Published in final edited form as:

Int J Comput Vis. 2009 March 1; 81(3): 331–355. doi:10.1007/s11263-008-0179-8.

Intrinsic Bayesian Active Contours for Extraction of Object Boundaries in Images

Shantanu H. Joshi and

Laboratory of Neuroimaging, University of California, Los Angeles, CA 90024, USA

Anuj Srivastava

Department of Statistics, Florida State University, Tallahassee, FL 32306, USA

Shantanu H. Joshi: sjoshi@loni.ucla.edu

Abstract

We present a framework for incorporating prior information about high-probability shapes in the process of contour extraction and object recognition in images. Here one studies shapes as elements of an infinite-dimensional, non-linear quotient space, and statistics of shapes are defined and computed intrinsically using differential geometry of this shape space. Prior models on shapes are constructed using probability distributions on tangent bundles of shape spaces. Similar to the past work on active contours, where curves are driven by vector fields based on image gradients and roughness penalties, we incorporate the prior shape knowledge in the form of vector fields on curves. Through experimental results, we demonstrate the use of prior shape models in the estimation of object boundaries, and their success in handling partial obscuration and missing data. Furthermore, we describe the use of this framework in shape-based object recognition or classification.

Keywords

Shape extraction; Segmentation; Bayesian shape extraction; Tangent PCA; Intrinsic shape analysis; Elastic shapes; Riemannian metric

1 Introduction

Appearances of objects in images can be characterized to a certain extent by shapes of their boundaries. Therefore, the task of extracting object boundaries in images may prove important in problems of detection, tracking, and recognition of objects using image data. In some applications this task has to be performed solely on the basis of images (data) while in others some a-priori information about the shape of the boundary (to be extracted) is available to the algorithm. Boundary extraction has been studied in detail for many years but mostly in the former situation. Focusing on the latter, we present a method for *representing, modeling, and incorporating prior information* about shapes of closed curves for use in the process of shape extraction. Examples of these situations include battlefield image analysis where the interest is restricted to certain military vehicles, or medical image analysis where one focuses on certain anatomical parts with known shape variability. Stated differently, we are not seeking “general purpose” segmentation algorithms; our inferences will correspond to curves that form boundaries of known and desired objects in images. From this perspective, our technique can be viewed as a preliminary step in shape-based inferences—classification and recognition of

objects—in images. Besides being application oriented, this perspective provides another advantage over the general-purpose segmentation algorithms: a criterion for evaluating boundary estimation performance. Since estimated boundaries are to be used for future inferences, performance in boundary estimation can be measured directly in terms of eventual object classification or recognition performance.

To highlight the novelty and contributions of this paper, we first discuss past and present methodologies in this area.

1.1 Current Methods for Curve Extraction

The seminal paper by Kass et al. (1988) formulated a snake evolving over a landscape of the intensity image, guided by its edges and corners. This model, in its original proposed form, had many limitations such as its limited capture range and preference to local solutions. As a remedy, Cohen and Cohen (1993) included a balloon force which expands the curve until it rests on edges. The Gradient Vector Flow (GVF) (Xu and Prince 1998b) and generalized GVF (GGVF) snakes (Xu and Prince 1998a), have gained popularity due to their ability to attract active contours toward object boundaries from large distances and move contours into object cavities. However, these snakes may fail to converge at weak edges, especially at the locations where strong and weak edges are in close proximity. Li et al. (2005) proposed an Edge Preserving GVF (EPGVF) to overcome this drawback; EPGVF prevents snakes from crossing over weak boundaries, while preserving the desired properties of GVF and GGVF. Caselles et al. (1997) proposed geodesic active contours to detect boundaries in images by searching for minimal-length curves in a space with a suitable Riemannian metric. Kichenassamy et al. (1995) formulated a geometric curve evolution approach based upon curvature flows in images. Both of these approaches use Euclidean curve-shortening flows that evolve a contour according to a normal vector field that is proportional to its curvature. A desirable property of these curvature flows, also called Euclidean heat flows, is the naturally occurring smoothing of the curve (Gage and Hamilton 1986; Grayson 1987). Additionally, curve evolution has also been accomplished using level-set representations, where curves are level sets of an appropriate function such as a signed-distance function (Caselles et al. 1997; Kichenassamy et al. 1995). While level-set representations allow for a change of topologies during curve evolution, an important strength of this approach, it is quite difficult to explicitly characterize shapes in this representation. Staib and Duncan (1992) use (parametric) deformable contours based on the elliptic Fourier decomposition of boundaries for image segmentation. In all these methods, the prior term, if used at all, regulates only the smoothness of the extracted curves; no specific shape information is incorporated in the extraction process.

1.2 Ideas in Shape Analysis

Shape is a characteristic that is invariant to rigid motion and uniform scaling. How can one represent, model, and incorporate a-priori information about high-probability shapes? A major effort in shape analysis has been on a finite-dimensional or “landmark-based” approach, where shapes are represented by a coarse, expert-driven sampling of objects (Dryden and Mardia 1998; Small 1996). This process however requires expert intervention as the automatic detection of landmarks is not straightforward. Additionally, since the analysis depends heavily on chosen landmarks, this approach is limited in its scope. Grenander's formulation (Grenander 1993; Miller and Younes 2001) considers shapes as points on infinite-dimensional manifolds, where the variations between the shapes are modeled by the action of diffeomorphism groups. Although this approach is relevant for modeling variations in full appearances of objects, i.e. both shapes and textures, its computational complexity makes it expensive for problems limited to curves and their shapes.

There are two general approaches for performing statistical analysis on nonlinear manifolds, including the shape manifolds: extrinsic and intrinsic. We briefly introduce these terms and refer the reader to Patrangenaru and Bhattacharya (2003) for a more detailed discussion.

Extrinsic Analysis—In this approach a manifold is embedded inside a larger vector space, the statistical analysis is performed in this larger space, and the final inference is projected back on to the manifold. For instance, sample means are computed in the larger space, using classical tools, and projected back to the manifold. While extrinsic analysis enjoys the advantage of being simple and computationally efficient, it is generally not the most natural. The main disadvantage is that solutions often depend on the choices of the ambient space and the embedding. As a simple example of this issue, consider a unit circle that can be embedded inside $\mathbb{R}^2$ in multiple ways where each embedding possibly leads to a different value of the extrinsic mean on the circle. The active shape model approach (Cootes et al. 1995) is an extrinsic approach that uses principal component analysis (PCA) of shapes, represented via vectors of landmarks, to model shape variability. Despite its simplicity and efficiency, this approach is limited because it ignores the nonlinear geometry of the underlying shape space. A similar idea is to use level sets (e.g. of signed-distance functions) to represent contours, and to compute PCA in the space of signed distance functions (Leventon et al. 2000).

Intrinsic Analysis—An intrinsic analysis is restricted to the manifold and uses quantities such as distances, probability measures, and moments, that are defined strictly on that manifold. For a Riemannian structure imposed on that manifold, the resulting distances between points are achieved by paths (geodesics) that lie completely on the manifold. In most formulations, shape spaces are quotient spaces of nonlinear manifolds, and one needs to apply tools derived from differential geometry to carry out intrinsic analysis on the underlying shape spaces. There has been a resurgence of ideas in studying shapes of closed, planar curves (Charpiat et al. 2005; Klassen et al. 2004; Michor and Mumford 2006; Mio et al. 2007; Shah 2005) that use differential geometry of manifolds. There has also been substantial progress in developing intrinsic (parametric) statistical models on such spaces (Srivastava et al. 2005a, 2005b).

1.3 Bayesian Shape Extraction

Several papers have presented approaches for extracting object boundaries using priors on either shapes or some shape related quantities. Grenander et al. (1991, 2000) used shape models learned from past observations to help improve shape estimation in future data. Cremers et al. (2003, 2006) and Rousson and Paragios (2002) have used priors in several formulations, including the level set formulation, to perform segmentation and object extraction.

There are several important criteria, such as invariance, shape metric, intrinsic nature, the use of higher moments, and the computational cost, in evaluating different Bayesian shape estimation methods. We motivate these criteria with a commonly used method. Consider the following way of comparing shapes of two closed curves, either directly or through their level sets. Let $\phi(\beta)$ denote the signed-distance function of a closed, planar curve β . Then, possible distances between shapes of two closed curves β_1 and β_2 are:

$$\begin{aligned} d(\beta_1, \beta_2) &= \min_{R \in SO(2), T \in \mathbb{R}^2, a \in \mathbb{R}_+} \|\phi(\beta_1) - \phi(R \cdot a\beta_2 + T)\|^2, \\ &\text{or} \min_{R \in SO(2), T \in \mathbb{R}^2, a \in \mathbb{R}_+} \|\beta_1 - (R \cdot a\beta_2 + T)\|^2 \end{aligned} \quad (1)$$

where R denotes a rotation, T denotes a translation, and a denotes a scale. These $\mathbb{L}^2$ -type norms are computed on ambient vector spaces and not on spaces of closed curves or spaces of signed-

distance functions. For example, in the case of signed-distance functions, $\|\phi_1 - \phi_2\|$ (or its version with the Heaviside step function) is often given by the $\mathbb{L}^2$ norm (Vemuri and Chen 2003; Cremers et al. 2002). With this example, we discuss different evaluation criteria for Bayesian shape estimation.

1. Level of Invariance to Similarity Transformations—Shape is traditionally considered invariant to the action of translation, rotation, parameterization, and scaling, and one would like a representation that removes as many of these groups as possible (ideally all of them) before analysis. The remaining groups have to be removed through computational means by posing them as optimization problems. In certain approaches, such as level-set methods or diffeomorphism based methods, it is difficult to build representations that are maximally invariant. For example, the formulation in (1) tends to be expensive because all the similarity transformations are still present in the representation and have to be removed using an optimization method.

2. Intrinsic or Extrinsic Analysis—Once the representation is chosen and the shape space is established, the ensuing shape analysis can be performed in either an extrinsic or intrinsic fashion. As mentioned earlier, while the intrinsic analysis is more natural with unique solutions, the extrinsic analysis is often the simpler one. A recent paper (Cremers et al. 2006) has obtained several desired invariances for shape analysis using level set representations, but it remains an extrinsic approach. One consequence is that in the resulting probability models, one will encounter functions that are invalid for representing shapes, with high probability. Another consequence is that shape statistics defined extrinsically will not be unique.

3. Choice of Shape Metric—While several metrics are available for comparing shapes, it is preferable to use one that lends to physical interpretations such as bending and stretching of curves. The elastic metric, defined later, is one such preferred metric as it provides a better matching of features through stretching, compression, and bending. However, despite its wide acceptance for shape comparisons, it has not been previously used in Bayesian shape estimation.

4. Use of Higher Moments—Most of the early work in shape priors involved terms of the type $\exp(-d^2)$, where d is a distance between shapes. However, a more interesting idea is to involve higher moments, such as the covariance, in building a shape prior. To our knowledge, covariances have seldom been used in past literature on shape extraction.

