Geometric Analysis of Constrained Curves for Image Understanding

Anuj Srivastava	Washington Mio	Eric Klassen	Xiuwen Liu
Department of Statistics	Department of Mathematics		Department of Computer Science
anuj@stat.fsu.edu	*mio@math.fsu.edu*	*klassen@math.fsu.edu*	*liux@cs.fsu.edu*

Florida State University
Tallahassee, FL 32306, USA

Abstract

We offer an overview a novel, comprehensive approach to the computational differential geometry of curves, with applications to shape analysis and discovery/recognition of objects in images. The main idea is to specify a space of curves with constraints suited to the application, and exploit the differential geometry of this space to solve optimization and inference problems. Applications of this approach include the statistical analysis of planar shapes, curve interpolations with elasticae, and the Bayesian discovery of contours of objects in noisy images.

1 Introduction

An important goal in image understanding is to detect, track and label objects of interest present in observed images. Imaged objects can be characterized in many ways: according to their colors, textures, shapes, movements, and locations. The past decade has seen significant advances in the modeling and analysis of pixel values or textures to characterize objects in images, albeit with limited success. On the other hand, planar curves that represent contours of objects have been studied independently for a long time. An emerging opinion in the vision community is that global features such as shapes of contours should also be taken into account for the successful detection and recognition of objects. A common approach to analyzing curves in images is to treat them as level sets of functions, and algorithms involving such *active contours* are governed usually by partial differential equations (PDEs) driven by appropriate data terms and smoothness penalties (see for example [12]). Regularized curve evolutions and region-based active contours offer alternatives in similar frameworks. This remarkable body of work contains various studies of curve evolution, each with relative strengths and drawbacks. Recent work of Charpiat et al. [1] provides another viewpoint on shape analysis using a different formulation.

In this paper, we present a novel framework for the algorithmic study of curves, their variations and statistics. In this approach, a fundamental element is a *space of constrained curves*, called SCC henceforth, with constraints appropriate to the application of interest. (As an example, we may be interested in analyzing closed curves and this provides a constraint.) We exploit the geometry of SCCs using elements such as tangents, normals, geodesics and gradient flows, to solve optimization and statistical inference problems for a variety of cost functions and probability densities. This framework differs from those employed in previous works on "geometry-driven flows" [10] in the sense that here both the geometry of the curves and the geometry of SCCs are utilized. The dynamics of active contours is described by vector fields on SCCs, thus reducing the evolution of curves to the study of **ordinary differential equations** (ODEs) on SCCs. It is important to emphasize that a SCC is usually a non-linear, infinite-dimensional manifold, and its elements are the individual curves of interest. Several interesting applications can be addressed in this formulation:

1. Efficient deformations between any two curves are generated by geodesic paths connecting the elements they represent in the SCC. Geodesic lengths also provide a natural metric for clustering curves (automated learning) according to their shapes.

2. The discovery of planar shapes in given noisy images can be treated as the problem of finding *maximum a-posteriori* points in SCCs.

3. Given a set of curves (or shapes), one can define the concepts of *mean* and *covariance* using geodesic paths, and thus develop statistical frameworks for studying shapes. Furthermore, one can define probabilities on a SCC to perform curve (or shape) classification via hypothesis testing.

4. Tracking time-varying curves or dynamic shapes can be studied as a problem of Bayesian nonlinear filtering on SCCs.

5. Completions of partially occluded contours using elasticae (or local elasticae) can be accomplished via gradient-based optimizations on a SCC.

Many of these problems have been studied in the past with elegant solutions presented in the literature (examples include [11, 13, 8, 2, 6]), but using significantly different ideas. We demonstrate the strength of the proposed framework by addressing the above-mentioned applications in a comprehensive and unified manner. This framework relates closely to the ideas presented in [15].

Given past achievements in PDE-based approaches to curve evolution, what is the need for newer frameworks? The study of the structure of SCCs provides new insights and solutions to problems involving dynamic contours and problems in quantitative shape analysis. Once the constraints are utilized in definitions of SCCs the resulting solutions automatically satisfy these constraints. It also complements existing methods of image processing and analysis well by providing new computational efficiencies. The main strength of this approach is its exploitation of the differential geometry of SCCs. For instance, a geodesic or gradient flow X_t of an energy function E (on a SCC) can be generated as solution of an ODE of the type

$$\frac{dX_t}{dt} = \Pi(\nabla E(X_t)) \, , \tag{1}$$

where Π denotes an appropriate projection onto the tangent bundle of that SCC. This is in contrast with the nonlinear PDE-based curve evolutions of past works. The geometry of SCCs also enables us to derive statistical elements: probability measures, means and covariances on SCCs; these quantities have rarely been treated in previous studies. In shape extraction, the main focus in past works has been on solving PDEs driven by image features under smoothness constraints, and not on the statistical analysis of shapes of curves. The use of geodesic paths, or piecewise geodesic paths, has also seen limited use in the past.

