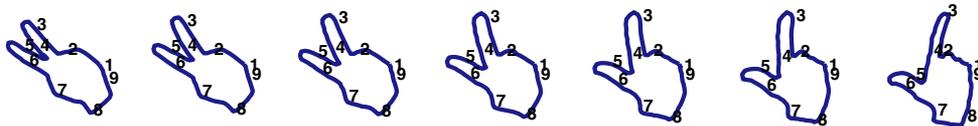


Figure 2. A geodesic in the pre-shape space of closed curves.

submanifold of \mathcal{H}_n . An explicit basis for the normal space at any point has been computed in (16).

4.3. Preliminary Illustrations

Figure 1 shows examples of geodesics in the pre-shape space \mathcal{A}_n computed with the methods described above. A similar illustration in the pre-shape space \mathcal{C}_n of closed curves is shown in Fig. 2. For each curve, the labels represent points associated with 9 uniformly sampled points over the “time” interval I . In each row, the initial and terminal frames display given pre-shapes to be interpolated. The labels illustrate correspondences obtained during geodesic morphing. The elasticity parameters $a = 0.1$ and $b = 1$ were used in our experiments and the curves were sampled at $n = 100$ points. A low tension coefficient was used to allow plenty of stretch elasticity. In this model, shape correspondences obtained during geodesic deformations look natural and intuitively correct. For example, the finger tips tend to be preserved, as well as some other geometric features shared by the curves. Comparisons with shape geodesics in other models are provided below.

Although we provide no theoretical assurance that the geodesics minimize length, extensive experimentation indicates that, for a vast collection of curves, the geodesics obtained are intuitively correct. However, further investigations along these lines are needed.

5. Matching and Interpolating Shapes

As pointed out in the Introduction, in many applications, a correspondence matching points between shapes to be compared is given. For example, in medical imaging, a correspondence between shapes is often established to preserve some landmark points. In such cases, it is desirable that shape metrics and morphing techniques be compatible with the given matching. The same applies to situations where shape correspondences are established using matching techniques that seek to optimally align elements such as velocity fields or curvature functions of contours, as in Cohen et al. (1992), Geiger et al. (1995), Tagare et al. (2002), and Sebastian et al. (2003). The goals of this section are:

- (i) to explain how geodesics in pre-shape spaces can be used to quantify divergence and interpolate shapes in the presence of a preferred matching;
- (ii) to investigate a variant of the dynamic programming shape matching strategy of Tagare et al. (2002), and Sebastian et al. (2003) using an energy functional that is more compatible with the elastic curve model adopted in this paper.

5.1. Interpolations and Metric

We present the details of the case of general plane curves without the closure condition, but closed curves can be treated in a similar manner. Let $\alpha_0, \alpha_1 : I \rightarrow \mathbb{R}^2$

be parametric curves. The time parameter t induces a correspondence between the curves, namely, $\alpha_0(t)$ is matched with $\alpha_1(t)$, for every t . Conversely, for smooth plane curves with no self-intersections, any diffeomorphic correspondence between them arises in this fashion. Moreover, if β_0 and β_1 are parameterizations of the same curves inducing the same correspondence as α_0 and α_1 , then the parameterizations differ by a diffeomorphism $\gamma: I \rightarrow I$; that is,

$$\beta_i(t) = \alpha_i(\gamma(t)),$$

for $i = 0, 1$. This motivates the following approach. Think of a correspondence ρ between two curves C_0 and C_1 as given by a choice of parameterizations α_0 and α_1 , and let $(\phi_0, \theta_0), (\phi_1, \theta_1) \in \mathcal{A}_n$ be the associated pre-shapes obtained after applying the normalizations discussed in Section 3. For a given choice of elasticity constants $a, b > 0$, define the distance between the curves, under the correspondence ρ , to be

$$d_\rho(C_0, C_1) = d_{a,b}((\phi_0, \theta_0), (\phi_1, \theta_1)), \quad (24)$$

where $d_{a,b}$ is the geodesic distance in pre-shape space. The length-minimizing geodesic between the pre-shapes gives a natural interpolation that is compatible with the given correspondence. Notice that the distance $d_\rho(C_0, C_1)$ is well-defined because any other choice of pre-shapes associated with the given correspondence can be obtained by the action of a diffeomorphism γ on $(\phi_0, \theta_0), (\phi_1, \theta_1)$ (see Eq. (3)) and the diffeomorphism group \mathcal{D}_I acts on \mathcal{A}_n by distance-preserving transformations, as explained in Section 2.1.3. A similar fact applies to geodesics. Let $(\phi(t, s), \theta(t, s))$ be a geodesic in \mathcal{A}_n between (ϕ_0, θ_0) and (ϕ_1, θ_1) , where t and s denote the curve and deformation parameters, respectively. Then, the corresponding geodesic between $(\phi_0, \theta_0) \cdot \gamma$ and $(\phi_1, \theta_1) \cdot \gamma$ is obtained by the action of γ ; that is, it is given by

$$(t, s) \mapsto (\phi(\gamma(t), s) + \log \gamma'(t), \theta(\gamma(t), s)).$$

As a consequence, in computations of distances and geodesics with respect to given curve correspondences, we can choose any parameterization for one of the curves and adjust the other to be compatible with the desired matching. In particular, we can assume that one of the curves is parameterized by arc length; i.e., $\phi_0 \equiv 0$.

