

Shape of Elastic Strings in Euclidean Space

Washington Mio · John C. Bowers · Xiuwen Liu

Received: 14 September 2007 / Accepted: 14 October 2008 / Published online: 15 November 2008
© Springer Science+Business Media, LLC 2008

Abstract We construct a 1-parameter family of geodesic shape metrics on a space of closed parametric curves in Euclidean space of any dimension. The curves are modeled on homogeneous elastic strings whose elasticity properties are described in terms of their tension and rigidity coefficients. As we change the elasticity properties, we obtain the various elastic models. The metrics are invariant under reparametrizations of the curves and induce metrics on shape space. Analysis of the geometry of the space of elastic strings and path spaces of elastic curves enables us to develop a computational model and algorithms for the estimation of geodesics and geodesic distances based on energy minimization. We also investigate a curve registration procedure that is employed in the estimation of shape distances and can be used as a general method for matching the geometric features of a family of curves. Several examples of geodesics are given and experiments are carried out to demonstrate the discriminative quality of the elastic metrics.

Keywords Shape analysis · Shape space · Shape geodesics · Elastic shapes · Shape manifold

W. Mio (✉)
Department of Mathematics, Florida State University,
Tallahassee, FL 32306-4510, USA
e-mail: mio@math.fsu.edu

J.C. Bowers · X. Liu
Department of Computer Science, Florida State University,
Tallahassee, FL 32306-4530, USA

1 Introduction

Much of the present interest in shapes of curves in Euclidean space stems from questions arising in computer vision, medical imaging and pattern recognition. Many studies of shapes have aimed at the design of discriminative shape descriptors, which are useful in retrieval and classification problems. A different viewpoint is based on the construction of shape spaces—in which “all” shapes are represented—equipped with metrics that attempt to quantify shape resemblance and dissimilarity. This provides a framework for shape analysis and inference, and the general philosophy is more in line with *Pattern Theory* (Grenander 1993; Mumford 2002). We seek to integrate these views by constructing shape spaces of curves equipped with families of geodesic metrics, whose choice can be tuned to a particular problem to enhance the discriminative qualities of the metric.

Our perception of shapes and the notions of shape similarity and divergence tend to be very contextual. This motivates the development of flexible models that are adaptable to different settings. A structured family of shape metrics, from which one may select or learn a metric that is best suited to a specific scenario, can be instrumental in the mathematical formulation of context dependence. Michor and Mumford offer an overview of the organization of some shape metrics in Michor and Mumford (2007). Largely motivated by this problem, we extend the 1-parameter family of *homogeneous elastic models* (HEM) for plane curves developed in Mio et al. (2007b) to geodesic metrics for curves in Euclidean space of any dimension. The shape metrics are given by geodesic distances calculated with respect to Riemannian structures on a manifold of elastic strings. As we change the elasticity properties of the strings, we obtain the different metrics in the family. The development of these

models and the investigation of their geometry for the design of effective computational strategies form the core of this paper. Even in the planar case, the formulation and computational models of this paper lead to algorithms that are more efficient and robust than those of Mio et al. (2007b).

Shape spaces of curves equipped with a variety of metrics have been investigated in several recent studies (Klassen et al. 2004; Michor and Mumford 2006; Mio et al. 2007b; Klassen and Srivastava 2006; Joshi et al. 2007; Michor et al. 2007). Two main types of shape representation have been adopted: curves as subspaces of \mathbb{R}^k and parametric presentations of curves. The most common view is that two curves have the same shape if they differ by the action of the group generated by rigid transformations and homotheties of \mathbb{R}^k . If parametrizations are used, shape is also preserved under reparameterizations by diffeomorphisms of the parameter space. The parametric models of this paper yield metrics that are invariant under rigid transformations, scale and reparameterizations, thus inducing metrics on shape space. We also develop a curve registration technique that allows more efficient estimations of shape distances and geodesics. The curve matching algorithm is implemented via dynamic programming and is a fully symmetric, multidimensional variant of those investigated in Tagare (1999), Sebastian et al. (2003), Zheng et al. (2005), Mio et al. (2007b). The technique can be employed not only to help estimate shape geodesics, but also as a tool to establish correspondences between curves in a given family.

To motivate the representation of curves adopted in the paper and explain the nature of the elastic metrics, we begin with a few remarks on some shape models of parametric plane curves developed in previous studies. We focus on the case of closed curves, as shapes of arcs can be treated with a simpler version of the same techniques. The Riemannian model of Klassen et al. (2004) employs a representation of plane curves via direction (or angle) functions with respect to the arc-length parameter. The model was first implemented using shooting methods. Later, energy minimization was used in Schmidt et al. (2006) to improve the computational efficiency, and an extension to curves in \mathbb{R}^k was studied in Klassen and Srivastava (2006). A drawback of this model is that the shape representation relies on the arc-length parameter so that geodesic deformations are not free to stretch or compress curves to match their geometric features. A geodesic deformation essentially gives the most efficient way of bending a shape into another respecting the arc-length parameter. In that model, curves are flexible, but have infinite tension and do not allow any tangential deformations. As a consequence, the resulting shape geodesics often do not yield natural or intuitive deformations. Thus, relaxation of the tension of the strings is very desirable for shape analysis. An early elastic model of shapes of plane curves was proposed by Younes (1998, 1999) and

provides a combination of bending and stretching elasticity. This model has been fully worked out in recent work by Michor et al. (2007). Mio et al. (2007b) introduced a continuous 1-parameter family of metrics that considers both the stretching and bending properties using a first-order representation of plane curves. A variant for plane curves based on (signed) curvature functions was subsequently studied by Shah (2006). More recently, Joshi et al. (2007) introduced a Riemannian pre-shape space of curves in \mathbb{R}^k based on a square-root representation of velocity fields, which turns out to be isometric to one of the models constructed in this paper. To describe the Riemannian models studied in Mio et al. (2007b) and motivate the present treatment, we first introduce some notation.

Let \mathbb{S}^1 be the unit circle in \mathbb{R}^2 centered at the origin. We express a point $z \in \mathbb{S}^1$ as $z = e^{js}$, where $s \in [0, 2\pi]$ and $j = \sqrt{-1}$. Given a parametric closed plane curve $\alpha: \mathbb{S}^1 \rightarrow \mathbb{R}^2$, we denote the velocity vector at z by $v(z) = \partial_s \alpha(e^{js})$. We often abuse notation and write $v(z) = \alpha'(s)$. Assuming that the curve is non-singular (that is, the velocity vector never vanishes), write

$$\alpha'(s) = e^{\varphi(s)} e^{j\theta(s)}, \quad (1)$$

where $\varphi(s)$ is the speed of the curve expressed in logarithmic scale and θ is a measurement of the angle that $\alpha'(s)$ makes with a horizontal axis. If we express an infinitesimal variation of (φ, θ) as (h, f) , the shape model of Mio et al. (2007b) is based on the Riemannian metric

$$\begin{aligned} \langle (h_1, f_1), (h_2, f_2) \rangle_{(\varphi, \theta)} &= a \int_0^{2\pi} h_1(s) h_2(s) e^{\varphi(s)} ds \\ &+ b \int_0^{2\pi} f_1(s) f_2(s) e^{\varphi(s)} ds, \quad (2) \end{aligned}$$

where $a, b > 0$ are parameters that can be interpreted as the tension and rigidity coefficients of the curves. Large values of a relative to b indicate that the curves offer much higher resistance to stretching or compression than to bending. If a is small as compared to b , the curves are tensile and compressible, but more rigid. Since the arc-length element of α is $d\ell = e^{\varphi(s)} ds$, (2) defines a weighted inner product with respect to $d\ell$. Multiplying a and b by a common factor simply scales the metric globally. Therefore, this construction essentially yields a 1-parameter family of metrics. To obtain a shape metric, various normalizations are made on (φ, θ) to ensure that the representation is invariant under shape preserving transformations. Reparameterizations via diffeomorphisms of the circle are taken into account, as well.

In this paper, we extend the elastic models of shapes of plane curves to closed curves in arbitrary Euclidean space. We devise computational strategies and develop algorithms to calculate geodesics and geodesic distances that apply equally well to all metrics in the family. Shooting

methods for the calculation of geodesics have been used in some previous studies of plane shapes (Klassen et al. 2004; Mio et al. 2007b). However, energy minimization is a more attractive alternative as it leads to more efficient and robust algorithms for the calculation of shape geodesics. Evidence supporting this fact is offered by the results of Schmidt et al. (2006), Klassen and Srivastava (2006), Joshi et al. (2007). Energy minimization leads to an improvement of computational efficiency over shooting methods by orders of magnitude. However, in contrast with the computational models of Klassen and Srivastava (2006), Joshi et al. (2007), the energy minimization algorithms of this paper have the added advantage that each step of the minimization process scales linearly with the dimension k of the ambient Euclidean space. Given two shapes, they are first connected by a path in the pre-shape manifold, which is gradually deformed to a geodesic following the negative gradient flow of the energy. A substantial part of this work is devoted to the investigation of the geometry of pre-shape manifolds of elastic strings and path spaces of pre-shapes to set up and integrate the evolution equations. From a computational standpoint, the payoff of this detour through geometry is a set of algorithms that are efficient and rather simple to implement. We illustrate the flexibility offered by the models through various examples, carry out shape retrieval experiments and compare the performance of the metrics with some previously reported results to demonstrate the ability of the elastic metrics to discriminate shapes.

Another important problem in shape analysis is to determine the particular features or regions that make two shapes to be perceived as similar or different. For example, in medical imaging, shape differences are often concentrated in particular areas, so it is important to have a localization tool to characterize and detect the regions where the main morphological differences occur. To quantify these local contributions, we resort to the geodesic deformation fields and introduce the notion of energy density functions that describe the distribution of the total (geodesic) deformation energy along the strings. We also examine the limit behavior of the HEM metrics as the tension coefficient $a \rightarrow \infty$. As expected, the ∞ -tension limit coincides with the arclength model of Klassen and Srivastava (2006). At the other extreme, with the full relaxation of the tension (that is, as $a \rightarrow 0$), the homogeneous elastic metrics degenerate and become singular in directions tangential to the curves. Extensions of the model to shapes of surfaces and other multi-dimensional objects, the investigation of models of inhomogeneous and anisotropic elastic shapes, the study of learning techniques to select the elasticity coefficients for a particular problem, and statistical modeling of shapes based on the elastic metrics will be considered in future work.

The paper is organized as follows. In Sect. 2, we construct the pre-shape manifold of closed elastic strings and

introduce the homogeneous Riemannian metrics. In Sect. 3, we show that the metrics are invariant under the action of the diffeomorphism group of the circle and thus induce metrics on shape space. Matching of a family of curves is discussed in Sect. 4. Path spaces and the energy functional are introduced in Sect. 5. This is followed by a discussion, in Sect. 6, of a procedure to find the closest pre-shape to a given curve using a Riemannian version of Newton's method. Section 7 brings all of these elements together in the calculation of pre-shape geodesics via energy minimization. Energy density functions are also introduced and several examples of geodesics are given. In Sect. 8, we discretize the model and provide pseudo-code for the algorithms. Experimental results are presented in Sect. 9 and the limit behavior of the homogeneous metrics is discussed in Sect. 10. We close the paper with a summary and discussion of future work.