Our framework has the advantage of having: (i) a maximal invariance to shape preserving transformations, (ii) an intrinsic shape analysis framework, (iii) an attractive elastic metric, and (iv) the ability to use covariance information in the shape prior. As a combination of these properties, the resulting solution is advantageous, although it may be computationally more complex than certain methods that employ non-elastic metrics or extrinsic solutions.

1.4 Our Approach: Bayesian Active Contours

The main contribution of this paper is to apply fully intrinsic as well as maximally invariant prior probability models on well-defined shape spaces to the problem of boundary extraction. A priori knowledge of shape assumes importance when the image quality is low, perhaps due to low contrast or excess clutter, or when parts of the object have been obscured. Our goal is to use intrinsic, geometric representations of shapes of continuous, closed, planar curves in imposing shape priors and using these priors in a Bayesian framework for shape extraction.

Numerous ideas proposing different shape representations of closed curves and their intrinsic analysis have been presented in the literature (Joshi et al. 2007b; Klassen et al. 2004; Michor

and Mumford 2006; Mio et al. 2007; Shah 2005; Yezzi and Mennucci 2005) and others referenced therein. The basic idea is to represent curves denoting object boundaries as parameterized functions (not necessarily by arc-length), appropriately constrained, resulting in a nonlinear manifold $\mathcal{C}$. Most of the similarity transformations are already removed, but to remove the remaining transformations one forms a quotient space $\mathcal{S} = \mathcal{C}/\mathcal{S}$, where $\mathcal{S}$ is the space of the remaining similarity transformations. Consequently, the shapes of closed curves are analyzed as elements of $\mathcal{S}$. The main tool in this analysis is the choice of a Riemannian metric and the construction of geodesic paths on $\mathcal{S}$ between shapes under that metric. While several metrics have been proposed in the literature, the *elastic* metric is gaining popularity lately. Under the elastic metric, the curves are matched to each other using an optimal combination of bending and stretching/compression. This metric was first proposed by Younes (1998), advanced by Mio et al. (2004, 2007), and more recently simplified by Joshi et al. (2007b). The difference between these papers is in the choice of representation of curves.

Using the elastic metric, one can develop a framework for statistical models by defining means, covariances, and other statistics in tangent spaces of $\mathcal{S}$ (Srivastava et al. 2005b; Vaillant et al. 2004). The problem of boundary extraction can then be posed as a maximum a-posteriori (MAP) estimation, and rewritten in terms of an energy formulation as:

$$\widehat{\beta} = \arg \min_{\beta} (E_{image}(\beta) + E_{smooth}(\beta) + E_{prior}(\beta)), \quad (2)$$

where the minimization is over all closed curves in the given image domain. To ensure a fully intrinsic formulation, all the terms will be independent of parameterizations of β . Additionally, E_{smooth} will be independent of translation and rotation, and E_{prior} will be independent of translation, rotation, and a global scale of β as well.

E_{prior} is the novel term representing prior energy associated with the shape of curve β , while E_{image} and E_{smooth} are standard terms that are used frequently for driving *active contours*. The addition of E_{prior} results in the formulation of *Bayesian active contours*. We term this approach Bayesian even though a formal prior probability distribution on $\mathcal{S}$ is not presented explicitly. Instead, we either sample from the shape prior or use its gradient or both. The proposed framework has the following advantages:

1. In addition to image information and roughness constraints, it incorporates a priori information about shapes in boundary extraction. This additional information is presented using change of variables in the form of deformation vector fields on curves for efficient energy minimization.
2. This approach does not require expert-generated landmarks, diffeomorphic embeddings, or level-set functions for shape analysis. Instead, it studies the set of shapes of curves directly. Since the computations here are performed only on curves, this analysis is also computationally efficient.
3. The prior term is fully intrinsic to the shape space. Metrics, shape statistics, and probability models are all defined and computed on a well-defined shape space or its tangent bundle. The metric used is the elastic metric for shape analysis.
4. The shape prior is fully invariant not only to the similarity transformations, but also to re-parameterizations of curves by changes in speed. Similarity transformations such as rotation, translation, and scaling are already removed and need not be estimated at each step. However, this representation introduces a variability, associated with re-parameterizations of curves, which is solved numerically.

The main *limitations* of our framework are presented here as follows. Firstly, it does not allow objects to split and merge during the evolution. Shape models are defined and applied to each closed curve individually. Secondly, this approach is restricted to 2D images and does not extend to analysis of surfaces in volumetric data (although surfaces can sometimes be represented as indexed collections of curves and then this framework applies (Samir et al. 2006)). This framework easily extends to shape estimation of curves in $\mathbb{R}^n$ (Joshi et al. 2007b). Lastly, as mentioned above, the representation of curves using parametric forms introduces an additional variability of re-parameterization that has to be removed using optimizations.

The remainder of this paper is organized as follows. Section 2 summarizes our approach for representing and analyzing shapes of closed curves and Sect. 3 presents two possibilities for forming a prior (probability) model on the shape space. Section 4 outlines different energy terms that form the posterior energy and presents an algorithm for finding MAP estimates of curves. Section 5 presents some experimental results that demonstrate the main ideas in the paper, while Sect. 6 shows an example from statistical decision theory that allows selection and testing of different prior shapes models on a given image data. We finish the paper with a section on conclusions and summary.

2 Elastic Shape Analysis

2.1 Representation

As mentioned earlier, elastic analysis has emerged as a dominant tool in statistical analysis of variations in shapes of curves. However, the formulation of this analysis is not unique and it depends on the chosen representation of curves. Let J denote the interval $[0, 2\pi]$. The following representations of a curve β have been proposed:

- Younes (1998) uses a function $\xi : J \mapsto \mathbb{S}^1 \subset \mathbb{C}$, given by $\xi = \frac{\dot{\beta}}{\|\dot{\beta}\|}$, with β viewed as a parameterized curve in the complex space $\mathbb{C}$, to represent a curve. For any (increasing) diffeomorphism of J , call it γ , and a complex function $r : J \mapsto \mathbb{S}^1 \subset \mathbb{C}$, the mapping $(r, \gamma) \circ \xi = r \cdot \xi(\gamma)$ defines a group action on the functions denoting curves. In particular, the component γ denotes the stretching and the component r denotes the bending of the original curve β .
- For a curve β , one can denote a deformation in its shape using a normal vector field on it and, further, specify this vector field as a real-valued function on J . That is, for a function $a : J \mapsto \mathbb{R}$, define a vector field $a(t)\mathbf{n}_{\beta(s)}$, where $\mathbf{n}_{\beta(s)}$ denotes a unit normal to the curve at $\beta(s)$. Using this representation, Michor and Mumford (2006) have studied shape analysis of closed curves under a variety of metrics, including the elastic metric.
- Mio et al. (2007) represent a closed curve β by a pair of functions: the log-speed function $\phi \equiv \log(\|\dot{\beta}\|)$ and the angle function $\theta \equiv \angle(\frac{\dot{\beta}}{\|\dot{\beta}\|})$, to represent and analyze shapes under the elastic metric.
- Joshi et al. (2007b) use a square-root representation, given by $\frac{\dot{\beta}}{\sqrt{\|\dot{\beta}\|}}$, that studies the shape of closed curves in $\mathbb{R}^n$ using the same elastic metric.

These different representations lead to different spaces, or manifolds, for analyzing shapes and the elastic shape analysis is performed by choosing an appropriate Riemannian structure on these manifolds. For the purpose of this paper, one can choose any of these representations and

metrics. We choose the third representation, i.e. the pair (ϕ, θ) , as it provides a direct interpretation of deformation between shapes as being elastic, a change in θ implies bending and a change in ϕ implies stretching/compression.

Next we summarize ideas from elastic shape analysis of planar closed curves; for details the interested reader is referred to Mio et al. (2007). Let $\beta : J \mapsto \mathbb{R}^2$ be a non-singular, parameterized curve and we are interested in analyzing its shape. We assume that the curve has been scaled to be of length 2π . We represent the velocity vector of the curve in complex polar coordinates, $\beta(t) = e^{\phi(t)} e^{j\theta(t)}$, where $\phi : J \rightarrow \mathbb{R}$ and $\theta : J \rightarrow \mathbb{R}$ are $\mathbb{L}^2$ functions, and $j = \sqrt{-1}$. ϕ is called the *log-speed* function of β and measures the rate of local stretching and compression, whereas θ is called the *angle* function and measures local bending. We will represent the shape of β with the pair $v \equiv (\phi, \theta)$, and denote by $\mathcal{H}$ the collection of all such pairs. In order to make the shape representation invariant to rigid motions and uniform scalings, we restrict shape representatives to pairs (ϕ, θ) satisfying the conditions

$$\mathcal{C} = \left\{ v \equiv (\phi, \theta) : \int_J e^{\phi(t)} dt = 2\pi, \frac{1}{2\pi} \int_J \theta(t) e^{\phi(t)} dt = \pi, \int_J e^{\phi(t)} e^{j\theta(t)} dt = 0 \right\} \subset \mathcal{H}, \quad (3)$$

where $\mathcal{C}$ is called the *pre-shape space*. The first constraint forces the length of β to be 2π , the second forces its orientation to be π and the last ensures that β is closed. Therefore, $\mathcal{C}$ contains all planar, closed curves with length 2π and orientations π . In other words, the set of valid shapes lying in $\mathcal{C}$ are defined as the inverse image $\mathcal{G}^{-1}(2\pi, \pi, 0, 0)$, where

$$\begin{aligned} \mathcal{G} &\equiv (\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3, \mathcal{G}_4) \\ &= \left(\int_J e^{\phi(t)} dt, \frac{1}{2\pi} \int_J \theta(t) e^{\phi(t)} dt, \int_J e^{\phi(t)} \cos \theta dt, \int_J e^{\phi(t)} \sin \theta dt \right). \end{aligned} \quad (4)$$

An arbitrary differentiable curve having a representation $(\phi, \theta) \in \mathcal{H}$, is not necessarily an element of $\mathcal{C}$. However, we can project it in $\mathcal{C}$ by moving perpendicular to the level sets of $\mathcal{G}$ as follows. Using the gradient of $\mathcal{G}$, we can specify the space normal to $\mathcal{G}$ at (ϕ, θ) :

$$N_{\phi, \theta}(\mathcal{G}) = \text{span} \left\{ \nabla \mathcal{G}_1 = \left(\frac{1}{a}, 0 \right), \nabla \mathcal{G}_2 = \left(\frac{\theta}{a}, \frac{1}{b} \right), \nabla \mathcal{G}_3 = \left(\frac{\cos(\theta)}{a}, \frac{-\sin(\theta)}{b} \right), \nabla \mathcal{G}_4 = \left(\frac{\sin(\theta)}{a}, \frac{\cos(\theta)}{b} \right) \right\}. \quad (5)$$

Then, using Algorithm 1 we can iteratively update (ϕ, θ) in a direction in $N_{\phi, \theta}(\mathcal{G})$ that moves $\mathcal{G}(\phi, \theta)$ towards $(2\pi, \pi, 0, 0)$ as fast as possible.