5.2. Shape Matching with Dynamic Programming

Consider curves represented by the pre-shapes $(\phi_0, \theta_0), (\phi_1, \theta_1)$. As explained in the previous section, to establish a correspondence between the curves, one may fix a parametrization of the first and consider all reparametrizations of the other. Thus, we assume that normalized pre-shapes initially have unit speed parameterizations; that is, $\phi_0 = \phi_1 \equiv 0$. Given elasticity constants $a, b > 0$ and a diffeomorphism $\gamma: I \rightarrow I$, our matching criterion will seek to minimize the functional

$$\begin{aligned} E(\gamma) &= \|(0, \theta_0) - (0, \theta_1) \cdot \gamma\|_{(0, \theta_0)}^2 \\ &= \int_I [a(\log \gamma'(t))^2 + b(\theta_0(t) - \theta_1 \circ \gamma(t))^2] dt, \end{aligned} \quad (25)$$

which quantifies the compatibility of $(0, \theta_0)$ and the reparameterized pre-shape $(0, \theta_1) \cdot \gamma$ from the viewpoint of (ϕ_0, θ_0) . More symmetric forms of the energy can be considered as in Tagare (1999). As before, each pre-shape is sampled uniformly at n points on the interval I ; in our experiments, we use $n = 100$. Consider the uniform $n \times n$ grid on the square $I \times I$ with grid points labeled (i, j) , $0 \leq i, j \leq n - 1$. As indicated in Fig. 3(a), diffeomorphisms of the interval I will be approximated by piecewise linear (PL) homeomorphisms whose graphs are PL paths on the square $I \times I$ from $(0, 0)$ to $(n - 1, n - 1)$, with each node a grid point. Note that all segments in such paths have positive slope. The matching problem is now reduced to finding the allowable path that minimizes (a discrete form of) the energy $E(\gamma)$. Dynamic programming, DP, is well suited for this problem since the cost associated

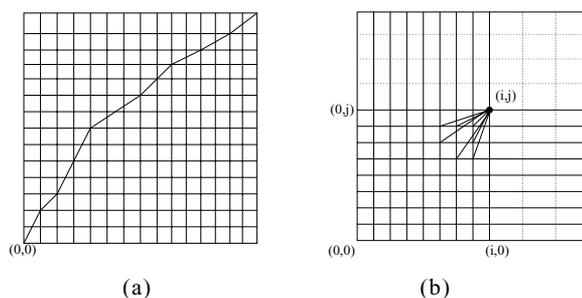


Figure 3. (a) A piecewise linear path on the square $I \times I$ approximating the graph of a diffeomorphism; (b) Restricting the possible slopes of the path at a node (i, j) .

with a path is additive over its segments. Thus, it is natural to introduce a localized form of the energy E over a segment. For $k < i$ and $l < j$, let $L(k, l; i, j)$ denote the line segment joining nodes (k, l) and (i, j) and let

$$E(k, l; i, j) = \int_{I_{ki}} [a(\log \gamma'_{klij}(t))^2 + b(\theta_0(t) - \theta_1 \circ \gamma_{klij}(t))^2] dt, \quad (26)$$

where $I_{ki} \subseteq I$ is the subinterval determined by the points indexed by k and i , and γ_{klij} is the linear diffeomorphism from I_{ki} to I_{lj} whose graph is $L(k, l; i, j)$.

To minimize the energy, in principle, one should consider all possible ways of reaching a node (i, j) through segments of the form $L(k, l; i, j)$, $k < i$ and $l < j$. For computational efficiency, in practice, we restrict the indexes (k, l) to a subset N_{ij} . A possible choice of N_{ij} is illustrated in Fig. 3(b). Define the minimum energy $H(i, j)$ needed to reach the node (i, j) , iteratively, as follows:

- (i) $H(0, 0) = 0$;
- (ii) $H(i, j) = E(\hat{k}, \hat{l}; i, j) + H(\hat{k}, \hat{l})$,

where

$$(\hat{k}, \hat{l}) = \underset{(k, l) \in N_{ij}}{\operatorname{argmin}} (E(k, l; i, j) + H(k, l)).$$

The energy $H(i, j)$ is computed sequentially, starting from $(0, 0)$ and increasing i, j until all allowable vertices are visited. For nodes that cannot be reached, we preassign $H(i, j) = \infty$. This is the case, e.g., for all $(i, 0)$ and $(0, j)$, $1 \leq i, j \leq n - 1$. Once $H(n - 1, n - 1)$ has been computed, we backtrack to find a path of minimum energy that represents a *PL* homeomorphism of the interval I . Further details on the calculation of $E(k, l; i, j)$ are presented in Appendix B.

