International Journal of Computer Vision 69(1), 137–144, 2006

DOI: 10.1007/s11263-006-6856-6

Contour Inferences for Image Understanding

WASHINGTON MIO
Department of Mathematics, Florida State University, Tallahassee, FL 32306
mio@math.fsu.edu

ANUJ SRIVASTAVA
Department of Statistics, Florida State University, Tallahassee, FL 32306
anuj@stat.fsu.edu

XIUWEN LIU
Department of Computer Science, Florida State University, Tallahassee, FL 32306
liux@cs.fsu.edu

Received April 1, 2004; Revised May 26, 2005; Accepted September 30, 2005

First online version published in April, 2006

Abstract. We present a new approach to the algorithmic study of planar curves, with applications to estimations of contours in images. We construct spaces of curves satisfying constraints suited to specific problems, exploit their geometric structure to quantify properties of contours, and solve optimization and inference problems. Applications include new algorithms for computing planar elasticae, with enhanced performance and speed, and geometric algorithms for the estimation of contours of partially occluded objects in images.

Keywords: curve interpolation, elastica, partial occlusions, shape completion

1. Introduction

The detection and recognition of objects in images are central problems in computer vision. Significant advances have been made in modeling appearances, but applications to recognition tasks only have found limited success. This is due, in part, to difficulties related to partial occlusions of objects and the high variability of observed pixel values. Thus, it is important to take into account additional global features of objects such as the shapes of their contours in order to improve the performance of image analysis algorithms. In this paper, we present a framework for the algorithmic study of planar curves and apply the methodology developed to the design of new algorithms for the completion of partially occluded contours of objects in images. Curves are treated as *continua* and discrete formulations are used only at the implementation level. Thus, representations of curves often involve infinitely many degrees of freedom, which lead us to the study of both finite and infinite-dimensional manifolds. In this paper, a fundamental element is a *space of constrained curves*, referred to as an SCC henceforth. As an example, if we are interested in analyzing closed contours that outline objects in images, the *closure condition* provides a constraint on curves.

A typical SCC is a *smooth manifold* equipped with a *Riemannian structure*; that is, a manifold with an inner product on each tangent space that varies smoothly along the space. We exploit such structures to develop

an algorithmic approach to gradient flows, and to solve optimization and statistical inference problems associated with a variety of cost functions and probability models on SCCs. The following problems will be addressed: (i) *completion of contours with elasticae*; (ii) *global estimation of contours using curves whose curvature functions are localized in the frequency domain*. Details are provided below.

Since the introduction of "snakes" in Kass et al. (1988), active contours have been the subject of numerous studies [see e.g. (Sethian, 1996; Chan et al., 2003)]. *Given the past success of PDE-based methods, why introduce new methodology?* The study of SCCs provides new insights and solutions to problems involving planar contours, as demonstrated here [see also (Klassen et al., 2004; Joshi et al., 2004; Mio et al., 2004b)]. An important new characteristic is that there is a well-defined, structured space of curves on which optimization and inference problems can be formulated and statistical models developed. In our approach, the dynamics of active contours is modeled on vector fields on SCCs. This treatment of curve evolution leads to algorithms with improved computational efficiency, as exemplified by our algorithm for computing planar *elasticae*.

We should also point out the main limitations of the proposed approach. One drawback is its inability to handle changes in topology, which is one of the key features of level-set methods. Secondly, it does not extend easily to the analysis of surfaces in 3D space or hypersurfaces in n-dimensional Euclidean space. Despite these limitations, the proposed methodology leads to powerful new geometric algorithms for the analysis of planar curves.

A word about the organization of the paper. In Section 2, we discuss representation of curves and the geometry of SCCs. In Section 3 we develop an algorithm to compute *elasticae* and apply it to the completion of contours in images. Section 4 is devoted to the study of *localization* of curvature and angle functions, which is applied to the resolution of partial occlusions in images in Section 5.

2. Representation of Planar Curves

We study planar curves presented in parametric form and traversed with constant speed. More precisely, we consider curves $\alpha\colon I \to \mathbb{R}^2$ parameterized over the unit interval $I = [0, 1]$ with $\|\alpha'(s)\|$ constant. Thus, if the length of α is L, then $\|\alpha'(s)\| = L$, for every $s \in I$. A continuous function $\theta\colon I \to \mathbb{R}$ is said to be an *angle function* for α if $\alpha'(s) = Le^{j\theta(s)}$, for every s, where $j = \sqrt{-1}$. Here, we are identifying $\mathbb{R}^2$ with the complex plane $\mathbb{C}$ in the usual manner. Notice that angle functions are only defined up to the addition of integer multiples of 2π. Moreover, they are invariant to scale and translations of the curve; the effect of a rotation is to add a constant to θ. The *rotation index* of α, which measures the total number of turns made by the tangent vector $\alpha'(s)$ as the curve α is traversed, is defined as $\iota(\alpha) = \frac{1}{2\pi}(\theta(1) - \theta(0))$ and is independent of the angle function chosen. The rotation index of a closed curve α with $\alpha'(0) = \alpha'(1)$ is an integer, with the sign depending on the orientation of the curve. If, in addition, α is a *simple* curve (i.e., with no self-intersections other than the end points) the rotation index is known to be ± 1 [see e.g. (Do Carmo, 1976)].