We should also point out the main limitations of the proposed framework. One drawback is that curve evolutions can not handle certain changes in topology, which is one of the key features of level-set methods; a SCC is purposely setup to not allow curves to branch into several components. Secondly, this idea does not extend easily to the analysis of surfaces in $\mathbb{R}^3$. Despite these limitations, the proposed methodology provides powerful algorithms for the analysis of planar curves as demonstrated by the examples presented later. Moreover, even in applications where branching appears to be essential, the proposed methods may be applicable with additional developments.

This paper is laid out as follows: Section 2 studies geometric representations of constrained curves with two examples of SCCs. A brief geometric analysis of these spaces is presented in Section 3. Section 4 provides an example of statistical analysis on a SCC, while Section 5 presents two applications of this framework: (i) completion of partially occluded objects in images, and (ii) discovery of shapes in noisy images using a Bayesian framework.

2 Geometric Representations of Constrained Curves

In this section we provide two examples $\mathcal{C}_1$ and $\mathcal{C}_2$ of SCCs on which optimization and inference problems will be solved later in the paper. In this paper we restrict mostly to curves in $\mathbb{R}^2$ although curves in $\mathbb{R}^3$ can be handled similarly. Let $\alpha : \mathbb{R} \mapsto \mathbb{R}^2$ denote the coordinate function of a curve parameterized by arc-length, i.e., satisfying $\|\dot{\alpha}(s)\| = 1$, for every s. A direction function $\theta(s)$ is a function satisfying $\dot{\alpha}(s) = e^{j\,\theta(s)}$, where $j = \sqrt{-1}$. θ captures the angle made by the velocity vector with the x-axis, and is defined up to the addition of integer multiples of 2π. The curvature function $\kappa(s) = \dot{\theta}(s)$ can also be used to represent a curve.

The choice of representation of curves will depend on the specific application. Furthermore, different constraints imposed on curves lead to different SCCs. Two examples suitable for the applications considered in this paper are discussed next.

1. **A Space of Closed Curves with Fixed Length**: Consider the problem of studying shapes of contours or silhouettes of imaged objects as closed, planar curves in $\mathbb{R}^2$, parameterized by arc length. Since shapes are invariant to rigid motions (rotations and translations) and uniform scaling, a shape representation should be insensitive to these transformations. Scaling can be resolved by fixing the length of α to be 2π, and translations by representing curves via their direction functions. Thus, we consider the space $\mathbb{L}^2$ of all square integrable functions $\theta\colon [0, 2\pi] \to \mathbb{R}$, with the usual inner product $\langle f, g\rangle = \int_0^{2\pi} f(s)g(s)\,ds$. To account for rotations and ambiguities on the choice of θ, we restrict direction functions to those having a fixed average, say, π. For α to be closed, it must satisfy the *closure condition* $\int_0^{2\pi} e^{j\theta(s)}\,ds = 0$. Thus, we represent curves by direction functions satisfying the average-π and closure conditions; we call this space of direction functions $\mathcal{D}_1$. Summarizing, $\mathcal{D}_1$ is the subspace of $\mathbb{L}^2$ consisting of all (direction) functions satisfying the constraints

$$\begin{aligned} \int_0^{2\pi} \theta(s)\,ds &= \pi \, ; \int_0^{2\pi} \cos(\theta(s))\,ds = 0 \, ; \\ \int_0^{2\pi} \sin(\theta(s))\,ds &= 0 \, . \end{aligned} \tag{2}$$

It is still possible to have multiple elements of $\mathcal{D}_1$ representing the same shape. This variability is due to the choice of the reference point ($s = 0$) along the curve. For $x \in \mathbb{S}^1$ and $\theta \in \mathcal{D}_1$, define $(x \cdot \theta)$ as a curve whose initial point ($s = 0$) is changed by a distance of x along the curve. We term this a re-parametrization of the curve. To remove the variability due to this re-parametrization group, define the quotient space $\mathcal{C}_1 \equiv \mathcal{D}_1/\mathbb{S}^1$ as the space of continuous, planar shapes.

2. **Space of Curves with Specified Boundary Conditions**: As another example, we consider curves $\alpha\colon I \to \mathbb{R}^2$ satisfying given boundary conditions to first order, where $I = [0, 1]$. Given points $p_0, p_1 \in \mathbb{R}^2$ with $\|p_1 - p_0\| < 1$ and angles $\theta_0, \theta_1 \in \mathbb{R}$, we are interested in curves α that admit angle functions θ satisfying $\alpha(i) = p_i$ and $\theta(i) = \theta_i$, for $i = 0, 1$.