Figure 4(a) illustrates this shape matching strategy applied to pre-shapes initially parameterized by arc length; i.e., $\phi_0 = \phi_1 \equiv 0$. We adopted a low value for the tension parameter a in order to allow plenty of stretch elasticity. As stretching is essentially not penalized, the matching diffeomorphism γ will attempt to align the angle functions θ_0 and $\theta_1 \circ \gamma$ as much as possible to minimize E . Figure 4(b) is an intensity plot of the energy H and the overlaid path represents the optimal diffeomorphism γ estimated using dynamic programming. Figure 4(c) illustrates the improved alignment of angle functions after the action of γ on θ_1 .

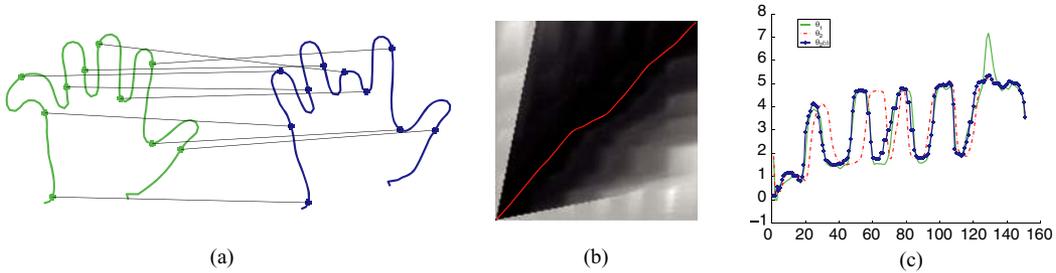


Figure 4. (a) Optimal correspondence between two shapes; (b) an intensity plot of the energy H and the graph of the matching diffeomorphism γ ; (c) a plot of the angle functions θ_0 , θ_1 and $\theta_1 \circ \gamma$.

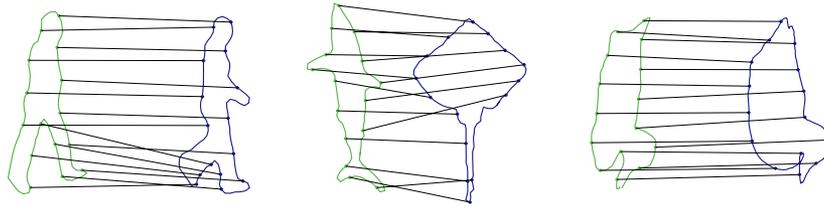


Figure 5. Optimal correspondences between closed curves.

The implementation for closed contours is similar, except that we also need to consider all possible choices of initial points. Computational strategies to manage this extra degree of freedom are discussed in Sebastian et al. (2003). Examples of correspondences for closed shapes are shown in Fig. 5.

5.3. Geodesics

We conclude this section with several examples of geodesics between closed shapes that respect optimal alignments obtained via the dynamic programming approach of Section 5.2. To illustrate the improvement over geodesics in the more rigid, bending-only shape model of Klassen et al. (2004) and the landmark model of Kendall (1984), we offer a few comparisons of these three types of geodesics in Fig. 6.

Figure 7 shows an application of elastic shape geodesics to echocardiography. The first and last panels show the *end diastolic* (ED) and *end systolic* (ES) frames, taken from the apical four-chamber view, during systole (the contracting part of the cardiac cycle). Overlaid on these frames are expert tracings of the

epicardial (as solid lines) and endocardial (as dashed lines) contours of the left ventricle. Among other things, cardiologists are interested in the temporal evolution of these contours. Geodesics in the pre-shape space of elastic curves with *DP* alignment were scaled and positioned appropriately to estimate the contours for intermediate frames. Images were acquired a rate of 30 frames per second; some of the frames are displayed in the figure. Modeling shapes on elastic curves that can be stretched, compressed and bent in different ways along the extension of the curve has a series of advantages. For example, one can observe that the apex of both endocardial and epicardial contours as well as sharp corners and edges are well aligned during the evolution.

6. Spaces of Elastic Shapes

Shape metrics and interpolations based on geodesics in pre-shape spaces and *DP* alignment, as discussed in Section 5, are very attractive from a computational standpoint. Extensive experimentation also indicates that matchings and interpolations so obtained tend to