Associated with α, there is a normalized unit-speed curve $\beta\colon I \to \mathbb{R}^2$ obtained by scaling α to have unit length; i.e., $\beta(s) = \alpha(s)/L$. If $\alpha'(s) = Le^{j\theta(s)}$, then $\beta'(s) = e^{j\theta(s)}$. Thus, any angle function for α is also an angle function for β and $\kappa(s) = \theta'(s)$ is the curvature function of β. Since curvature scales inversely, the curvature of α at s is $\kappa(s)/L$.

In applications involving elastic energies, we can only allow curves with square integrable curvature functions. Thus, we only consider angle functions $\theta\colon I \to \mathbb{R}$ that can be expressed as integrals of square integrable functions. In more technical terms, we assume that angle functions are *absolutely continuous* [see e.g. (Royden, 1988)] with square integrable first derivatives. The space of all such functions equipped with the inner product, $\langle f, g\rangle_1 = f(0)g(0) + \int_0^1 f'(s)g'(s)\,ds$, is denoted $\mathbb{H}^1$. This inner product, introduced in Palais (1963), is a variant of the Sobolev inner product $\int_0^1 f(s)g(s)\,ds + \int_0^1 f'(s)g'(s)\,ds$.

2.1. A Manifold of Constrained Curves

We begin our discussion of SCCs by considering the space of all planar curves with square integrable curvature function, fixed rotation index, and given first-order boundary conditions. We prescribe the endpoints $p_0, p_1 \in \mathbb{R}^2$ and angles $\theta_0, \theta_1 \in \mathbb{R}$ that determine the tangent directions at the endpoints. Our goal is to construct an SCC, denoted $\mathcal{A}$, consisting of all curves α of rotation index $(\theta_1 - \theta_0)/2\pi$ satisfying $\alpha(i) = p_i$ and $\alpha'(i)/\|\alpha'(i)\| = e^{j\theta_i}$, for $i = 0, 1$.

Using a logarithmic scale for the length, write $L = e^{\ell}$. If θ is an angle function for α, the velocity vector can be written as $\alpha'(s) = e^{\ell}e^{j\theta(s)}$. Thus, curves

of length L satisfying $\alpha(0) = p_0$ can be expressed as $\alpha(s) = p_0 + e^{\ell} \int_0^s e^{j\theta(\sigma)}\, d\sigma$. Let $d = (d_1, d_2) = p_1 - p_0$ be the desired total displacement of α. Then, the condition $\alpha(1) = p_1$ can be written as $e^{\ell} \int_0^1 e^{j\theta(s)}\, ds = d_1 + j d_2$, or equivalently, as

$$e^{\ell} \int_0^1 \cos\theta(s)\, ds = d_1 \quad \text{and} \quad e^{\ell} \int_0^1 \sin\theta(s)\, ds = d_2. \tag{1}$$

Hence, there is a one-to-one correspondence between elements of $\mathcal{A}$ and pairs $(\ell, \theta) \in \mathbb{R} \times \mathbb{H}^1$ satisfying $\theta(0) = \theta_0$, $\theta(1) = \theta_1$ and (1). These four constraints define a codimension-4 submanifold of $\mathbb{R} \times \mathbb{H}^1$; i.e., a submanifold whose normal space at each point is four dimensional. We are assuming that the ambient space $\mathbb{R} \times \mathbb{H}^1$ is equipped with the inner product $\langle (v, f), (w, g) \rangle = v \cdot w + \langle f, g \rangle_1$, where $(v, f), (w, g) \in \mathbb{R} \times \mathbb{H}^1$.

Before proceeding with the calculations, we describe the space $\mathcal{A}$ as a level set, as this will shed some light on its geometric structure. Define a function $F\colon \mathbb{R} \times \mathbb{H}^1 \to \mathbb{R}^4$ by:

$$\begin{aligned} F^1(\ell, \theta) &= e^{\ell} \int_0^1 \cos\theta(s)\, ds; \\ F^2(\ell, \theta) &= e^{\ell} \int_0^1 \sin\theta(s)\, ds; \\ F^3(\ell, \theta) &= \theta(0); \quad F^4(\ell, \theta) = \theta(1). \end{aligned} \tag{2}$$

Then, $\mathcal{A}$ is the level set $F^{-1}(d_1, d_2, \theta_0, \theta_1)$. Note that F^3 and F^4 are linear functionals on $\mathbb{R} \times \mathbb{H}^1$ and can be expressed in terms of the inner product $\langle\, , \rangle$ as $F^3(\ell, \theta) = \langle (\ell, \theta), (0, 1) \rangle$ and $F^4(\ell, \theta) = \langle (\ell, \theta), (0, 1 + s) \rangle$, where 1 and s denote the constant function 1 and the function $s \mapsto s$, respectively.