   Curves with $\alpha(0) = p_0$ are determined by their direction functions via the expression $\alpha(s) = p_0 + \int_0^s e^{j\theta(u)}\, du$. For these curves, the conditions above can be rephrased as

$$\theta(0) = \theta_0, \quad \theta(1) = \theta_1, \text{ and } \int_0^1 e^{j\theta(s)}\, ds = d,$$

   where $d = p_1 - p_0$ is the total displacement of α. This last condition ensures that the end point of the curve α is p_1. We consider the vector space $\mathbb{H}^1$ of all absolutely continuous functions $\theta\colon I \to \mathbb{R}$ with square integrable derivative, equipped with the inner product $\langle f, g\rangle = f(0)g(0) + \int_0^1 \dot{f}(s)\dot{g}(s)\, ds$. Here, we use the space $\mathbb{H}^1$ instead of $\mathbb{L}^2$ because we wish to be able to control the values of θ at the end points. The space $\mathcal{C}_2$ consists of all functions in $\mathbb{H}^1$ satisfying the three conditions above.

3 Geometries of SCCs

The main idea in the proposed framework is to use the geometric structure of SCCs to solve optimization and statistical inference problems on these spaces. This approach often leads to simple formulations of these problems and to more efficient vision algorithms. Thus, we must study issues related to the differential geometry and topology of those SCCs. In this paper we restrict to the tangent and normal bundles, and geodesic flows on these spaces.

3.1 Tangents and Normals to SCCs

There are two main reasons for studying the tangential and normal structures: (i) to compute the gradient of the restriction of a functional on $\mathbb{L}^2$ ($\mathbb{H}^1$, resp.) to a SCC, we can first compute the gradient on the linear space $\mathbb{L}^2$ ($\mathbb{H}^1$, resp.) in which the SCC is contained (usually, a simpler task), and then *subtract the normal components* to obtain the component that is tangent to the SCC; (ii) we wish to employ iterative numerical methods in the simulation of geodesic and gradient flows; at each step in the iteration, we first flow in the linear space $\mathbb{L}^2$ ($\mathbb{H}^1$, resp.) using standard methods, and then project the new point back onto the SCC using our knowledge of the normal structure, as discussed in Section 3.2. The tangent spaces on these two SCCs are described next.

1. **Case 1**: For technical reasons, it is convenient to reduce optimization and inference problems on $\mathcal{C}_1$ to problems on the manifold $\mathcal{D}_1$, so we study the latter. It is difficult to specify the tangent spaces to $\mathcal{D}_1$ directly, because they are infinite-dimensional. When working with finitely many constraints, as is the case here, it is easier to describe the space of normals to $\mathcal{D}_1$ in $\mathbb{L}^2$ instead. It can be shown that a vector $f \in \mathbb{L}^2$ is tangent to $\mathcal{D}_1$ at θ if and only if f is orthogonal to the subspace spanned by $\{1, \sin\theta, \cos\theta\}$. Hence, these three functions span the normal space to $\mathcal{D}_1$ at θ. Implicitly, the tangent space is given as:

$$T_\theta(\mathcal{D}_1) = \{f \in \mathbb{L}^2 | f \perp \text{span}\{1, \cos\theta, \sin\theta\}\}\,.$$

   Thus, the projection Π in Eqn. 1 can be specified by subtracting from a function (in $\mathbb{L}^2$) its projection onto the space spanned by these three elements.

2. **Case 2**: Similar to Case 1, one can specify the four-dimensional space of normals to $\mathcal{C}_2$ inside $\mathbb{H}^1$ as the space spanned by $\{1, s, \varepsilon_1, \varepsilon_2\}$, where $\varepsilon_1, \varepsilon_2\colon I \to \mathbb{R}$ are characterized by $\ddot{\varepsilon}_1 = \cos\theta$, $\varepsilon_1(0) = \dot{\varepsilon}_1(0) = 0$, and $\ddot{\varepsilon}_2 = \sin\theta$, $\varepsilon_2(0) = \dot{\varepsilon}_2(0) = 0$ [7]. Thus,

$$T_\theta(\mathcal{C}_2) = \{f \in \mathbb{L}^2 | f \perp \text{span}\{1, s, \varepsilon_1, \varepsilon_2\}\}\,.$$

3.2 Geodesics Connecting Closed Curves

We first describe the computation of geodesics (or, one-parameter flows) in $\mathcal{D}_1$ with prescribed initial conditions. The intricate geometry of $\mathcal{D}_1$ disallows explicit analytic expressions. Therefore, we adopt an iterative strategy, where in each step, we first flow infinitesimally in the prescribed tangent direction in the space $\mathbb{L}^2$, and then project the end point of the path to $\mathcal{D}_1$. Next, we parallel transport the velocity vector to the new point by projecting the previous velocity orthogonally onto the tanget space of $\mathcal{D}_1$ at the new point. Again, this is done by subtracting normal components. The simplest implementation is to use Euler's method in $\mathbb{L}^2$, i.e., to move in each step along short straight line segments in $\mathbb{L}^2$ in the prescribed direction, and then project the path back onto $\mathcal{D}_1$. Details of this numerical construction of geodesics are provided in [5].