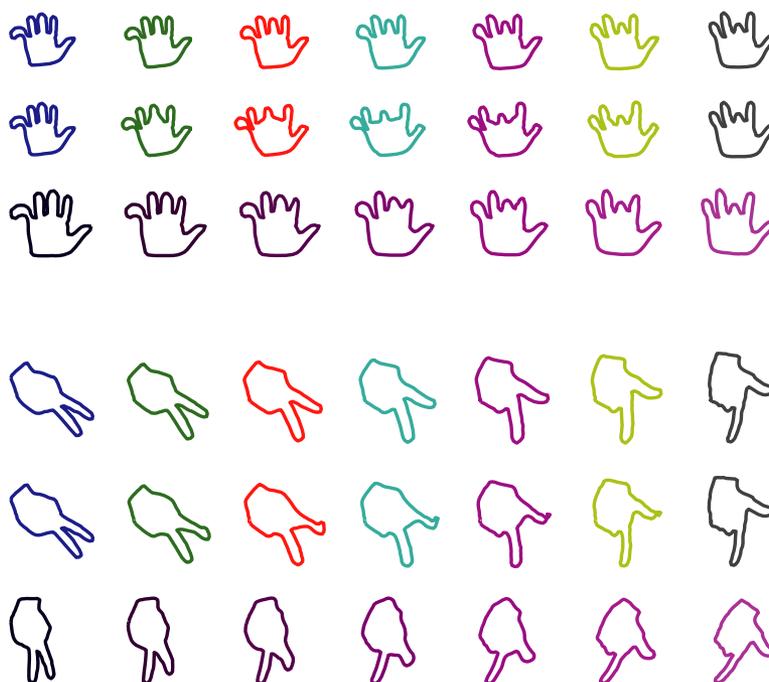


Figure 6. Geodesics in the fully elastic, bending-only and Kendall models. In each group, the first row shows a geodesic between elastic shapes with low tension coefficient, aligned with dynamic programming techniques. For comparison purposes, the second and third rows in each group display geodesics in the bending-only and Procrustes models, respectively.

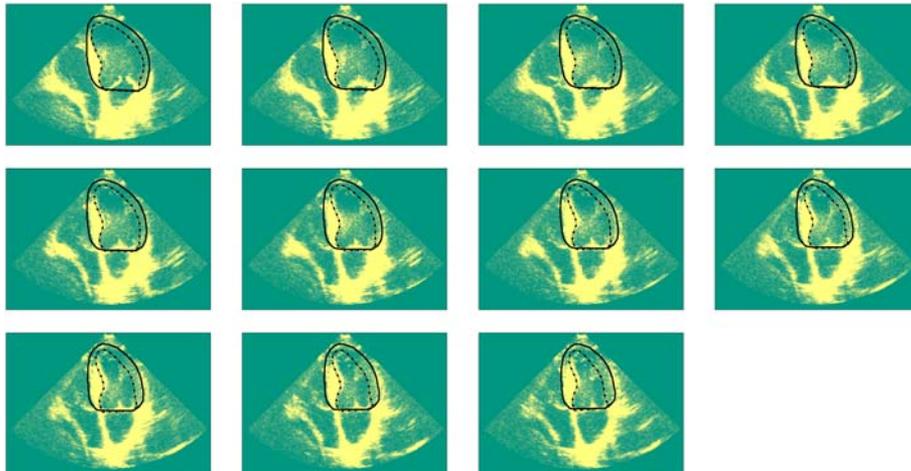


Figure 7. Interpolation using an elastic shape geodesic of the expert-drawn epicardial and endocardial contours of the left ventricle at end diastole and end systole to estimate their evolution during systole.

be intuitively correct; further evidence of this fact will be provided below. Thus, in practice, this is the model we propose to adopt. The framework that we have developed also allow us to define *spaces of shapes* as quotient spaces of pre-shape spaces under the action of reparameterizations and compute distances and interpolations without assuming any preferred shape correspondences; as a matter of fact, finding an optimal correspondence becomes an intrinsic part of the problem of defining geodesic distance in shape space. However, the additional computational costs are significant. Our experiments indicate that geodesics with *DP* alignment offer accurate approximations to actual geodesics in shape space. For completeness, we briefly outline the construction of shape spaces, metrics, and morphing techniques.

6.1. Shape of Planar Arcs

The shape of a planar arc admits multiple representative pre-shapes, as defined in Section 3.1, due to all possible reparameterizations of a curve. Thus, we define the shape space \mathcal{S} of planar arcs as the quotient space of \mathcal{A} by the action of \mathcal{D}_I , that is, $\mathcal{S} = \mathcal{A}/\mathcal{D}_I$. Define a (pseudo) metric on \mathcal{S} by

$$d(s_1, s_2) = \min_{(\phi_1, \theta_1), (\phi_2, \theta_2)} d((\phi_1, \theta_1), (\phi_2, \theta_2)),$$

where the minimum is taken over all pre-shapes (ϕ_1, θ_1) and (ϕ_2, θ_2) representing s_1 and s_2 , respectively. Let

$(0, \theta_i^*)$ be the representative of s_i , $i = 1, 2$, parameterized by arc length. Then, any pre-shape representing s_i can be written as $(0, \theta_i^*) \cdot \gamma$, with $\gamma \in \mathcal{D}_I$. Moreover, since \mathcal{D}_I acts on \mathcal{A} by isometries, we can fix the representative of s_1 to be $(0, \theta_1^*)$ and take the minimum only over (ϕ_2, θ_2) . Thus, the distance in \mathcal{S} can be expressed as

$$d(s_1, s_2) = \inf_{\gamma \in \mathcal{D}_I} d((0, \theta_1^*), (0, \theta_2^*) \cdot \gamma).$$

The cost function (22) used in the construction of geodesics in pre-shape spaces is now modified to

$$E(h, f; \gamma) = \left\| \exp_{(0, \theta_0^*)}(h, f) - (0, \theta_1^*) \cdot \gamma \right\|_{(0, \theta_0^*)}^2. \quad (27)$$

to include the action of γ . A path realizing the geodesic distance has the property that it is orthogonal to the orbit of the diffeomorphism group at any point. Thus, in the construction of geodesics, the search for the appropriate direction (h, f) to shoot a geodesic from $(0, \theta_0^*)$ is restricted to those tangent to \mathcal{A}_n and perpendicular to the orbit of \mathcal{D}_I at $(0, \theta_0^*)$ in \mathcal{H}_n . The minimization of E is carried out iteratively alternating between the variables (h, f) and γ . The optimization over γ is done with dynamic programming as in Section 5.2. Gradient methods are utilized for the optimization over (h, f) .

