

Intrinsic Bayesian Active Contours for Extraction of Object Boundaries in Images *

Shantanu H. Joshi¹ and Anuj Srivastava²

¹ Department of Electrical Engineering, Florida State University, Tallahassee, FL, 32310 USA,

² Department of Statistics, Florida State University, Tallahassee, FL, 32306 USA

The date of receipt and acceptance will be inserted by the editor

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 probability models are constructed on the tangent bundle of shape space. Past work on boundary extraction has used active curves driven by vector fields that were based on image gradients and roughness penalties. The proposed method incorporates *a priori* knowledge of shapes in the form of gradient fields in addition to the previously used image vector fields. Through experimental results, we demonstrate the use of prior shape models in 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.

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. In some applications, this task has to be performed solely on the basis of image data, while in others some prior information about the shape 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 interesting objects 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 interesting objects in images. From this perspective, our technique can be viewed as a tool for 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 further inferences, performance in boundary estimation can be measured directly in terms of eventual object classification or recognition performance.

* This paper was presented in part at the Asian Conference on Computer Vision, January 2006

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

1.1 Current Methods for Curve Extraction

The seminal paper by Kass et al [17] initiated the area of curve extraction; it considered an energy minimizing spline or 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. To remedy them, Cohen et al [5] improved the model by including a balloon force which expands the curve till it rests on edges. The Gradient Vector Flow (GVF) [36] and generalized GVF (GGVF) snakes [35] proposed by Xu and Prince 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 a strong and a weak edge are in close proximity. Li et al [21] proposed an Edge Preserving Gradient Vector Flow (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 [2] 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. [18] formulated a geometric curve evolution approach based upon curvature flows in images. Both of these approaches use Euclidean curve-shortening flows that evolve the contour according to its curvature in the normal direction. A desirable property of these curvature flows, or Euclidean heat flows, is the naturally occurring smoothing of the curve [11,12]. 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 [2,18]. While level-set representations allow for 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 [30] 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 expected 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 [10,27]. This process however requires an 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 [13,23] considers shapes as points on infinite-dimensional manifolds, where the variations between the shapes are modeled by the action of diffeomorphism groups [13]. 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. 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. However, to impose statistical models, certain approaches often embed shapes in larger Euclidean spaces and analyze them as elements of vector spaces. For example, active shape model (ASM) [6] uses principal component analysis (PCA) of shapes, represented via vectors of landmarks, to model shape variability. Despite its simplicity and efficiency, this approach is extrinsic and 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 [20]. More recently, there has been a resurgence of ideas in studying shapes of close, planar curves [19, 24, 22, 26, 37, 3] that use differential geometry of the manifolds of such curves to develop tools for shape analysis. An important ongoing discussion in this literature is the choice of Riemannian metric for use in different applications. On the other hand, there has been some progress in developing parametric statistical models intrinsically on these spaces [29, 28].

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. [14, 1] used shape models learned from past observations to help improve shape estimation in future data. Cremers et al. [7, 8] have used priors in several formulations to perform segmentation and object extraction, with good results. A very important consideration in shape analysis is invariance of shapes to similarity transformations. In certain approaches, such as level-set methods or diffeomorphism based methods, it is difficult to build representations that are fully invariant. Consequently, any use of shape priors involves additional minimizations on similarity groups. To understand this issue, consider the problem 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 **extrinsic** distances ([33, 9]) between shapes of two closed curves β_1 and β_2 are:

$$d(\beta_1, \beta_2) = \operatorname{argmin}_{R \in SO(2), T \in \mathbb{R}^2, a \in \mathbb{R}_+} \|\phi(\beta_1) - \phi(R \cdot a\beta_2 + T)\|^2, \quad \text{or} \quad \operatorname{argmin}_{R \in SO(2), T \in \mathbb{R}^2, a \in \mathbb{R}_+} \|\beta_1 - (R \cdot a\beta_2 + T)\|^2$$

where R denotes a rotation, T denotes a translation, and a denotes a scale. Such extrinsic shape analysis tends to be expensive because of these additional optimizations. Moreover, the norms are often computed on ambient Euclidean spaces and not on spaces of closed curves or spaces of signed-distance functions. For example, in case of signed-distance functions, if $\|\phi_1 - \phi_2\|$ (or its version with the Heaviside step function) is used to define a distance, it is difficult to show existence of a path on the space of signed-distance functions that actually achieves this distance. This fact is important for defining and computing statistics intrinsically on shape spaces [25]. A recent paper [8] has successfully achieved several desired invariances for shape analysis using level set representations. Still, it is not fully intrinsic for two reasons: the distances are computed on an ambient space of all functions (not just signed distance functions) and not all similarity transformations are removed. One consequence is that in the resulting probability models one will encounter, with high probability, functions that are invalid for representing shapes. Another consequence is that shape statistics defined extrinsically will not be unique; they will depend upon: (i) the ambient space and its metric, and (ii) the embedding.

1.4 Our Approach: Bayesian Active Contours

The main contribution of this paper is to apply fully intrinsic as well as fully 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 geometric, intrinsic representations of shapes of continuous, closed, planar curves in imposing shape priors and using these priors in a Bayesian framework for shape extraction. To represent and analyze shapes of closed curves, we use the approach presented in [19, 24], although any of the recent ideas [22, 26, 37] are also applicable. 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} . To remove desired similarity transformations, one forms a quotient space $\mathcal{S} = \mathcal{C}/S$, where S is the space of similarity transformations, and shapes of closed curves are analyzed as elements of \mathcal{S} . The main tool in this analysis is the construction of geodesic paths on \mathcal{S} between any two given shapes.

Using geodesic constructions, one can develop a framework for statistical models by defining means, covariances, and other statistics in tangent spaces of \mathcal{S} [29,32]. Finally, the problem of boundary extraction is posed as maximum a-posteriori (MAP) estimation, which in terms of energies is rewritten as:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} (E_{image}(\beta) + E_{smooth}(\beta) + E_{prior}(\beta)) , \quad (1)$$

where the minimization is over all closed curves in the given image domain. E_{prior} is the novel term representing prior energy associated with shape of curve β . To ensure intrinsic formulation, all the terms will be independent of re-parameterizations of β . Furthermore, E_{smooth} will be independent of translation and rotation, and E_{prior} will be independent of translation, rotation, and global scale of β . The first two terms have been used previously in the literature for driving *active contours*. Adding the third term results in *Bayesian active contours*.

The proposed framework has the following advantages:

1. In addition to image information and roughness constraints, it incorporates an *a priori* information about shapes in boundary extraction. This additional information is presented in the form of deformation vector fields on curves for efficient energy minimization.
2. This approach does not require expert-generated landmarks, diffeomorphic embedding, or level-set function for shape analysis. Instead, it studies full curves of shapes. 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.
4. 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.

One limitation of this approach, as presented here, is that 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.