2.2. *The Normal Bundle of $\mathcal{A}$*

We exhibit an explicit basis for the normal space to the level sets of F at any regular point $(\ell, \theta) \in \mathbb{R} \times \mathbb{H}^1$; that is, at any point where the derivative $dF\colon T_{(\ell,\theta)}\mathcal{A} \to \mathbb{R}^4$ is an onto mapping. Here, $T_{(\ell,\theta)}\mathcal{A}$ is the tangent space of $\mathcal{A}$ at the point (ℓ, θ). The derivative of F at (ℓ, θ) in the direction of the vector $(v, f) \in \mathbb{R} \times \mathbb{H}^1$ is given by:

$$\begin{aligned} dF^1(v, f) &= \langle (v, f), (F^1(\ell, \theta), e^{\ell} h_1) \rangle; \\ dF^2(v, f) &= \langle (v, f), (F^2(\ell, \theta), e^{\ell} h_2) \rangle; \\ dF^3(v, f) &= \langle (v, f), (0, 1) \rangle; \\ dF^4(v, f) &= \langle (v, f), (0, 1 + s) \rangle, \end{aligned} \tag{3}$$

where $h_1, h_2\colon I \to \mathbb{R}$ are characterized by $h_1''(s) = \sin\theta(s)$, $h_1(0) = h_1'(0) = 0$ and $h_2''(s) = -\cos\theta(s)$, $h_2(0) = h_2'(0) = 0$. This means that the gradient of F^i, $1 \le i \le 4$, is:

$$\begin{aligned} &\nabla F^1(\ell, \theta) = (F^1(\ell, \theta), e^{\ell} h_1); \\ &\nabla F^2(\ell, \theta) = (F^2(\ell, \theta), e^{\ell} h_2); \\ &\nabla F^3(\ell, \theta) = (0, 1); \quad \nabla F^4(\ell, \theta) = (0, 1 + s). \end{aligned} \tag{4}$$

Hence, a vector (v, f) is tangent to the level set of F at (ℓ, θ) if and only if it is orthogonal to the subspace of $\mathbb{R} \times \mathbb{H}^1$ spanned by $\{\nabla F^i(\ell, \theta), 1 \le i \le 4\}$. We thus have a basis of the normal space to the level set of F at any regular point (ℓ, θ). An orthonormal basis can be obtained using the Gram-Schmidt process.

In applications, the calculation of normal structures will be useful in the computation of the gradient of energy functionals on $\mathcal{A}$. One can usually proceed in two steps: first, calculate the gradient of the energy as a functional on $\mathbb{R} \times \mathbb{H}^1$, often an easier task. Then, subtract the normal component to $\mathcal{A}$ to obtain the gradient as a functional on $\mathcal{A}$.

2.3. *Projection onto $\mathcal{A}$*

In numerical simulations of the flow associated with vector fields on $\mathcal{A}$, we use a variant of Euler's method, as follows. Infinitesimally, first flow in the ambient linear space $\mathbb{R} \times \mathbb{H}^1$ in the direction of the vector field. The new point typically falls off of $\mathcal{A}$ due to its nonlinearity, so we devise a mechanism to project the point back onto $\mathcal{A}$ in order to iterate the procedure. Higher-order integration methods can be modified similarly.

Let $(\ell, \theta) \in \mathbb{R} \times \mathbb{H}^1$. The residual vector $r(\ell, \theta) = (d_1, d_2, \theta_0, \theta_1) - F(\ell, \theta) \in \mathbb{R}^4$ is zero precisely at points $(\ell, \theta) \in \mathcal{A}$. Thus, $\varepsilon(\ell, \theta) = \|r(\ell, \theta)\|^2$ gives a measurement of how far off (ℓ, θ) is from $\mathcal{A}$ and achieves its minimum value 0 at points in $\mathcal{A}$. We project (ℓ, θ) onto $\mathcal{A}$ by following the negative gradient flow of the functional ε. The gradient of ε is given by $\nabla\varepsilon(\ell, \theta) = -\sum_{i=1}^4 r_i(\ell, \theta) \nabla F^i(\ell, \theta)$, with $r = (r_1, r_2, r_3, r_4)$ and ∇F^i as in (4). Hence, the projection of a point $(\ell_*, \theta_*) \in \mathbb{R} \times \mathbb{H}^1$ onto $\mathcal{A}$ is obtained by asymptotically following the solution of the initial

value problem

$$(\dot{\ell}, \dot{\theta}) = \sum_{i=1}^{4} r_i(\ell, \theta)\, \nabla F^i(\ell, \theta), \qquad \ell(0) = \ell_*,\ \theta(0) = \theta_*. \tag{5}$$

Alternatively, a version of Newton's method can be used to search for the zeros of ϵ, as in Klassen et al. (2004) and Mio et al. (2004a).

