APPLICATIONS OF PLANAR SHAPE ANALYSIS TO IMAGE-BASED INFERENCES

*Anuj Srivastava**, *Shantanu Joshi***, *David Kaziska**, *and David Wilson****

* Department of Statistics, Florida State University, Tallahassee, FL
** Department of Eletrical Engineering, Florida State University, Tallahassee, FL
*** Department of Mathematics, University of Florida, Gainesville, FL

ABSTRACT

We describe an approach for statistical analysis of shapes of closed curves using tools from differential geometry. This approach uses geodesic paths to define a metric on shape space, that is used to compare shapes, to compute intrinsic statistics for a set of shapes, and to define probability models on shape spaces. We demonstrate this approach using: (i) interpolation of heart-wall boundaries in echocardiographic image sequences and (ii) a study of shapes of human silhouettes in infrared surveillance images.

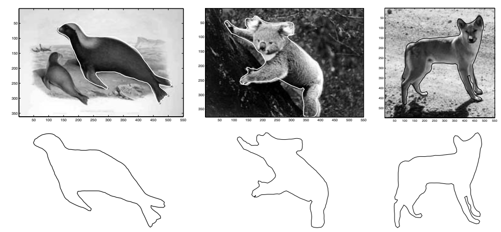

Fig. 1. Analysis of shapes of objects' boundaries in images can help in computer vision tasks such as object recognition.

1. INTRODUCTION

Detection, extraction and recognition of objects in an image is an important area of research. Objects can be characterized using a variety of features: textures, edges, boundaries, colors, motion, shapes, locations, etc. Shape often provides an important clue for determining how an object appears in an image. For example, we have displayed the images of three animals in the top panels of Figure 1. The lower panels show the silhouettes of these animals in the corresponding images. It is easy to see that the shapes of these silhouettes can help shortlist, or even identify, the animals present in these images. Tools for shape analysis can prove important in several applications including medical image analysis, human surveillance, military target recognition, fingerprint analysis, space exploration, and underwater search.

A significant part of the past efforts has been restricted to "landmark-based" analysis, where shapes are represented by a coarse, discrete sampling of the object contours [1]. A recent approach [2] considers the shapes of **continuous**, closed curves in $\mathbb{R}^2$. In this paper we describe two applications of this idea: First, we look at a problem in ecocardiographic image analysis where shapes of epicardial and endocardial boundaries are studied to determine the extent and progression of disease in a patient's heart. We focus on the specific problem of interpolating these boundaries in image sequences when an expert provides contours for the first and last frames in the sequence. Secondly, we will present an application involving human surveillance with a goal of detecting humans in low-quality night-vision (infrared) images. Our goal here is to build a statistical model to capture human shapes.

The rest of this chapter is organized as follows. In Section 2 we summarize past work on differential-geometric representation of shapes. In the next two sections we describe two applications of this approach.

2. A FRAMEWORK FOR PLANAR SHAPE ANALYSIS

The basic idea presented in [2] is to identify a space of closed curves, remove shape-preserving transformations from it, impose a Riemannian structure on it, and treat the resulting quotient space as the shape space.

1. **A Geometric Representation of Shapes**: Consider the boundaries or silhouettes of the imaged objects as closed, planar curves in $\mathbb{R}^2$, parameterized by the arc length. Denote by $\theta(s)$ the angle made by the velocity vector with the positive x-axis, as a function of arc length s. Coordinate function $\alpha(s)$ relates to the angle function $\theta(s)$ according to $\dot{\alpha}(s) = e^{j\,\theta(s)}$, $j = \sqrt{-1}$, with an example in Figure 2. We choose angle functions to represent and analyze shapes. To

This research was supported in part by the grants NSF (FRG) DMS-0101429, NSF (ACT) DMS-0345242, and ARO W911NF-04-01-0268.