Although we have removed rotation, translation and scaling variabilities, many elements of $\mathcal{C}$ can still have the same shape. These elements result from re-parameterizing the same curve and, thus, need to be removed from the representation. Let Γ be the set of all diffeomorphisms of a unit circle $\mathbb{S}^1$ to itself. This set can be identified with the set of all re-parameterizations of all closed curves, using the right composition action, $(\beta, \gamma_0) = \beta(\gamma_0)$. Any element $\gamma_0 \in \Gamma$ can be written as $\gamma_0 = r \cdot \gamma$ where r is an element of $\mathbb{S}^1$, γ is the set of all origin-preserving diffeomorphisms of $\mathbb{S}^1$ (call it $\mathcal{D}$), and $\cdot$ is the group operation in Γ . Another way to write this is $\Gamma = \mathbb{S}^1 \cdot \mathcal{D}$. Stated differently, the re-parameterizations of a closed curve can be broken into two components: (i) The change in the placement of origin on the curve is represented by the action of a unit circle $\mathbb{S}^1$ on $\mathcal{C}$, according to $r \cdot (\phi(s), \theta(s)) = (\phi(s - r), \theta((s - r)_{\text{mod } 2\pi} - r))$, for $r \in \mathbb{S}^1$. (ii) The change in speed of a curve in $\mathcal{C}$, while maintaining orientation and origin, can

be written as an action of subgroup $\mathcal{D}$ on $\mathcal{C}$: for a $\gamma \in \mathcal{D}$, the action is given by $(\phi, \theta) \cdot \gamma = (\phi \circ \gamma + \log \dot{\gamma}, \theta \circ \gamma)$. Using this structure, we will divide the analysis of Γ into these two distinct subgroups $\mathbb{S}^1$ and $\mathcal{D}$. This greatly helps the computational search over Γ as existing efficient algorithms, such as dynamic programming, can be utilized. As an example, Fig. 1(a) shows an arbitrary arc-length parameterized curve, where the sampling of points around the curve is shown for clarity. Its ϕ and θ functions are shown in Fig. 1(b) and (c). As expected the ϕ function is identically zero. This curve is acted upon by a monotonically increasing, continuous reparameterization function γ (shown in (d)) to yield a new log-speed and an angle function (shown in (e) and (f)). The reparameterized curve is shown in Fig. 1(g). The space of all reparameterizations of a shape in $\mathcal{C}$ is thus given by $\mathbb{S}^1 \times \mathcal{D}$. The elastic shape space is now defined as the quotient space $\mathcal{S} = \mathcal{C}/(\mathbb{S}^1 \times \mathcal{D})$. For any $(\phi, \theta) \in \mathcal{C}$, its equivalence class is given by:

$$[(\phi, \theta)] = \{(r \cdot (\phi, \theta)) \cdot \gamma | r \in \mathbb{S}^1, \gamma \in \mathcal{D}\}. \quad (6)$$

This set is also called the *orbit* of (ϕ, θ) each orbit in $\mathcal{C}$ represents a unique shape. Analysis of shapes is performed on the set $\mathcal{S}$ of such orbits.

Shape versus Nuisance Variables—Note that the pair $v \equiv (\phi, \theta)$ represents the shape of β and, thus, ignores its placement, orientation, parameterization and scale. These latter variables are called *nuisance variables* in shapes analysis, as they do not contribute to the shape of a curve. However, in extraction of boundaries from images, matching shapes of curves, or even in visual displays of a shape, these nuisance variables become important. In Table 1, we enumerate the shape and the nuisance variables separately.

We summarize the role of these variables used in different parts of the analysis:

- The shape variables (ϕ, θ) and the matching variables $x_m \equiv (r, \gamma)$ play a role in shape analysis and are used in computation of E_{prior} and its gradient.
- The length variable l completely defines the term E_{smooth} , while some additional variables (ϕ, θ, p, τ) are used to compute the gradient of E_{smooth} .
- The shape variables (ϕ, θ) and the nuisance (appearance) variables $x_a \equiv (p, \tau, l)$ are involved in computation of E_{image} and its gradients. This set of variables is also called the *full appearance variables*.
- There are two places where $\mathbb{S}^1$ enters in our representation: the orientation $\tau \in \mathbb{S}^1$ and the placement of origin $r \in \mathbb{S}^1$. However, they are not completely independent. A change in r , while keeping $\tau = \pi$, results in a rotation of the curve. Hence, we need to keep either an orientation parameter or a placement of origin, but not both, to capture the rotation variability. We will choose one according to the convenience: for rotating curves in a plane we will vary τ while keeping r fixed and for optimally matching any two shapes we will vary r while keeping $\tau = \pi$.

We will use $x = [x_a, x_m]$ to denote the whole set of nuisance variables, $x_a \in \mathcal{N}_a \equiv (\mathbb{R}_+ \times \mathbb{S}^1 \times \mathbb{R}^2)$ and $x_m \in \mathcal{N}_m \equiv (\mathbb{S}^1 \cdot \mathcal{D})$, and the full nuisance space is given by $\mathcal{N} \equiv \mathcal{N}_a \times \mathcal{N}_m$. The curve β is determined both by its shape and its nuisance variables $\beta \equiv [(\phi, \theta), x] \in \mathcal{B} \equiv \mathcal{S} \times \mathcal{N}$, according to:

$$\beta(s) = p + l \int_0^s e^{\tilde{\varphi}(s)} e^{j(\tilde{\theta}(s) + \tau)} ds, \\ (\tilde{\varphi}, \tilde{\theta}) = (r \cdot (\varphi, \theta)) \cdot \gamma \quad (7)$$

where $x = (p, l, \tau, r, \gamma) \in \mathcal{N}$. To emphasize dependence of β on shape v and nuisance variables x , we will often use an explicit notation $\beta \equiv \beta(v, x)$, where $v = (\phi, \theta)$.

2.2 Geodesics

Geodesics are important in defining and computing statistics of shapes. To specify geodesics, we need to impose a Riemannian structure on the tangent space of shapes first. The tangent space at any point of $\mathcal{H}$ is $\mathcal{H}$ itself, while the tangent spaces to $\mathcal{C}$ and $\mathcal{S}$ are naturally more restricted. We start by making $\mathcal{H}$ a Riemannian manifold by imposing the following metric. For a point $(\phi, \theta) \in \mathcal{H}$, let h_i and f_i , $i = 1, 2$ be two small perturbations of ϕ and θ each, respectively. Then, for arbitrary $a, b > 0$, define

$$\langle (h_1, f_1), (h_2, f_2) \rangle_{(\varphi, \theta)} = a \int_{\mathcal{J}} h_1(s) h_2(s) e^{\varphi(s)} ds + b \int_{\mathcal{J}} f_1(s) f_2(s) e^{\varphi(s)} ds. \quad (8)$$

The parameters a and b control the *tension* and *rigidity* in the metric. A higher ratio of a/b favors bending, while a lower value favors stretching/compression.

According to (5), the tangent space $T_{(\phi, \theta)}(\mathcal{C})$ is the set of all pairs in $\mathcal{H}$ that are orthogonal to the set:

$$\text{span} \left\{ \left(\frac{1}{a}, 0 \right), \left(\frac{\theta}{a}, \frac{1}{b} \right), \left(\frac{\cos \theta}{a}, \frac{-\sin \theta}{b} \right), \left(\frac{\sin \theta}{a}, \frac{\cos \theta}{b} \right) \right\}$$

with orthogonality defined using the inner product in (8). So, one can project an arbitrary pair $(h, f) \in \mathcal{H}$ into $T_{(\phi, \theta)}(\mathcal{C})$ using:

$$(h, f)_{proj} \equiv (h, f) - \sum_{i=1}^4 \langle (h, f), \mathbf{e}_i \rangle_{(\varphi, \theta)} \mathbf{e}_i, \quad (9)$$

where $\mathbf{e}_1, \dots, \mathbf{e}_4$ are the orthonormal basis elements of $N_{(\phi, \theta)}(\mathcal{C})$. Furthermore, the tangent space $T_{[(\phi, \theta)]}(\mathcal{S})$ is the set of all elements of $T_{(\phi, \theta)}(\mathcal{C})$ that are orthogonal to the space $T_{(\phi, \theta)}([(\phi, \theta)])$ (with $[(\phi, \theta)]$ as defined in (6)).

With the metric given in (8), the geodesics in $\mathcal{S}$ are well defined and the remaining task is to compute geodesics for any two elements, say $[(\phi_1, \theta_1)]$ and $[(\phi_2, \theta_2)]$, of $\mathcal{S}$. There are two prominent numerical methods for finding geodesics on nonlinear manifolds when analytical expressions are not available.

1. Shooting Method (Klassen et al. 2004)—Here one starts by constructing a geodesic starting from $[(\phi_1, \theta_1)]$ in an arbitrary initial direction $(h, g) \in T_{[(\phi_1, \theta_1)]}(\mathcal{S})$ and follows it for unit time. Then, one optimizes over (h, g) so that the $\mathbb{L}^2$ norm, between the end point of the geodesic and $[(\phi_2, \theta_2)]$, becomes zero. Qualitatively, this corresponds to *shooting* geodesics

from the first shape so as to hit the second shape in unit time. The readers are referred to the paper (Mio et al. 2007) for more details on the shooting method.

We emphasize that this approach requires an optimal matching of the two shapes before or during the calculation of a geodesic. In other words, we need to solve an optimization problem over $\mathcal{S}^1 \times \mathcal{D}$ so that the parameterization of the second shape best matches a fixed (usually arc-length) parameterization of the first shape. In our implementation, the search over $\mathcal{S}^1$ is performed exhaustively while the search over $\mathcal{D}$ is performed using the dynamic programming (see Mio et al. 2007).

2. Path Straightening Method (Klassen and Srivastava 2006; Joshi et al. 2007a, 2007b)—The main idea here is to connect the two shapes $[(\phi_1, \theta_1)]$ and $[(\phi_2, \theta_2)]$ by an arbitrary initial path in $\mathcal{S}$ and to iteratively straighten it until the path results in a geodesic. The straightening is performed by perturbing the path along the gradient of an energy. An interesting point is that the gradient is known analytically and the perturbation is computationally efficient. In general, path straightening is faster and more stable of the two methods.

For the purpose of this paper, either method can be used for finding geodesic paths in $\mathcal{S}$ and we have chosen the former. Once a geodesic is found, it can be completely characterized by a starting point and a starting direction. Let $\Psi_t(v, g)$ denote a geodesic starting from $v \in \mathcal{S}$ in the direction g , and parameterized by time t . The value of a geodesic at $t = 1$ is also known as the *exponential map*, i.e. $\exp_v: T_v(\mathcal{S}) \mapsto \mathcal{S}$ is such that $\exp_v(g) = \Psi_1(v, g)$. Given two shapes v_1 and v_2 in $\mathcal{S}$, computing a geodesic between them results in an optimal direction $g \in T_{v_1}(\mathcal{S})$, such that $\Psi_1(v_1, g) = v_2$. This also provides the value of a matching variable x_m ($r \in \mathbb{S}^1$ and $\gamma \in \mathcal{D}$) for each point along the geodesic. That is, one obtains an optimal value of the matching variable x_m that matches v_2 to v_1 . When needed, we will express this availability of x_m as $\Psi_1(v_1, g; v_2, x_m)$ in this paper.