A one-parameter flow can be specified by an initial condition $\theta \in \mathcal{D}_1$ and a direction $f \in T_\theta(\mathcal{D}_1)$, the space of all tangent directions at θ. We will denote the corresponding flow by $\Psi(\theta, t, f)$, where t is the time parameter. The technique just described allows us to compute Ψ numerically.

Next, we focus on the problem of finding a geodesic path between any two given shapes θ_1, $\theta_2 \in \mathcal{D}_1$. The only remaining issue is to find that appropriate direction $f \in T_{\theta_1}(\mathcal{D}_1)$ such that a geodesic from θ_1 in that direction passes through θ_2 at time $t = 1$. In other words, the problem is to solve for an $f \in T_{\theta_1}(\mathcal{D}_1)$ such that $\Psi(\theta_1, 0, f) = \theta_1$ and $\Psi(\theta_1, 1, f) = \theta_2$. One can treat the search for this direction as an optimization problem over the tangent space $T_{\theta_1}(\mathcal{D}_1)$. The cost to be minimized is given by the functional $H[f] = \|\Psi(\theta_1, 1, f) - \theta_2\|^2$, and we are looking for that $f \in T_{\theta_1}(\mathcal{C}_1)$ for which: (i) $H[f]$ is zero, and (ii) $\|f\|$ is minimum among all such tangents. Since the space $T_{\theta_1}(\mathcal{D}_1)$ is infinite dimensional, this optimization is not straightforward. However, since $f \in \mathbb{L}^2$, it has a Fourier decomposition, and we can solve the optimization problem over a finite number of Fourier coefficients. For any two shapes $\theta_1, \theta_2 \in \mathcal{D}_1$, we have used a shooting method to find the optimal f [5]. The basic idea is to choose an initial direction f specified by its Fourier coefficients and then use a gradient search to minimize H as a function of the Fourier coefficients.

Finally, to find the shortest path between two shapes in $\mathcal{C}_1$, we compute the shortest geodesic connecting representatives of the given shapes in $\mathcal{D}_1$. This is a simple numerical problem, because $\mathcal{C}_1$ is the quotient of $\mathcal{D}_1$ by the 1-dimensional re-parametrization group $\mathbb{S}^1$.

Shown in Figure 1 are two examples of geodesic paths in $\mathcal{C}_1$ connecting given shapes. Drawn in between are shapes corresponding to equally spaced points along the geodesic paths. Similar ideas can be used to find geodesics on the space $\mathcal{C}_2$.

4 Statistical Models on SCCs

Using the case of $\mathcal{C}_1$, we illustrate statistical modeling of the constrained curves. Algorithms for finding geodesic paths on SCCs allow us to compute means and covariances in these spaces. We adopt a notion of mean known as the *intrinsic mean* or the *Karcher mean* ([4]) that is quite natural in our geometric framework. Let $d(\_,\_)$ be the shortest-path metric on $\mathcal{C}_1$. To calculate the Karcher mean of shapes $\{\theta_1, \ldots, \theta_n\}$ in $\mathcal{C}_1$, define a function $V : \mathcal{C}_1 \to \mathbb{R}$ by $V(\theta) = \sum_{i=1}^n d(\theta, \theta_i)^2$. Then, define the *Karcher mean* of the given shapes to be any point $\mu \in \mathcal{C}_1$ for which $V(\mu)$ is a local minimum. In the case of Euclidean spaces this definition agrees with the usual definition $\mu = \frac{1}{n}\sum_{i=1}^n p_i$. Since $\mathcal{C}_1$ is complete, the intrinsic mean as defined above always exists. However, there may be collections of shapes for which μ is not unique.

We now review an iterative algorithm given in [6] for finding a Karcher mean of given shapes.

Algorithm 1 *Set $k = 0$. Choose some time increment $\epsilon \leq \frac{1}{n}$. Choose a point $\mu_0 \in \mathcal{C}_1$ as an initial guess of the mean. (For example, one could just take $\mu_0 = \theta_1$.)*

1. *For each $i = 1, \ldots, n$ choose the tangent vector $f_i \in T_{\mu_k}(\mathcal{C}_1)$ which is tangent to the shortest geodesic from μ_k to θ_i, and whose norm is equal to the length of this shortest geodesic. The vector $g = \sum_{i=1}^n f_i$ is equal to (-2) times the gradient at μ_k of the function $V : \mathcal{C}_1 \to \mathbb{R}$ which we defined above.*

2. *Flow for time ϵ along the geodesic which starts at μ_k and has velocity vector g. Call the point where you end up μ_{k+1}, i.e. $\mu_{k+1} = \Psi(\mu_k, \epsilon, g)$.*

3. *If not converged, set $k = k + 1$, and go to Step 1.*

A similar algorithm and convergence results for a (finite-dimensional) landmark-based representation of shapes are described in [6].