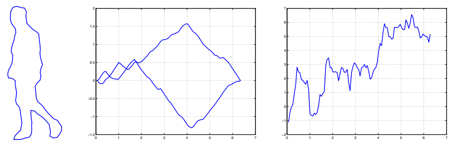

Fig. 2. A closed curve (left panel), its coordinate functions α (second panel) and its angle function θ (third panel).

build shape-invariant representations, one restricts to the set $\mathcal{C} = \{\theta \in \theta_0 + \mathbb{L}^2 | \frac{1}{2\pi}\int_0^{2\pi} \theta(s)ds = \pi, \int_0^{2\pi} e^{j\theta(s)}ds = 0\}$, where $\theta_0(s) = s$. The first constraint removes the orientation variability from these representations, while the second constraint ensures that the shapes are closed. This does allow self-intersecting curves. Furthermore, to remove the reparametrization group (relating to different placements of the origin), define the quotient space $\mathcal{S} \equiv \mathcal{C}/\mathbb{S}^1$ as the space of continuous, planar shapes, where $\mathbb{S}^1$ denotes the unit circle in $\mathbb{R}^2$. $\mathcal{C}$ is called the *pre-shape space* and $\mathcal{S}$ is called the *shape space*. It should be noted that angle functions representing shapes are actually analyzed discretely, in view of the discrete (pixellated) images containing those shapes.

2. **Geodesic Paths Between Shapes**: An important tool in a Riemannian analysis of shapes is to construct geodesic paths between arbitrary shapes. Klassen et al. [2] approximate geodesics on $\mathcal{S}$ by successively drawing infinitesimal line segments in $\mathbb{L}^2$ and projecting them onto $\mathcal{S}$. For any two shapes $\theta_1, \theta_2 \in \mathcal{S}$, one uses a *shooting method* to construct a geodesic between them. The geodesic metric is $\langle g_1, g_2 \rangle = \int_0^{2\pi} g_1(s)g_2(s)ds$ on the tangent space of $\mathcal{S}$. The resulting geodesic between any two shapes is the path that uses **minimum energy to bend one shape into the other**. Shown in Figure 3 are two examples of geodesic paths connecting the two end shapes. We will use the notation $\Psi_t(\theta, g)$ for a geodesic path starting from $\theta \in \mathcal{S}$, in the direction $g \in T_\theta(\mathcal{S})$, as a function of time t. Here $T_\theta(\mathcal{S})$ denotes the space of functions tangents to $\mathcal{S}$ at the point θ. If $g \in T_{\theta_1}(\mathcal{S})$ is the shooting direction to reach θ_2 in unit time from θ_1, then the following holds: $\Psi_0(\theta_1, g) = \theta_1$, $\Psi_1(\theta_1, g) = \theta_2$, and $\dot{\Psi}_0(\theta_1, g) = g$. The length of this geodesic is given by $d(\theta_1, \theta_2) = \sqrt{\langle g, g \rangle}$.

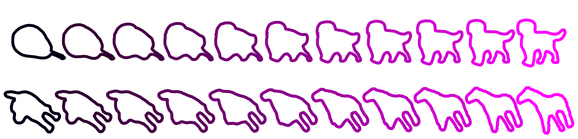

Fig. 3. Examples of geodsic paths in $\mathcal{S}$.

3. **Mean Shape in** $\mathcal{S}$: For a collection $\theta_1, \ldots, \theta_n$ in $\mathcal{S}$, and $d(\theta_i, \theta_j)$ the geodesic length between θ_i and θ_j, the Karcher mean is defined as the element $\mu \in \mathcal{S}$ that minimizes the quantity $\sum_{i=1}^n d(\theta, \theta_i)^2$. A gradient-based, iterative algorithm for computing the Karcher mean is particularized to $\mathcal{S}$ in [2].

This approach provides a comprehensive framework for a statistical analysis of planar shapes. In the next two sections, we present some applications of this framework to problems of practical interest.

3. INTERPOLATION OF SHAPES IN ECHOCARDIOGRAPHIC IMAGE-SEQUENCES

Shape analysis continues to play a major role in medical diagnostics using non-invasive imaging. Shapes and shape variations of anatomical parts are often important factors in deciding normality/abnormality of imaged patients. For example, the two images displayed in Figure 4 were acquired as the end diastolic (ED) and end systolic (ES) frames from a sequence of echocardiographic images during systole, taken from the apical four chamber view. Superimposed on both images are expert tracings of the epicardial (solid lines) and endocardial borders (broken lines) of the left ventricle of the heart. From these four borders, indices of cardiac health, including chamber area, fractional area change, and wall thickness, can be easily computed.