Algorithm 1 (Projection onto $\mathcal{A}$). Given a pair $(\ell, \theta) \in \mathbb{R} \times \mathbb{H}^1$, this algorithm projects it onto $\mathcal{A}$, for a given choice of d_1, d_2, θ_0 and θ_1. Let $\epsilon, \delta > 0$ be small.

1. Compute $F(\ell, \theta)$ using Eq. (2) and the residual vector $r(\ell, \theta) = (d_1, d_2, \theta_0, \theta_1) - F(\ell, \theta)$. If $\|r(\ell, \theta)\| < \epsilon$, then stop. Else, continue.
2. Use Eqs. (4) and (6) to compute $\nabla F^i(\ell, \theta)$, $1 \leq i \leq 4$, and $(\dot{\ell}, \dot{\theta})$.
3. Update $(\ell, \theta) = (\ell, \theta) + \delta(\dot{\ell}, \dot{\theta})$. Return to Step 1.

3. Scale-Invariant Planar Elasticae

The *elastica model* for the estimation of partially hidden contours in images was introduced in Horn (1983), but the study of the elastic energy of curves dates back to Euler (1744). A stochastic interpretation of elasticae as the most likely curves according to a probability model based on length and (bending) elastic energy was given in Mumford (1994). In Weiss (1988) and Bruckstein and Netravali (1990), it was argued that scale-invariant models are more natural in vision problems and the notion of *scale-invariant elastica* was introduced. Further studies of elasticae and associated stochastic completion fields were carried out in Williams and Jacobs (1997), but calculations were computationally very intensive. Algorithmically efficient approximations were investigated in Sharon et al. (2000) under the simplifying assumption that 2D elasticae could be described as graphs of functions. More recently, geometric algorithms to compute elasticae in general Euclidean spaces were developed in Mio et al. (2004a). As an application of SCCs, in this paper, we present an algorithm to compute planar elasticae. By restricting the analysis to the 2D case, not only we simplify the algorithm of Mio et al. (2004a), but also improve computational efficiency. Completions with *Euler spirals* have been studied in Kimia et al. (2003).

3.1. The Elastic Energy Flow

Let $(\ell, \theta) \in \mathcal{A}$ represent a curve $\alpha: I \to \mathbb{R}^2$ of length $L = e^{\ell}$ and let $\kappa(s) = \theta'(s)$. The (bending) *elastic energy* of α is the integral of the square of the curvature function κ/L with respect to the arc-length parameter. Since the arc-length element is $L\,ds$, the elastic energy of α can be expressed as $\frac{1}{2L}\int_0^1 \kappa^2(s)\,ds$. The *scale-invariant elastic energy* of α (Bruckstein and Netravali, 1990) is given by

$$E_{si}(\ell, \theta) = \frac{1}{2}\int_0^1 \kappa^2(s)\,ds. \tag{6}$$

The main goal of this section is to describe an algorithm to calculate curves of least scale-invariant elastic energy with prescribed rotation index and first-order boundary conditions; these curves are known as *scale-invariant elasticae*. Equivalently, we wish to minimize the functional E_{si} on the manifold $\mathcal{A}$. From (6), the differential of E_{si} (as a functional on $\mathbb{R} \times \mathbb{H}^1$) at the point (ℓ, θ) evaluated in the direction (v, f) is $dE_{si}(v, f) = \int_0^1 f'(s)\theta'(s)\,ds = \langle (v, f), (0, \theta - \theta_0) \rangle$. Thus, $\nabla_{\mathbb{R}\times\mathbb{H}^1} E_{si}(\ell, \theta) = (0, \theta - \theta_0)$. Using the calculation of normal spaces to $\mathcal{A}$ carried out in Section 2.2, we project this vector orthogonally onto the tangent space $T_{(\ell,\theta)}\mathcal{A}$ to obtain $\nabla_{\mathcal{A}} E_{si}(\ell, \theta)$. Flow lines of the negative gradient vector field on $\mathcal{A}$ associated with E_{si} approach scale-invariant elasticae, asymptotically.

3.2. Algorithms and Experimental Results

Algorithm 2 (Finding Elasticae). Given boundary conditions $(d_1, d_2; \theta_0, \theta_1)$, find the corresponding scale-invariant elastica. Initialize (ℓ, θ) with a straight line segment between the end points.