The remainder of this paper is organized as follows. Section 2 summarizes a framework for representing and analyzing shapes of closed curves. Section 3 explains the main idea of this paper – an intrinsic probability measure on the shape space \mathcal{S} . It details two candidates for prior shape models and points out basic differences between them. Section 4 outlines different energy terms that form the posterior energy and the algorithm for finding MAP estimates of curves. Section 5 presents experimental results that demonstrate the main ideas in the paper. Section 6 shows an idea for Bayesian inference that allows selection and testing of different prior shapes models on a given image data followed with conclusion.

2 Elastic Shape Representations

In this section we summarize ideas from *elastic shape analysis* of planar closed curves; for details the interested reader is referred to [24]. Let β be a parameterized curve of interest of length l , and define a re-scaled version $\alpha = 2\pi\beta/l$. We will assume $\alpha: [0, 2\pi] \rightarrow \mathbb{R}^2$ to be a non-singular, parametric curve in the sense that $\dot{\alpha}(s) \neq 0, \forall s \in [0, 2\pi]$. Define the velocity vector of the curve as $\dot{\alpha}(s) = e^{\phi(s)} e^{j\theta(s)}$, where $\phi: [0, 2\pi] \rightarrow \mathbb{R}$ and $\theta: [0, 2\pi] \rightarrow \mathbb{R}$ are smooth, and $j = \sqrt{-1}$. ϕ is called the *speed* function of α and measures the rate of stretching and compression, whereas θ is the angle made by $\dot{\alpha}(s)$ with the X -axis and measures bending. We will represent α via the pair $\nu \equiv (\phi, \theta)$, 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

$$\mathcal{C} = \{(\phi, \theta) : \int_0^{2\pi} e^{\phi(t)} dt = 2\pi, \frac{1}{2\pi} \int_0^{2\pi} \theta(t) e^{\phi(t)} dt = \pi, \int_0^{2\pi} e^{\phi(t)} e^{j\theta(t)} dt = 0\} \subset \mathcal{H}, \quad (2)$$




Fig. 1 Shown from left is the original shape, its ϕ and θ representation, an arbitrary time warp γ , ϕ and θ functions after the transformation, and the re-parameterized shape.

where \mathcal{C} is called the *pre-shape space*.

Note that the pair (ϕ, θ) represents the shape of β and, thus, ignores its placement, orientation, and scale. These variables are called **nuisance variables** in the study of shapes, as they do not contribute to the analysis of shapes. However, in extraction of boundaries from images, these nuisance variables become important and also need to be estimated from data. We will use x to denote these nuisance variables, $x \in \mathcal{N} \equiv (\mathbb{S}^1 \times \mathbb{R}^2 \times \mathbb{R}_+)$. The full contour β is determined both by its shape and its nuisance variables $\beta \equiv ((\phi, \theta), x) \in \mathcal{B} \equiv \mathcal{C} \times \mathcal{N}$, according to: $\beta(t) = x_2 + x_3 \int_0^t e^{\phi(s)} e^{j(\theta(s) + x_1)} ds$, where $x_2 \in \mathbb{R}^2$, $x_3 \in \mathbb{R}_+$, and $x_1 \in \mathbb{S}^1$ denote the translation, scale, and rotation components of x , respectively. To emphasize dependence of β on shape ν and nuisance variables x , we will often use the notation $\beta = \beta(\nu, x)$, where $\nu = (\phi, \theta)$.

Geodesics are important in defining and computing statistics of shapes. To specify a geodesic, we use the following Riemannian metric: Given $(\phi, \theta) \in \mathcal{C}$, let h_i and f_i , $i = 1, 2$ be tangent to \mathcal{C} at (ϕ, θ) . For $a, b > 0$, define

$$\langle (h_1, f_1), (h_2, f_2) \rangle_{(\phi, \theta)} = a \int_0^{2\pi} h_1(s) h_2(s) e^{\phi(s)} ds + b \int_0^{2\pi} f_1(s) f_2(s) e^{\phi(s)} ds. \quad (3)$$

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. With this metric, geodesics are completely characterized by a starting point and a starting direction. Let $\Psi_t(\nu, g)$ denote a geodesic starting from $\nu \in \mathcal{C}$ in the direction g , and parameterized by time t , and let $T_\nu(\mathcal{C})$ denote the vector space of all elements tangent to \mathcal{C} at ν . The value of a geodesic at $t = 1$ is also known as the exponential map, i.e. $\exp_\nu : T_\nu(\mathcal{C}) \mapsto \mathcal{C}$, such that $\exp_\nu(g) = \Psi_1(\nu, g)$. Given two shapes ν_1 and ν_2 in \mathcal{C} , computing a geodesic between them involves finding an optimal direction $g \in T_{\nu_1}(\mathcal{C})$, such that $\Psi_1(\nu_1, g) = \nu_2$. Similar to ideas presented in [19], we use a shooting method to construct geodesic paths in \mathcal{C} . The shooting method solves for g such that a geodesic from ν_1 in the direction g reaches the target shape ν_2 in unit time.

So far we have looked at geodesics in \mathcal{C} but we need to remove some (shape) invariance present in elements of \mathcal{C} , as follows. Note that the following re-parameterizing actions do not change the shape of a closed curve: (i) The change of origin on a shape in \mathcal{C} is represented by the action of a unit circle \mathbb{S}^1 on a shape (ϕ, θ) , according to $r \cdot (\phi(s), \theta(s)) = (\phi(s - r), \theta(s - r) - r)$, for $r \in \mathbb{S}^1$. (ii) Traversing a shape in \mathcal{C} with different velocities, while maintaining orientation, constitutes variable-speed re-parameterizations of the same shape. Such re-parameterizations define a group action \mathcal{D} of diffeomorphisms, $\gamma \in \mathcal{D}$ on \mathcal{C} , given by $(\phi, \theta) \circ \gamma = (\phi \circ \gamma + \log \dot{\gamma}, \theta \circ \gamma)$, $\gamma : [0, 2\pi] \rightarrow [0, 2\pi]$, $\dot{\gamma}(s) > 0, \forall s \in [0, 2\pi]$. Figure 1 shows an example of shape representation and illustrates the action of the re-parameterization group. The space of all (shape preserving) re-parameterizations of a shape in \mathcal{C} is thus given by $\mathbb{S}^1 \times \mathcal{D}$. We define the elastic shape space as an equivalence class given by $\mathcal{S} = \mathcal{C}/(\mathbb{S}^1 \times \mathcal{D})$, and construct geodesics in \mathcal{S} by appropriately restricting geodesics in \mathcal{C} . For details, please refer to [24]. Shown in Figure 2 is an example of a geodesic path.



Fig. 2 Example of a geodesic path between the end shapes using elastic representation.




Fig. 3 From left, (a) A shape ν , (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 (ν, x) , where $E_{prior} = d_g(\nu, \mu)^2$.