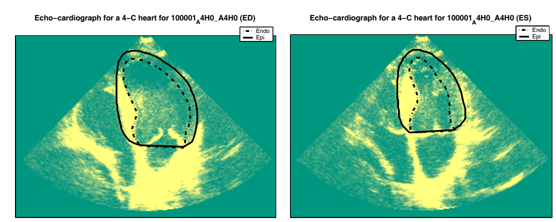

Fig. 4. Expert generated boundaries, denoting epicardial (solid lines) and endocardial (broken lines) borders, drawn over ED (left) and ES (right) frames of an echocardiographic image sequence.

A major goal in echocardiographic image analysis has been to develop and implement automated methods for computing these two sets of borders as well as the sets of borders for the 10-12 image frames that are typically acquired between ED and ES. Different aspects of past efforts include both the construction of geometric figures to model the shape of the heart as well as validation. While it is a rare cardiologist who is willing to submit to the tedium of drawing borders for all image frames between ED and ES, a few will in fact agree to draw borders for the first and last frames. Since the heart walls may exhibit diskinetic (i.e irregular) motion patterns during systole, the tracking of these borders may be important in a diagnosis. Our goal is to estimate epicardial and endocardial boundaries in the intermediate frames given the boundaries at the ED and ES frames.

A closed contour α has two sets of descriptors associated with it: a shape descriptor denoted by $\theta \in \mathcal{S}$ and a vector $z \in \mathcal{Z}$ of nuisance variables such as position, orientation, and scale. In our approach, interpolation between two closed curves is performed via interpolations between their shapes and nuisance components, respectively. The interpolation of shape is obtained using geodesic paths, while that of the nuisance components is obtained using linear methods. The implication is to use a Riemannian structure on the joint space of shape and nuisance variables, with a Euclidean metric on the nuisance space $\mathcal{Z}$. Let $\alpha_1 = (\theta_1, z_1)$ and $\alpha_2 = (\theta_2, z_2)$ be the two closed curves, and our goal is to find a path $\Phi : [0, 1] \mapsto \mathcal{S} \times \mathcal{Z}$ such that $\Phi_0 = (\theta_1, z_1)$ and $\Phi_1 = (\theta_2, z_2)$. For example, in Figure 4, the endocardial boundary (broken curves) of the ED and ES frames can form α_1 and α_2, respectively. Alternatively, one can treat the epicardial boundaries (solid curves) of ED and ES frames as α_1 and α_2 as well. The different components are interpolated as follows:

1. **Shape Component**: Given the two shapes θ_1 and θ_2 in $\mathcal{S}$, we use the shooting method to find the geodesic that starts from the first and reaches the other in unit time. This results in the flow $\Psi_t(\theta_1, g)$ such that $\Psi_0(\theta_1, g)$ = θ_1 and $\Psi_1(\theta_1, g) = \theta_2$. This also results in a reparametrization of θ_2 such that the origins (points where $s = 0$) on the two curves are now registered. With a slight abuse of notation we will also call the new curve θ_2. Let a shape along this path be given by $\theta_t = \Psi_t(\theta_1, g)$. Since the path θ_t lies in $\mathcal{S}$, the average value of θ_t for all t is π.

2. **Translation**: If p_1, p_2 represent the locations of the initial points on the two curves, i.e. $p_i = \alpha_i(0)$, $i = 1, 2$, then the linear interpolation between them is given by $p(t) = (1 - t)p_1 + tp_2$.

3. **Orientation**: For a closed curve α_i, the average orientation is defined by $\phi_i = \frac{1}{2\pi} \int_0^{2\pi} \frac{1}{j} \log(\dot{\alpha}_i(s))ds$, $i = 1, 2$, $j = \sqrt{-1}$. Given ϕ_1 and ϕ_2, a linear interpolation between them is $\phi(t) = (1 - t)\phi_2 + t\tilde{\phi}_2$, where $\tilde{\phi}_2 = \text{argmin}_{\phi \in \{\phi_2 - 2\pi, \phi_2, \phi_2 + 2\pi\}} |\phi - \phi_1|$.