1. Project (ℓ, θ) onto $\mathcal{A}$ using Algorithm 1. Compute $e_i \equiv \nabla F^i(\ell, \theta)$, for $i = 1, \ldots, 4$, according to Eq. (4). Transform $\{e_1, e_2, e_3, e_4\}$ into an orthonormal basis using the Gram-Schmidt process.
2. Project the gradient vector $(0, \theta - \theta_0)$ orthogonally onto the tangent space using: $(d\ell, d\theta) \equiv (0, \theta - \theta_0) - \sum_{i=1}^{4} \langle (0, \theta - \theta_0), e_i \rangle$. If $\|(d\ell, d\theta)\| < \epsilon$, stop. Else, continue.
3. For a step size $\delta > 0$, perform the update: $(\ell, \theta) = (\ell, \theta) + \delta(d\ell, d\theta)$. Return to Step 1.

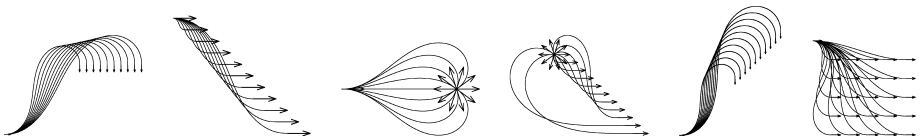

Figure 1. Examples of scale-invariant planar elasticae.

Figure 2. Examples of partial occlusions in natural images resolved with scale-invariant elasticae are shown as solid lines. For comparison, dashed lines show hand-drawn completions.

Figure 1 shows several examples of planar, scale-invariant elasticae generated using Algorithm 2. In our current implementation, we compute an average of 50 elasticae per second on a Pentium IV processor.

Examples of partial occlusions of contours in natural images resolved with scale-invariant elasticae are shown in Fig. 2. The curves drawn as solid lines are elasticae computed with the algorithm developed in this section. For comparison purposes, dashed lines show hand-drawn contours.

4. Localization of Curvature and Angle Functions

In this section, we study an SCC of closed curves with fixed length L, rotation index 1, localized curvature functions, and prescribed first-order initial conditions. By *localization*, we mean that curvature functions will be restricted to some subspace $W \subseteq \mathbb{L}^2$ such as that spanned by certain dominant harmonics. We carry out experiments with localization in the frequency domain, but the discussion is more general.

4.1. V-Local Curves

All curves will have the same initial point $p_0 \in \mathbb{R}^2$ and initial value $\theta(0) = \theta_0$ of the angle function. Thus, angle functions will be of the form $\theta(s) = \theta_0 + \varphi(s)$, where $\varphi: I \to \mathbb{R}$ satisfies $\varphi(0) = 0$. In Section 2, curves were represented by pairs $(\ell, \theta) \in \mathbb{R} \times \mathbb{H}^1$. For simplicity, we assume that the length $L = e^{\ell}$ is fixed, so we just need the variable θ, or equivalently, φ. Thus, the curve associated with $\theta \in \mathbb{H}^1$ is given by $\alpha(s) = p_0 + L \int_0^s e^{j\theta(t)}\, dt$.

We wish to localize curvature functions and consequently angle functions, as well. What subspaces are we going to restrict angle functions to? This is equivalent to asking what restrictions to impose on φ. We formulate these restrictions in terms of the $\mathbb{H}^1$ inner product.

Since $\langle \varphi, 1 \rangle_1 = \varphi(0)$, functions $\varphi \in \mathbb{H}^1$ satisfying $\varphi(0) = 0$ are those orthogonal to constant functions. They can be written as $\varphi(s) = \int_0^s \varphi'(\xi)\, d\xi$. Hence, we localize angle functions as follows: pick a finite-dimensional subspace $W \subset \mathbb{L}^2$, with orthonormal basis $\{\sigma_1, \ldots, \sigma_n\}$; this is to be interpreted as localization of the curvature function $\kappa(s) = \theta'(s) = \varphi'(s)$ to W. If

$$\varphi_i(s) = \int_0^s \sigma_i(t)\, dt, \tag{7}$$

we restrict φ to the subspace $V = \text{span}\,\{\varphi_1, \ldots, \varphi_n\} \subset \mathbb{H}^1$. We only consider localization to finite-dimensional subspaces, but the construction can be extended to other (closed) subspaces.

Example. Let $W \subset \mathbb{L}^2$ be the subspace spanned by the truncated Fourier basis

$$\{1, \sqrt{2}\cos(2\pi s), \ldots, \sqrt{2}\cos(2k\pi s), \sqrt{2}\sin(2\pi s), \ldots, \sqrt{2}\sin(2k\pi s)\}.$$

If $\{\sigma_i\}_{i=1}^{2k+1}$ are these basis elements, use (7) to obtain a basis $\{\varphi_i\}_{i=1}^{2k+1}$ of V. This choice of W corresponds to localizing curvature functions to the first $2k + 1$ harmonics.