Specifically, the task of calculating a geodesic between a given pair of shapes (ϕ_1, θ_1) and (ϕ_2, θ_2) , involves finding a tangent vector, (h, f) such that the geodesic starting from (ϕ_1, θ_1) reaches (ϕ_2, θ_2) in unit time. For the pre-shape space, this amounts to solving an optimization problem that involves the exponential map as follows,

$$(\widehat{h}, \widehat{f}) = \arg \min_{(h, f)} \|\exp(h, f) - v_2\|_{v_1}^2. \quad (10)$$

In the process of finding the geodesic, any update to the vector (h, f) , may result in the terminal point being located outside the manifold $\mathcal{C}$. This point is projected back on the manifold using a mechanism that projects arbitrary points in $\mathcal{H}$ on the pre-shape space $\mathcal{C}$, using Algorithm 1. This projection can be constructed by iterating the point in the direction given by the basis vectors of the normal space of $\mathcal{C}$. The initial tangent vector is then parallel-transported to the new point along the infinitesimal geodesic. A first order infinitesimal update to the current estimate of the vector (h, f) is given as follows,

$$(\widetilde{h}, \widetilde{f}) = (h, f) - \delta \left(\frac{h^2}{2} - \frac{b}{2a} f^2, hf \right). \quad (11)$$

This is also the equation of the parallel transport and is derived from the system of differential equations of the geodesics approximated in $\mathcal{C}$. Shown in Fig. 2 is an example of a geodesic path under the elastic metric for $a = 1$ and $b = 1$.

3 Construction of Prior Shape Models

We will specify a prior model on shape using E_{prior} that in some circumstances can be seen as the negative log of a prior density. Although shape spaces are infinite-dimensional, prior densities are defined on some finite-dimensional subset using standard reference measures (Grenander 1981). In other cases E_{prior} is simply treated as an energy term without connecting it to a prior density. The two important items are: the form of E_{prior} and the gradient of E_{prior} with respect to the curve β . The latter is needed because we will use a variational approach to solve for the optimal curve. The gradient will be conveniently in the form of a vector field on β , that will drive its evolution in time. We emphasize that E_{prior} being a shape prior depends only on the shape of β and not on its nuisance variables. On any nonlinear manifold there are two general ideas for imposing probability measures: (i) specify a probability measure on the manifold itself, and (ii) specify a measure on its tangent bundle and transfer it onto the manifold. We will study both the cases, starting with the former.

3.1 Squared-Distance as Prior Energy

A simple choice for prior energy is $E_{prior} = d(v, \mu)^2$ where d denotes the geodesic distance connecting the shape $v \equiv [(\phi, \theta)]$ with a prior mean $\mu \in \mathcal{S}$. Since the geodesic distance is invariant to all the similarity transformations, including the action of the re-parameterization group, E_{prior} is also invariant to these transformations. We proceed with this term with a word of caution. One would like to view the negative exponent of this E_{prior} as being proportional to a probability density on $\mathcal{S}$, with respect to some underlying invariant measure. The difficulty, however, comes from infinite-dimensionality of $\mathcal{S}$ that may result in an unbounded normalization constant. Therefore, it is better to treat E_{prior} as an energy term rather than being proportional to $-\log(\text{prior density})$. The choice of the underlying invariant measure that results in an invariant estimation is another difficult issue (Jermyn 2005) that we will skip.

Let v and x be the shape and the nuisance variables of β , respectively; we can write $\beta = \beta(v, x)$. Since the prior depends only on v , the choice of x is completely arbitrary here. In the larger optimization over β , x is determined by other terms in E_{total} .

Proposition 1—The gradient of E_{prior} with respect to β , $\nabla_{\beta} E_{prior}$, is approximated by the vector field:

$$\nabla_{\beta} E_{prior}(s) \approx \frac{1}{\varepsilon} (\beta(\Psi_{\varepsilon}(v, g; \mu, x_m), x_a)(s) - \beta(v, x_a)(s)), \quad \text{for } \varepsilon > 0 \text{ small}, \quad (12)$$

where

- $\Psi_t(v, g; \mu, x_m)$ is a geodesic path from v to μ having the optimal direction g and the optimal matching $x_m \in \mathcal{N}_m$.
- $\beta(\Psi_{\varepsilon}(v, g; \mu, x_m), x_a)$ is the closed curve that results when the shape component of β is perturbed along the tangent direction g by an amount ε . The appearance nuisance variable x_a remains unchanged.

Proof: Let $g \in T_v(\mathcal{S})$ be such that $\Psi_1(v, g) = \mu$. Then, since $E_{prior} = d(v, \mu)^2$ or $\|g\|^2$, the gradient of E_{prior} with respect to v , using the metric given in (8), is simply g . In other words, v should be changed along the geodesic, connecting v to μ , to maximally decrease E_{prior} . How should the curve β be changed? The relationship between β and v is given by (7). Let x_a be the set of nuisance appearance variables of β . Then, the change in β , if v follows the geodesic along

g , is given by $\beta(\Psi_t(v, g; \mu, x_m), x_a)$. The derivative with respect to t gives the desired vector field, and is approximated numerically using (12).

Figure 3 shows a sample shape v , a prior mean μ , and a prior vector field $\nabla_{\beta} E_{prior}$ induced on it.

While this E_{prior} provides a mechanism for introducing some prior information in the process of boundary extraction, it does not permit any more additional structure. For instance, this E_{prior} is isotropic and it cannot incorporate information about certain directions being more important than others in terms of capturing shape variability. To allow for more structure, we have to utilize the second idea of including covariances in the prior.

3.2 TPCA in Shape Space: Normal Energy

Since the tangent space at any point on the manifold is a vector space, it has been the preferred domain for imposing probabilities for statistical analysis of shapes. Depending upon the representations used in shape analysis there are several references for this idea: for landmark-based shape analysis (Dryden and Mardia 1998), shape analysis of non-elastic (Srivastava et al. 2005b) and elastic curves (Srivastava et al. 2005a), diffeomorphism of image domains (Vaillant et al. 2004), and for set-theoretic shape analysis using Hausdorff type metrics (Charpiat et al. 2005). In case of functional representations, it is also common to impose a multivariate normal density on a finite-dimensional subspace of the tangent space with respect to the underlying Lebesgue measure. Although such a probability density can be pushed onto the manifold using the exponential (or some other) map, it is difficult to do so analytically due to the nonlinear nature of such mappings. Therefore, one restricts to considering probability densities on tangent spaces rather than on S . In the implementation, however, one can generate (random) samples in tangent spaces and map them freely to S using the exponential map. We will take this approach to define a E_{prior} that allows more structure than the squared-distance energy.

The basic idea is to estimate a mean shape μ for an object class, and to use the tangent space $T_{\mu}(S)$, or rather a finite-dimensional subspace of $T_{\mu}(S)$ obtained using PCA, for imposing probabilities. This approach has been termed *Tangent PCA* (TPCA) (Srivastava et al. 2005b). The mean μ is taken to be the Karcher mean as defined previously (Srivastava et al. 2005b, 2005a). Figure 4 shows an example of a Karcher mean of some observed shapes. The actual use of TPCA here is slightly different from the past usage and is described in a few steps here. Firstly, it relies on estimated versions of tangent spaces and their principal subspaces, and secondly, it results in a direction for modifying the current shape v rather than just sampling from a probability distribution.

TPCA Approach

1. Find mean shape and sample directions: One starts with an ensemble of k observed (in this case, training) shapes $\{v_i, i = 1, \dots, k\}$, and computes their Karcher mean $\mu \in S$. This results in, for each v_i , a tangent vector $g_i \in T_{\mu}(S)$ such that $\Psi_1(\mu, g_i) = v_i$.

2. Form observed space: The next step is to form an orthonormal basis for the set $\{g_i, i = 1, \dots, k\}$ with respect to the inner product given in (8) using the Gram-Schmidt process. Let the new basis elements be $\{f_i, i = 1, \dots, m\}$ and define $M \equiv \text{span}(f_1, f_2, \dots, f_m) \subset T_{\mu}(S)$ with $m \leq k$.

Denote the projection of each tangent g_i into M by $y_i \in \mathbb{R}^m$, where $y_{ij} = \langle g_i, f_j \rangle_{\mu}$ for $j = 1, 2, \dots, m$.

3. Find a principal subspace: To obtain computational efficiency, it is often useful to restrict to a principal subspace of M as follows. Perform PCA of the observed set $\{y_i, i = 1, \dots, k\}$ by finding their principal subspace of size $n \leq m$; n is chosen by the user. Denote $N \subset M \subset T_\mu(S)$ as the principal subspace of observed tangent vectors. We will use $y \mapsto U^T y$ as the mapping from M into N , where $U \in \mathbb{R}^{m \times n}$ is an orthogonal basis of N .

4. Impose a probability density on N : One can impose one of several probability densities on N as a prior. For example, a Gaussian model for individual principal coefficients, with mean zero and variances given by observed singular values, is a simple choice. Another model that imposes a different mixture of Gaussians for each principal component, independent of each other, has been found to perform better than a simple Gaussian model and a nonparametric model (Srivastava et al. 2005a). For the simplicity of presentation, we take the first idea and form a multivariate normal density on the subspace M of $T_\mu(S)$.

Shape Prior and Its Gradient

1. Shape prior: Let $g \in T_\mu(S)$ be a tangent vector such that $\exp_\mu(g) = v$ and let y be the projection of g into M . Now define a prior energy as

$$\tilde{E}_{prior}(y) = \frac{1}{2} y^T (UK^{-1}U^T)y + \frac{1}{2\delta^2} \|y - UU^T y\|^2. \quad (13)$$

$K \in \mathbb{R}^{n \times n}$ is the sample covariance matrix associated with past observations on N and δ is a positive number chosen to be less than the smallest singular value of K .

2. Gradient of shape prior: The gradient $\nabla_y \tilde{E}_{prior} = Ay$, where $A = UK^{-1}U^T + \frac{1}{\delta^2}(\mathbf{I}_m - UU^T)$. Recall that y is simply an expression of g in the basis $\{f_j, j = 1, \dots, m\}$. Therefore, one can write $\tilde{E}_{prior}$ and its gradient as a function of g by simply lifting them to M , using a change of variables $y \mapsto \sum_{j=1}^m y_j f_j$. For instance, $E_{prior}(g) = \tilde{E}_{prior}(\{ \langle g, f_j \rangle_\mu, j = 1, \dots, m \})$, and its derivative $\nabla_g E_{prior} = \sum_{j=1}^m (\nabla_y \tilde{E}_{prior})_j f_j$.