2 Pre-Shape Space of Closed Curves

We study the shapes of parametric curves $\alpha: \mathbb{S}^1 \rightarrow \mathbb{R}^k$. A point $z \in \mathbb{S}^1$ is often expressed as $z = e^{js}$, $s \in [0, 2\pi]$. For a mapping F defined on \mathbb{S}^1 , we abuse notation and frequently write $F(z) = F(s)$. We also express the velocity vector $\partial_s \alpha(e^{js})$ simply as $\alpha'(s)$. Throughout the paper, we assume that the curves are non-singular; that is, $\alpha'(s) \neq 0$, for every s . Let \mathbb{S}^{k-1} be the unit sphere in \mathbb{R}^k centered at the origin. The simplest generalization of the representation of plane curves used in Mio et al. (2007b) is the log-polar representation of the velocity field of α given by the pair (φ, ν) , where $\varphi: \mathbb{S}^1 \rightarrow \mathbb{R}$ and $\nu: \mathbb{S}^1 \rightarrow \mathbb{S}^{k-1}$ are given by $\varphi(s) = \log \|\alpha'(s)\|$ and $\nu(s) = \alpha'(s)/\|\alpha'(s)\|$. Thus, $\alpha'(s) = e^{\varphi(s)} \nu(s)$. φ is the *modular component* of the velocity in logarithmic scale and ν is the *direction field* or *tangent indicatrix*. This leads us to consider the space M formed by all such pairs equipped with various Riemannian structures to be described below. Note that a pair (φ, ν) determines a parametric curve α up to translations. The curve is given by

$$\alpha(s) = x_0 + \int_0^s e^{\varphi(\zeta)} \nu(\zeta) d\zeta, \quad (3)$$

where $x_0 \in \mathbb{R}^k$ is arbitrary.

It is often convenient to view M as a subspace of the vector space N of pairs of mappings (φ, ν) , $\varphi: \mathbb{S}^1 \rightarrow \mathbb{R}$ and $\nu: \mathbb{S}^1 \rightarrow \mathbb{R}^k$. The difference between M and N is that, in the space N , ν is not restricted to take values in \mathbb{S}^{k-1} . Infinitesimal variations of (φ, ν) in N (or equivalently, tangent vectors to N at (φ, ν)) are given by pairs (h, w) , where $h: \mathbb{S}^1 \rightarrow \mathbb{R}$ and $w: \mathbb{S}^1 \rightarrow \mathbb{R}^k$. Tangent vectors to M at $(\varphi, \nu) \in M$ are those that satisfy the constraint $w(s) \cdot \nu(s) = 0$, for every $s \in \mathbb{S}^1$, which ensures that $w(s)$ be tangent to \mathbb{S}^{k-1} at $\nu(s)$.

The representation of a parametric curve α via the pair (φ, ν) is clearly invariant under translations since it is based

on the velocity field of α . Invariance under scale will be achieved by fixing the length to be, say, 2π . Since $\|\alpha'(s)\| = e^{\varphi(s)}$, this condition may be expressed as $\int_0^{2\pi} e^{\varphi(s)} ds = 2\pi$. Moreover, as $\alpha'(s) = e^{\varphi(s)}v(s)$, the pair $(\varphi, v) \in M$ represents a closed curve if and only if $\int_0^{2\pi} v(s)e^{\varphi(s)} ds = 0$. Thus, our interest is in the submanifold P of M consisting of all pairs satisfying these two constraints, which we refer to as the *pre-shape space* of closed curves. Thus, we have the following hierarchy of spaces:

$$P \subset M \subset N, \tag{4}$$

where P is the space of primary interest, while M and N will facilitate the geometric analysis of P . More formally, let $\ell : N \rightarrow \mathbb{R}$ and $\delta : N \rightarrow \mathbb{R}^k$ be the *length* and *displacement* functionals defined as

$$\begin{aligned} \ell(\varphi, v) &= \int_0^{2\pi} e^{\varphi(s)} ds \quad \text{and} \\ \delta(\varphi, v) &= \int_0^{2\pi} v(s)e^{\varphi(s)} ds, \end{aligned} \tag{5}$$

respectively. Then, the pre-shape space may be expressed as

$$P = M \cap \ell^{-1}(2\pi) \cap \delta^{-1}(0). \tag{6}$$

Now, we introduce Riemannian structures on N that generalize (2) to curves in Euclidean space of any dimension. Let $a, b > 0$ be parameters representing the tension and rigidity coefficients of the strings. Define

$$\begin{aligned} &\langle (h_1, w_1), (h_2, w_2) \rangle_{(\varphi, v)} \\ &= a \int_0^{2\pi} h_1(s)h_2(s)e^{\varphi(s)} ds \\ &\quad + b \int_0^{2\pi} (w_1(s) \cdot w_2(s))e^{\varphi(s)} ds. \end{aligned} \tag{7}$$

The elasticity coefficients a and b have been omitted from the left-hand side of (7) to avoid cumbersome notation.

Given pre-shapes $p_i = (\varphi_i, v_i)$, $i = 0, 1$, we define the *pre-shape distance* $d(p_0, p_1)$ as the geodesic distance in P with respect to the Riemannian structure on P induced by (7). If we multiply a and b by a common factor, we simply scale the metric globally. Hence, we often assume that $a + b = 1$ and essentially have a 1-parameter family of pre-shape metrics.

3 Shape Spaces

The pre-shape representation of a curve, introduced in Sect. 2, is invariant under translations and the scale has been fixed by normalizing the lengths of all curves to be 2π . However, the representation is sensitive to reparametrizations and rotations. In this section, we first examine the effect of these transformations on (φ, v) .

3.1 The Right Action of the Diffeomorphism Group

If $\alpha : \mathbb{S}^1 \rightarrow \mathbb{R}^k$ is a curve and $\gamma : \mathbb{S}^1 \rightarrow \mathbb{S}^1$ is a diffeomorphism, the reparametrization of α by γ is the curve α_γ given by $\alpha_\gamma(s) = \alpha(\gamma(s))$. The velocity of α_γ at s is $\alpha'_\gamma(s) = \alpha'(\gamma(s))\|\partial_s \gamma(s)\|$. Thus, if (φ, v) is the pair that represents α , the curve α_γ is represented by $(\varphi \circ \gamma + \log \|\partial_s \gamma\|, v \circ \gamma)$, where \circ denotes composition of mappings. This suggests that we define a right action of the diffeomorphism group \mathcal{D} of the circle on N by

$$(\varphi, v) \cdot \gamma = (\varphi \circ \gamma + \log \|\partial_s \gamma\|, v \circ \gamma). \tag{8}$$

A simple change-of-variables argument shows that the diffeomorphism group \mathcal{D} acts by isometries on N , so that

$$d((\varphi_0, v_0), (\varphi_1, v_1)) = d((\varphi_0, v_0) \cdot \gamma, (\varphi_1, v_1) \cdot \gamma), \tag{9}$$

for any $\gamma \in \mathcal{D}$. The invariance of d under γ implies that the distance does not depend on the particular parametrizations chosen, only on the point correspondences they induce since parametrizations that induce the same correspondences are those that differ by the action of some γ . Note that the action is also compatible with the hierarchy (4) of spaces in the sense that it preserves the subspaces P and M . In other words, if $(\varphi, v) \in P$ or M , then $(\varphi, v) \cdot \gamma \in P$ or M , respectively.

3.2 The Left Action of the Rotation Group

The action of the rotation group $SO(k)$ on \mathbb{R}^k induces a left-action of $SO(k)$ on N by isometries, as follows:

$$U \cdot (\varphi, v) = (\varphi, Uv), \tag{10}$$

where $U \in SO(k)$ and $(\varphi, v) \in N$. The action is trivial on the modular component φ since rotations do not affect the speed of a parametric curve. In some applications, one may wish to consider the action of the full orthogonal group $O(k)$ in order to include orientation-reversing orthogonal transformations such as reflections. The $SO(k)$ action is also compatible with the hierarchy (4).

3.3 Geodesic Shape Spaces

We define the shape space \mathcal{S} of closed curves as the (double) quotient space of P under the actions of $SO(k)$ and \mathcal{D} ; that is,

$$\mathcal{S} = SO(k) \backslash P / \mathcal{D}. \tag{11}$$

The orbit of a pre-shape (φ, v) under this two-sided action is given by

$$\begin{aligned} \mathcal{O}(\varphi, \nu) &= \{U \cdot (\varphi, \nu) \cdot \gamma \mid U \in SO(k) \text{ and } \gamma \in \mathcal{D}\} \\ &= \{(\varphi \circ \gamma + \log \|\partial_s \gamma\|, U\nu \circ \gamma) \mid U \in SO(k), \gamma \in \mathcal{D}\} \end{aligned} \tag{12}$$

and each element of \mathcal{S} can be identified with an orbit in P . If $s_0, s_1 \in \mathcal{S}$, we define the geodesic shape distance by

$$d(s_0, s_1) = \inf_{\substack{(\varphi_0, \nu_0) \\ (\varphi_1, \nu_1)}} d((\varphi_0, \nu_0), (\varphi_1, \nu_1)), \tag{13}$$

where $(\varphi_i, \nu_i), i = 0, 1$, range over the orbit associated with s_i . Thus, for each choice of a and b , we have constructed a geodesic shape space of closed homogeneous elastic curves, which we refer to as the HEM(a, b) model. Since both $SO(k)$ and \mathcal{D} act on P by isometries, to calculate the distance, one may fix any pre-shape (φ_0, ν_0) representing s_0 and take the infimum only over the orbit of a representative (φ_1, ν_1) ; that is,

$$d(s_0, s_1) = \inf_{\substack{\gamma \in \mathcal{D} \\ U \in SO(k)}} d((\varphi_0, \nu_0), U \cdot (\varphi_1, \nu_1) \cdot \gamma). \tag{14}$$

4 Correspondences and Parametrizations

In this section, we address the elastic registration problem for curves in \mathbb{R}^k . The techniques can be used as a general tool for establishing point correspondences for a given family of curves and also for the estimation of a diffeomorphism γ and an orthogonal transformation U that will allow us to approximate the shape distance defined in (14). To obtain accurate estimates of the shape distance, the registration criterion should be as compatible as possible with the geodesic metrics in pre-shape space and not just seek to match some geometric features of the curves.

Curve matching based on velocity fields or curvature functions have been investigated by many authors, primarily in the context of plane curves (see e.g. Cohen et al. 1992; Geiger et al. 1995; Tagare 1999; Tagare et al. 2002; Sebastian et al. 2003; Zheng et al. 2005; Mio et al. 2007b). We utilize a variant of these models, which take the elasticity coefficients into account and can be applied to curves in arbitrary Euclidean spaces. We base correspondences on elastic alignment of velocity fields, as a first-order representation is more compatible with the pre-shape metrics proposed in this paper. Since velocity fields are not invariant under rotations of a curve, we will have to optimize alignments over rotations as well. This is similar to the sensitivity of the proposed pre-shape metrics to rotational alignment, as discussed in Sect. 3.3.