3 Construction of Prior Shape Models

The main contribution of this paper is the use of intrinsic prior probability models on \mathcal{S} in extracting boundaries in images. We will specify the prior using an energy E_{prior} that sometimes can be seen as negative log of a prior density. For our purpose, 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. It should be noted 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 possible spaces for imposing probability measures: the manifold itself and its tangent bundle. We will study both the cases, starting with the former.

3.1 Squared Distance as Prior Energy

A simple choice for prior energy is $E_{prior} = d(\nu, \mu)^2$ where d denotes the geodesic distance connecting the shape $\nu \equiv (\phi, \theta)$ with the prior mean μ . Since the geodesic distance is invariant to all similarity transformations, including the re-parametrization 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 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 $-\log(\text{prior density})$. The choice of the underlying invariant measure that results in an invariant estimation will be another difficult issue that we will avoid [16].

The gradient of E_{prior} with respect to ν is given by $\nabla_{\nu} E_{prior} = g \in T_{\nu}(\mathcal{S})$, such that $\Psi_1(\nu, g) = \mu$. What we really need is $\nabla_{\beta} E_{prior}$, the gradient of E_{prior} with respect to the curve β . This is a vector field on the curve β , that can be more easily approximated numerically than analytically. Let ν and x be the shape and nuisance variables of β , respectively; we can write $\beta = \beta(\nu, x)$. Define a vector field:

$$\nabla_{\beta} E_{prior}(s) = (\beta(\Psi_{\epsilon}(\nu, g), x)(s) - \beta(\nu, x)(s)), \quad \text{for } \epsilon > 0 \text{ small.} \quad (4)$$

$\beta(\Psi_{\epsilon}(\nu, g), x)$ is the closed curve that results when the shape component of $\beta(\nu, x)$ is perturbed along the tangent direction g by an amount ϵ . The size (norm) of this field is not important as we will multiply this term with an arbitrary scalar later on. Figure 3 shows a sample shape ν , 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 form of E_{prior} cannot incorporate any information about observed directions containing the most variation in a shape class around μ . To allow for more structure, we will utilize the second idea of imposing probability measures on tangent spaces of \mathcal{S} .

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 [10, 32, 29, 28]. It is also common to impose a multivariate normal density on 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 \mathcal{S} . In the implementation, however, one can generate (random) samples in tangent spaces and map them freely to \mathcal{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 broad idea is to estimate a mean shape μ for an object class, and to use the tangent space $T_\mu(\mathcal{S})$, or rather a finite-dimensional subspace of $T_\mu(\mathcal{S})$ obtained using PCA, for imposing probabilities. This approach has been termed *tangent* PCA (TPCA). The mean μ is taken to be the Karcher mean [29, 28]. Figure 4 shows an example of a Karcher mean of some observed shapes. Algorithm 1 outlines the main computational steps for the TPCA approach.

Algorithm 1 TPCA in shape space \mathcal{S}

- 1: For an ensemble of k training shapes $\{\nu_i, i = 1, \dots, k\}$, compute their Karcher mean μ .
For each ν_i , compute a tangent vector $g_i \in T_\mu(\mathcal{S})$ such that $\Psi_1(\mu, g_i) = \nu_i$.
 - 2: Compute an orthonormal basis for the set $\{g_i\}$ with respect to the inner product given in Eqn. 3 using the Gram-Schmidt process. Let the new basis elements be $\{f_i, i = 1, \dots, m\}$ and define $M \equiv \text{span}(f_1, f_2, \dots, f_m) \subset T_\mu(\mathcal{S})$ with $m \leq k$.
 - 3: 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 $j = 1, 2, \dots, m$.
 - 4: 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 $N \subset M \subset T_\mu(\mathcal{S})$ 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 .
-

One can impose one of the 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 [28]. For simplicity of presentation we take the first idea and form a multivariate normal density on the subspace M of $T_\mu(\mathcal{S})$. Let $g \in T_\mu(\mathcal{S})$ be a tangent vector such that $\exp_\mu(g) = \nu$ and let y be the projection of g into M . Now define a prior energy as

$$\tilde{E}_{prior}(y) = \frac{1}{2} y^T (UK^{-1}U^T)y + \frac{1}{2\delta^2} \|y - UU^T y\|^2. \quad (5)$$

$K \in \mathbb{R}^{n \times n}$ 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 . The gradient $\nabla_y \tilde{E}_{prior} = Ay$, where $A = UK^{-1}U^T + \frac{1}{\delta^2}(I_m - UU^T)$. 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 a simple change of variables $y \mapsto (g \equiv \sum_{i=1}^m y_i f_i)$. For instance, $E_{prior}(g) = \tilde{E}_{prior}(\{\langle g, f_j \rangle_\mu, j = 1, \dots, m\})$, and its derivative $\nabla_g E_{prior} = \sum_{j=1}^m (\nabla_y \tilde{E}_{prior})_j f_j$.

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 ν , using the exponential map $g \mapsto \nu = \exp_\mu(g)$. However, writing these terms explicitly will introduce burdensome notation due to complicated change of variables under \exp_μ . 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:

$$\nabla_\beta E_{prior}(s) = (\beta(\Psi_1(\mu, g + \epsilon \nabla_g E_{prior}), x)(s) - \beta(\nu, x)(s)), \quad \text{for } \epsilon > 0 \text{ small}, \quad (6)$$




Fig. 4 Left: Sample shapes from a cluster. Right: Mean shape and singular values of covariance in $T_\mu \mathcal{S}$.

keeping in mind that $\Psi_1(\mu, g) = \nu$. 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 Figure 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 Figure 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 ν . Since the distance-squared term is simply the squared length of a tangent vector, the gradient process will simply follow this straight line from ν towards μ in the tangent space $T_\mu(\mathcal{S})$. The normal energy term however will form a curved path from ν to μ favoring lower energy shapes (under the normal energy). Both these paths are shown in Figure 5(a) and successive shapes along these paths are shown in Figure 5(b). Together these plots demonstrate that gradient of normal energy will find low energy shapes more quickly and closer to ν 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 Figure 15 using a synthetic image.) For 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 ν is constrained to be unit length. As expected, the limiting shape is the shape along the dominant eigenvector of K . As 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). We can see that even a small variation from μ along random directions results in a 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 better represent the variability of the observed ensemble 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 , $\nu(\beta)$. 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, translations, and scales, though in the future, more informative priors can be pursued.