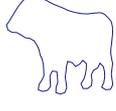
| Sample Shapes | | | | Elastic | Bending Only |
|---|---|---|---|--|---|
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |

Figure 8. Mean shapes calculated using the elastic model and the bending-only model.

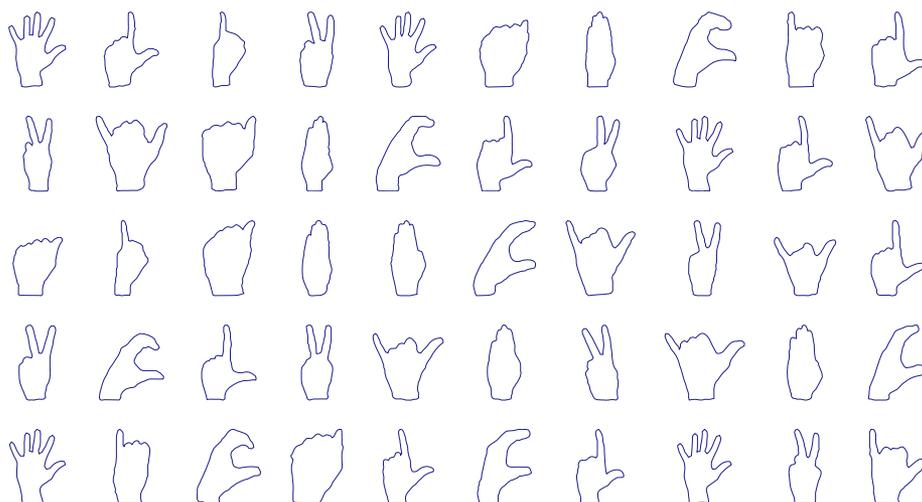


Figure 9. A database of 50 shapes of hands at different poses.

6.2. Closed Shapes

Geodesic distance and morphing techniques for closed shapes can be treated in an almost identical manner, with the extra closure condition enforced.

7. Mean Shapes

Observations of the shape of an object or shapes associated with a family of similar objects typically exhibit significant variations. For example, contours of objects

in images are subject to variations due to pose, perspective, noise, and partial occlusions. Hence, an important goal in the algorithmic study of shapes is to develop tools for a statistical treatment of shapes. Tangent-space representation is becoming a standard approach to the analysis of data on Riemannian manifolds (Dryden and Mardia, 1998). The idea is to define and calculate the sample mean and lift the data to the tangent space at the mean via the inverse exponential map. The data points are now represented as points on a vector space equipped with an inner product where standard tools of data analysis can be applied. For example, estimation of covariance in the tangent-space representation has been studied in many different contexts (see e.g. (Dryden and Mardia, 1998; Srivastava et al., 2005; Vailliant et al., 2004)). Thus, in this paper, we focus on the most basic notion needed, that of mean elastic shape. We adopt the notion of Fréchet mean shape that

has been previously used in the landmark and bending-only models (Dryden and Mardia, 1998; Klassen et al., 2004). These are special cases of a more general notion of mean on a Riemannian manifold studied by Karcher (1977). We present a formulation for closed shapes, but shapes of elastic arcs can be treated similarly.

Let s_1, \dots, s_k be a collection of closed shapes. For a choice of elasticity constants $a, b > 0$, define the scatter of the collection with respect to a shape s to be

$$V(s) = \frac{1}{2} \sum_{i=1}^k d_{a,b}^2(s, s_i), \quad (28)$$

where $d_{a,b}$ denotes geodesic distance. A *Fréchet mean* of the collection is defined as a shape that is a (local) minimum of V . In practice, we approximate $d_{a,b}(s, s_i)$ by the geodesic distance in pre-shape space after a *DP* alignment, as discussed in Section 5. To search for a Fréchet mean of the collection, we adopt a gradient-type strategy. For data on a Riemannian manifold, the (negative) gradient of the scatter functional V is known to be given by