3. Parallel transport: This can be carried one step further, i.e. we can write the posterior energy in terms of the curve β , or rather its shape v , using the exponential map $g \mapsto v = \exp_\mu(g)$. However, writing these terms explicitly will introduce burdensome notation due to complicated change of variables under $\exp_\mu$. Taking a numerical approach, we parallel transport $\nabla_g E_{prior}$ along the geodesic connecting μ with v in S . Since we already have a geodesic path from μ to v (from the step on Shape Prior), we can use it here. Using Algorithm 2 we construct a vector field on this geodesic whose covariant derivative along that geodesic is zero and it equals $\nabla_g E_{prior}$ on μ . Its value at v , call it $h \in T_v(S)$, provides us with the desired transport.

4. Gradient field on curve: Our interest is in writing the gradient of E_{prior} with respect to the shape of β , we will do so using the numerical approximation (similar to Proposition 1):

$$\nabla_\beta E_{prior}(s) \approx \frac{1}{\varepsilon} (\beta(\Psi_1(v, \varepsilon h; \mu, x_m), x_a)(s) - \beta(v, x)(s)), \quad \text{for } \varepsilon > 0 \text{ small.} \quad (14)$$

keeping in mind that $\Psi_1(\mu, g) = v$. As earlier, $\nabla_\beta E_{prior}$ is a vector field on the curve β which can be used to update β directly.

3.3 Comparison of the Two Prior Terms

We have proposed two prior energy terms on β —a squared-distance term on $\mathcal{S}$ and a multivariate normal term on the tangent bundle $T\mathcal{S}$ —and a natural question is to compare their effectiveness in incorporating shape priors. Intuitively, the normal term should provide better results as it additionally contains the directions of major variability in a particular shape class. This idea is further demonstrated pictorially in Fig. 5 where shape evolutions under the (negative) gradients of the two energies are presented.

In this particular example, we use tank shapes and their mean μ shown in Fig. 4. Using the TPCA approach, we compute the dominant subspace N and a covariance matrix K from the observations. Next, we plot shape evolutions under the gradient process with respect to both the energy terms, starting at an arbitrary shape v . Since the distance-squared term is simply the squared length of a tangent vector, the gradient process will simply follow this straight line from v towards μ in the tangent space $T_\mu(\mathcal{S})$. The normal energy term however will form a curved path from v to μ favoring lower energy shapes (under the normal energy). Both these paths are shown in Fig. 5(a) and successive shapes along these paths are shown in Fig. 5(b). Together these plots demonstrate that gradient of the normal energy will find low energy shapes in fewer iterations and closer to v than the gradient of squared-distance term. One important consequence is that if the shape to be discovered is not close to μ but still a low (normal) energy shape, then the squared-distance term may not be useful. However, the normal energy term will get close to it by definition. (This has been demonstrated later in Figs. 15 and 18 using both real and synthetic images.) For an additional insight in this discussion, we also plot the evolution according to gradient of the Rayleigh coefficient of K , i.e. the normal term with an additional constraint that the distance between μ and v is constrained to be unit length. As expected, the limiting shape is the shape along the dominant eigenvector of K . For another perspective, we present variation of shapes along different tangent directions at μ . We plot the shapes $\exp_\mu(k\sqrt{\sigma_i}v_i)$, for the first three dominant eigen directions and three random directions perpendicular to M , by varying k in Fig. 6(a). In these experiments the norms of v_{rand} are set to one for a fair comparison. We can see that even a small variation from μ along random directions results in a significant loss of structure. Level sets of squared-distance energy are circles while that of normal energy are ellipses. Shown in Fig. 6(b) are shapes along these two sets in a two-dimensional subspace spanned by the dominant direction v_1 and a random direction in $M^\perp$. The shapes along the ellipse represent the variability of the observed ensemble better and thus, normal energy is a better candidate for the shape prior.

So far we have discussed prior models for the shape component in β , $v(\beta)$. Another interesting question is, what prior model can be chosen for the nuisance variable x ? In this paper, we assume an uniform prior on rotations and translations though, in the future, more informative priors can be pursued.

4 Posterior Energy for Boundary Extraction

Boundary extraction typically involves minimizing an energy functional on the space of closed curves. The posterior energy functional usually involves the image information and a smoothing criterion, whereas this paper adds a term on shape prior, resulting in:

$$E_{total}(\beta) = \lambda_1 E_{image}(\beta) + \lambda_2 E_{smooth}(\beta) + \lambda_3 E_{prior}(\beta), \quad (15)$$

where λ_1 , λ_2 , and λ_3 are arbitrary positive constants. The term E_{image} forces the curve β to pass through areas of highest intensity changes, i.e. edges and ridges in the image, whereas E_{smooth} maintains the smoothness of β by penalizing the length of β . The prior term E_{prior} forces

the shape of β to be of high probability under the chosen prior shape model. The resulting $\hat{\beta}$ is a combination of these three forces depending upon the relative weights $\lambda_i > 0$. The optimal curve is given by,

$$\hat{\beta} = \arg \min_{\beta \in \mathcal{B}} E_{total}(\beta), \quad (16)$$

where the minimization is performed over all closed curves in the given image domain. We will use a gradient approach to solve for $\hat{\beta}$.

In the next subsections, we describe the remaining terms, E_{image} and E_{smooth} , and calculations of their gradients with respect to β .

4.1 Image Energy

There exists a large literature on the problem of extracting curves from 2D images. Since our goal is to illustrate the influence of a shape prior, any such method can be chosen as a starting point as long as it allows shape analysis of underlying curves. Here we select methods that represent contours as parametric curves explicitly.

Most segmentation methods involving curves, depend on intensity or edge information to evolve over the image. Kass et al. (1988) suggest the use of various external forces like simple image intensity, edge energy functional or termination functionals using the curvature of level lines in an image: $I(x, y)$, $-|\nabla I(x, y)|^2$ or $-\|\nabla(G_\sigma(x, y) * I(x, y))\|^2$. Caselles et al. (1997) also

suggest the energy function: $\frac{1}{1+|\nabla I(x, y)|^i}$, $i = 1$ or 2 . These image-derived forces, separately or when combined together, emphasize the detection of lines, edges, ridges or corners and serve as a stopping criterion for the curve evolution. In order to reduce sensitivity to initialization, and to smooth these image-driven vector fields, Xu et al. (1998a, 1998b) present an energy directly in terms of vector fields $\mathbf{v}(x, y) = (v_x(x, y), v_y(x, y))$: $\int g(|\nabla f|)|\nabla \mathbf{v}|^2 + h(|\nabla f|)(|\mathbf{v} - \nabla f|^2) dx dy$, where f is one of the energy functions suggested by Kass et al. The papers (Xu and Price 1998a, 1998b) describe a numerical approach to solving for $\mathbf{v}$ using an iterative approach. However, both GVF and GGVF may fail to stop the snake at weak edges. Li et al. (2005) proposed an edge preserving gradient vector flow (EPGVF) to overcome this drawback of GVF and GGVF, by minimizing the energy:

$$\mathcal{E}_{ext}(\mathbf{v}) = \int g(|\nabla f|)|\nabla \mathbf{v}|^2 + h(|\nabla f|) \times (\mu |J_{\mathbf{v}} p|^2 + |\mathbf{v} - \nabla f|^2) dx dy. \quad (17)$$

This definition is taken verbatim from (6) (Li et al. 2005), where f and g weighting functions, and p is a normalized vector field that is orthogonal to the gradient vector field of the image. Here $J_{\mathbf{v}}$ is the Jacobian of the vector field $\mathbf{v}$ and $J_{\mathbf{v}}(\mathbf{p})$ denotes the matrix-vector multiplication,

$$\mathbf{p} = \left[\frac{-I_y}{\sqrt{I_x^2 + I_y^2}}, \frac{I_x}{\sqrt{I_x^2 + I_y^2}} \right]$$
 where This term $\mu |J_{\mathbf{v}} p|^2$ in (17) has been added to ensure smoothing of vector fields in directions tangential to dominant edges, while keeping them sharp along the normals. The EPGVF accentuates weak boundaries, while retaining the desired properties of GVF and GGVF. One obtains the optimal vector field, $\mathbf{v}^* = \arg \min_{\mathbf{v}} \mathcal{E}_{ext}(\mathbf{v})$ using a numerical iterative approach (Li et al. 2005). From now onwards we will only consider this optimal vector field $\mathbf{v}^*$, and to simplify notation we will denote it by $\mathbf{v}$. Using this external force field, we can

now define a dynamic energy $E_{image}(\beta(s))$, that changes with every step of the evolution of the curve. This energy E_{image} is given by,

$$E_{image}(\beta) = \int_{\beta} |v(p)|^2 dp. \quad (18)$$

Proposition 2—The energy term E_{image} is invariant to re-parameterizations of β .

Proof: For any parameterized curve β , E_{image} is given by $\int_0^l |v(\beta(s))|^2 (|\dot{\beta}(s)|) ds$. Replacing β by $\tilde{\beta} = \beta(\gamma)$ for any $\gamma \in \mathcal{D}$, we obtain:

$$E_{image}(\tilde{\beta}) = \int_0^l |v(\beta(\gamma(s)))|^2 (|\dot{\beta}(\gamma(s))|) \dot{\gamma}(s) ds$$

which, by a change of variable $t = \gamma(s)$, becomes same as $E_{image}(\beta)$.

Although one can derive the gradient of E_{image} with respect to β under arbitrary parameterization of β , we simplify the gradient by assuming arc-length parameterization of β , i.e. $|\dot{\beta}(s)| = 1$. The resulting gradient is $\nabla_{\beta} E_{image}(s) = J_{\beta}(\mathbf{v})(\beta(s)) \mathbf{v}(\beta(s))$, where $J_{\beta}(\mathbf{v})$ is the Jacobian of $\mathbf{v}$ with respect to (x, y) , evaluated at $\beta(s)$, given by

$$\left[\begin{array}{cc} \frac{\partial v_x}{\partial x} & \frac{\partial v_x}{\partial y} \\ \frac{\partial v_y}{\partial x} & \frac{\partial v_y}{\partial y} \end{array} \right] \bigg|_{\beta(s)}.$$

Figure 7 shows some examples of the EPGVF, and vector fields induced on curves. The top row shows some images with some arbitrary closed curves. The middle row shows EPGVF for the full images, and the bottom row shows $\nabla_{\beta} E_{image}$ as the vector field restricted to the curve β .

4.2 Smoothing Energy and Length Penalty

A desirable property of the curve evolution is that smoothness of the curve be preserved or even enhanced during extraction. All parametric and free-form active contour models (Kass et al. 1988; Xu and Prince 1998b) enforce smoothness by including a regularization term that penalizes the first and second order derivatives of the curve. While such an approach usually depends on the parameterization, we follow a more intrinsic approach (Caselles et al. 1997; Kichenassamy et al. 1995). We consider an arc-length parameterized curve β . Define the smoothing energy functional as:

$$E_{smooth}(\beta) = \int_0^l \langle \dot{\beta}(s), \dot{\beta}(s) \rangle^{1/2} ds.$$

E_{smooth} is simply the length of the curve β and is naturally invariant to any re-parameterization of β . We include E_{smooth} in the total energy, despite having a separate E_{prior} on the shape of β , because E_{smooth} provides a prior on the length of β while E_{prior} does not. Shape and length are two different variables and are governed by separate terms.