4.1 Elastic Correspondence

We begin with two curves C_1 and C_2 viewed as submanifolds of \mathbb{R}^k . We use parametrizations $\alpha, \beta: \mathbb{S}^1 \rightarrow \mathbb{R}^k$ of C_1 and C_2 , respectively, to describe a correspondence. For each s , $\alpha(s)$ and $\beta(s)$ are to be viewed as matching points. Let $\gamma: \mathbb{S}^1 \rightarrow \mathbb{S}^1$ be an orientation-preserving diffeomorphism. Recall that the reparametrization of β by γ is the curve $s \mapsto \beta(\gamma(s))$, which is denoted β_γ . Clearly, for any diffeomorphism γ_1 , the correspondence between C_1 and C_2 established by α and β is the same as that given by the pair α_{γ_1} and β_{γ_1} . Thus, one may fix a parametrization of C_1 and only consider reparametrizations of C_2 . Throughout our discussion of curve registration, α and β will be constant speed parametrizations and we consider correspondences induced by α and β_γ . Let $v(s) = \alpha'(s)/\|\alpha'(s)\|$ and $w(s) = \beta'(s)/\|\beta'(s)\|$ be the direction fields of the curves C_1 and C_2 . For an orientation-preserving diffeomorphism γ , consider the functional

$$\begin{aligned} G_1(\gamma; \alpha, \beta) &= a \int_0^{2\pi} \log^2 \|\gamma'(s)\| ds \\ &\quad + b \int_0^{2\pi} \|v(s) - w(\gamma(s))\|^2 \sqrt{1 + \|\gamma'(s)\|^2} ds, \end{aligned} \tag{15}$$

with $a, b > 0$. The first term quantifies the stretching energy associated with γ on a logarithmic scale. The second term measures the discrepancy of the velocity fields of α and β_γ , which is calculated with respect to the arc-length element $d\tau = \sqrt{1 + \|\gamma'(s)\|^2} ds$ of the graph $\Gamma = \{(z, \gamma(z)) : z \in \mathbb{S}^1\}$ of γ . The cost function G_1 can be re-expressed more symmetrically as

$$\begin{aligned} G_1(\gamma; \alpha, \beta) &= a \int_0^{2\pi} \log^2 \|\gamma'(s)\| ds \\ &\quad + b \int_\Gamma \|v(p_1(\tau)) - w(p_2(\tau))\|^2 d\tau, \end{aligned} \tag{16}$$

where p_1 and p_2 are the projections onto the first and second coordinates, respectively. G_1 is symmetric in the sense that $G_1(\gamma; \alpha, \beta) = G_1(\gamma^{-1}; \beta, \alpha)$. For curve registration, we seek $\hat{\gamma}$ that minimizes G_1 . In implementations, we approximate diffeomorphisms with piecewise linear homeomorphisms. Computational strategies to estimate $\hat{\gamma}$ for energy functionals such as G_1 using dynamic programming (DP) were studied in Sebastian et al. (2003) for both open and closed curves. As velocity fields are not invariant under orthogonal transformations, we will combine this elastic alignment procedure with an optimization over orthogonal maps or rotations.

4.2 Rotational Alignment

We now discuss the optimal rotational alignment of the direction fields of α and β_γ . If $U \in SO(k)$, the direction field of the rotated curve $U \circ \beta_\gamma$ is $s \mapsto Uw(\gamma(s))$. The goal is to minimize

$$G_2(U; \alpha, \beta, \gamma) = \int_\Gamma \|v(p_1(\tau)) - Uw(p_2(\tau))\|^2 d\tau. \quad (17)$$

To include orientation reversing mappings, the minimization should be carried out over the full orthogonal group $O(k)$. This optimization problem is similar to that encountered in Procrustes alignment of shapes (Kendall 1984). We give an explicit description of the solution in the discrete case, but the continuous version is similar (Mio et al. 2007a). Let $\zeta : [0, 2\pi] \rightarrow \Gamma$ be a constant speed parametrization of Γ . Sample the fields $v(p_1(\zeta(s)))$ and $w(p_2(\zeta(s)))$ at n uniformly spaced point $s_1, \dots, s_n \in [0, 2\pi]$ to obtain unit vectors $v_j, w_j, 1 \leq j \leq n$. Let V and W be the $k \times n$ matrices whose columns are v_j and w_j . If $VW^T = V_1 \Sigma V_2^T$ is a singular value decomposition of VW^T , the solution of (17) over $O(k)$ is given by $\hat{U} = V_1 V_2^T$. To optimize only over rotations, one may use the following variant. If $\det(V_1 V_2^T) < 0$, let the least eigenvalue of the nonnegative diagonal matrix Σ occur at the j th column. Then, we change the sign of the j th column of V_1 and proceed as before.

4.3 Full Correspondence

The full alignment problem takes into account both the action of orthogonal transformations and reparametrizations. Starting with constant speed parametrizations α and β of C_1 and C_2 , the goal is to minimize the energy functional

$$G(\gamma, U; \alpha, \beta) = a \int_0^{2\pi} \|\log \gamma'(s)\|^2 ds + b \int_\Gamma \|v(p_1(\tau)) - Uw(p_2(\tau))\|^2 d\tau, \quad (18)$$

over the pair (γ, U) . Starting with the variable U , we approach this optimization problem alternating over U and γ . We initialize the search with the diffeomorphism γ as the identity map, so that $\|\gamma'(s)\| = 1, \forall s$. Note that, for a fixed

γ , the minimization of G over U reduces to the minimization of the functional G_2 of Sect. 4.2, which has a closed-form solution. The γ -step was discussed in Sect. 4.1.

If $(\hat{U}, \hat{\gamma})$ represents an optimal pair, then we use the correspondence induced by $\alpha(s)$ and $\beta^*(s) = \hat{U} \circ \beta(\hat{\gamma}(s))$ to compare the shapes of C_1 and C_2 . As observed earlier, for any γ , the point correspondence between C_1 and C_2 induced by α and β^* is identical to that induced by α_γ and β_γ^* , for any diffeomorphism γ . However, the discrete representation of β^* may lead to regions where C_2 may be highly under-sampled or over-sampled, due to the distortions introduced by $\hat{\gamma}$, while the sampling of C_1 is uniform. In implementations, this may produce noticeable errors, so that it is desirable to reparameterize both curves to distribute the distortions more equitably. This is discussed next.

4.4 Balanced Parametrizations

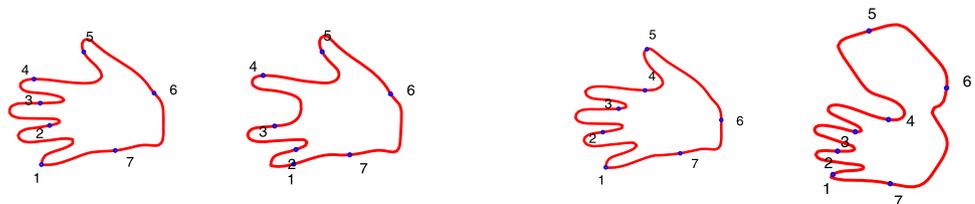
Let $\alpha_j : \mathbb{S}^1 \rightarrow \mathbb{R}^k, 1 \leq j \leq n$, be constant-speed parametrizations of a family of curves C_j , and let $\gamma_j : \mathbb{S}^1 \rightarrow \mathbb{S}^1$ be diffeomorphisms. Consider the parametrization of C_j given by $\alpha_{\gamma_j}(s) = \alpha_j(\gamma_j(s))$. Suppose we sample α_{γ_j} using a uniform grid on the interval $[0, 2\pi]$. Then, C_j will be under-sampled where stretching occurs (near points where $\|\gamma_j'(s)\| > 1$) and over-sampled near compression points. Our goal is to reparameterize all curves so that the correspondences induced by α_{γ_j} are maintained and the sampling distortion becomes more uniformly distributed. The problem posed in Sect. 4.3 is the special case where $n = 2, \gamma_1$ is the identity map, and $\gamma_2 = \hat{\gamma}$.

Let $\Lambda = \{(\gamma_1(s), \dots, \gamma_n(s)), s \in [0, 2\pi]\} \subset \mathbb{S}^1 \times \dots \times \mathbb{S}^1$ and $\lambda : \mathbb{S}^1 \rightarrow \Lambda$ be a constant-speed parametrization of Λ . To make the choice of λ unique, we impose the condition $\lambda(0) = (\gamma_1(0), \dots, \gamma_n(0))$. Let $p_j : \Lambda \rightarrow \mathbb{S}^1$ be the projection onto the j th coordinate. Define $\lambda_j : \mathbb{S}^1 \rightarrow \mathbb{S}^1$ by $\lambda_j = p_j \circ \lambda$. Then, one can verify that the parametrizations $\alpha_{\gamma_j \circ \lambda_j}, 1 \leq j \leq n$, preserve the correspondences and yield more uniform samplings upon discretization. The implementation details are similar to those in Sebastian et al. (2003), Mio et al. (2007b).

4.5 Examples of Shape Correspondences

Figure 1 shows two examples of shape correspondences obtained with the method just described. For each pair, dynamic programming was used to align the velocity fields and

Fig. 1 Matching pairs of plane shapes from the LEMS database



then balanced parametrizations were extracted as described in Sect. 4.4. The shapes used are from the LEMS database. In each case, 200 points were used to represent the contours and correspondences were obtained with elasticity coefficients $a = 0.1$ and $b = 0.9$. Some corresponding points are highlighted for visualization purposes. Even in the second example, where the thumb is subject to a severe deformation, the use of a low tension coefficient allows us to obtain fairly natural correspondences on other parts of the contour. A similar example for curves in 3D space is shown in Fig. 2; the curves were extracted from the contour surface of a cow. Figure 3 shows the results of the shape matching technique applied to two families of shapes. The leaf data used is from the Swedish Museum of Natural History and the jets from the LEMS database. In each case, all curves were aligned to the first using dynamic programming and then a balanced set of parametrizations was constructed for the entire set.

5 Path Spaces and the Energy Functional

Given pre-shapes $(\Phi_0, \mathcal{V}_0), (\Phi_1, \mathcal{V}_1) \in P$, our next goal is to construct a geodesic connecting them in P . As remarked

in the Introduction, it is well established that energy minimization leads to more efficient algorithms, so this is the strategy adopted. The idea is to begin with a path in P from (Φ_0, \mathcal{V}_0) to (Φ_1, \mathcal{V}_1) and gradually deform it to a geodesic following the negative (Riemannian) gradient flow of the energy. Implicit in this statement is the assumption that a path space, where the energy is defined, equipped with a Riemannian structure has been constructed. We now introduce the various path spaces that arise in our calculation of geodesics.