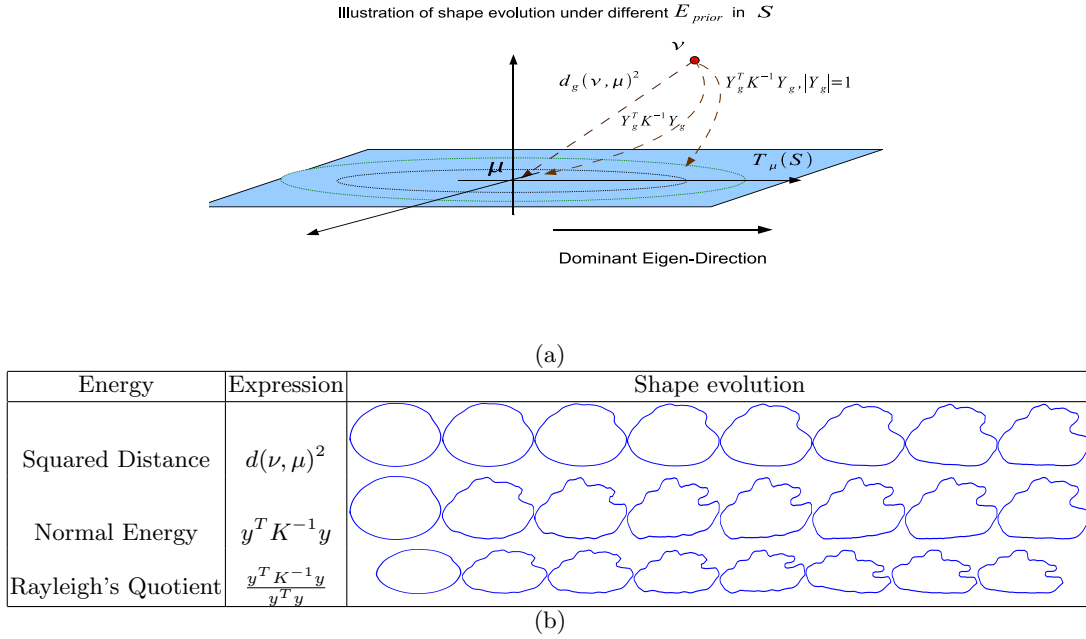


Fig. 5 (a) Cartoon diagram of shape evolution under the squared distance, normal energy, and Rayleigh quotient. (b) Successive shapes along gradient paths of the three energies.

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 gradient information and a smoothing criterion, and this paper adds a term on shape prior, resulting in:

$$E_{total}(\beta) = \lambda_1 E_{image}(\beta) + \lambda_2 E_{smooth}(\beta) + \lambda_3 E_{prior}(\beta), \quad (7)$$

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 . 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,

$$\hat{\beta} = \operatorname{argmin}_{\beta \in \mathcal{B}} E_{total}(\beta), \quad (8)$$

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

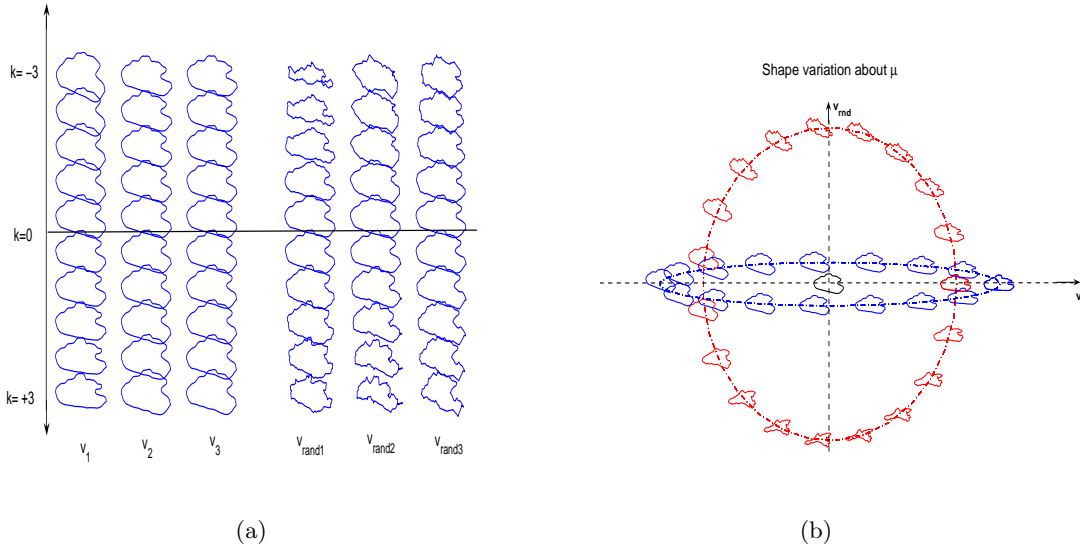


Fig. 6 (a) Eigen variation (columnwise) along dominant vectors v_1, v_2, v_3 and some random directions $v_{rand} \in M^\perp \subset T_\mu(S)$. (b) Shape variation about the mean μ (center) for the squared distance (circle), and the normal (ellipse) energy.

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. [17] 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. [2] also suggest the energy function: $\frac{1}{1+|\nabla I(x, y)|^p}$, $p = 1$ 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. [36, 35] present an energy directly in terms of vector fields $\mathbf{v}(x, y) = (u(x, y), v(x, y))$: $\int g(|\nabla f|)|\nabla \mathbf{v}|^2 + h(|\nabla f|)(|\mathbf{v} - \nabla f|^2) dx dy$, where f is one of the energy functions suggested by Kass et al. The papers [36, 35] 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 [21] proposed an edge preserving gradient vector flow (EPGVF) to overcome this drawback of GVF and GGVF, by minimizing the energy:

$$\mathcal{E}_{ext}(\mathbf{v}) = \int g(|\nabla f|)|\nabla \mathbf{v}|^2 + h(|\nabla f|)(\mu|J_{\mathbf{v}}p|^2 + |\mathbf{v} - \nabla f|^2) dx dy. \quad (9)$$

The term $\mu|J_{\mathbf{v}}p|^2$ in Eqn. 9 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, $\mathbf{v}^* = \operatorname{argmin}_{\mathbf{v}} \mathcal{E}_{ext}(\mathbf{v})$ using a numerical iterative approach [21]. From now onwards we will only consider this optimal vector field \mathbf{v}^* , and to simply 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,

$$E_{image}(\beta) = \int_0^l |\mathbf{v}(\beta(s))|^2 (|\dot{\beta}(s)|) ds \quad (10)$$

Proposition 1 *The energy term E_{image} is invariant to the action of \mathcal{D} on \mathcal{C} . In other words, $E_{image}(\beta)$ does not change with the re-parameterizations of β .*

Proof Replacing β by $\tilde{\beta} = \beta(\gamma)$, for any $\gamma \in \mathcal{D}$, we obtain:

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

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 parametrization of β , we simplify the gradient by assuming arc-length parametrization 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 v with respect to β , given by $\begin{bmatrix} \frac{\partial v_x}{\partial x} & \frac{\partial v_x}{\partial y} \\ \frac{\partial v_y}{\partial x} & \frac{\partial v_y}{\partial y} \end{bmatrix}$. 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 .