For the mean μ, let $T_\mu(\mathcal{C}_1) \subset \mathbb{L}^2$ be the space of all tangents to the shape space at μ. Let a tangent element $f \in T_\mu(\mathcal{C}_1)$ be represented by its Fourier expansion:

$$f(s) \approx \sum_{k=1}^{m}(a_k \cos(ks) + b_k \sin(ks))$$

for a large positive integer m. Using the identification $f \approx \mathbf{a} = \{a_k, b_k\} \in \mathbb{R}^{2m-1}$, one can define a probability distribution on the tangent vectors in an approximate fashion. We will model $\mathbf{a}$ as multivariate normal with mean $\mathbf{0}$ and covariance $K \in \mathbb{R}^{(2m-1)\times(2m-1)}$. We will term $\sigma_0^2 = \mathrm{trace}(K)$ the *dispersion* of a shape model. Estimation of K from observed shapes follows the standard procedure. One can easily sample tangent vectors from this multivariate normal, and use the geodesic calculation to generate sample shapes. Shown in Figure 2 is an illustration of this idea. The left nine panels show the actual observed shapes from which the mean and covariance are calculated. The mean is shown in the middle image and the nine random shapes generated under the Gaussian model are shown in the right panels.

5 Applications

In this section, we describe some applications that involve the solution of optimization or inference problems on SCCs with geometric tools.

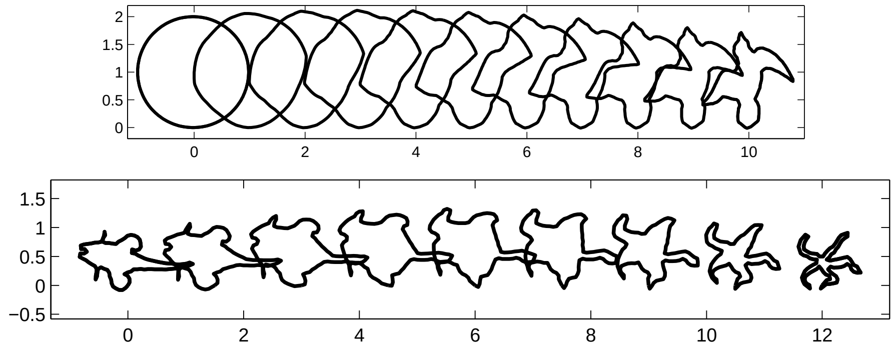

Figure 1: Examples of evolving one shape into another via a geodesic path. Leftmost shape is θ_1, rightmost curves are θ_2, and intermediate shapes are equispaced points along the geodesic.

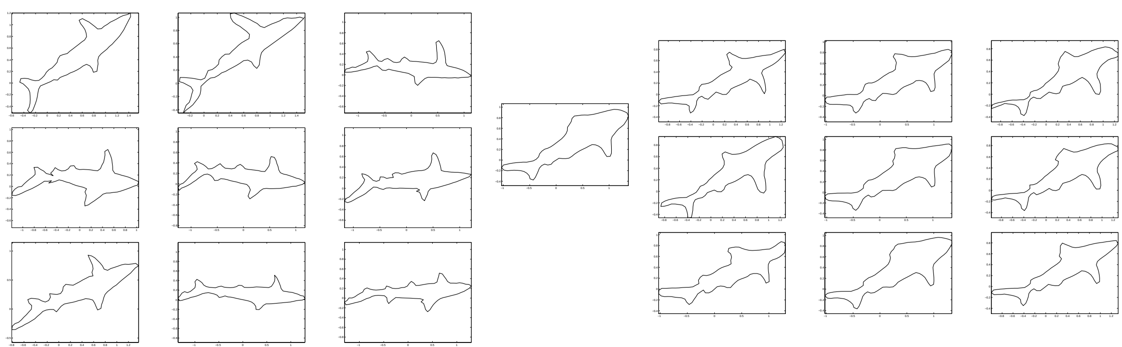

Figure 2: Learning shape models: For the nine observed shark shapes shown in left, the middle panel shows the mean shapes, and the right panels show nine random samples the Gaussian model.

5.1 Completion of Curves Using Constrained Elastic Curves

In the problem of recognizing objects in given images, the extraction and use of edges present in the images play an important role. If the objects of interest are partially obscured by other objects, an important task is to interpolate between the observed edges to complete contours. The geometry near the end points of the observed edges provide the boundary points p_0, p_1, and the boundary angles θ_0, θ_1. We complete the missing edge using an *elastica* satisfying these boundary conditions. This problem has been studied by several other researchers as well; see for example [13, 8]. For simplicity, we only consider the case of interpolations with elasticae with a given fixed length, normalized to be 1.