$$-\nabla V(s) = \sum_{i=1}^k v_i, \quad (29)$$

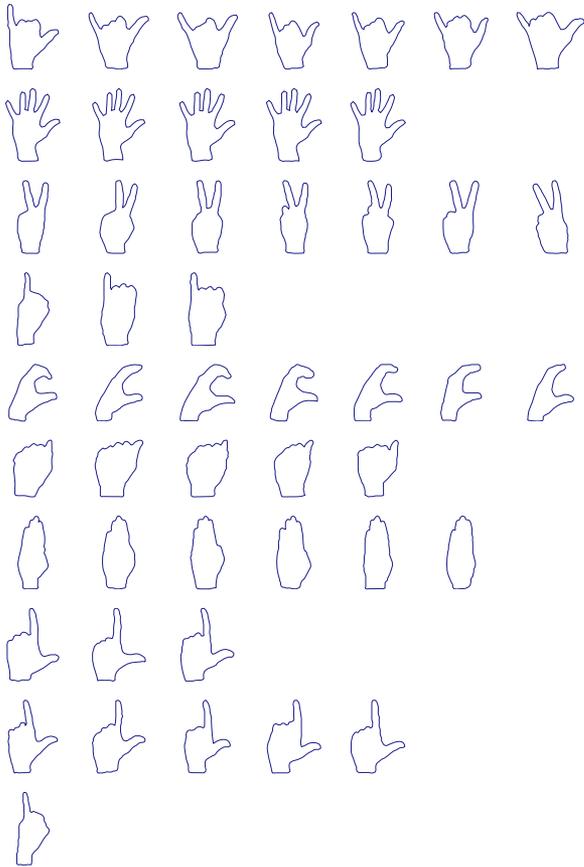


Figure 10. Ten clusters obtained with a variant of the *k-Means Algorithm* applied to the shapes shown in Fig. 9 using the geodesic distance with *DP* alignment as metric.

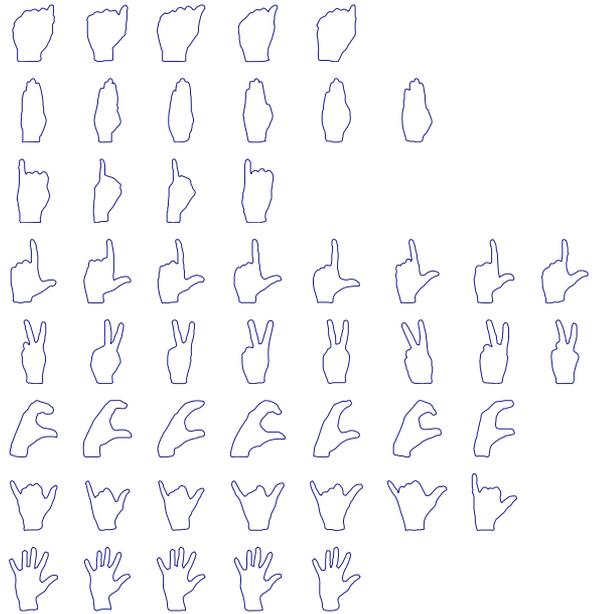


Figure 11. Eight clusters obtained using a hierarchical clustering algorithm using the elastic geodesic distance with *DP* alignment as metric.

where v_i is the initial velocity vector of the geodesic that runs from s to s_i in unit time (Karcher, 1977); that is, $\exp_s(v_i) = s_i$. Motivated by this observation, initialize the search with s as one of the shapes in the given collection and choose pre-shapes $p = (\phi, \theta)$ and $p_i = (\phi_i, \theta_i)$ representing s and s_i , $1 \leq i \leq k$, with each p_i aligned with p via dynamic programming. Find the geodesics in the pre-shape manifold \mathcal{C}_n from p to p_i and calculate their initial velocities $v_i = (h_i, f_i)$. Then, infinitesimally displace p in \mathcal{C}_n along a geodesic from p in the direction $v = \sum_{i=1}^k v_i$ (see Section 4.1). The updated shape s is the shape associated with the new pre-shape p . The process is iterated until the norm of v falls below a set threshold thus yielding an estimation of a mean shape.

Figure 8 shows several examples of mean shapes computed with the algorithmic procedure just described. For comparison purposes, we also display the mean shape in the more rigid bending-only model of Klassen et al. (2004). The examples illustrate well the fact that the finer shape correspondences built-in to the elastic model leads to more natural mean shapes. On the first three rows, unlike the bending-only means, the elastic means preserve the main features shared by the shapes. Similarly, on the fourth row, more of the

common sharp features of the sample shapes are retained by the elastic mean.

8. Clustering

Alongside statistical modeling of shapes, clustering techniques are of basic importance in applications involving large collections of shapes. In Srivastava et al. (2005), a variant of the classical *k-Means Algorithm* was developed and implemented for clustering shapes using the bending-only model. The MCMC techniques adopted in that paper for the construction of clusters apply to other shape metrics since the shape model is only used to compute the pairwise distances between the sample shapes. Thus, here, we just offer illustrations of results obtained by replacing the shape metric with the elastic geodesic distance with *DP* alignment. Figure 9 shows a collection of 50 shapes of hands at various different poses. Clusters obtained with the *k-Means Algorithm* with $k = 10$ are displayed in Fig. 10. The results of another clustering experiment with the elastic geodesic distance with *DP* alignment are shown in Fig. 11. Here, the 50 sample shapes were grouped into 8 clusters using a hierarchical clustering algorithm.