5.1 Path Spaces

Let $I = [0, 1]$. A path in M is given by a pair of mappings $\varphi: \mathbb{S}^1 \times I \rightarrow \mathbb{R}$ and $v: \mathbb{S}^1 \times I \rightarrow \mathbb{S}^{k-1}$. We denote the path by $(\varphi_t, v_t), t \in I$, where $\varphi_t(s) = \varphi(s, t)$ and $v_t(s) = v(s, t)$. We think of t as a time (or deformation) parameter and s as the curve parameter. Let Y be the space of all continuous paths in M (with square integrable derivative). A tangent vector to Y at (φ_t, v_t) , which can be interpreted as the “direction” of an infinitesimal deformation of (φ_t, v_t) , is represented by a pair (h_t, w_t) , with $h_t: \mathbb{S}^1 \rightarrow \mathbb{R}$ and $w_t: \mathbb{S}^1 \rightarrow \mathbb{R}^k$, with the property that $w_t(s) \cdot v_t(s) = 0$, for each t and s . This last condition ensures that $w_t(s)$ is tangent to the unit sphere \mathbb{S}^{k-1} at $v_t(s)$. Define a Riemannian structure on Y by the inner product

$$\begin{aligned} & \langle (h_t^1, w_t^1), (h_t^2, w_t^2) \rangle_{(\varphi_t, v_t)} \\ &= \langle (h_0^1, w_0^1), (h_0^2, w_0^2) \rangle_{(\varphi_0, v_0)} \\ &+ \int_0^1 \langle D_t(h_t^1, w_t^1), D_t(h_t^2, w_t^2) \rangle_{(\varphi_t, v_t)} dt, \end{aligned} \tag{19}$$



Fig. 2 Matching the curves highlighted on the contour surface of a cow. Color-coded correspondences are shown on the right

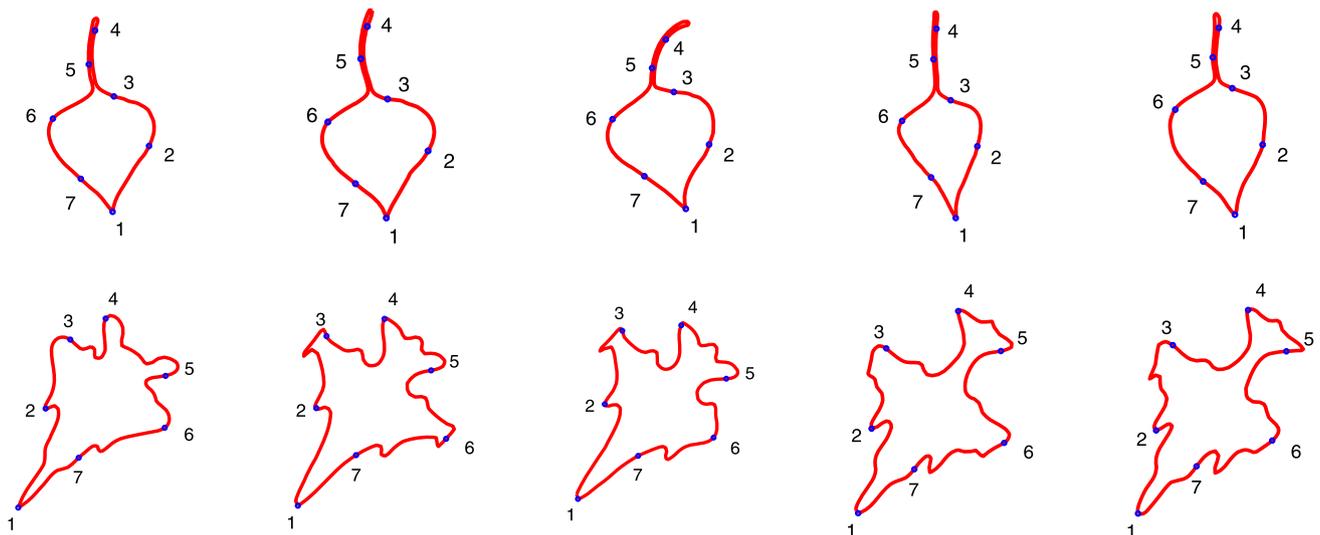


Fig. 3 Row 1 shows the alignment of the silhouettes of 5 leaves with a balanced set of parametrizations. Row 2 displays the result of a similar experiment with the contours of 5 jets

where D_t denotes covariant differentiation in M along the path. This type of metric was introduced by Palais (1963). In computer vision, variants of the metric have been used in curve evolution based on gradient methods by Mio et al. (2004) for the calculation of elastica, by Sundaramoorthi et al. (2007) in the study of Sobolev active contours, and by Klassen and Srivastava (2006) in shape analysis. We shall consider the following submanifolds of Y :

- (i) the space $Z_M \subset Y$ of paths in M satisfying the boundary conditions $(\varphi_0, \nu_0) = (\Phi_0, \mathcal{V}_0)$ and $(\varphi_1, \nu_1) = (\Phi_1, \mathcal{V}_1)$;
- (ii) the space $Z_P \subset Z_M$ of paths in the pre-shape space P satisfying the boundary conditions described in (i); a path in Z_P has the property that each $(\varphi_t, \nu_t), t \in I$, also satisfies the length and closure constraints that define pre-shapes.

We thus have the following nested sequence of path spaces: $Z_P \subset Z_M \subset Y$. Our goal is to find a minimal energy path in Z_P , but we exploit this hierarchy in the development of our algorithms.

5.2 The Energy of a Path

On the path space Y , define the energy functional $E: Y \rightarrow \mathbb{R}$ by

$$E(\varphi_t, \nu_t) = \frac{1}{2} \int_0^1 \langle (\partial_t \varphi_t, \partial_t \nu_t), (\partial_t \varphi_t, \partial_t \nu_t) \rangle_{(\varphi_t, \nu_t)} dt. \quad (20)$$

A pre-shape geodesic between $(\Phi_0, \mathcal{V}_0), (\Phi_1, \mathcal{V}_1) \in P$ is a path $(\varphi_t, \nu_t) \in Z_P$, which is a critical point of the energy E restricted to Z_P . We are particularly interested in paths of minimal energy since they represent minimal length geodesics.

6 Pre-Shape Projection

Our energy minimization algorithm will rely on a projection Π that maps a pair $(\varphi, \nu) \in M$ (near P) to the nearest pre-shape in P . If α is a curve associated with (φ, ν) , as described in (3), Π turns α into a closed curve and normalizes its length to be 2π . In the construction of geodesics, we shall use Π to map a path in M obtained as an infinitesimal deformation of a path in P back to P . Thus, our main interest is in the projection of pairs (φ, ν) that lie in a small vicinity of P in M . We employ a Riemannian version of Newton’s method to search for the pre-shape closest to (φ, ν) .

A pair $(\varphi, \nu) \in M$ represents a pre-shape if and only if $\ell(\varphi, \nu) = 2\pi$ and $\delta(\varphi, \nu) = 0 \in \mathbb{R}^k$, where ℓ and $\delta = (\delta_1, \dots, \delta_k)$ are the length and displacement functionals defined in (5). Consider the $k + 1$ residual functions

$$\rho_0(\varphi, \nu) = 2\pi - \ell(\varphi, \nu) \quad \text{and} \quad \rho_j(\varphi, \nu) = -\delta_j(\varphi, \nu), \quad (21)$$

$1 \leq j \leq k$, whose simultaneous vanishing is equivalent to (φ, ν) being a pre-shape. Letting

$$H(\varphi, \nu) = \frac{1}{2} \rho_0^2(\varphi, \nu) + \frac{1}{2} \sum_{j=1}^k \rho_j^2(\varphi, \nu), \quad (22)$$

the pair (φ, ν) is a pre-shape if and only if $H(\varphi, \nu) = 0$. We employ Newton’s method on the manifold M to find the nearest zero of H , thereby projecting (φ, ν) onto P . This requires the computation of the gradient of H relative to the Riemannian metric determined by the elasticity coefficients a, b . As shown in Appendix A, the gradient of H at (φ, ν) , as a functional on N , is the mapping $\nabla_N H: \mathbb{S}^1 \rightarrow \mathbb{R} \times \mathbb{R}^k$ given by

$$\nabla_N H(s) = -\rho_0(\varphi, \nu) \nabla \ell(s) - \sum_{j=1}^k \rho_j(\varphi, \nu) \nabla \delta_j(s), \quad (23)$$

where

$$\nabla \ell(s) = \left(\frac{1}{a}, 0 \right) \quad \text{and} \quad \nabla \delta_j(s) = \left(\frac{v_j(s)}{a}, \frac{e_j(s)}{b} \right). \quad (24)$$

Here, $e_j: \mathbb{S}^1 \rightarrow \mathbb{R}^k$ is the constant function $e_j(s) = e_j$, where $\{e_1, \dots, e_k\}$ is the canonical basis of \mathbb{R}^k , 0 denotes the constant function $0 \in \mathbb{R}^k$, and the unit vector $\nu(s)$ is written as $\nu(s) = (\nu_1(s), \dots, \nu_k(s))$. We have dropped (φ, ν) from the notation of gradients to keep the expressions more manageable. The gradient $\nabla_M H$ of H as a functional on M can be obtained by orthogonally projecting $\nabla_N H$ onto the tangent space to M at (φ, ν) . If we write the modular and directional components as

$$\begin{aligned} \nabla_N H(s) &= (\nabla_N^\varphi H(s), \nabla_N^\nu H(s)) \quad \text{and} \\ \nabla_M H(s) &= (\nabla_M^\varphi H(s), \nabla_M^\nu H(s)), \end{aligned} \quad (25)$$

the projection amounts to making $\nabla_N^\nu H(s)$ orthogonal to $\nu(s)$ in \mathbb{R}^k , for each s . Therefore,

$$\begin{aligned} &(\nabla_M^\varphi H(s), \nabla_M^\nu H(s)) \\ &= (\nabla_N^\varphi H(s), \nabla_N^\nu H(s) - [\nabla_N^\nu H(s) \cdot \nu(s)]\nu(s)). \end{aligned} \quad (26)$$

In each step of Newton’s method, the geodesic update of the directional component is performed along great circles in \mathbb{S}^{k-1} to ensure that each $\nu(s)$ remains a unit vector. Letting

$$\epsilon(\varphi, \nu) = \frac{H(\varphi, \nu)}{\langle \nabla_M H, \nabla_M H \rangle_{(\varphi, \nu)}} \quad \text{and} \quad g(s) = \nabla_M^\nu H(s), \quad (27)$$

the (spherical) update rule is

$$\begin{cases} \varphi_{\text{new}}(s) = \varphi(s) - \epsilon \nabla_M^\varphi H(s) \\ \nu_{\text{new}}(s) = \cos(\epsilon \|g(s)\|) \nu(s) - \sin(\epsilon \|g(s)\|) \frac{g(s)}{\|g(s)\|}, \end{cases} \quad (28)$$

if $g(s) \neq 0$. Otherwise, $\nu_{\text{new}}(s) = \nu(s)$.

7 Geodesics in Pre-Shape Space

To calculate geodesics in pre-shape space, we resort to a gradient search for paths of minimal energy.

7.1 Initialization

To initialize the energy minimization process, we use the following construction. Let $\alpha_0, \alpha_1: \mathbb{S}^1 \rightarrow \mathbb{R}^k$ be paths associated with the given pre-shapes, as described in (3). We linearly interpolate the curves to obtain a 1-parameter family of curves $\alpha_t, 0 \leq t \leq 1$. If the velocity fields of these curves vanish at some points, we gently deform the family to make all curves non-singular and then scale each α_t to have length 2π . The pre-shapes (φ_t, ν_t) associated with $\alpha_t, t \in I$, yield the desired path.

7.2 Covariant Integration and Parallel Transport

In order to describe the calculation of the gradient of the energy E at (φ_t, ν_t) , we first discuss covariant integration in M of vector fields along the path (φ_t, ν_t) . Let (f_t, x_t) be a vector field along a path $(\varphi_t, \nu_t) \in Y$, which is tangential to M for each $t \in I$. This means that $x_t(s) \cdot \nu_t(s) = 0$ is satisfied for every s, t . As shown in Appendix B, a vector field (F_t, X_t) along (φ_t, ν_t) is tangential to M and represents a *covariant integral* of (f_t, x_t) if and only if it satisfies the system of differential equations

$$\begin{cases} \partial_t F_t(s) = f_t(s) - \frac{1}{2} \partial_t \varphi_t(s) F_t(s) \\ \quad + \frac{1}{2} \frac{b}{a} [X_t(s) \cdot \partial_t \nu_t(s)] \\ \partial_t X_t(s) = x_t(s) - \frac{1}{2} (X_t(s) \partial_t \varphi_t(s) + F_t(s) \partial_t \nu_t(s)) \\ \quad - [X_t(s) \cdot \partial_t \nu_t(s)] \nu_t(s). \end{cases} \tag{29}$$

In the calculation of geodesics, we will integrate this system numerically with prescribed initial conditions (F_0, X_0) . In the special case where the field (f_t, x_t) is identically zero, the integral field (F_t, X_t) is the *parallel transport* of (F_0, X_0) along (φ_t, ν_t) .