We restrict our attention to V-local angle functions; that is, those of the form

$$\theta(s) = \theta_0 + x_1\varphi_1(s) + \cdots + x_n\varphi_n(s), \tag{8}$$

where $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$. Thus, curves α of fixed length $L = e^{\ell}$, localized angle functions and given first-order initial conditions can be represented by a vector $x \in \mathbb{R}^n$ via $\alpha(s) = p_0 + L\int_0^s e^{j\left(\theta_0 + \sum_{i=1}^n x_i\varphi_i(\tau)\right)}\,d\tau$. If $f = \sum_{i=1}^n x_i\varphi_i(s)$ and $g = \sum_{i=1}^n y_i\varphi_i(s)$, then $\langle f, g\rangle_1 = x_1y_1 + \cdots + x_ny_n$. Thus, the inner product $\langle f, g\rangle_1$ can be expressed in terms of x and y as the standard inner product $x \cdot y = x_1\,y_1 + \cdots + x_n\,y_n$ in $\mathbb{R}^n$.

4.2. A Manifold of Localized Closed Curves

We are interested in closed curves with rotation index 1. Therefore, analogous to the constraints on pairs (ℓ, θ) imposed by the function F, defined in (3), we consider the map $G\colon \mathbb{R}^n \to \mathbb{R}^3$ given by:

$$G^1(x) = \int_0^1 \cos\theta(s)\,ds\,; \quad G^2(x) = \int_0^1 \sin\theta(s)\,ds\,;$$
$$G^3(x) = \theta(1) - \theta_0\,, \tag{9}$$

with θ as in (8). The relevant SCC of V-local closed curves is the iso-set $\mathcal{H}_V = G^{-1}(0, 0, 2\pi)$. For almost all choices of V, $\mathcal{H}_V$ is a smooth $(n-3)$-dimensional submanifold of $\mathbb{R}^n$. As in Section 2, a computation of the normal structure to level sets of the mapping G is needed for projections of points and vector fields in the ambient space $\mathbb{R}^n$ onto $\mathcal{H}_V$. The mechanism is identical to that discussed earlier, so we just provide expressions for ∇G:

$$\nabla G^1(x) = -\left(\int_0^1 \varphi_1(s)\sin\theta(s)\,ds, \ldots, \int_0^1 \varphi_n(s)\sin\theta(s)\,ds\right);$$
$$\nabla G^2(x) = \left(\int_0^1 \varphi_1(s)\cos\theta(s)\,ds, \ldots, \int_0^1 \varphi_n(s)\cos\theta(s)\,ds\right);$$
$$\nabla G^3(x) = (\varphi_1(1), \ldots, \varphi_n(1)). \tag{10}$$

Algorithm 3. [*Projection onto* $\mathcal{H}_V$]. Given a point $x \in \mathbb{R}^n$, this algorithm projects it onto $\mathcal{H}_V$. Choose $\epsilon, \delta > 0$ small.

1. Compute G(x) according to Eq. (9) and the residual vector $r(x) = (0, 0, 2\pi) - G(x)$. If $\|r(x)\| < \epsilon$, then stop. Else, continue.
2. Compute $\nabla^i G(x)$, for $i = 1, \ldots, 3$, according to Eq. (10) and let $dx = \sum_{i=1}^3 r_i(x)\nabla G^i(x)$.
3. Perform the update: $x = x + \delta dx$. Return to Step 1.

Applications of this algorithm to contour inference in images will be discussed in the next section.

5. Discovery of Contours in Images

The resolution of partial occlusions of contours with elasticae only yields good results in situations where the hidden parts are small or not very rich in geometric features. This is because elasticae only take first-order boundary information into account and are designed to minimize the bending elastic energy. If examples of likely shapes are known, algorithms can be trained to use this information, e.g., as in Cootes et al. (1995); Klassen et al. (2004). However, if additional contextual knowledge is unavailable, the problem of discovering hidden shapes is very challenging. In such generality, it might be even difficult to develop meaningful quantitative criteria to measure the "goodness" of a proposed completion.

We propose to use localization, in a Bayesian framework, to produce completions that take into consideration more information about the global geometry of the visible parts. The basic idea is to *restrict the search to curves whose curvature (or angle) functions contain just enough harmonics to capture the geometry of the visible portions*, so that observed geometric patterns will propagate to the parts to be estimated without high energy costs. Examples will illustrate the fact that, in the presence of periodic patterns, this technique leads to improved contour estimations.

5.1. Bayesian Estimation of Contours

To simplify the discussion, we assume that the visible portion of the contour of an imaged object has been extracted and consists of a single arc. Hence, we work under the hypothesis that the data is presented as a parametric curve γ. Also, we suppose that the length of

the part of the contour to be estimated is known. This implicitly assumes that, in practice, some additional contextual knowledge is available.