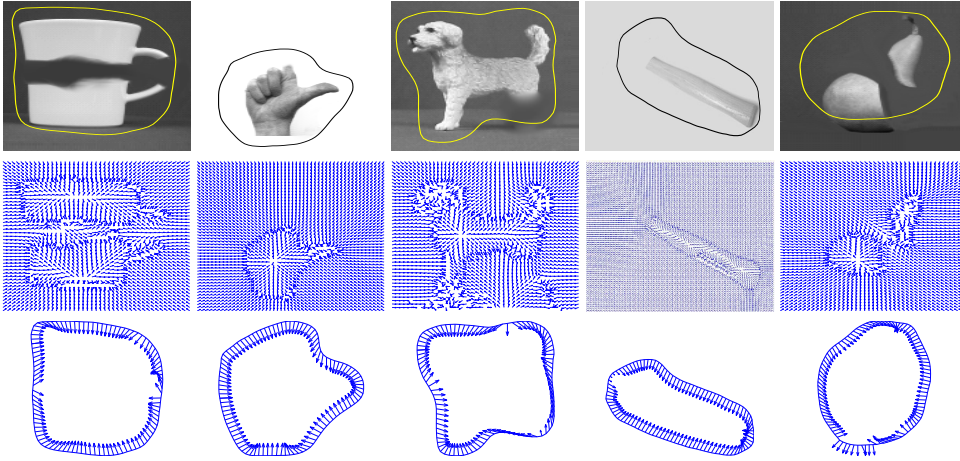


Fig. 7 Top row shows arbitrarily initialized curves. Middle row shows the image gradient field \mathbf{v} . Bottom row shows the induced vector field ($\nabla_{\beta} E_{image}$) on the curves.

4.2 Smoothing Energy

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 [17,36] 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 parametrization, we follow a more intrinsic approach following [2,18]. We consider an arc-length parameterized curve β . Define the smoothing energy functional as:

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




Fig. 8 Evolution of β under $\nabla_{\beta} E_{smooth}$. All curves are scaled to the same length for display.

It is easy to see that E_{smooth} is invariant to any re-parametrization of β . As shown in many articles, e.g. [18], the gradient of E_{smooth} with respect to β is given by $-\frac{d}{ds}(\frac{\dot{\beta}(s)}{|\dot{\beta}(s)|})$. 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 parametrization resulting in $\|\dot{\beta}(s)\| = 1$ and

$$\nabla_{\beta} E_{smooth}(\beta) = -\ddot{\beta}(s) = -\kappa(s) n(s), \quad (11)$$

where, κ is the curvature along β and $n(s)$ is the inward normal unit vector to β at s . It has been shown [11,12] 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, Figure 8 shows the evolution of a curve under $\nabla_{\beta} E_{smooth}$. All the curves have been scaled for a better display.

From the three different energy terms, we can write the total gradient vector field as,

$$\nabla E_{total}(\beta) = \lambda_1 J_{\beta}(\mathbf{v})(\beta(s)) \mathbf{v}(\beta(s)) - \lambda_2 \kappa(s) n(s) + \lambda_3 \nabla_{\beta} E_{prior}(s), \quad (12)$$

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

$$\frac{dX(t)}{dt} = -\nabla E_{total}(X(t)), \quad X(0) \in \mathcal{B}, \quad (13)$$

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.

5 Experimental Results

In this section we describe 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, in addition to that, we have a prior knowledge about the shape of 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. In this section we present some examples.

Even though $E_{total}(\beta)$ is invariant to re-parametrization of β , the gradient term simplifies if the arc-length parametrization is assumed. Therefore, we need to re-sample β uniformly after every iteration to ensure arc-length parametrization. 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

$$(\phi, \theta) \circ \gamma, \quad \text{where} \quad \gamma(t) = \int_0^t e^{-\phi(\gamma(s))} ds. \quad (14)$$

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

1. **Example 1: (Visible Spectrum Images):** In this case we use images of a tank, shown in Figure 9, to demonstrate Bayesian active curves under a shape prior. The right panel 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 Figure 9 are the evolutions of $X(t)$, with $(\lambda_3 > 0)$ and without the prior term $(\lambda_3 = 0)$. In the absence of prior, the algorithm is misled by a partial obscuration of the object in I . In contrast, the Bayesian solution seems a better estimate of the tank outline as compared to when the final result only depends on the image data.




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 .

2. **Example 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 [34, 4] include 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 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, and use these intermediate curves as initial conditions for boundary extraction using minimization of E_{total} . Figure 10 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. **Example 3: (Natural and Synthetic Images):** Next we present boundary extraction results for a few images from the ETH-80 dataset, as well as some natural partially obscured images.




Fig. 10 Evolution of β under $\nabla E_{total}(\beta)$ with $\lambda_3 = 0$ and $\lambda_3 > 0$. The last panel shows the prior mean μ .

Each row in Figure 11 shows curve evolutions with and without the E_{prior} term; the upper row is for $\lambda_3 = 0$ and the lower row is for $\lambda_3 > 0$.

4. **Experiment 4: (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 12 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, eventually for large λ_3 the most probable shape is simply μ .
5. **Experiment 5: (Locally Random Search):** So far the estimated boundary has been obtained by implementing a gradient flow that typically yields locally optimal solutions. Moreover, depending upon the initial condition, the method may converge rather slowly to the desired answer. Algorithm 2 presents a stochastic method that uses localized random search for curve extraction. Figure 13 shows curve evolutions for arbitrarily initializations with and without the effect of the random vector field. Figure 14 shows the evolution of the energy E_{total} for both cases. It is observed that the random algorithm requires lesser iterations to converge although we have not established such a result analytically. In this particular experiment, we set $\lambda_3 = 0$ as the prior term is not important for this issue.
6. **Experiment 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. Fig. 5(b) illustrates the difference in evolution when the squared distance and the normal energy are used. We compare the effect of these two terms in an image-based experiment involving only primitive shapes such as circles and ellipses. The top panel of Fig. 15 shows some sample shapes and their Karcher mean μ . A multivariate normal distribution is imposed on the tangent space at the mean shape using Algorithm 1. As expected, the variation in this class (Fig. 15) is mostly along a single, dominant eigenvector. 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. 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.