7.3 The Gradient of the Energy

Given a path $(\varphi_t, \nu_t) \in Z_P$, we first calculate the gradient of E at (φ_t, ν_t) as a functional on the path space Y . For this purpose, we consider a variation $(\varphi_t(s; \mu), \nu_t(s; \mu))$ of (φ_t, ν_t) in Y along a direction (h_t, w_t) , which is defined on a small interval $-\epsilon < \mu < \epsilon$. This means that $(\varphi_t(s; 0), \nu_t(s; 0)) = (\varphi_t, \nu_t)$, the path $(\varphi_t(_ ; \mu), \nu_t(_ ; \mu)) \in Y$ for each fixed μ ,

$$\begin{aligned} h_t(s) &= \frac{\partial}{\partial \mu} \varphi_t(s; \mu) \Big|_{\mu=0} \quad \text{and} \\ w_t(s) &= \frac{\partial}{\partial \mu} \nu_t(s; \mu) \Big|_{\mu=0}. \end{aligned} \tag{30}$$

Differentiating (20) at $\mu = 0$, we obtain

$$\begin{aligned} dE_{(\varphi_t, \nu_t)}(h_t, w_t) &= \int_0^1 \left\langle D_\mu(\partial_t \varphi_t, \partial_t \nu_t) \Big|_{\mu=0}, (\partial_t \varphi_t, \partial_t \nu_t) \right\rangle_{(\varphi_t, \nu_t)} dt \\ &= \int_0^1 \langle D_t(h_t, w_t), (\partial_t \varphi_t, \partial_t \nu_t) \rangle_{(\varphi_t, \nu_t)} dt. \end{aligned} \tag{31}$$

In (29), we set $f_t = \partial_t \varphi_t$ and $x_t = \partial_t \nu_t$, and integrate the differential equation with initial condition $(F_0, X_0) = (0, 0)$ to get a vector field (F_t^1, X_t^1) along the path (φ_t, ν_t) . Then, using (19), we may rewrite (31) as

$$\begin{aligned} dE_{(\varphi_t, \nu_t)}(h_t, w_t) &= \int_0^1 \left\langle D_t(h_t, w_t), D_t(F_t^1, X_t^1) \right\rangle_{(\varphi_t, \nu_t)} dt \\ &= \langle (h_t, w_t), (F_t^1, X_t^1) \rangle_{(\varphi_t, \nu_t)}. \end{aligned} \tag{32}$$

Thus, the gradient of E at (φ_t, ν_t) as a functional on Y is given by

$$\nabla_Y E(\varphi_t, \nu_t) = (F_t^1, X_t^1). \tag{33}$$

To obtain $\nabla_{Z_M} E(\varphi_t, \nu_t)$, we project $\nabla_Y E(\varphi_t, \nu_t)$ orthogonally onto the tangent space of Z_M with respect to the Palais inner product. Since the space Z_M is obtained from Y by imposing the boundary conditions $(\varphi_0, \nu_0) = (\Phi_0, \mathcal{V}_0)$ and $(\varphi_1, \nu_1) = (\Phi_1, \mathcal{V}_1)$ on a path (φ_t, ν_t) , a tangent vector to Y at (φ_t, ν_t) is tangential to Z_M if and only if it vanishes at $t = 0$ and $t = 1$. Moreover, as explained in Appendix C, the orthogonal complement of the tangent space of Z_M in the tangent space of Y at (φ_t, ν_t) is formed by the covariantly linear fields (that is, fields whose second covariant derivatives vanish) in M along (φ_t, ν_t) . By construction, the field (F_t^1, X_t^1) vanishes at $t = 0$. Thus, to orthogonally project (F_t^1, X_t^1) onto the tangent space of Z_M , we simply need to subtract from (F_t^1, X_t^1) the covariantly linear field that vanishes at $t = 0$ and coincides with (F_t^1, X_t^1) at $t = 1$. Again, we resort to covariant integration, this time applied to the reverse of the path (φ_t, ν_t) . We first construct a parallel field along the reverse path with initial condition (F_1^1, X_1^1) . Reversing the path again, after integration, we obtain a parallel field (G_t, Y_t) along (φ_t, ν_t) whose value at $t = 1$ is (F_1^1, X_1^1) . The field (tG_t, tY_t) is covariantly linear with the desired properties. Therefore, the gradient is given by

$$\nabla_{Z_M} E(\varphi_t, \nu_t) = \nabla_Y E(\varphi_t, \nu_t) - (tG_t, tY_t). \tag{34}$$

Our goal is to minimize E on the path space Z_P . The direct calculation of $\nabla_{Z_P} E(\varphi_t, \nu_t)$ is a possible approach, however, the computation is costly for large values of k . For this reason, we resort to an alternative strategy, which replaces the gradient descent in Z_P with its counterpart in

Z_M followed by the closest-point projection onto Z_P . Experiments with the two strategies for $k = 2$ indicate that the number of iterations needed for the calculation of geodesics is approximately the same and the results are nearly identical. Thus, the proposed strategy is adopted.

7.4 Pre-Shape Geodesics

We now present an algorithm to calculate a pre-shape geodesic from (Φ_0, \mathcal{V}_0) to (Φ_1, \mathcal{V}_1) . Let $\epsilon, \delta > 0$ be small real numbers.

- (i) Initialize the search with a path (φ_t, ν_t) in Z_P , which can be constructed, for example, as described in Sect. 7.1.
- (ii) Let $(f_t, x_t) = (\partial_t \varphi_t, \partial_t \nu_t)$. Using (29), integrate this field covariantly along the path (φ_t, ν_t) with zero initial condition. According to (33), the integral field (F_t, X_t) gives the gradient $\nabla_Y E(\varphi_t, \nu_t)$.
- (iii) Using (29), calculate the parallel transport of (F_1, X_1) along the reverse of the path (φ_t, ν_t) . Reverse the path and the parallel field again to obtain a parallel field (G_t, Y_t) along (φ_t, ν_t) . By (34), the Z_M -gradient of E is given by $\nabla_{Z_M} E(\varphi_t, \nu_t) = \nabla_Y E(\varphi_t, \nu_t) - (tG_t, tY_t)$.
- (iv) Write the modular and directional components of the gradient as $\nabla_{Z_M} E = (\nabla_{Z_M}^\varphi E, \nabla_{Z_M}^\nu E)$. Update (φ_t, ν_t) as a path in Z_M according to

$$\begin{aligned} \varphi_t^* &= \varphi_t - \epsilon \nabla_{Z_M}^\varphi E(\varphi_t, \nu_t); \\ \nu_t^*(s) &= \nu_t(s), \quad \text{if } \nabla_{Z_M}^\nu E(\varphi_t, \nu_t)(s) = 0; \\ \nu_t(s) &= \cos(\epsilon \eta(t, s)) \nu_t(s) - \sin(\epsilon \eta(t, s)) A(t, s), \\ &\text{otherwise.} \end{aligned} \tag{35}$$

Here, $\eta(t, s) = \|\nabla_{Z_M}^\nu E(\varphi_t, \nu_t)(s)\|$ and $A(t, s) = \nabla_{Z_M}^\nu E(\varphi_t, \nu_t)(s) / \eta(t, s)$, where the norm is Euclidean. Note that the update of ν_t takes place along great circles ensuring that ν_t^* is a unit vector.

- (v) Project each (φ_t^*, ν_t^*) onto the pre-shape space P , as discussed in Sect. 6, to obtain a path $(\varphi_t, \nu_t)_{\text{new}}$ in P .
- (vi) Iterate the process until $\|(\varphi_t, \nu_t)_{\text{new}} - (\varphi_t, \nu_t)\|_{(\varphi_t, \nu_t)} < \delta$; that is, until the update is δ -small as measured by the Palais norm.

7.5 Energy Density

A (parametric) pre-shape geodesic $(\varphi_t(s), \nu_t(s))$, $0 \leq s \leq 2\pi$, $0 \leq t \leq 1$, is traversed with constant speed ω , where ω is the length of the geodesic. The energy of the path is given by

$$\begin{aligned} E &= \int_0^{2\pi} \int_0^1 [a(\partial_t \varphi(s; t))^2 + b\|\partial_t \nu(s; t)\|^2] e^{\varphi_t(s)} dt ds \\ &= \omega^2. \end{aligned} \tag{36}$$

Thus, we define the energy density function $\rho: \mathbb{S}^1 \rightarrow \mathbb{R}$ by

$$\rho(s) = \frac{1}{\omega^2} \int_0^1 (a(\partial_t \varphi(s; t))^2 + b\|\partial_t \nu(s; t)\|^2) e^{\varphi_t(s)} dt, \tag{37}$$

which quantifies the local contribution at s to the total energy of the geodesic path. The density function allows us to trace back the regions that exhibit the highest shape resemblance or dissimilarity as measured by the elastic shape metric. Note that ρ has been normalized so that $\int_0^{2\pi} \rho(s) ds = 1$. If desired, one can further decompose the energy into tension and bending components to quantify separately the local contributions due to stretching (or compression) and bending.

7.6 Examples of Geodesics

Figure 4 shows two examples of geodesics between shapes from the LEMS database, which are calculated with elasticity coefficients $a = 0.1$ and $b = 0.9$ and 200 points on

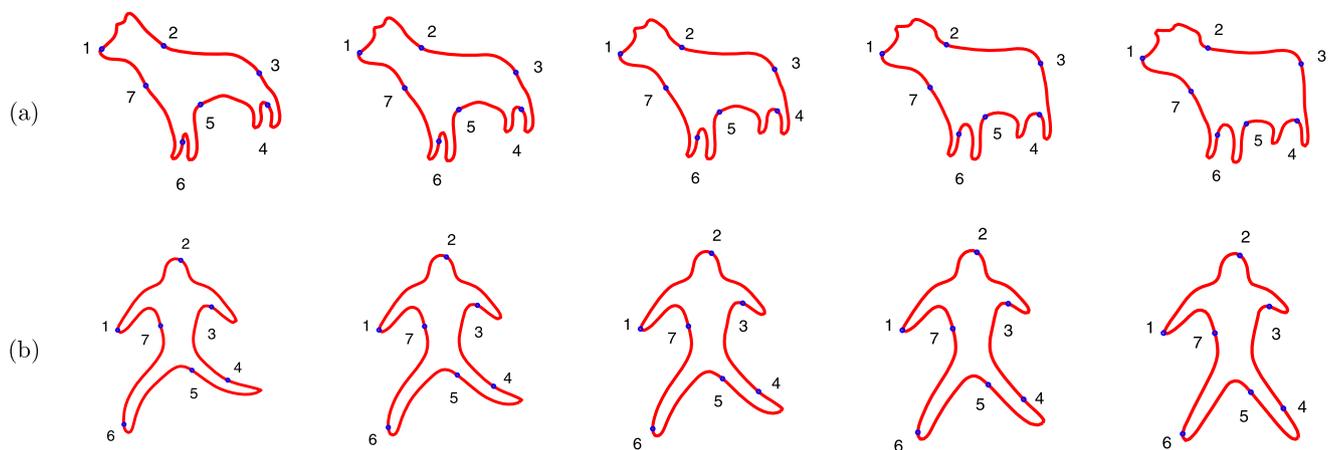


Fig. 4 Examples of geodesic interpolations between shapes from the LEMS database calculated with elasticity coefficients $a = 0.1$ and $b = 0.9$

each contour. On each row, the first and last images depict the given data and the intermediate shapes represent several stages of the geodesic deformation. Figure 5 shows the evolution of the energy during the construction of the geodesics and the energy density functions associated with the geodesics. We divided the interval $[0, 2\pi]$ into seven equal parts, as illustrated in Fig. 4, and calculated the average value of the density function in each of these intervals, which are shown in the bar graphs. The first bin corresponds to the arc 1–2 and continues sequentially to bin 7, which is associated with the arc 7–1. For the geodesic between the dog and the cow, the highest values occur along the arcs 3–4 and 4–5, as intuitively expected since these are the regions where deformation due to bending is most pronounced. For the second geodesic, the peak value occurs along the leg where the most significant bending takes place. Although there is significant shrinking along one of the arms, the density function has low value along the arc 2–3 reflecting the fact that the tension coefficient used is relatively small so that compression can be achieved with low energy expenditure. Similar examples for curves in \mathbb{R}^3 are shown in Fig. 6. The knot data used was obtained from The KnotPlot Site developed by R. Scharein.