Given an open curve $\gamma\colon[0, L_0] \to \mathbb{R}^2$, $L_0 > 0$, with unit-speed parameterization, we would like to find an "optimal completion" of γ to a closed curve α of rotation index 1 and length $L > L_0$. For the problem to be interesting, we assume that $L > \|\gamma(L_0) - \gamma(0)\| + L_0$. We let $\theta_\gamma\colon[0, L_0] \to \mathbb{R}$ be an angle function for the curve γ, $\theta_0 = \theta_\gamma(0)$, and $p_0 = \gamma(0)$.

We take a Bayesian approach to the estimation of contours. We adopt a model having the scale-invariant elastic energy $E(x) = \frac{1}{2}\sum_{i=1}^{n} x_i^2$ as prior energy and data likelihood energy $F(\gamma \mid x) = \frac{1}{2}\int_0^{\zeta_0} \big(\theta(t) - \theta_\gamma(Lt)\big)^2\, dt$, where $\zeta_0 = L_0/L$. The functional $F(\gamma|x)$ quantifies the consistency of θ_γ with the angle function represented by $x \in \mathbb{R}^n$ over the interval $[0, L_0]$. For each λ, $0 < \lambda < 1$, the posterior energy is proportional to the functional $E_\lambda(x|\gamma) = \lambda\, F(\gamma|x) + (1 - \lambda)E(x)$. We refer to a *maximum-a-posteriori* (MAP) estimation of the contour according to E_λ as a V-local λ-*completion* of γ. Equivalently, a λ-completion of γ is a curve represented by $x \in \mathcal{H}_V \subset \mathbb{R}^n$ that minimizes the posterior energy E_λ on $\mathcal{H}_V$.

The search for V-local λ-completions requires that we solve an optimization problem on $\mathcal{H}_V$, which we approach with gradient methods. To simulate the negative gradient flow of E_λ on $\mathcal{H}_V$, we need to compute the partial derivatives of E_λ. For each $x \in \mathbb{R}^n$, let $Q_x(s) = \theta(s) - \theta_\gamma(Ls)$, $0 \le s \le \zeta_0 = L_0/L$. Then,

$$\frac{\partial F}{\partial x_i} = \int_0^{\zeta_0} Q_x(s)\varphi_i(s)\, ds \qquad \text{and} \qquad \frac{\partial E}{\partial x_i} = x_i. \tag{11}$$

Thus,

$$\nabla_{\mathbb{R}^n} E_\lambda(x) = \lambda\left(\frac{\partial F}{\partial x_1}, \ldots, \frac{\partial F}{\partial x_n}\right) + (1 - \lambda)\,(x_1, \ldots, x_n). \tag{12}$$

Projecting this vector orthogonally onto the tangent space $T_x\mathcal{H}_V$, we obtain $\nabla_{\mathcal{H}_V} E_\lambda(x)$. Flow lines of the negative gradient field on $\mathcal{H}_V$ associated with E_λ approach λ-completions of γ asymptotically.

5.2. *Algorithms and Experimental Results*

Algorithm 4. (Finding V-Local Completions of Contours). Given the visible curve γ and a basis $\{\varphi_i\}$ of V, find a local closed curve that minimizes the posterior energy. L_0 is the length of γ and L is the chosen length of its completion. Initialize $x \in \mathbb{R}^n$.

1. Project x onto $\mathcal{H}_V$ using Algorithm 3. Compute $e_i \equiv \nabla H^i(x)$, for $i = 1, 2, 3$, and make them orthonormal in $\mathbb{R}^n$ using the Gram-Schmidt procedure.

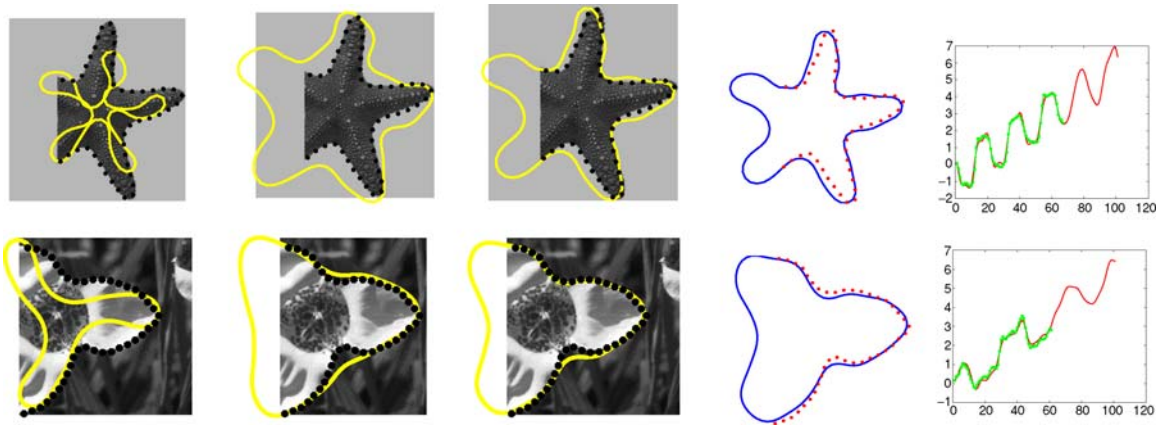

Figure 3. Curve evolution during MAP estimations of partially occluded contours.