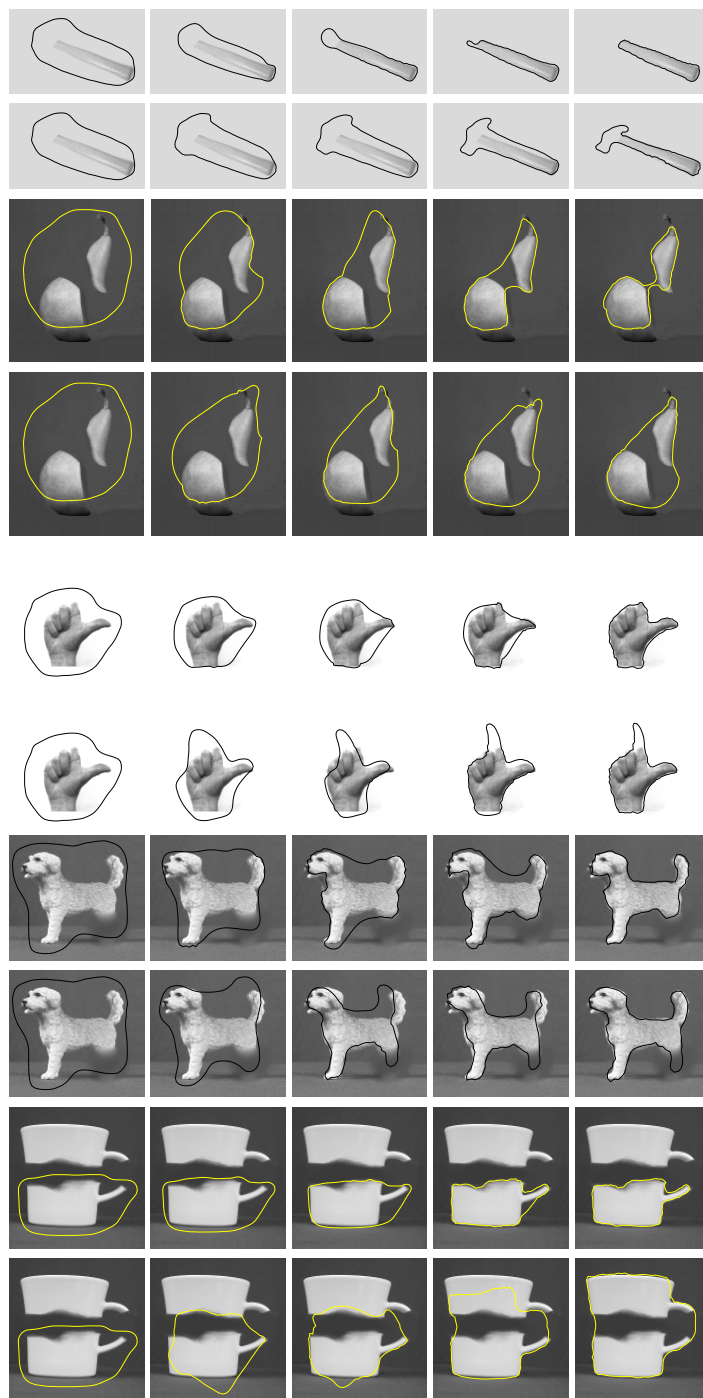


Fig. 11 Odd rows: Evolution of β under $\nabla E_{total}(\beta)$ with $\lambda_3 = 0$. Even rows: Evolution of β under all three terms including E_{prior} .

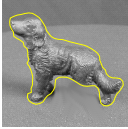
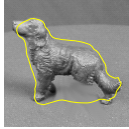
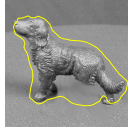
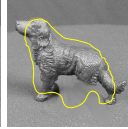
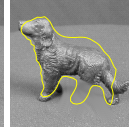
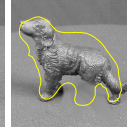
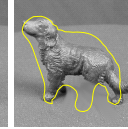
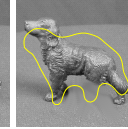
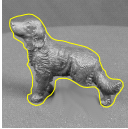
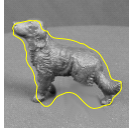
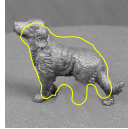
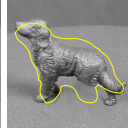
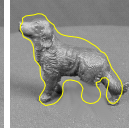
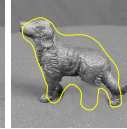
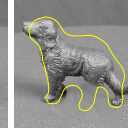
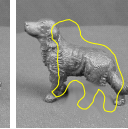
E_{prior}	Final Extracted Curves for different values of λ_3							
Squared-Distance								
Normal Energy								
λ_3	0	0.1	0.2	0.25	0.3	0.4	0.6	0.9

Fig. 12 Final extracted curves under varying effects of Gaussian, and the squared distance prior energy.

Algorithm 2 Localized Random Search using vector fields

Initialize curve β on the image I . Let the shape of β be ν .

while not converged **do**

 Calculate $\mathbf{v}_{image}(\beta)$, and $\mathbf{v}_{smooth}(\beta)$

 Generate $\mathbf{v}_{rand}(\beta)$ as follows:

a: Let a tangent vector $g = \sum a_n \cos nx + \sum b_n \sin nx$

 where $a_n, b_n \sim N(0, \sigma), n = 0, 1, \dots, N$

b: Let the randomly perturbed shape be $\nu^* = \exp_\nu(\bar{0}, g)$

c: Calculate $\beta^* = (\nu^*, x)$, where $x \in (\mathbb{S}^1 \times \mathbb{R}^2 \times \mathbb{R}_+)$ is the pose.

d: Now compute the random vector field as $\mathbf{v}_{rand}(\beta) = \lambda_4(\beta - \beta^*)$

 Obtain β_{new} from β , and $\nabla E_{total}(\beta) = \lambda_1 \mathbf{v}_{image} + \lambda_2 \mathbf{v}_{smooth} + \lambda_3 \mathbf{v}_{prior} + \lambda_4 \mathbf{v}_{rand}$

if $E_{total}(\beta_{new}) > E_{total}(\beta)$ **then**

$\beta = \beta_{new}$

end if

end while

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 prior can be used in 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 prior 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 maximum-a-posteriori (MAP) framework. In other words, Bayesian approach suggests searching over all possible shape models and selects the model

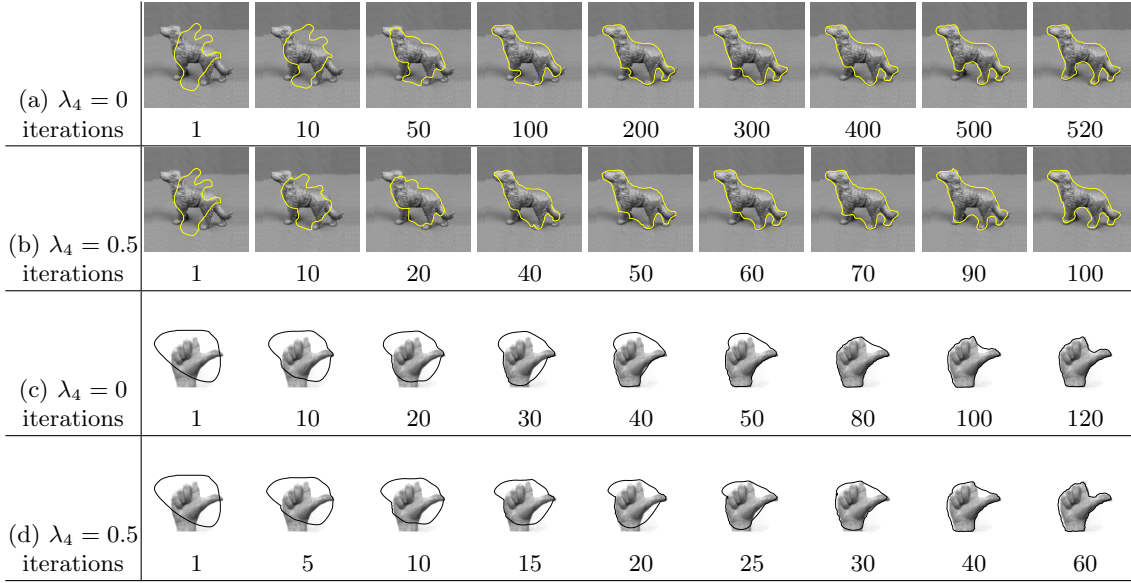


Fig. 13 Odd rows: Evolution in absence of random vector field. Even rows: Evolution under a random vector field. In this experiment $\lambda_3 = 0$ for all cases.