8 Computational Model and Algorithms

In the discrete model, we use polygonal representations of curves with $n + 1$ vertices A_0, \dots, A_n , with the regularity assumption that the edges $e_i = A_i - A_{i-1}$, $1 \leq i \leq n$, are

nondegenerate. For closed curves, $A_0 = A_n$. If $r_i = \log \|e_i\|$ and $v_i = e_i^T / \|e_i\|$ are the modular and the (transpose of the) directional components of the i th edge, we let

$$r = \begin{bmatrix} r_1 \\ \vdots \\ r_n \end{bmatrix} \quad \text{and} \quad v = \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} v_{11} & \dots & v_{1k} \\ \vdots & \ddots & \vdots \\ v_{n1} & \dots & v_{nk} \end{bmatrix}. \quad (38)$$

The pair (r, v) gives a discrete analogue of $(\varphi, \nu) \in M$. More generally, pairs (r, v) without the requirement that each v_i have unit Euclidean norm give discrete analogues of elements of N . Similarly, a tangent vector to M or N is represented by a pair (h, w) , where $h \in \mathbb{R}^n$ and w is an $n \times k$ matrix. In this representation, the inner product (7) becomes

$$\begin{aligned} \langle (h, w), (h^*, w^*) \rangle_{(r,v)} \\ = a \sum_{i=1}^n h_i h_i^* e^{r_i} + b \sum_{i=1}^n (w_i \cdot w_i^*) e^{r_i}. \end{aligned} \quad (39)$$

To discretize a path (φ_t, ν_t) , we sample the interval $I = [0, 1]$ at $m + 1$ uniformly distributed points, so that a path becomes a sequence $(r(0), v(0)), \dots, (r(m), v(m))$. The discrete analogues of the length and displacement functionals, introduced in (5), are

$$\ell(r, v) = \sum_{i=1}^n e^{r_i} \quad \text{and} \quad \delta_j(r, v) = \sum_{i=1}^n v_{ij} e^{r_i}, \quad (40)$$

and the residual functions defined in (21) become $\rho_0(r, v) = 2\pi - \ell(r, v)$ and $\rho_j(r, v) = -\delta_j(r, v)$, $1 \leq j \leq k$. The rep-

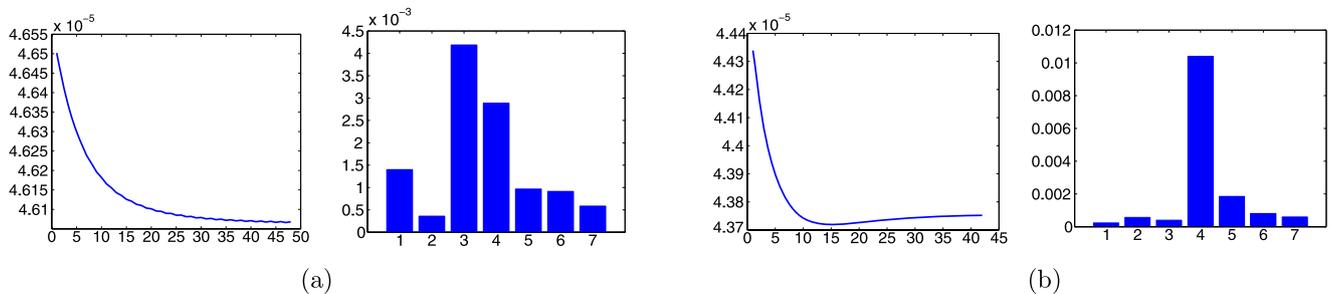
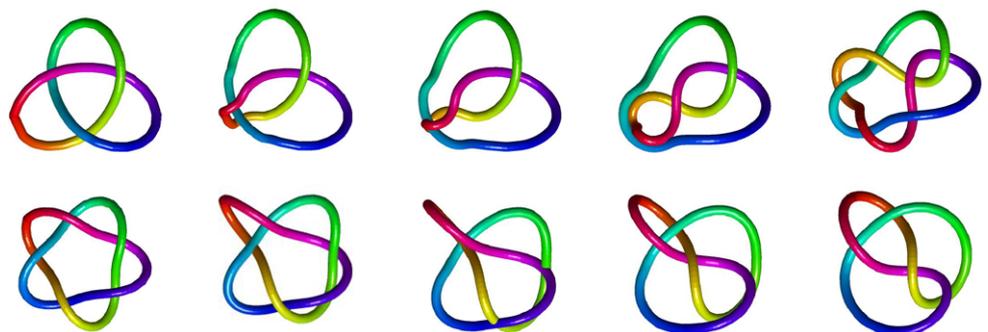


Fig. 5 Plots of the evolution of the energy during the gradient search and the energy density functions for the geodesics in Fig. 4

Fig. 6 Two examples of shape geodesics in 3D space calculated with elasticity coefficients $a = 0.1$ and $b = 0.9$ using data from The KnotPlot Site



resentation of the gradient vectors (24) take the form

$$\begin{aligned} \nabla \ell(r, v) &= \left(\begin{bmatrix} 1/a \\ \vdots \\ 1/a \end{bmatrix}, \begin{bmatrix} 0 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 0 \end{bmatrix} \right) \text{ and} \\ \nabla \delta_j(r, v) &= \left(\begin{bmatrix} v_{1j}/a \\ \vdots \\ v_{nj}/a \end{bmatrix}, \begin{bmatrix} 0 & \dots & 1/b & \dots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \dots & 1/b & \dots & 0 \end{bmatrix} \right), \end{aligned} \tag{41}$$

where the entries $1/b$ appear on the j th column.

8.1 The Projection Algorithm

Let (r, v) be as above and $\delta > 0$ a small real number. The goal is to find the closest pair with the property that $\rho_j(r, v) = 0$, for $0 \leq j \leq k$.

Algorithm 8.1.1 (Pre-Shape Projection)

1. Calculate $H(r, v) = \frac{1}{2} \sum_{j=0}^k \rho_j^2(r, v)$.
2. Using (42), calculate

$$\begin{aligned} \nabla_N H(r, v) &= -\rho_0(r, v) \nabla \ell(r, v) \\ &\quad - \sum_{j=1}^k \rho_j(r, v) \nabla \delta_j(r, v). \end{aligned}$$

Let $\nabla_N^r H(r, v)$ and $\nabla_N^v H(r, v)$ be the modular and directional components of $\nabla_N H(r, v)$, respectively.

3. Set $\nabla_M^r H(r, v) = \nabla_N^r H(r, v)$.
4. Let q_i be the i th row of $\nabla_N^v H(r, v)$. Calculate $q_i^* = q_i - (q_i v_i^T) v_i$ and let $\nabla_M^v H(r, v)$ be the $n \times k$ matrix whose i th row is q_i^* .
5. Set

$$\nabla_M H(r, v) = ((\nabla_M^r H(r, v), \nabla_M^v H(r, v))).$$

6. Let $\epsilon(r, v) = H(r, v) / \langle \nabla_M H(r, v), \nabla_M H(r, v) \rangle_{(r,v)}$.
7. Update (r, v) spherically, as follows:

$$\begin{aligned} r &= r - \epsilon(r, v) \nabla_M^r H(r, v) \quad \text{and} \\ v_i &= \cos(\epsilon(r, v) \|q_i^*\|) v_i - \sin(\epsilon(r, v) \|q_i^*\|) \frac{q_i^*}{\|q_i^*\|}. \end{aligned}$$

8. Iterate until $\|\nabla_M H(r, v)\|_{(r,v)} < \delta$.

8.2 The Covariant Integration Algorithm

Let $(r(t), v(t))$, $t = 0, \dots, m$, be the discrete representation of a path in M and let $(f(t), x(t))$ be a vector field tangential to M along the path. This means that $x_i(t) v_i(t)^T = 0$, for every $0 \leq t \leq m$ and $1 \leq i \leq n$. We now present a first-order algorithm to estimate the covariant integral of $(f(t), x(t))$ with initial conditions (F_0, X_0) .

Algorithm 8.2.1 (Covariant Integration)

1. Set $t = 0$ and let $F(t) = F_0$ and $X(t) = X_0$. While $t < m$, do:
2. For $i = 1, \dots, n$, calculate

$$\begin{aligned} \Delta r_i(t) &= m(r_i(t+1) - r_i(t)) \quad \text{and} \\ \Delta v_i(t) &= m(v_i(t+1) - v_i(t)). \end{aligned}$$

Set

$$\Delta_i F(t) = f_i(t) - \frac{1}{2} \Delta r_i(t) F_i(t) + \frac{b}{2a} X_i(t) (\Delta v_i(t))^T$$

and

$$\begin{aligned} \Delta_i X(t) &= x_i(t) - \frac{1}{2} [X_i(t) (\Delta r_i(t))^T + F_i(t) \Delta v_i(t)] \\ &\quad - [X_i(t) (\Delta v_i(t))^T] v_i(t). \end{aligned}$$

3. Let $\Delta F(t) = [\Delta F_1(t) \dots \Delta F_n(t)]^T \in \mathbb{R}^n$ and $\Delta X(t)$ be the $n \times k$ matrix whose i th row is $\Delta_i X(t)$.
4. Set

$$\begin{aligned} F(t+1) &= F(t) + \frac{1}{m} \Delta F(t) \quad \text{and} \\ X(t+1) &= X(t) + \frac{1}{m} \Delta X(t). \end{aligned}$$

5. Set $t = t + 1$.

8.3 Computation of Geodesics

Our goal is to construct a geodesic path between two discrete pre-shapes (R_0, V_0) and (R_1, V_1) , following the strategy described in Sect. 7.4.

Algorithm 8.3.1 (Pre-Shape Geodesics)

Let (R_0, V_0) and (R_1, V_1) be discrete pre-shapes and $\epsilon, \delta > 0$ small real numbers.