1. **Elasticae for Curve Completion**: In our notation, the elastic energy E of $\theta\colon I \to \mathbb{R}$ can be expressed as

$$E(\theta) = \frac{1}{2}\int_0^1 \dot{\theta}(s)^2\, ds. \tag{3}$$

We are interested in finding the critical points of E restricted to $\mathcal{C}_2$ using a gradient search method. Curves represented by these critical points are known as *elasticae*. Let $\theta^*(s) = \theta(s) - \theta_0$, and $f\colon I \to \mathbb{R}$ be the function obtained by projecting θ^* onto the tangent space of $\mathcal{C}_2$ at θ. Then, $\nabla_{\mathcal{C}_2} E = f$, i.e., f is the gradient of $E\colon \mathcal{C}_2 \to \mathbb{R}$ at θ. The flow lines of the negative gradient field $-\nabla_{\mathcal{C}_2} E$ on $\mathcal{C}_2$ approach elasticae asymptotically. Flows of this type that seek to minimize the elastic energy efficiently are known as *curve straightening flows*.

Shown in the top panels of Figure 3 are some examples of elasticae in $\mathbb{R}^2$ and $\mathbb{R}^3$ for given boundary conditions p_0, p_1, θ_0, and θ_1 (depicted via arrows). The bottom row shows an example of using a variant known as scale-invariant elasticae for completing missing edges of a partially occluded object. Object in the left panel is obscured artificially, and the boundaries of the visible parts are used to find interpolating elasticae that are shown in white lines in the right panel.

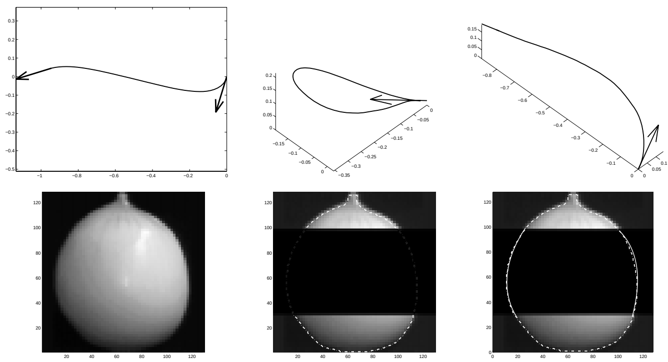

Figure 3: Upper row: solid lines show the elastica between given points and directions shown by arrows. Lower two rows: objects in left images are obscured and the end points of the visible edges are used to find elastica, shown in the right panels.

2. **Using Local Harmonics to Constrain Curves**: So far, the task of completing curves has been based solely on first-order boundary data. It seems logical to use more information from the visible portions. Our idea is to consider a subspace V of $\mathbb{H}^1$ associated with the dominant lower harmonics of the direction function of the visible portions of the edge, and restrict the search for completions to the space $\mathcal{C}_2^V = V \cap \mathcal{C}_2$. The energy now consists of two terms: E_1 is the same as the elastic energy defined in item 1, while E_2 gives a measurement of the similarity between the predicted and observed curves.

$$E_2(\theta) = \int_0^l |\theta(s) - \theta^*(s)|^2 ds \ ,$$

where l and θ^* are the length and the direction function of the visible curve, respectively. We have derived analytical expressions for the gradients of these two energies on $\mathcal{C}_2^V$, and have utilized them to search for the completion with least total energy. The edge completion can be complemented with a study of textures. Statistics of the texture of visible regions can used to predict the pixels values in the augmented region [14].

Shown in Figure 4 are simple examples computed using these ideas. The left panels show images of partially obscured objects and our goal is to predict the missing pieces of these objects. We solve this prediction problem in two steps using: (i) curve completion, and (ii) texture growth. For curve completion, we extract the boundaries of the visible portions, extract dominant harmonics, and use these harmonics in the prediction of closed contours . The middle panels show the visible portions in marked lines and the corresponding optimal completions in plain lines. Texture growth is used to produce the final results displayed on the right panels of each row.

5.2 Bayesian Discovery of Objects in Images

An important application of these shape analysis tools in the discovery of partially occluded objects in noisy images. Given some prior knowledge of their shapes, how can we incorporate it in the search for the objects? We define a prior probability distribution on an appropriate curve space, and use a Bayesian framework to infer new shapes. The curve space $\mathcal{C}$ used here is a simple variation of $\mathcal{C}_1$ that takes position, orientation, and scale into consideration. Elements $\gamma \in \mathcal{C}$ can (up to reparametrizations) be represented as pairs $\gamma = (x, \theta)$, where x is a finite-dimensional variable encoding initial position, initial velocity and scale, and $\theta \in \mathcal{C}_1$. We start by deriving a posterior density on $\mathcal{C}$.