As shown in many articles, e.g. Kichenassamy et al. (1995), the gradient of E_{smooth} with respect

to β is given by $-\frac{d}{ds}\left(\frac{\dot{\beta}(s)}{|\dot{\beta}(s)|}\right)$. This term results from calculus of variations and the fact that other term drops out because the variation can be assumed to be zero at the boundary points. (In other words, since a translation of β does not change E_{smooth} , the variation is taken in such a way that a point on β remains fixed.) One can derive the gradient more explicitly by choosing a proper metric on spaces of vector fields on curves. We, however, seek simplification by assuming arc-length parameterization resulting in $\|\dot{\beta}(s)\|=1$ and

$$\nabla_{\beta} E_{smooth}(\beta) = -\ddot{\beta}(s) = -\kappa(s)n(s), \quad (19)$$

where, κ is the curvature along β and $n(s)$ is the inward normal unit vector to β at s . It has been shown (Gage and Hamilton 1986; Grayson 1987) that these curve-shortening or Euclidean heat flows automatically lead to smoothing of the curve. Furthermore they cause the curve to shrink and become convex, and if this process is continued, the curve then turns into a circle and ultimately collapses to a point.

As an example, Fig. 8 shows the evolution of a curve under: (i) $\nabla_{\beta} E_{smooth}$ alone (top row) and (ii) under a combination of $\nabla_{\beta} E_{smooth}$ and $\nabla_{\beta} E_{prior}$ (bottom row). The prior energy is the squared geodesic distance to the mean shape in this case and the mean shape is the final shape in the bottom row. In this figure, all the curves have been scaled to the same size for better display. As expected, under the effect of the smoothing vector field, the length of the curve diminishes gradually, while the curve becomes circular in shape. However when the curve is additionally influenced by the prior vector field, the shape of the curve gradually converges to the prior mean shape, while it's length gradually reduces.

4.3 Total Energy

For the total energy given in (15), the gradient vector field is a sum on individual gradients:

$$\nabla E_{total}(\beta)(s) = \lambda_1 J_{\beta}(v)(\beta(s))v(\beta(s)) - \lambda_2 \kappa(s)n(s) + \lambda_3 \nabla_{\beta} E_{prior}(s), \quad (20)$$

and contour estimation is obtained as a solution of the differential equation:

$$\frac{dX(t)}{dt} = -\nabla E_{total}(X(t)), \quad X(0) \in B, \quad (21)$$

with $\hat{\beta} = \lim_{t \rightarrow \infty} X(t)$. It is difficult to establish asymptotic properties of $X(t)$ analytically and we resort to experimental results to analyze this method.

Computational Cost—The computational cost in updating β is in evaluating the three vector fields on β at each iteration. While the other two terms are standard, we discuss the computational cost of the gradient of E_{prior} . The main cost in this step comes from the computation of a geodesic path between the current shape v and the prior mean μ . Depending upon the choice of shape representation, shape metric, and the method for finding geodesic, this cost varies a lot. For instance, using the bending metric (Klassen et al. 2004) and the path straightening approach (Klassen and Srivastava 2006), one can compute approximately 50 geodesics per second for the shapes and resolutions studied in this paper. For the elastic metric

(Mio et al. 2007), the cost of similar computation increases to result in approximately one geodesic per second. We believe that these times for computing geodesics in shape spaces of curves are much smaller than those obtained using level sets or region-based shape analysis.

5 Experimental Results

In this section we deibe experimental results on estimating boundaries of objects using real and simulated image data. Bayesian shape extraction is most effective in situations where the image quality is low, or when the image consists of partially obscured or occluded objects and, most importantly, when we have a prior knowledge about the shape of the object contained in the region. This is often times the case in medical image analysis using ultrasound images, echocardiographs or low-resolution PET scans, where signal to noise ratio is rather small and object to background contrast is not sufficient for image-based object extraction. Similar low-contrast images appear in surveillance applications using infrared cameras. Using past observations of those objects: anatomical parts or human subjects, one can develop a shape prior and use it in order to compensate for missing data.

Even though $E_{total}(\beta)$ is invariant to re-parameterization of β , the computation of the gradient term simplifies if the arc-length parameterization is assumed. Therefore, we need to re-sample β uniformly after every iteration to ensure arc-length parameterization. This also helps in stabilizing computations by ensuring that samples along β do not coalesce. If the shape of β is given by the pair $(\phi, \theta) \in \mathcal{C}$, one can re-sample it into an arc-length parameterized curve using

$$(\phi, \theta) \circ \gamma, \quad \text{where } \gamma(t) = \int_0^t e^{-\varphi(s)} ds. \quad (22)$$

The warping function γ can be computed numerically, and applied to (ϕ, θ) to obtain an arc-length parameterized version of β .

1. Visible Spectrum Images

Firstly, we present two examples where we use images of a tank shown in Fig. 9, to demonstrate Bayesian active curves under a shape prior using the squared-distance energy. The rightmost panel in each case shows the prior mean shape μ used in this experiment and the prior energy used here was the squared-distance term. The goal is to find a MAP estimate of the object shape present in I . Shown in Fig. 9 are the evolutions of $X(t)$, without the prior term ($\lambda_3 = 0$) and with ($\lambda_3 > 0$). In the absence of prior, the algorithm is misled by a partial obscuration of the object in I . In contrast, the Bayesian solution gives a better estimate of the tank boundary because of a strong prior on the shape of the curve. We emphasize that the scale, position, and orientation of β , along with its shape, are estimated by the algorithm automatically.

Secondly, we consider the effect of using different prior shape averages for Bayesian extraction by comparing the evolution of curves for a fixed, arbitrary test image. In this experiment, the prior mean shape was computed using an independent set of silhouettes from IR tank images, shown later in Fig. 19. Comparisons of curve evolution under the effect of different prior mean shapes are shown in Fig. 10.

2. Ultrasound Images

This experiment involves a set of cardiac images taken using an ultrasound equipment. Each set consists of a sequence of images beginning with the ED (End-Diastole) frame and ending with the ES (End-Systole) frame. An important goal in echocardiographic image analysis is to develop tools for extracting cardiac boundaries in approximately 10–15 image frames that are

typically acquired between ED and ES. Different aspects of past efforts on this dataset (Wilson et al. 2000; Chen et al. 2002) have focused on both the construction of geometric figures to model the shape of the heart as well as validation of extracted contours. Given a manual tracing of the boundaries for the ED and ES frames, we aim to extract epicardial and endocardial boundaries in the intermediate frames. We first compute a geodesic interpolation of the curves between the first and the last frame (Joshi et al. 2005), and use these intermediate curves as initial conditions for boundary extraction using minimization of E_{total} . Figure 11 shows extraction results for the cardiac boundary for a specific frame in two different image sequences. In this experiment, we have used the squared-distance energy as E_{prior} .

3. Shape Estimation from Incomplete Images

In this example, we consider the task of completing shapes of dinosaur's bones using images of bones with missing pieces (left column of Fig. 12). As a prior information, we are given some pictures of complete bones from animals of the same species (second column). For each case, we compute the mean shape μ of the prior observations and use the distance squared prior on shapes for Bayesian shape estimation. The estimated shapes $\hat{\beta}$ are shown in the third column of Fig. 12. Since there is no ground truth available, we use manual segmentations by an expert (right column) for comparison. The visual closeness of the estimated shapes with expert drawings shows the utility of prior information in the shape estimation process. It is also interesting to note some differences in the two results. In each of the first two rows, the Bayes' estimate shows an elongation on the right side whereas the expert's sketch does not have a major change on the right corner.

4. Natural and Synthetic Images

Next we present boundary extraction results for a few images from the ETH-80 dataset, as well as some partially obscured image consisting of a hand gesture. Each row in Fig. 13 shows curve evolutions with and without the E_{prior} term; the upper row is for $\lambda_3 = 0$ and the lower row is for $\lambda_3 > 0$.

5. Results for varying λ_3

At each step of the evolution, the curve is influenced by the relative weights of the terms λ_1 , λ_2 , and λ_3 . The appropriate values for these weights are mostly application dependent. For noisy images lacking in reliable edge information, it is desirable that the final segmentation be largely governed by the prior term, thereby emphasizing λ_3 . Figure 14 shows the changes in the final extracted curve by varying λ_3 while keeping λ_1 and λ_2 fixed, for both types of E_{prior} —the squared-distance energy and the normal energy. It is noticed that an increasing contribution by the prior vector field reduces the effect of object boundary edges on the final segmentation, and eventually for large λ_3 the most probable shape is simply μ .

6. Squared-distance versus normal energy

Bayesian shape extraction is expected to benefit from the use of advanced statistical models, but different choices of E_{prior} can lead to different results. The Fig. 5(b) illustrates the difference in evolution when the squared-distance and the normal energies are minimized using a gradient approach. We first compare the effects of these two energy terms using primitive shapes, such as circles and ellipses. The top panel of Fig. 15 shows some circles and ellipses, and their Karcher mean μ . A multivariate normal distribution is imposed on the tangent space at the mean shape using the TPCA approach. As expected, the variation in this class (Fig. 15) is mostly along a single, dominant eigenvector with the eigenvalue of $\sigma_1 = 0.2399$, with the next largest eigenvalue of $\sigma_2 = 0.0015$. Next we use the gradient of E_{total} , with E_{prior} being one of the two types, to extract the boundary in a binary image. The actual boundary in this image is a noisy version of an ellipse, a shape that occurs often in the observed set. We also include two

results of curve extraction of noisy versions of elliptical objects with different aspect ratios. From the estimated boundaries in Fig. 15, it is observed, as expected, that the estimated curve for the normal term is better than that for the squared-distance prior. The explanation for this improvement is closely related to the discussion presented with Fig. 5 earlier.

Next, we perform a similar experiment, this time with a shape prior model derived from real image observations. The data consists of a sequence of infrared images resulting from the rotation of a tank mounted on a pedestal at various elevations. We are interested in a subset of this data that corresponds roughly to an azimuthal rotation of the tank from 200° to 315° at a fixed elevation. This amounts to a change in the angle between the gun and the body of the tank at different perspectives. Figure 16 shows a few examples out of the total number of 108 of such images.

The TPCA model is constructed by first computing the mean shape shown in Fig. 17 and then projecting the tangent vectors from all the 108 observations at the mean shape. Figure 17 also shows the decay in the eigenvalues corresponding to each principal component as well as the eigen-shape variation along the dominant eigendirections.

Next, we compare the curve evolution under both the squared-distance and the normal energy on a test image in Fig. 18. The final extracted curve for the squared-distance model closely resembles the mean shape, while converging on the boundaries of the tank image. It is noticed that there is a mismatch in the extracted shape w.r.t. the location of the turret in the image. This problem is avoided in the case of the normal model where the covariance information affects the curve evolution such that the misalignment between the turrets of the extracted curve and the tank image is reduced.