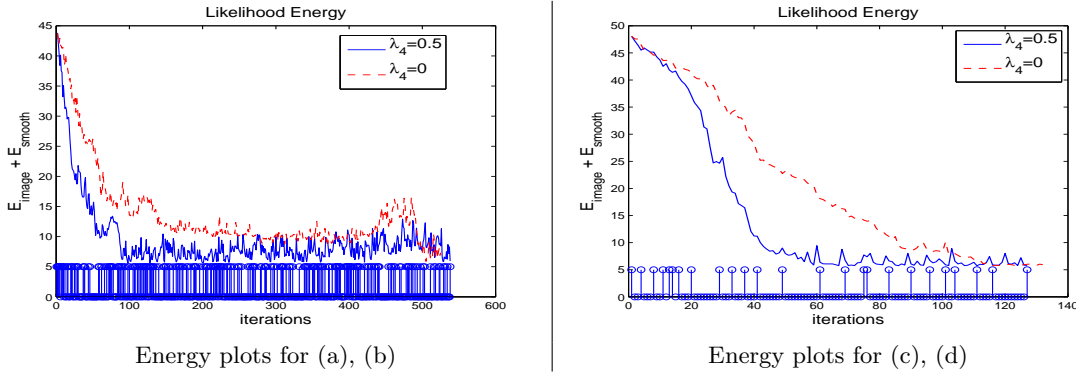


Fig. 14 Plots of $\lambda_1 E_{image} + \lambda_2 E_{smooth}$ for evolutions given in Figure 13. The overlaid stem plot indicates a '1' for the iteration that accepted the random vector field.

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 shapes of different poses of the same object, or both. For each class $q \in \mathcal{Q}$, assume that we have a prior shape model P_q on \mathcal{S} and a prior probability $P(q)$ of that class. In case of Gaussian prior models, this implies that we have a different mean μ_q , principal space $N_q \subset T_{\mu_q}(\mathcal{S})$, 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

$$P(q|I) = \frac{P(q) \int_{\mathcal{B}} P(I|\beta) P(\beta|q) d\beta}{\sum_q (P(q) \int_{\mathcal{B}} P(I|\beta) P(\beta|q) d\beta)} = \frac{P(q) \int_{\mathcal{B}} e^{-E_{total}^q(\beta)} d\beta}{\sum_q (P(q) \int_{\mathcal{B}} e^{-E_{total}^q(\beta)} d\beta)},$$

where $E_{total}^q = \lambda_1 E_{image} + \lambda_2 E_{smooth} + \lambda_3 E_{prior}^q$. Object class is selected by integrating out all possible contours with respect to the class-dependent shape models. Since it is difficult to perform

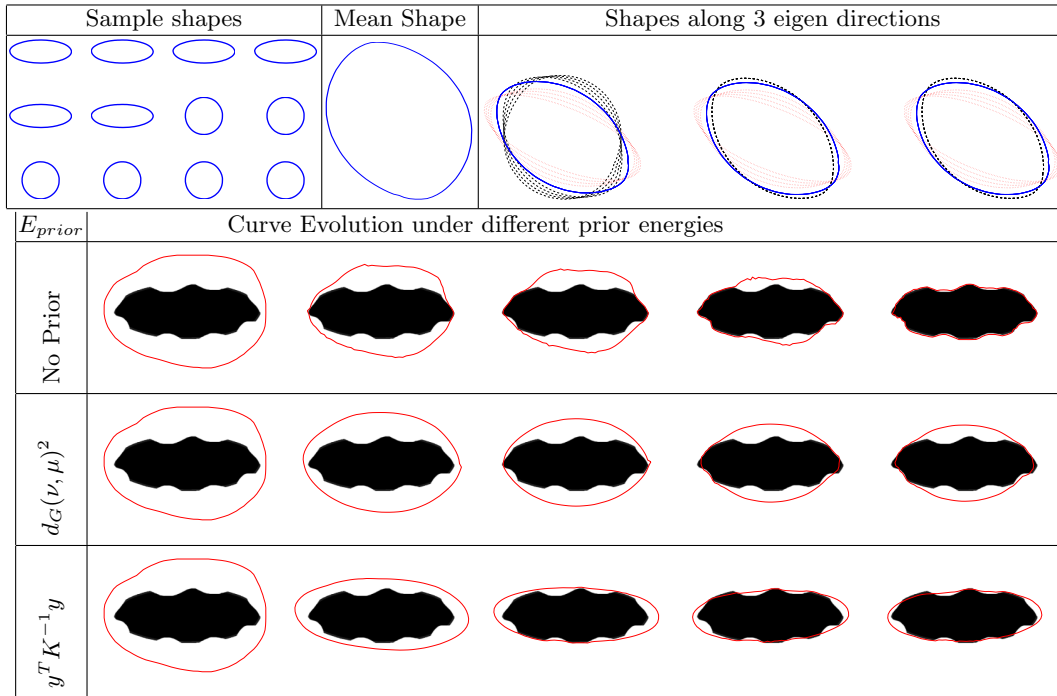


Fig. 15 Top panel: Sample shapes consisting of ellipses and circles, their mean shape, and their eigen-variation. Bottom panel: Curve evolution under two different E_{prior} : squared-distance and normal energy.

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 [15]. A coarse approximation is to evaluate the integrands at $\hat{\beta}$, the MAP estimates. That is, choose:

$$\hat{q} = \operatorname{argmax}_{q \in \mathcal{Q}} \left(P(q) e^{-E_{total}(\hat{\beta}^q)} \right), \quad (15)$$

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 [31], and is often termed as *Bayesian model selection*.