1. **Image Likelihood**: The likelihood function can be described as follows: Let $D \subset \mathbb{R}^2$ be the image domain and $I\colon D \to \mathbb{R}^+$ be an image. A (simple) closed curve γ in D divides the image domain into a region $D_i(\gamma)$ inside the curve, and a region $D_o(\gamma)$ outside. Let P_i be a probability model for the pixel values inside the curve, and P_o be a model for pixels outside. For example, for a noisy two-phase image, one can choose P_i and P_o to be Gaussian distributions with different means.

   For a given image I, the likelihood that γ is present in it is proportional to $e^{-\frac{1}{\sigma_1^2} H(\theta, I)}$, where H is given by:

$$-\log\Big(\int_{D_i(\gamma)} P_i(I(x))dx + \int_{D_o(\gamma)} P_o(I(x))dx\Big) \, . \quad (4)$$

   This image model can be modified with energies such as the magnitude square of the image gradient [9], or information-theoretic entropy based terms [3].

2. **Prior Density**: We choose a "Gaussian" probability density as a prior. Let θ_0 represent a mean shape, σ_0^2

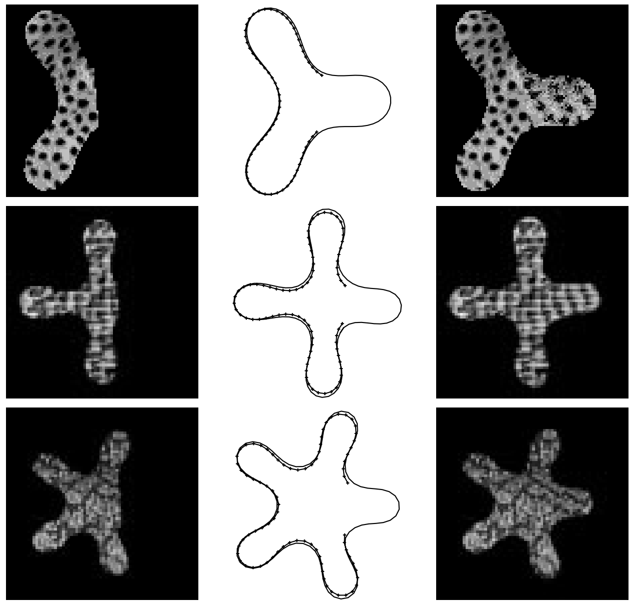

Figure 4: Left panels: given images of obscured objects. Middle panels: edges extracted from visible parts (marked lines) are used to find optimal curve completions (solid lines). Right panels: Texture statistics from the visible portions are used to grow pixels in new regions.

be the shape dispersion, and $\theta(\gamma) \in \mathcal{C}_1$ be the shape associated with γ. Then, define the prior density by $\mu(\gamma) \propto e^{-\frac{1}{\sigma_0^2} d(\theta(\gamma),\theta_0)^2} e^{-\frac{P(x)}{\sigma_2^2}}$, where $d(\cdot,\cdot)$ is the geodesic distance on $\mathcal{C}_1$ discussed earlier and P is a prior energy on the space of nuisance variables.

The posterior density is given by:

$$\mu(\gamma|I) = \frac{1}{Z} e^{-\frac{1}{\sigma_0^2} d(\theta(\gamma),\theta_0)^2 - \frac{1}{\sigma_2^2} P(x) - \frac{1}{\sigma_1^2} H(\gamma, I)} .$$

We have used a gradient approach to find the MAP estimate of γ for a given image. For an initial condition $\gamma_0 = (x_0, \theta_0)$, let $\Psi(\gamma_0, t, w, f)$ be the geodesic flow from γ_0 with initial velocity (w, f). Here, w is the component tangential to the "nuisance" variable x and $f \in T_{\theta_0}(\mathcal{C}_1)$. Using the fact that $d(\theta, \theta_0) = \|f\|$, we rewrite the posterior energy as:

$$E[w,f] = \frac{1}{\sigma_1^2} H(\Psi(\gamma_0, 1, w, f), I) + \frac{1}{\sigma_0^2}\|f\|^2 + \frac{1}{\sigma_2^2} P(x(w)) .$$

Now using a Fourier decomposition of f, we use a gradient process to minimize the posterior energy.

Shown in Figure 5 is an illustration of this Bayesian shape extraction. Top left panel shows the true shape embedded in the observed images, and top right shows the prior mean shape associated with θ_0. In the lower two rows, the left panels show the observed images and the remaining panels show MAP estimates of θ under an increasing influence of the prior (going from left to right). Successful discovery of the hidden shape despite partial obscuration (in the last row) emphasizes the need and power of a Bayesian approach in such problems.

6 Conclusions

We have presented an overview of an ambitious framework to solve optimization and inference problems on spaces of constrained curves (SCC). The main idea is to exploit the differential geometry of these Riemannian manifolds to obtain simpler solutions as compared to those obtained with PDE-based methods. Using two examples of SCCs, we have presented some applications of this framework in image understanding. In particular, these ideas lead to a novel statistical theory of shapes of planar objects with powerful tools for shape analysis.