4. **Scale**: If ρ_1 and ρ_2 are the lengths of the curves α_1 and α_2, then a linear interpolation on the lengths is simply $\rho(t) = (1 - t)\rho_1 + t\rho_2$.

Using these different components, the resulting geodesic on the space of closed curves is given by $\{\Phi_t : t \in [0, 1]\}$ where:

$$\Phi_t(s) = p(t) + \rho(t) \int_0^s \exp(j(\theta_t(\tau) - \pi + \phi(t)))d\tau \ .$$

Shown in Figure 5 is a sequence of 11 image frames for the same patient as displayed in Figure 4. Again, each image frame has a set of epicardial and endocardial borders overlaid on the image. In Figure 5, borders in the first and last frames have been traced by an expert, while the borders on the intermediate frames have been generated using the path Φ_t, one each for epicardial and endocardial boundaries. In view of the geodesic paths in $\mathcal{S}$ relating to the minimum bending energy, this method provides a smoother interpolation for the endocardial borders, as compared to a direct linear interpolation of coordinates.

We foresee a number of uses for this idea. First, this method could be included in an acquisition system so that if an expert traces sets of borders at ED and ES, then the borders for the intermediate frames can be generated automatically. As a future extension, one might modify the proposed interpolation to include image information. That is, formulate a boundary-value problem in $\mathcal{S}$ that seeks an optimal path under an image-based energy function, while fixing the expert generated boundaries as the end points.

4. STUDY OF HUMAN SILHOUETTES IN INFRARED IMAGES

There is a great interest in detection and recognition of humans using static images and video sequences. Night-vision cameras, or infrared cameras, have been found important in human detection and tracking, especially in surveillance and security environments. These cameras capture emissivity, or thermal states, of the imaged objects, and are largely invariant to ambient illumination. In this section, we investigate the problem of building statistical shape models for human silhouettes, for use in future shape detection.

Using a hand-held Raytheon Pro250 IR camera, we have hand-generated a database of human silhouettes. Shown in Figure 6 are some examples: the top panels show three IR images and the bottom panels show the corresponding hand-extracted human silhouettes. Furthermore, the database has been partitioned into clusters of similar shapes. These clusters correspond to front views with legs appearing together, side views with legs apart, side views with leg together, etc, and an example cluster is shown in Figure 7.

Our goal is to "train" probability models by assuming that elements in the same cluster are samples from the same probability model. These models can then be used for future Bayesian discoveries of shapes or for classification of new shapes. To train a probability model amounts to estimating a probability density function on the shape space $\mathcal{S}$, a task that is rather difficult to perform precisely. The two main difficulties are: nonlinearity and infinite-dimensionality of $\mathcal{S}$, and they are handled here as follows. $\mathcal{S}$ is a nonlinear manifold, so we impose a probability density on a tangent space instead. For a mean shape $\mu \in \mathcal{S}$, $T_\mu(\mathcal{S}) \subset \mathbb{L}^2$, is a

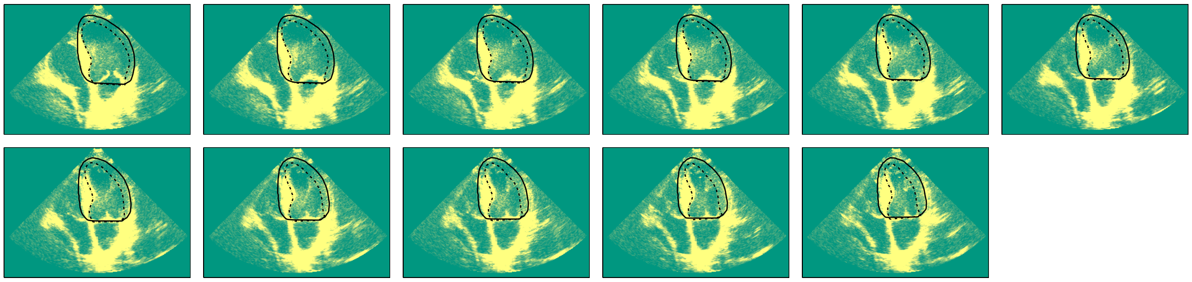

Fig. 5. Interpolated shapes using geodesic paths in shape space.