1. Initialize the search with a path $(r(t), v(t))$, $t = 0, \dots, m$, where each $(r(t), v(t))$ is a pre-shape. This can be done, for example, as described in Sect. 7.1.
2. For each $0 \leq t \leq m - 1$, calculate

$$\begin{aligned} \Delta r(t) &= m(r(t+1) - r(t)) \quad \text{and} \\ \Delta v(t) &= m(v(t+1) - v(t)). \end{aligned}$$

3. Set

$$f(t) = \Delta r(t) \quad \text{and} \quad x(t) = \Delta v(t),$$

for $0 \leq t \leq m - 1$, $f(m) = 0$, and $x(m) = 0$. Using Algorithm 8.2, integrate the field $(f(t), x(t))$ covariantly

Fig. 8 Shape preprocessing: (a) original data; (b) segmented contour; (c) smoothed contour

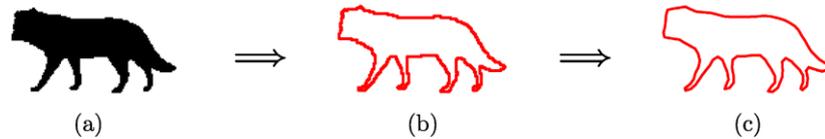


Table 1 Results of shape retrieval experiments with the LEMS-99 database shown in Fig. 7. The results for SC, GM, SE and IDSC-DP were reported in Ling and Jacobs (2007)

Method	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
SC	97	91	88	85	84	77	75	66	56	37
GM	99	97	99	98	96	96	94	83	75	48
SE	99	99	99	98	98	97	96	95	93	82
IDSC-DP	99	99	99	98	98	97	97	98	94	79
HEM(0.25, 0.75)	99	99	99	99	99	97	96	92	80	73
HEM(0.35, 0.65)	99	99	99	99	98	97	97	91	84	75
HEM(0.40, 0.60)	99	99	99	99	99	96	95	94	86	74

Table 2 Results of leaf classification experiments with the HEM metric for two different values of the elasticity coefficients and comparison with the results reported in Söderkvist (2001), Ling and Jacobs (2007)

Method	SO	FO	SC-DP	MDS SC-DP	IDSC-DP	SPTC-DP	HEM (0.15, 0.85)	HEM (0.25, 0.75)
Perf.	82%	89.6%	88.12%	95.33%	94.13%	95.33%	96.67%	96.53%

same elasticity coefficients employed for alignment. Subsequently, using the same parametrizations, we also carried out retrieval experiments with other values of a and b . For each shape, the ten closest shapes were retrieved and ranked 1–10 according to the increasing distance to the query shape. For each i , $1 \leq i \leq 10$, the total number of i th retrievals that fall in the correct category is shown in Table 1. For each choice of a and b , a total of 4,851 geodesic distances were calculated. For comparison, the table also shows results reported in the literature obtained with the following methods: (i) shape context (SC) of Belongie et al. (2002); (ii) shock edit (SE) developed in Sebastian et al. (2004); (iii) the generative model (GM) of Tu and Yuille (2004); (iv) shape context based on the inner-distance implemented with dynamic programming (IDSC-DP) (Ling and Jacobs 2007). Overall, the results obtained are somewhat comparable to the best previously reported for this data set. Note that the HEM metrics tend to perform well at the top retrievals. We argue that this is an important property of a shape metric. In applications involving large data sets, instead of computing the distance to all shapes in a database, it is much more likely that the metric will be used for retrieval in conjunction with classification and labeling based on a training set. For example, one could use the shape metric and the k -nearest neighbor classifier to first identify the class to which a shape belongs and then retrieve from that class based on shape distance. In such more realistic scenarios, metrics for which the top retrievals are correct will exhibit better performance.

9.2 Swedish Foliage

The data set used in this experiment is from a project at Linköping University and the Swedish Museum of Natural History (Söderkvist 2001); it consists of images of 75 leaves from each of 15 different species of Swedish trees for a total of 1,125 samples. We used segmented contours made publicly available by the authors of Ling and Jacobs (2007). Twenty five samples from each species were used as training shapes and the remaining 750 contours were used for testing our shape metrics. As in the previous experiment, the contours were aligned using the algorithm described in Sect. 4 and re-sampled with 100 points. Table 2 shows the performance of the homogeneous elastic metrics in shape recognition using the nearest neighbor classifier and two different choices of elasticity coefficients. For comparison purposes, we include several previously reported results. The label SO refers to the experiments carried out by O. Söderkvist (2001) using a combination of several classical shape descriptors. The results obtained with Fourier descriptors (FO), shape context with dynamic programming (SC-DP), MDS with shape context and dynamic programming (MDS-SC-DP), inner-distance shape context with dynamic programming (IDSC-DP), and shortest path texture context (SPTC-DP) were reported in Ling and Jacobs (2007). Note that SPTC is not a single cue method as it integrates shape and texture features.

10 The Infinite-Tension Limit

For a fixed rigidity coefficient b , we sketch an argument that the limit of the HEM(a, b) models, as the tension coefficient $a \rightarrow \infty$, is the arc-length parametrization model of Klassen and Srivastava (2006), up to a scaling factor. A geodesic deformation in shape space is realized by a geodesic in the pre-shape manifold P that is always orthogonal to the orbits of the action of diffeomorphisms and rotations. We investigate in more detail the orthogonality with respect to the orbit of the diffeomorphism group \mathcal{D} of the circle at (φ, ν) . We first characterize vectors that are tangent to the orbit of \mathcal{D} at (φ, ν) .

The action of a diffeomorphism γ on (φ, ν) is described in (8). Given a 1-parameter family $\gamma(\cdot, \mu)$ of diffeomorphisms, $-\epsilon < \mu < \epsilon$, satisfying $\gamma(s, 0) = s$, let $g(s) = \gamma_\mu(s, 0)$. Differentiating (8) at $\mu = 0$, we see that tangent vectors to the orbit of diffeomorphisms at (φ, ν) are those of the form $(d\varphi \circ g + \partial_s g \cdot \partial_s \gamma, d\nu \circ g)$. In particular, if $\varphi = 0$, tangent vectors to orbits are of the form $(\partial_s g \cdot \partial_s \gamma, d\nu \circ g)$. Note that $\partial_s \gamma$ is simply the counterclockwise unit tangent field $\partial/\partial s$ on \mathbb{S}^1 . If we write, $\partial_s g = B(s)\partial/\partial s$, where $B(s)$ is a scalar field, then $\partial_s g \cdot \partial_s \gamma = B(s)$. Thus, tangent vectors to orbits at $(0, \nu)$ are those of the form $(B, d\nu \circ g)$ and the orthogonality of (h, w) to the orbit of \mathcal{D} may be expressed as

$$\begin{aligned} \langle (h, w), (B, d\nu \circ g) \rangle_{(0, \nu)} \\ = a \int_0^{2\pi} h(s)B(s) ds \\ + b \int_0^{2\pi} w(s) \cdot d\nu(g(s)) ds = 0, \end{aligned} \quad (42)$$

for any g of the form $g(s) = \gamma_\mu(s, 0)$. Equivalently,

$$\int_0^{2\pi} h(s)B(s) ds + \frac{b}{a} \int_0^{2\pi} w(s) \cdot d\nu(g(s)) ds = 0. \quad (43)$$

If b is fixed, in the limit as $a \rightarrow \infty$, we get

$$\int_0^{2\pi} h(s)B(s) ds = 0. \quad (44)$$

Since (44) holds for every B , it follows that $h = 0$. The same can be shown at any (φ, ν) using the fact that diffeomorphisms act on M by isometries. Therefore, in the limit as $a \rightarrow \infty$, orthogonality of a path (φ_t, ν_t) to orbits implies that $\partial_t \varphi_t = 0$. If $\varphi_0(s) = 0$, then $\varphi_t(s) = 0$, for every t . In other words, if the initial curve is parameterized by arc-length, it remains so for every $t \in I$. For a fixed b , although the Riemannian structure becomes singular in the limit, it is non-singular on the subspace of pairs of the form $(0, \nu)$ and coincides with (a multiple of) the standard \mathbb{L}^2 metric. Thus, up to a scaling factor, we obtain the arc-length model at the ∞ -tension limit.

11 Summary and Discussion

On the pre-shape space of closed parametric curves in Euclidean space \mathbb{R}^k , we constructed a family of Riemannian metrics indexed by the elastic tension and rigidity coefficients. Curves were represented by their velocity fields in log-polar coordinates and the elastic metrics were defined so as to account for the stretching and bending properties of the strings. The metrics are invariant under the action of the diffeomorphism group of the parameter space \mathbb{S}^1 and induce metrics on shape space. We studied the geometric properties of the pre-shape manifold and various path spaces of elastic strings to develop a computational model and algorithms to calculate geodesics using energy minimization. The shape metrics are related to the elastic metrics for plane shapes studied in Mio et al. (2007b), but even in that case, the shape representation and the computational strategies are more robust and efficient. We also developed a shape registration method, implemented via dynamic programming, that allows us to estimate shape distances more efficiently and can be used as a general curve registration technique. Although the elastic shape distance is a global measure of shape dissimilarity, energy density functions were introduced to identify the regions where shape similarities and differences tend to be most pronounced. This type of local-global geometric analysis of shapes is important in applications in which one needs to trace back the main sources of morphological differences. Several examples of geodesics were given and experiments were carried out to demonstrate the ability of the elastic metrics to discern and classify shapes, as well as to characterize the relevance of shapes of curves in high dimensional space in pattern recognition. This paper provides a framework and a set of basic tools for modeling the shapes of curves. An important element in our approach is the flexibility to adjust the shape metric to a particular problem to address the problem of context dependence typically encountered in shape analysis. This raises the problem of developing criteria to select the most appropriate elasticity coefficients for a given application. A possible approach, assuming that training data is available, is to use a modified version of the cost function employed in Linear Discriminant Analysis (LDA) to select a shape metric that maximizes the ratio of inter-class to within-class scatter. Other criteria used in machine learning and dimension reduction algorithms can also be adapted to the selection of shape metrics and will be investigated in future work. We also plan to study extensions of the techniques to shapes of surfaces, other compact Riemannian manifolds, and finite simplicial complexes. These higher dimensional analogues of the curve models developed in this paper are needed in several different settings, for example, in computational anatomy to model normal morphological variations and pathological changes in anatomy. We also will further investigate the use

of curves to obtain sparse representations of surfaces and higher dimensional objects for tasks such as object classification and recognition.

Acknowledgements This work was supported in part by the National Science Foundation, grants CCF-0514743 and DMS-0713012. In our experiments, we used data from the LEMS shape database compiled by B. Kimia, the leaf database from a project at Linköping University and the Swedish Museum of Natural History, segmented leaf contours made publicly available by Ling and Jacobs, knot data from The KnotPlot Site developed by R. Scharein, and shape models provided courtesy of MPII by the AIM@Shape Shape Repository.