2. Compute the gradient vector $\nabla_{\mathbb{R}^n} E_\lambda(x)$ using Eqn. (12) and project it onto the tangent space using: $dx \equiv \nabla_{\mathbb{R}^n} E_\lambda(x) - \sum_{i=1}^{3} \langle \nabla_{\mathbb{R}^n} E_\lambda(x), e_i \rangle e_i$. If $\|dx\| < \epsilon$, stop. Else, continue.
3. For a step size $\delta > 0$, perform the update: $x = x + \delta\, dx$. Return to Step 1.

Figure 3 shows contour estimations with this algorithm. On the first row, the part of the contour of the starfish shown as a dotted line was hand extracted and represents the curve γ. To obtain a coarse estimation of the subspace V, we truncated the Fourier expansion of θ_γ keeping only the dominant harmonics and used the same fundamental harmonics to choose V. The first three panels show the evolution of the curve during the gradient search and the fourth panel shows the MAP estimation of the contour. The fifth panel on shows a plot of the angle function of γ and that of the estimated contour. The result of a similar Bayesian contour estimation experiment with an image of a flower is shown on the second row.

6. Summary and Conclusions

We investigated spaces of planar curves satisfying various linear and non-linear constraints, as well as methods for solving optimization and inference problems on these spaces. The methodology developed was applied to the design of new algorithms for computing planar elasticae, and to the Bayesian estimation of hidden contours in images. We concluded that planar elasticae can be computed efficiently without some restrictive assumptions adopted in the past, and that frequency localization techniques allows us to better infer occluded contours exhibiting periodic patterns.

Acknowledgments

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

References

Bruckstein, A. and Netravali, N. 1990. On minimal energy trajectories. *Computer Vision, Graphics and Image Processing*, 49:283–296.

Chan, T., Shen, J., and Vese, L. 2003. Variational PDE models in image processing. *Notices Amer. Math. Soc.*, 50:14–26.

Cootes, T.F., Cooper, D., Taylor, C.J., and Graham, J. 1995. Active shape models—their training and application. *Computer Vision and Image Understanding*, 61:38–59.

Do Carmo, M.P. 1976. *Differential Geometry of Curves and Surfaces*, Prentice Hall, Inc.

Euler, L. 1744. Methodus inveniendi lineas curvas maximi minimive proprietate gaudentes, sive solutio problematis isoperimetrici lattisimo sensu accepti. *Bousquet, Lausannae e Genevae, E65A. O. O. Ser. I*, 24.

Horn, B. 1983. The curve of least elastic energy. *ACM Trans. Math. Software*, 9:441–460.

Joshi, S., Srivastava, A., Mio, W., and Liu, X. 2004. Hierarchical organization of shapes for efficient retrieval. In *Proc. ECCV 2004, LNCS*, Prague, Czech Republic, pp. 570–581.

Kass, M., Witkin, A., and Terzopoulos, D. 1988. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331.

Kimia, B., Frankel, I., and Popescu, A. 2003. Euler spiral for shape completion. *International Journal of Computer Vision*, 54:159–182.

Klassen, E., Srivastava, A., Mio, W., and Joshi, S. 2004. Analysis of planar shapes using geodesic paths on shape spaces. *Trans. Pattern Analysis and Machine Intelligence*, 26:372–383.

Mio, W., Srivastava, A., and Klassen, E. 2004a. Interpolations with elastica in euclidean spaces. *Quarterly of Applied Mathematics*, 62:359–378.

Mio, W., Srivastava, A., and Liu, X. 2004b. Learning and bayesian shape extraction for object recognition. In *Proc. ECCV 2004, LNCS*, Prague, Czech Republic, pp. 62–73.

Mumford, D. 1994. *Elastica and Computer Vision*, Springer, New York. pp. 491–506.

Palais, R.S. 1963. Morse theory on hilbert manifolds. *Topology*, 2:299–340.

Royden, H. 1988. *Real Analaysis*, Prentice Hall.

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

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

Weiss, I. 1988. 3D shape representation by contours. *Computer Vision, Graphics and Image Processing*, 41:80–100.

Williams, L. and Jacobs, D. 1997. Local parallel computation of stochastic completion fields. *Neural Computation*, 9:837–858.