So far the estimated boundary has been obtained by implementing a gradient flow that typically yields locally optimal solutions. This is largely due to the sensitivity of gradient descent type optimizations to initializations. In such cases, one may employ stochastic optimization techniques to achieve global minima.

6 Statistical Decision Theory

There are some important questions that arise immediately for this framework: (i) In case we do not have any a priori expectation on the class of shapes, which prior model can be used for Bayesian active contour framework? (ii) How to evaluate performance of such an active contour approach? (iii) How to choose λ s for weighting different energy terms appropriately?

We suggest the use of statistical decision theory to address these and similar questions, as described next.

6.1 Bayesian Model Selection

What is the correct choice of the prior model for estimation of contours from a given image? In many situations, the presence of contextual information may help reduce possible hypotheses to a small number, but not necessarily to a single shape model. Consider for example an image taken from a battlefield containing military targets. In addition to the possibility that one of the several target types—tank, truck, jeep, hummer, etc.—may be present in that image, the shape model for even a single target will differ depending upon its pose. Rather than fixing an a priori shape model, a Bayesian approach suggests evaluating all possible models under the MAP framework. In other words, under the Bayesian approach one searches over all possible shape models and selects the model that best explains the image data. This also results in Bayesian object recognition as the selection of a shape class being equivalent to object recognition.

Let $\mathcal{Q}$ denote a discrete, finite set of shape models associated with a possible object present in an image. Different models can come from shapes of different objects, or different planar shapes resulting from different poses of the same three-dimensional object, or both. For each class $q \in \mathcal{Q}$, assume that we have a prior shape model and a prior probability $P(q)$ of that class. In case of the normal energy prior, this implies that we have a different mean μ_q , principal space $N_q \subset T_{\mu_q}(S)$, and a covariance K_q for each class. A MAP estimate of the target class, given an image I , is given by: $\hat{q} = \operatorname{argmax}_{q \in \mathcal{Q}} P(q|I)$, where

$$\begin{aligned} P(q|I) &= \frac{P(q) \int_{\mathcal{B}} P(I|\beta) P(\beta|q) d\beta}{\sum_{q \in \mathcal{Q}} (P(q) \int_{\mathcal{B}} P(I|\beta) P(\beta|q) d\beta)} \\ &= \frac{P(q) \int_{\mathcal{B}} e^{-E_{total}^q(\beta)} d\beta}{\sum_{q \in \mathcal{Q}} (P(q) \int_{\mathcal{B}} e^{-E_{total}^q(\beta)} d\beta)}, \end{aligned}$$

where $E_{total}^q = \lambda_1 E_{image} + \lambda_2 E_{smooth} + \lambda_3 E_{prior}^q$. Since it is difficult to perform this integration analytically, we look for some computational approximations. One idea is to use importance sampling, i.e. sample from the prior and use the likelihood term as weights. Another possibility is to use Laplace's approximation of the nuisance integral using asymptotics (Grenander et al. 2000). A coarse approximation is to evaluate the integrands at $\hat{\beta}$, the MAP estimates. That is, choose:

$$\hat{q} = \operatorname{argmax}_{q \in \mathcal{Q}} (P(q) e^{-E_{total}^q(\hat{\beta}_q)}), \quad (23)$$

where $\hat{\beta}_q$ is the MAP curve estimate for the model q . This process is akin to inferring from multiple hypotheses using generalized likelihood ratio tests (Van Trees 1971), and is often termed as *Bayesian model selection*.

We conducted an experiment that uses real image data to generate prior shape models for tanks. Figure 19(a) shows a schematic of a battle tank rotated at 3° steps in azimuth while keeping a fixed elevation. The sample infrared images corresponding to different side views of the tank are shown in Fig. 19(c). Figure 19(b) shows the automatically extracted shapes and their sample means grouped together in 9 quadrants for one half-circle. Due to viewing symmetry (the profile of a 3D object is same at antipodal viewing directions), we construct prior shape models for half of a circle, and reflect the shapes to get the models for the other half. In this experiment, the shape model comprises of the mean shape μ_q corresponding to each quadrant $\{q_i \in \mathcal{Q}\}$, $i = 1, \dots, 9, 1_r, \dots, 9_r$, where q_{ir} stands for the reflected shape for q_i . Further, we assume an uniform prior for $P(q)$ over all classes. For each of these quadrants, the MAP curve estimate,

$\hat{\beta}_q = \operatorname{argmin}_{\beta \in \mathcal{B}} E_{total}^q(\beta)$, is obtained. Based on these extracted curves, we use (23) to maximize the probability of selecting the best model $\hat{q}$. Figure 20 shows the extracted curves $\hat{\beta}_q$ overlaid on a partially obscured test image I . Figure 21 plots the discrete posterior energy in (15) against each of the tested prior models, $q_1, \dots, q_{9_r}$. It is observed that q_5 correctly results in the smallest E_{total} amongst all models, followed by q_6 , and q_4 . This method holds promise in situations where a large number of candidate prior models need to be evaluated.

6.2 Performance Evaluation Metrics

How should one evaluate the performance of this framework or any other boundary extraction algorithm? In a general segmentation framework this question is quite difficult to address.

However, in our narrow framework where extracted curves correspond to boundaries of objects, and shapes of objects relate to their classes, one can use the resulting classification performance for analyzing extraction algorithm.

If $\hat{\beta}$ is the result of our MAP estimation in an image I , we can use this curve to classify the object in that image. This can be done as suggested above using a Bayesian decision or using a more general classifier, e.g. SVM, nearest neighbor, etc. If we denote the classification performance on a test dataset to be F , then we claim that F can also be used to quantify performance of the original curve extraction algorithm.

In our framework, the estimate $\hat{\beta}$ and therefore F will actually be a function of λ_1 , λ_2 , and λ_3 . This suggests a technique for selecting these parameters which have been selected quite arbitrarily thus far. One can maximize F over these parameters on a validation set to select their values. If one seeks a fully Bayesian framework, then it would be better to use some non-informative priors on λ s themselves, e.g. a gamma density with large shape and scale parameters. We have left this study for future work.

7 Conclusion

This paper presents an intrinsic approach for constructing prior shape models for object extraction. Extraction is based on curve evolution under the influence of image, smoothing and prior vector fields. The strength of this work is the incorporation of object priors using a geometric shape model for improved segmentation. The shape prior is fully invariant not only to the similarity transformations, but also to re-parameterizations of curves by changes in speed. The last step comes at an additional computational cost for removing the re-parameterization group from the shape representation. The prior shape information drastically improves the object extractions under various obscuration to objects in images. The effectiveness of the shape prior in curve extraction is demonstrated using both the squared-distance energy as well as the normal energy using higher-order covariance information. This is evident from results on a variety of images including synthetic, as well as real data consisting of military, biological, and medical images.

Acknowledgments

We thank Prof. David Wilson of University of Florida for providing us the ultrasound images, Dr. Richard Sims of USAMRDEC for the tank data, and Albert Prieto-Marquez of Department of Biological Sciences, Florida State University for the use of the hadrosaurid bone images. Additionally, we thank Dr. Ian Jermyn of INRIA, Sophia Antipolis, France for several interesting discussions. This research was supported in part by the Army Research Office awards W911NF-04-01-0268 and W911-NF-04-1-0113, the Air Force Office of Scientific Research award FA9550-06-1-0324, and an Innovation Alliance Award from the Northrop Grumman Corporation.

References

- Amit Y, Grenander U, Piccioni M. Structural image restoration through deformable templates. *Journal of the American Statistical Association* 1991;86(414):376–387.
- Caselles V, Kimmel R, Sapiro G. Geodesic active contours. *International Journal of Computer Vision* 1997;22(1):61–79.
- Charpiat, G.; Faugeras, O.; Keriven, R.; Maurel, P. Approximations of shape metrics and application to shape warping and empirical shape statistics. In: Krim, H.; Yezzi, A., editors. *Statistics and analysis of shapes*. Berlin: Springer; 2005. p. 363–395.
- Chen Y, Tagare H, Thiruvankadam S, Huang D, Wilson D, Gopinath K, Briggs R, Geiser E. Using prior shapes in geometric active contours in a variational framework. *International Journal of Computer Vision* 2002;50(3):315–328.
- Cohen LD, Cohen I. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1993;15(11):1131–1147.

- Cootes TF, Taylor CJ, Cooper DH, Graham J. Active shape models,—their training and application. *Computer Vision and Image Understanding* 1995;61(1):38–59.
- Cremers D, Tischhäuser F, Weickert J, Schnörr C. Diffusion snakes: introducing statistical shape knowledge into the Mumford–Shah functional. *International Journal of Computer Vision* 2002;50(3):295–313.
- Cremers D, Kohlberger T, Schnörr C. Shape statistics in kernel space for variational image segmentation. *Pattern Recognition* 2003;36(9):1929–1943.
- Cremers D, Osher SJ, Soatto S. Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *International Journal of Computer Vision* 2006;69(3):335–351.
- Dryden, IL.; Mardia, KV. *Statistical shape analysis*. New York: Wiley; 1998.
- Gage M, Hamilton RST. The heat equation shrinking convex plane curves. *Journal of Differential Geometry* 1986;23:69–96.
- Grayson M. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry* 1987;26:285–314.
- Grenander, U. *Abstract inference*. New York: Wiley; 1981.
- Grenander, U. *General pattern theory*. London: Oxford University Press; 1993.
- Grenander, U.; Chow, Y.; Keenan, DM. *Hands: a pattern theoretic study of biological shapes*. Berlin: Springer; 1991.
- Grenander U, Srivastava A, Miller M. Asymptotic performance analysis of Bayesian target recognition. *IEEE Transactions on Information Theory* 2000;46:1658–1665.
- Jermyn I. Invariant Bayesian estimation on manifolds. *Annals of Statistics* 2005;33(2):583–605.
- Joshi SH, Srivastava A, Mio W. Elastic shape models for interpolations of curves in image sequences. *IPMI* 2005:541–552.
- Joshi, SH.; Klassen, E.; Srivastava, A.; Jermyn, IH. Removing shape-preserving transformations in square-root elastic (sre) framework for shape analysis of curves. In: Yuille, A., et al., editors. *LNCS: Vol 4679 Proceedings of energy minimization methods in computer vision and pattern recognition (EMMCVPR)*. Berlin: Springer; 2007a. p. 387–398.
- Joshi, SH.; Klassen, E.; Srivastava, A.; Jermyn, IH. An efficient representation for computing geodesics between n-dimensional elastic shapes. *IEEE conference on computer vision and pattern recognition (CVPR)*; June 2007; 2007b.
- Kass M, Witkin A, Terzopoulos D. Snakes: Active contour models. *International Journal of Computer Vision* 1988;1(4):321–331.
- Kichenassamy S, Kumar A, Olver P, Tannenbaum A, Yezzi A. Gradient flows and geometric active contour models. *Proc international conf computer vision* 1995:810–815.
- Klassen E, Srivastava A. Geodesics between 3D closed curves using path straightening. *European conference on computer vision* 2006;1:95–106.
- Klassen E, Srivastava A, Mio W, Joshi SH. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Transactions Pattern Analysis and Machine Intelligence* 2004;26(3):372–383.
- Leventon M, Eric W, Grimson L, Faugeras O. Statistical shape influence in geodesic active contours. *Proc IEEE conf comp vision and pattern recognition* 2000:316–323.
- Li C, Liu CJ, Fox MD. Segmentation of external force field for automatic initialization and splitting of snakes. *Pattern Recognition* 2005;38(11):1947–1960.
- Michor PW, Mumford D. Riemannian geometries on spaces of plane curves. *Journal of the European Mathematical Society* 2006;8:1–48.
- Miller MI, Younes L. Group actions, homeomorphisms, and matching: a general framework. *International Journal of Computer Vision* 2001;41(1/2):61–84.
- Mio W, Srivastava A. Elastic-string models for representation and analysis of planar shapes. *Proc IEEE conf comp vision and pattern recognition* 2004:10–15.
- Mio W, Srivastava A, Joshi SH. On shape of plane elastic curves. *International Journal of Computer Vision* 2007;73(3):307–324.
- Patrangenaru V, Bhattacharya R. Large sample theory of intrinsic and extrinsic sample means on manifolds. *Annals of Statistics* 2003;31(1):1–29.