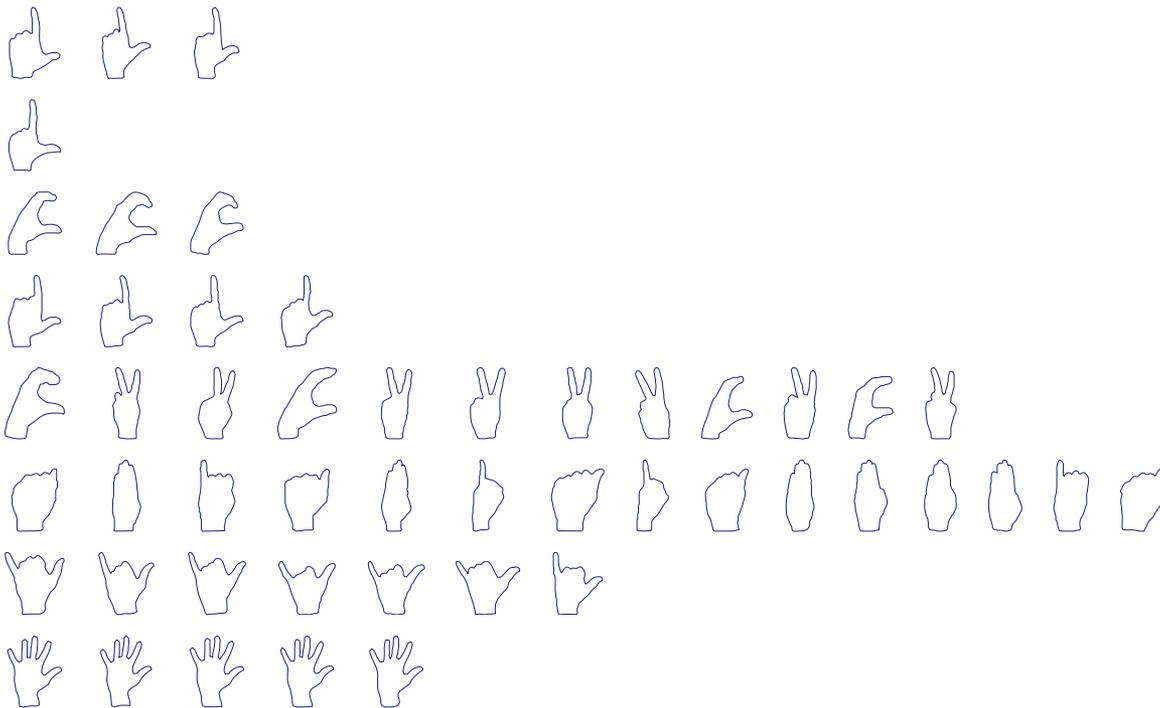


Figure 12. Eight clusters obtained using a hierarchical clustering algorithm with the bending-only geodesic distance.

Starting with 50 clusters, each consisting of one of the given shapes, clusters were combined successively using the nearest neighbor as merging criterion. For comparison purposes, in Fig. 12 we show clusters obtained with the bending-only model. This experiment illustrates well the fact that the elastic model produces more natural groupings of shapes.

9. Summary and Concluding Remarks

Motivated by problems arising in computer vision and image understanding, we developed a new treatment of shapes of plane curves, where shapes are modeled on elastic strings that can stretched, compressed and bent at varying rates along their extensions. Shape spaces were constructed equipped with metrics that reflect elastic properties of the curves, and geodesics were used to quantify shape divergence and morph shapes. Approximations of elastic shape geodesics using dynamic programming alignment and geodesics in pre-shapes were presented that make the model practically feasible. The techniques developed were applied to statistical modeling of families of shapes and to the design of shape clustering algorithms. A distinctive advantage of the elastic shape model presented in this paper is the more natural and intuitive shape correspondences associated with shape distances and geodesics making it more suitable for various applications. Several comparisons with existing models were offered to illustrate this point.

Appendix

A. Geodesics in \mathcal{H}_n

Given elasticity constants $a, b > 0$ and $(x, y) \in \mathbb{R}^2$, define an inner product by

$$\langle (v_1, v_2), (w_1, w_2) \rangle_{(x,y)} = ae^x v_1 w_1 + be^x v_2 w_2. \quad (30)$$

The plane \mathbb{R}^2 equipped with the Riemannian structure induced by $\langle \cdot, \cdot \rangle_{(x,y)}$ will be denoted $\mathbb{L}_{a,b}$. From (6), the space \mathcal{H}_n can be expressed as the Cartesian product of n copies of $\mathbb{L}_{a,b}$; that is,

$$\mathcal{H}_n \cong \prod_{i=1}^n \mathbb{L}_{a,b}.$$

Thus, to calculate geodesics in \mathcal{H}_n , it suffices to derive the differential equation for geodesics in $\mathbb{L}_{a,b}$. The metric tensor is given in x - y coordinates by

$$g_{11}(x, y) = a e^x, \quad g_{12}(x, y) = g_{21}(x, y) = 0, \\ g_{22}(x, y) = b e^x,$$

and the Christoffel symbols by

$$\Gamma_{11}^1 = \Gamma_{12}^2 = \Gamma_{21}^2 = \frac{1}{2}, \quad \Gamma_{11}^2 = \Gamma_{12}^1 = \Gamma_{21}^1 = \Gamma_{22}^2 = 0, \\ \Gamma_{22}^1 = -\frac{b}{2a}.$$

This implies that a curve $\alpha(t) = (x(t), y(t))$ is a geodesic if and only if it satisfies the system

$$x'' + \frac{1}{2}(x')^2 - \frac{b}{2a}(y')^2 = 0 \quad \text{and} \quad y'' + x'y' = 0 \quad (31)$$

of second-order differential equations (see e.g. do Carmo (1994)).