References

[1] G. Charpiat, O. Faugeras, and R. Keriven. Approximations of shape metrics and application to shape warping and empirical shape statistics. *Institute Nationale de receherche en informatique et en automatique*, No. 4820, May, 2003.

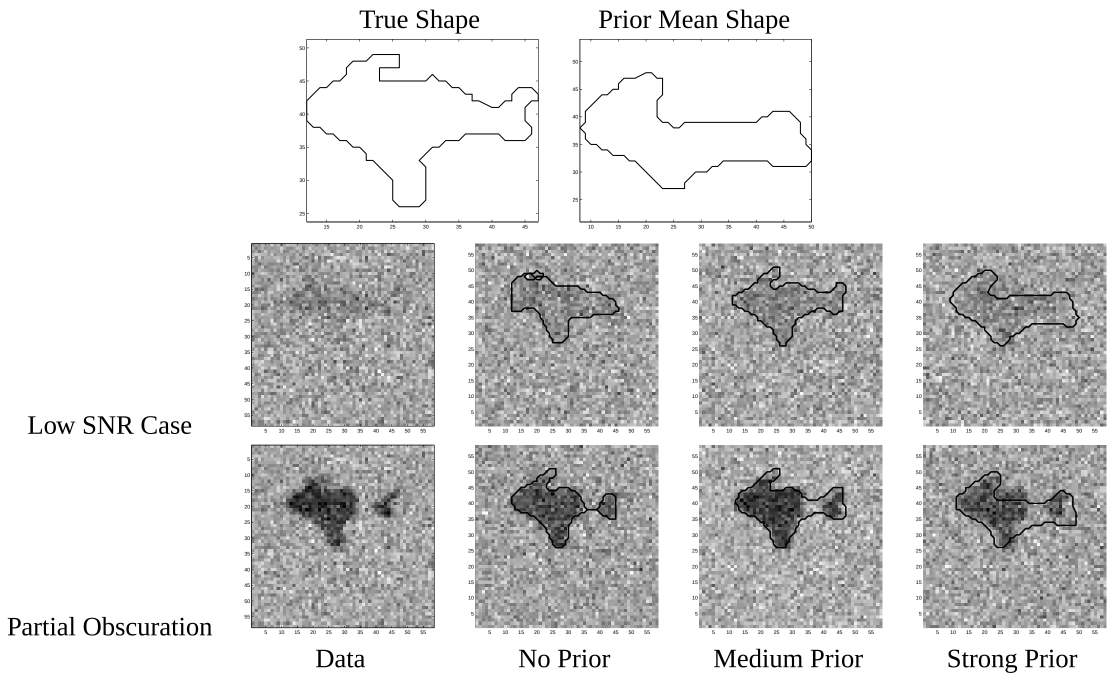

Figure 5: Top panels: left shows the true shape present in the images and right shows the prior mean. Middle and bottom panels: left shows the noisy data wand next three show MAP shape estimates (superimposed) with increasing influence of the prior from left to right.

[2] N. Duta, M. Sonka, and A. K. Jain. Learning shape models from examples using automatic shape clustering and Procrustes analysis. *Lecture Notes in Computer Science*, 1999.

[3] A. Ben Hamza and H. Krim. Entropic image registration and segmentation. In *Proceedings of EMMCVPR*, 2003, to appear.

[4] H. Karcher. Riemann center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*, 30:509–541, 1977.

[5] E. Klassen, A. Srivastava, and W. Mio. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Pattern Analysis and Machiner Intelligence*, in review, 2003.

[6] H. Le. Locating frechet means with application to shape spaces. *Advances in Applied Probability*, 33(2):324–338, 2001.

[7] W. Mio, A. Srivastava, and E. Klassen. Interpolation by elastica in euclidean spaces. *Quarterly of Applied Mathematics*, to appear, 2003.

[8] D. Mumford. *Elastica and computer vision*, pages 491–506. Springer, New York, 1994.

[9] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems, 1989.

[10] Editor: B. Romeny. *Geometry Driven Diffusions in Computer Vision*. Kluwer, 1994.

[11] T. B. Sebastian, P. N. Klein, and B. B. Kimia. On aligning curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(1):116–125, 2003.

[12] J. Sethian. *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Material Science*. Cambridge University Press, 1996.

[13] E. Sharon, A. Brandt, and R. Basri. Completion energies and scale. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(10):1117–1131, 2000.

[14] Y. N. Wu, S. C. Zhu, and X. Liu. Equivalence of julesz ensembles and frame models. *International Journal of Computer Vision*, 38(3):247–265, 2000.

[15] L. Younes. Optimal matching between shapes via elastic deformations. *Journal of Image and Vision Computing*, 17(5/6):381–389, 1999.