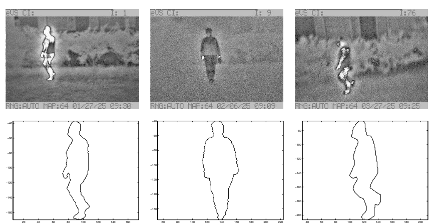

Fig. 6. Top panels: Examples of infrared images of human subjects. Bottom panels: hand extracted boundaries for analyzing shapes of human silhouettes.

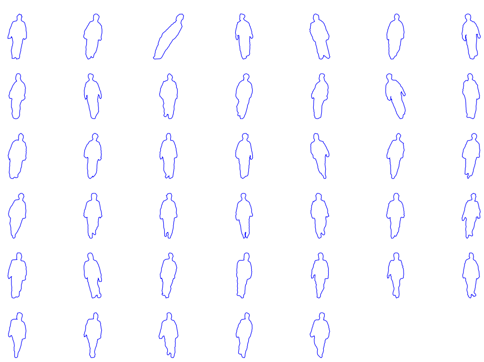

Fig. 7. An example of a cluster of human silhouettes.

vector space and more conventional statistics applies. Next, we approximate a tangent function g by a finite-dimensional vector, e.g. a vector of Fourier coefficients, and thus characterize a probability distribution on $T_\mu(\mathcal{S})$ as that on a finite-dimensional vector space. Let a tangent element $g \in T_\mu(\mathcal{S})$ be represented by its approximation: $g(s) = \sum_{i=1}^{m} x_i e_i(s)$, where $\{e_i\}$ is a complete orthonormal basis of $T_\mu(\mathcal{S})$ and m is a large positive integer. Using the identification $g \equiv \mathbf{x} = \{x_i\} \in \mathbb{R}^m$, one can define a probability distribution on elements of $T_\mu(\mathcal{S})$ via one on $\mathbb{R}^m$. The simplest model is a multivariate normal probability imposed as follows. Using principal component analysis (PCA) of the elements of $\mathbf{x}$, determine variances of the principal coefficients, and impose independent Gaussian models on these coefficients with zero means and estimated variances. This imposes a probability model on $T_\mu(\mathcal{S})$, and through the exponential map ($\exp_\mu : T_\mu(\mathcal{S}) \mapsto \mathcal{S}$ defined by $\exp_\mu(g) = \psi_1(\mu, g)$) leads to a probability model on $\mathcal{S}$. We term this model "Tangent PCA" or TPCA.

Consider the set of 40 human silhouettes displayed in Figure 7. For each observed shape θ_i, we compute a tangent vector g_i, such that $\Psi_1(\mu, g_i) = \theta_i$. Using TPCA model we obtain a normal probability model on the tangent space $T_\mu(\mathcal{S})$. Shown in Figure 8 are the mean μ (leftmost figure) and eight random shapes generated by this probability model.

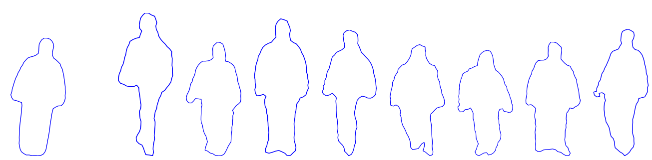

Fig. 8. Mean shape (left) and eight random shapes.

5. SUMMARY

We have presented a differential-geometric framework for statistical analysis of shapes, and have demonstrated this framework using two applications: medical image analysis and human surveillance.

6. REFERENCES

[1] I. L. Dryden and K. V. Mardia, *Statistical Shape Analysis*, John Wiley & Son, 1998.

[2] E. Klassen, A. Srivastava, W. Mio, and S. Joshi, "Analysis of planar shapes using geodesic paths on shape spaces," *IEEE Pattern Analysis and Machine Intelligence*, vol. 26, no. 3, pp. 372–383, March, 2004.