B. Calculation of E

Let $0 = t_0 < t_1 < \dots < t_{n-1} = 1$ be uniformly spaced. Then, the node labeled (i, j) in Section 5.2 has coordinates (t_i, t_j) . To describe the calculation of $E(k, l; i, j)$ defined in (26), we first discuss the action of γ_{klij} on a function θ defined on the interval $I_{ki} = [t_k, t_i]$. For $k \leq r \leq i$, we use the notation $\theta(r)$ for $\theta(t_r)$. The only approximation needed in the calculation of the action comes from the fact that $\gamma_{klij}(r)$ may not fall on the grid, in which case, we use a linear interpolation, as follows: if $\gamma_{klij}(r) \in [s_j, s_{j+1}]$, write

$$\gamma_{klij}(r) = \xi s_j + (1 - \xi) s_{j+1}, \quad (32)$$

with $0 \leq \xi \leq 1$. Then, $\theta \circ \gamma_{klij}(r)$ is calculated using the corresponding convex combination of $\theta(j)$ and $\theta(j+1)$; that is,

$$\theta \circ \gamma_{klij}(r) = \xi \theta(j) + (1 - \xi) \theta(j+1). \quad (33)$$

The derivative γ'_{klij} is constant since γ_{klij} is linear. Then, the integral in (26) is discretized using standard methods.

Acknowledgments

This work was supported in part by National Science Foundation grants CCF-0514743 and DMS-0101429, and Army Research Office grant W911NF-04-01-0268. The echocardiographic images used in the paper are courtesy of David Wilson at the University of Florida.

References

- Beg, M.F., Miller, M.I., Troune, A., and Younes, L. 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61:139–157.
- Bookstein, F.L. 1986. Size and shape spaces for landmark data in two dimensions. *Statistical Science*, 1:181–242.
- Cohen, I., Ayache, N., and Sulger, P. 1992. Tracking points on deformable objects using curvature information. In *Lecture Notes in Computer Science*, vol. 588.
- Cootes, T.F., Taylor, C.J., Cooper, D.H., and Graham, J. 1995. Active shape models: Their training and applications. *Computer Vision and Image Understanding*, 61:38–59.
- do Carmo, Manfredo, 1994. *Riemannian Geometry*. Birkhauser.
- Dryden, I.L. and Mardia, K.V. 1998. *Statistical Shape Analysis*. John Wiley & Son.
- Geiger, D., Gupta, A., Costa, L.A., and Vlontzos, J. 1995. Dynamic programming for detecting, tracking and matching elastic contours. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(3):294–302.
- Grenander, U. 1993. *General Pattern Theory*. Oxford University Press.
- Karcher, H. 1977. Riemann center of mass and mollifier smoothing. *Comm. Pure Appl. Math.*, 30:509–541.
- Kendall, D.G. 1984. Shape manifolds, Procrustean metrics and complex projective spaces. *Bulletin of London Mathematical Society*, 16:81–121.
- Klassen, E., Srivastava, A., Mio, W., and Joshi, S. 2004. Analysis of planar shapes using geodesic paths on shape manifolds. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 26:372–383.
- Le, H.L. and Kendall, D.G. 1993. The Riemannian structure of euclidean shape spaces: a novel environment for statistics. *Annals of Statistics*, 21(3):1225–1271.
- Mennucci, A. and Yezzi, A. 2004. Metrics in the space of curves. Technical report.
- Michor, P. and Mumford, D. Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.* in press.
- Sebastian, T.B., Klein, P.N., and Kimia, B.B. 2003. On aligning curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(1):116–125.
- Sharon, E. and Mumford, D. 2004. 2D-Shape analysis using conformal mappings. In *Proceedings of IEEE Conference on Computer Vision*, pp. 350–357.
- Srivastava, A., Joshi, S., Mio, W., and Liu, X. 2005. Statistical shape analysis: Clustering, learning and testing. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 27:590–602.
- Tagare, H.D. 1999. Shape-based non-rigid correspondence with applications to heart motion analysis. *IEEE Trans. on Medical Imaging*, 8(7):570–579.
- Tagare, H.D., O’Shea, D., and Groisser, D. 2002. Non-rigid shape comparison of plane curves in images. *Journal of Mathematical Imaging and Vision*, 16:57–68.
- Thompson, D.W. 1992. *On Growth and Form: The Complete Revised Edition*. Dover.
- Vailliant, M., Miller, M.I., and Younes, L. 2004. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23:161–169.
- Younes, L. 1998. Computable elastic distance between shapes. *SIAM Journal of Applied Mathematics*, 58:565–586.
- Younes, L. 1999. Optimal matching between shapes via elastic deformations. *Journal of Image and Vision Computing*, 17(5/6):381–389.