- Rousson M, Paragios N. Shape priors for level set representations. European conference in computer vision (ECCV) 2002;II:78–93.
- Samir C, Srivastava A, Daoudi M. Three-dimensional face recognition using shapes of facial curves. IEEE Transactions on Pattern Analysis and Pattern Recognition 2006;28(11):1858–1863.
- Shah, J. H^0 -type Riemannian metrics on spaces of planar curves. 2005. Preprint <http://arxiv.org/abs/math.DG/0510192>
- Small, CG. The statistical theory of shape. Berlin: Springer; 1996.
- Srivastava, A.; Jain, A.; Joshi, S.; Kaziska, D. Statistical shape models using elastic-string representations. In: Narayanan, PJ., editor. LNCS: Vol 3851. Proceedings of Asian conference on computer vision (ACCV). Berlin: Springer; 2005a. p. 612–621.
- Srivastava A, Joshi SH, Mio W, Liu X. Statistical shape analysis: Clustering, learning and testing. IEEE Transactions Pattern Analysis and Machine Intelligence 2005b;27(4):590–602.
- Staib LH, Duncan JS. Boundary finding with parametrically deformable models. IEEE Transactions Pattern Analysis and Machine Intelligence 1992;14(11):1061–1075.
- Van Trees, HL. Detection, estimation, and modulation theory (Vol I). New York: Wiley; 1971.
- Vaillant M, Miller MI, Younes L, Troune A. Statistics on diffeomorphisms via tangent space representations. NeuroImage 2004;23:161–169.
- Vemuri, B.; Chen, Y. Joint image registration and segmentation. In: Osher, S.; Paragios, N., editors. Geometric level set methods in imaging, vision, and graphics. Berlin: Springer; 2003. p. 251–269.
- Wilson D, Geiser E, Larocca J. Automated analysis of echocardiographic apical 4-chamber images. Proc international society for optical eng in math modeling, estimation, and imaging 2000:128–139.
- Xu C, Prince J. Generalized gradient vector flow external force for active contours. Signal Processing 1998a;71(2):131–139.
- Xu C, Prince J. Snakes, shapes, and gradient vector flow. IEEE Transactions Image Processing 1998b;7(3):359–369.
- Yezzi AJ, Mennucci A. Conformal metrics and true “gradient flows” for curves. Proc IEEE international conf computer vision 2005:913–919.
- Younes L. Computable elastic distance between shapes. SIAM Journal of Applied Mathematics 1998;58(2):565–586.

Fig. 1.

(a) A curve, (b) it's log-speed function ϕ , and (c) angle function θ . (d) An arbitrary warping function γ acting on the curve, (e) the new log-speed function $(\phi \circ \gamma + \log \gamma')$, and (f) the new angle function $(\theta \circ \gamma)$. (g) The reparameterized curve

Fig. 2.

Example of a geodesic path between the end shapes using elastic representation for $a = 1$ and $b = 1$

Fig. 3.

From left, **(a)** a shape v , **(b)** the prior mean shape μ , **(c)** a geodesic between them, and **(d)** the vector field $\nabla_{\beta} E_{prior}$ (arrows) overlaid on the curve (v, x) , where $E_{prior} = d_g(v, \mu)^2$

Fig. 4.
Left: Some samples from a shape population. *Right:* Mean shape and singular values of covariance in $T_{\mu}S$

Fig. 5. (a) Cartoon diagram of shape evolution under the squared distance, normal energy, and Rayleigh quotient in the space $T_\mu(S)$. (b) Successive shapes along gradient paths of the three energies

Fig. 6.

(a) Eigen variation (*column wise*) along dominant vectors v_1, v_2, v_3 and some random directions $v_{rand} \in M^\perp \subset T_\mu(S)$. (b) Shape variation about the mean μ (*center*) for the squared-distance (*circle*), and the normal (ellipse) energy

Fig. 7.

Top row shows arbitrarily initialized curves. *Middle row* shows the image gradient field $\mathbf{v}$. *Bottom row* shows the induced vector field $(\nabla_{\beta} E_{image})$ on the curves

Fig. 8.

Top row shows the evolution of β under $\nabla_{\beta} E_{smooth}$. Bottom row shows the evolution of β under $\nabla_{\beta} E_{smooth} + \nabla_{\beta} E_{prior}$. All the curves have been scaled to the same length for display, while the evolutions of their actual lengths are shown in the rightmost panels in each row

Fig. 9.

Evolution of β under $\nabla E_{total}(\beta)$ with $\lambda_3 = 0$ and $\lambda_3 > 0$. The *last panel* shows the prior mean μ

Fig. 10. Evolution of β under $\nabla E_{total}(\beta)$ with $\lambda_3 > 0$ for the same image, but for different prior mean shapes. The *last panel* shows the prior mean μ in each row

Fig. 11. Evolution of β under $\nabla E_{total}(\beta)$ with $\lambda_3 = 0$ and $\lambda_3 > 0$. The last panel shows the prior mean μ

Fig. 12.

Bayes' estimates of bone shapes. In *each row*, the *left panel* shows the image I , the *second panel* shows the training shapes used to build a prior, the *third panel* shows the prior mean shape μ , the *fourth panel* shows β^* estimated by the algorithm, and the *last panel* shows an expert's hand-drawn estimation

Fig. 13.
Odd rows: Evolution of β under $\nabla E_{total}(\beta)$ with $\lambda_3 = 0$. *Even rows:* Evolution of β under all three terms including E_{prior}

Fig. 14. Final extracted curves under varying effects of Normal, and the squared-distance prior energy

Fig. 15.

Top panel: Sample shapes consisting of ellipses and circles, their mean shape, and shape variations along two dominant eigenvectors. *Bottom panel:* Curve evolution under two different E_{prior} : squared-distance and normal energy. The last two rows show a corrupted image of an ellipse with different aspect ratios

Fig. 16.

A few examples of IR images of tanks for a sweep of 207° to 312° . *First column* shows the elevation of the tank

Fig. 17.
Left panel: Average shape of 108 observations of tanks silhouettes shown on the *top*, with the decay in the energy of the singular values shown on the *bottom*. *Right panel:* Eigen-shapes corresponding to 1'st, 2'nd and 3'rd principal components, with the mean shape shown by a *solid line*. The eigenvalues are shown at the *bottom*

Fig. 18.
Examples of curve evolution under the Squared-Distance and the Normal energy

Fig. 19.

(a) Experimental setup for image captures of different views. (b) Shapes for each quadrant q_i , and their mean shapes. (c) Examples of infrared images of a tank at different poses

Fig. 20.

Each vertical section shows a quadrant number, the extracted curves $\hat{\beta}_q$ from I for respective quadrant, and the prior shape model used for that quadrant

Fig. 21.
 E_{total}^q plotted against prior models $q_i \in \mathcal{Q}$

Table 1
Shape, and nuisance variables associated with appearance

Shape Variables			Nuisance Variables		
Symbol	Definition	Space	Symbol	Definition	Space
ϕ	$\log(\frac{2\pi}{l} \ \beta\)$, log-speed function	$\mathbb{L}^2(J)$	1. Appearance Variables		
θ	$\angle(\frac{\dot{\beta}}{\ \beta\ })$, angle function	$\mathbb{L}^2(J)$	l	$\int \ \beta(s)\ ds$, length of the curve	$\mathbf{R}_+$
			τ	$\int \theta(s) e^{ik(s)} ds$, orientation of the curves	$\mathbf{S}^1$
			p	$\beta(0)$, placement of the curve	$\mathbf{R}^2$
			2. Matching Variables		
			r	Placement of origin on the curve	$\mathbf{S}^1$
			γ	Origin-preserving diffeomorphism of J	$\mathcal{D}$

Algorithm 1

Projection of a curve $(\phi, \theta) \in \mathcal{H}$ into $\mathcal{C}$.

```

1:  Let  $n = 4$ . Initialize  $l(\phi, \theta)_i = \mathbf{1}_n$ . Let  $\varepsilon > 0$ .
2:  while  $\|l(\phi, \theta)\| > \varepsilon$  do
3:    Compute  $l(\phi, \theta)_i = -\mathcal{G}_i(\phi, \theta)$ ,  $i = 1, \dots, n$ .
4:    Calculate the Jacobian matrix,  $J_{i,j} = \langle \nabla \mathcal{G}_i(\phi, \theta), \nabla \mathcal{G}_j(\phi, \theta) \rangle$ 
5:    Solve the equation  $J(\phi, \theta)x^T = l^T(\phi, \theta)$  for  $x$ .
6:    Update  $(\varphi, \theta) = (\varphi, \theta) + \sum_{i=1}^n x_i \nabla \mathcal{G}_i(\theta) \delta$ ,  $\delta > 0$ .
7:  end while

```

Algorithm 2

Parallel Transport of A Tangent Vector $g \in T_{\mu}(S)$ to $T_{\nu}(S)$

-
- 1: Denote the geodesic from μ to ν by ψ_t , such that $\psi_0 = \mu$ and $\psi_1 = \nu$.
 - 2: Choose a small $\delta > 0$ to divide the interval $[0, 1]$ in intervals of size δ . Compute the length of g according to $l_g = \sqrt{\langle g, g \rangle} \mu$.
 - 3: Update the tangent vector (h, f) as $(\tilde{h}, \tilde{f}) = (h, f) - \delta(\frac{h^2}{2} - \frac{b}{2a}f^2, hf)$.
 - 4: Project $(\tilde{h}, \tilde{f})$ onto $T_{\psi_{(t+1)\delta}}(S)$ using (9) and re-scale it be of length l_g . Repeat this for each point along the geodesic.
-