Appendix A: The Gradient of H

Let $(\varphi(s; \mu), \nu(s; \mu))$, $\mu \in (-\epsilon, \epsilon)$, be a variation of (φ, ν) and let $(h, w) = (\partial_\mu \varphi(s; \mu), \partial_\mu \nu(s; \mu))|_{\mu=0}$ represent a tangent vector. We write the components of w as $w = (w_1, \dots, w_k)$. Differentiating (22) at $\mu = 0$ and using (5), we obtain

$$\begin{aligned}
 dH_{\varphi,\nu}(h, w) &= -\rho_0(\varphi, \nu) \int_0^{2\pi} h(s)e^{\varphi(s)} ds \\
 &\quad - \sum_{j=1}^k \rho_j(\varphi, \nu) \int_0^{2\pi} \nu(s)h(s)e^{\varphi(s)} ds \\
 &\quad - \sum_{j=1}^k \rho_j(\varphi, \nu) \int_0^{2\pi} w_j(s)e^{\varphi(s)} ds \\
 &= -\rho_0(\varphi, \nu) \int_0^{2\pi} a \frac{h(s)}{a} e^{\varphi(s)} ds \\
 &\quad - \sum_{j=1}^k \rho_j(\varphi, \nu) \int_0^{2\pi} a \frac{\nu_j(s)h(s)}{a} e^{\varphi(s)} ds \\
 &\quad - \sum_{j=1}^k \rho_j(\varphi, \nu) \int_0^{2\pi} b \frac{w_j(s)}{b} e^{\varphi(s)} ds \\
 &= -\rho_0(\varphi, \nu) \left\langle \left(\frac{1}{a}, 0 \right), (h, w) \right\rangle_{(\varphi,\nu)} \\
 &\quad - \sum_{j=1}^k \rho_j(\varphi, \nu) \left\langle \left(\frac{\nu_j}{a}, \frac{e_j}{b} \right), (h, w) \right\rangle_{(\varphi,\nu)}. \tag{45}
 \end{aligned}$$

Hence, $\nabla_N H$ can be calculated as claimed in (23).

Appendix B: Covariant Integration

The Riemannian structure on N given by (7) is defined by an integration over the unit circle. For each fixed s , if we set $r = \varphi(s) \in \mathbb{R}$ and $y = (y_1, \dots, y_k) = \nu(s) \in \mathbb{R}^k$, the integrand is

obtained from the Riemannian metric on $\mathbb{R} \times \mathbb{R}^k$ given at (r, y) by

$$\langle (h_1, w_1), (h_2, w_2) \rangle_{(r,y)} = a(h_1 h_2)e^r + b(w_1 \cdot w_2)e^r. \tag{46}$$

Thus, to derive the differential equation that governs covariant integration along a path in N , it suffices to derive the corresponding differential equation for $\mathbb{R} \times \mathbb{R}^k$ with respect to (46). Using the subscript 0 to identify the r -coordinate and the subscript i for the y_i -coordinate, the metric tensor on $\mathbb{R} \times \mathbb{R}^k$ is given by $g_{00}(r, x) = a e^r$, $g_{ii}(r, x) = b e^r$, for $1 \leq i \leq k$, and $g_{ij}(r, x) = 0$, otherwise. The Christoffel symbols of the Levi-Civita connection are

$$\Gamma_{00}^0 = \Gamma_{0i}^i = \Gamma_{i0}^i = \frac{1}{2}, \quad \Gamma_{ii}^0 = -\frac{b}{2a},$$

$1 \leq i \leq k$, and zero otherwise. Therefore, the covariant derivative of a vector field (F_t, X_t) along a path (φ_t, ν_t) in N is given by (see e.g. do Carmo 1994)

$$\begin{cases} D_t^N F_t(s) = \partial_t F_t(s) + \frac{1}{2} \partial_t \varphi_t(s) F_t(s) \\ \quad - \frac{1}{2} \frac{b}{a} [X_t(s) \cdot \partial_t \nu_t(s)] \\ D_t^N X_t(s) = \partial_t X_t(s) + \frac{1}{2} (X_t \partial_t \varphi_t + F_t \partial_t \nu_t), \end{cases} \tag{47}$$

where the superscript N on the left hand side just highlights the fact that covariant differentiation takes place in the Riemannian manifold N . Fields that are tangential to the submanifold $M \subset N$ are those that satisfy the additional orthogonality condition $F_t(s) \cdot \nu_t(s) = 0$, for every s and t . Thus, if (f_t, x_t) and (F_t, X_t) are tangential to M along a path in M , we can rephrase the condition that $(D_t^M F_t, D_t^M X_t) = (f_t, x_t)$ as

$$D_t^N F_t(s) = f_t(s) \quad \text{and} \quad D_t^N X_t(s) = x_t(s) - \tau_t(s) \nu_t(s),$$

where τ_t is a scalar field to be determined. Substituting in (47), we obtain

$$\begin{cases} \partial_t F_t(s) = f_t(s) + \frac{1}{2} \partial_t \varphi_t(s) F_t(s) \\ \quad - \frac{1}{2} \frac{b}{a} [X_t(s) \cdot \partial_t \nu_t(s)] \\ \partial_t X_t(s) = x_t(s) + \frac{1}{2} (X_t(s) \partial_t \varphi_t(s) \\ \quad + F_t(s) \partial_t \nu_t(s)) + \tau_t(s) \nu_t(s). \end{cases} \tag{48}$$

Differentiating $X_t(s) \cdot \nu_t(s) = 0$, it follows that $\partial_t X_t(s) \cdot \nu_t(s) = -X_t(s) \cdot \partial_t \nu_t(s)$. From the expression for $\partial_t X_t(s)$ obtained in (48), we get $\tau_t(s) = -X_t(s) \cdot \partial_t \nu_t(s)$, where we used the facts that $X_t(s) \cdot \nu_s = 0$ and $\partial_t \nu_t(s) \cdot \nu_t(s) = 0$. Substituting this value of $\tau_t(s)$ in (48) yields (29). It is easy to check that, for any $t \in I$, a solution of (29) is actually tangential to M .

Appendix C: Paths with Fixed Boundary

In Sect. 7.3, we used the fact that the orthogonal complement of the tangent space of the path space Z_M in the tangent space of Y at $(\varphi_t, \nu_t) \in Z_M$, with respect to the Palais inner product, consists of the covariantly linear fields along (φ_t, ν_t) . This follows from a simple integration-by-parts argument that holds in general Riemannian manifolds, as we recall next.

Let Q be the manifold of continuous paths $\alpha: I \rightarrow \mathcal{M}$ (with square integrable derivative) in a Riemannian manifold \mathcal{M} , and $Q_1 \subset Q$ the submanifold of paths satisfying the boundary conditions $\alpha(0) = p_0$ and $\alpha(1) = p_1$, where $p_0, p_1 \in \mathcal{M}$. If $\alpha \in Q_1$, a vector field $h(t)$, $t \in I$, along α represents a tangent vector in $T_\alpha Q_1$ if and only if the field vanishes at the end points, as can be easily seen from the constraints defining Q_1 . If $r(t)$ is another tangent vector, then the Palais inner product is given by

$$\begin{aligned} \langle h, r \rangle_\alpha &= \langle h(0), r(0) \rangle_{p_0} + \int_0^1 \langle D_t h(t), D_t r(t) \rangle_{\alpha(t)} dt \\ &= \int_0^1 \langle D_t h(t), D_t r(t) \rangle_{\alpha(t)} dt \\ &= \langle h(t), D_t r(t) \rangle_{\alpha(t)} \Big|_0^1 - \int_0^1 \langle h(t), D_t^2 r(t) \rangle_{\alpha(t)} dt \\ &= - \int_0^1 \langle h(t), D_t^2 r(t) \rangle_{\alpha(t)} dt. \end{aligned}$$

The last expression vanishes for every $h \in T_\alpha Q_1$ if and only if $D_t^2 r = 0$. Hence, r is orthogonal to $T_\alpha Q_1$ if and only if r is covariantly linear.

References

- Belongie, S., Malik, J., & Puzicha, J. (2002). Shape matching and object recognition using shape context. *PAMI*, 24, 509–522.
- Cohen, I., Ayache, N., & Sulger, P. (1992). Tracking points on deformable objects using curvature information. In *Lecture notes in computer science* (vol. 588). Berlin: Springer.
- do Carmo, M. P. (1994). *Riemannian geometry*. Basel: Birkhauser.
- Geiger, D., Gupta, A., Costa, L. A., & Vlontzos, J. (1995). Dynamic programming for detecting, tracking and matching elastic contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(3), 294–302.
- Grenander, U. (1993). *General pattern theory*. Oxford: Oxford University Press.
- Joshi, S., Klassen, E., Srivastava, A., & Jermyn, I. (2007). An efficient representation for computing geodesics between n -dimensional elastic shapes. In *IEEE conference on computer vision and pattern recognition*.
- Kendall, D. G. (1984). Shape manifolds, Procrustean metrics and complex projective spaces. *Bulletin of London Mathematical Society*, 16, 81–121.
- Klassen, E., & Srivastava, A. (2006). Geodesics between 3D closed curves using path straightening. In *European conference on computer vision (ECCV)*.
- Klassen, E., Srivastava, A., Mio, W., & Joshi, S. (2004). Analysis of planar shapes using geodesic paths on shape manifolds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26, 372–383.
- Ling, H., & Jacobs, D. (2007). Shape classification using the inner distance. *PAMI*, 29(2), 286–299.
- Michor, P., & Mumford, D. (2006). Riemannian geometries on spaces of plane curves. *Journal of the European Mathematical Society*, 8, 1–48.
- Michor, P., & Mumford, D. (2007). An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. *Applied and Computational Harmonic Analysis*, 23, 74–113.
- Michor, P., Mumford, D., Shah, J., & Younes, L. (2007). A metric on shape space with explicit geodesics. arXiv:0706.4299v1.
- Mio, W., Bowers, J. C., Hurdal, M. K., & Liu, X. (2007a). Modeling brain anatomy with 3D arrangements of curves. In *IEEE 11th international conference on computer vision* (pp. 1–8).
- Mio, W., Srivastava, A., & Joshi, S. (2007b). On shape of plane elastic curves. *International Journal of Computer Vision*, 73(3), 307–324.
- Mio, W., Srivastava, A., & Klassen, E. (2004). Interpolations with elasticae in Euclidean spaces. *Quarterly of Applied Mathematics*, 62, 359–378.
- Mumford, D. (2002). Pattern theory: The mathematics of perception. In *Proc. of the international congress of mathematicians*. Beijing, China.
- Palais, R. S. (1963). Morse theory on Hilbert manifolds. *Topology*, 2, 299–340.
- Schmidt, F. R., Clausen, M., & Cremers, D. (2006). Shape matching by variational computation of geodesics on a manifold. In *LNCs: Vol. 4174. Pattern recognition (Proc. DAGM)* (pp. 142–151). Berlin: Springer.
- Sebastian, T. B., Klein, P. N., & Kimia, B. B. (2003). On aligning curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(1), 116–125.
- Sebastian, T. B., Klein, P. N., & Kimia, B. B. (2004). Recognition of shapes by editing their shock graphs. *PAMI*, 26(5), 550–571.
- Shah, J. (2006). An H^2 type Riemannian on the space of planar curves. In *Workshop on the mathematical foundations of computational anatomy (MICCAI)*.
- Söderkvist, O. (2001). *Computer vision classification of leaves from Swedish trees*. Master's thesis, Linköping University.
- Sundaramoorthi, G., Yezzi, A., & Mennucci, A. (2007). Sobolev active contours. *The International Journal of Computer Vision*, 73, 345–366.
- Tagare, H. D. (1999). Shape-based non-rigid correspondence with applications to heart motion analysis. *IEEE Transactions on Medical Imaging*, 8(7), 570–579.
- Tagare, H. D., O'Shea, D., & Groisser, D. (2002). Non-rigid shape comparison of plane curves in images. *Journal of Mathematical Imaging and Vision*, 16, 57–68.
- Tu, Z., & Yuille, A. (2004). Shape matching and recognition using generative models and informative features. In *European conference on computer vision (ECCV)* (pp. 195–209).
- 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.
- Zheng, X., Chen, Y., Groisser, D., & Wilson, D. (2005). Some new results on non-rigid correspondence and classification of curves. In *EMMCVPR* (pp. 473–489).