This is demonstrated in the following experiment that uses real data to generate prior shape models for tanks. Figure 16(a) shows a schematic of a battle tank rotated at a 3° resolution around its central axis at a fixed elevation. The sample infrared images corresponding to different views of the tank are shown in Figure 16(c). Figure 16(b) shows the automatically extracted shapes and their sample means grouped together in 9 quadrants for one hemisphere. For simplicity, we reflect the mean shapes of quadrants in one hemisphere to obtain the model mean shapes in the other hemisphere. In this experiment, the shape model comprises of the mean shape μ_q corresponding to each quadrant $\{q_i \in \mathcal{Q}\}, i = 1, \dots, 9, 1_r, \dots, 9_r$. 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 Eqn. 15 to maximize the probability of selecting the best model \hat{q} . Figure 17 shows the extracted curves $\hat{\beta}_q$ overlaid on a partially obscured test image I . Figure 18 plots the discrete posterior energy in Eqn. 7 against each of the tested prior models, q_1, \dots, q_{9r} . 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.

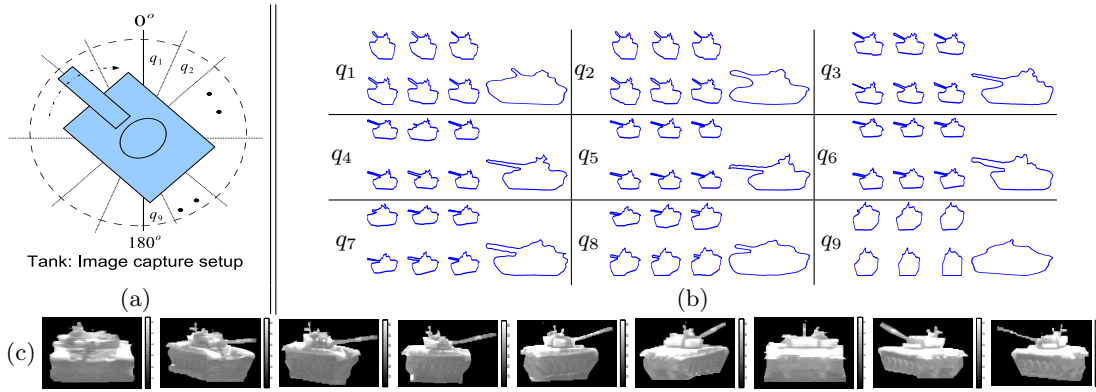


Fig. 16 (a) Experimental setup for image captures of different views. (b) Shapes for each quadrant q_i , and their mean shapes. (c) Examples of infrared images of a tank at different poses.

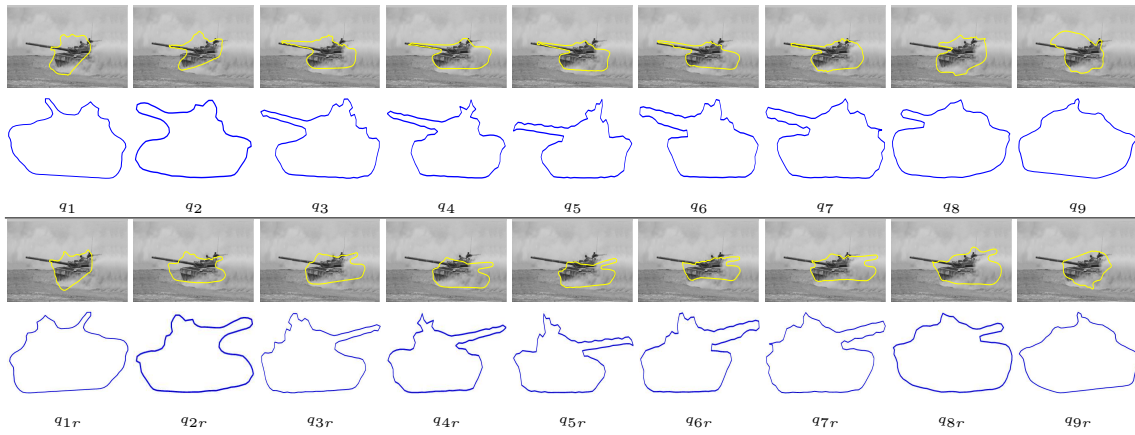


Fig. 17 Extracted curves $\hat{\beta}_q$ from I for respective prior models q_1, q_2, \dots, q_{9r} displayed below each image.




Fig. 18 E_{total}^q plotted against prior models $q_i \in \mathcal{Q}$.

6.2 Performance Evaluation Metrics

How should one evaluate the performance of this 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.

Let $\hat{\beta}$ be the result of our MAP estimation in an image I , and one uses 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. Denote the classification performance on a test dataset to be F . 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 prior shape information drastically improves the object extraction under various image obscurations. This is evident from various results on synthetic data as well as successful segmentation of cardiac boundaries in real data including video and ultrasound 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 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, and the Air Force Office of Scientific Research award FA9550-06-1-0324.

References

1. Y. Amit, U. Grenander, and M. Piccioni. Structural image restoration through deformable templates. *J. Amer. Statist. Assoc.*, 86(414):376–387, 1991.
2. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Intl. Journal of Computer Vision*, 22(1):61–79, 1997.
3. G. Charpiat, O. Faugeras, R. Keriven, and P. Maurel. Approximations of shape metrics and application to shape warping and empirical shape statistics. In H. Krim and A. Yezzi, editors, *Statistics and Analysis of Shapes*, pages 363–395. Springer, 2005.
4. Y. Chen, H. Tagare, S. Thiruvankadam, D. Huang, D. Wilson, K. Gopinath, R. Briggs., and M. Geiser. Using prior shapes in geometric active contours in a variational framework. *Intl. Journal of Computer Vision*, 50(3):315–328, 2002.
5. L. D. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 15(11):1131–1147, 1993.
6. T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models, – their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
7. D. Cremers, T. Kohlberger, and C. Schnörr. Shape Statistics in Kernel Space for Variational Image Segmentation. *Pattern Recognition*, 36(9):1929–1943, 2003.
8. D. Cremers, S. J. Osher, and S. Soatto. Kernel Density Estimation and Intrinsic Alignment for Shape Priors in Level Set Segmentation. *Intl. Journal of Computer Vision*, 69(3):335–351, 2006.
9. D. Cremers, F. Tischhäuser, J. Weickert, and C. Schnörr. Diffusion Snakes: Introducing Statistical Shape Knowledge into the Mumford–Shah functional. *Intl. Journal of Computer Vision*, 50(3):295–313, 2002.

10. I. L. Dryden and K.V. Mardia. *Statistical Shape Analysis*. John Wiley & Son, 1998.
11. M. Gage and R. S. T. Hamilton. The heat equation shrinking convex plane curves. *J. Differential Geometry*, 23:69–96, 1986.
12. M. Grayson. The heat equation shrinks embedded plane curves to round points. *J. Differential Geometry*, 26:285–314, 1987.
13. U. Grenander. *General Pattern Theory*. Oxford University Press, 1993.
14. U. Grenander, Y. Chow, and D. M. Keenan. *Hands: A Pattern Theoretic Study of Biological Shapes*. Springer-Verlag, 1991.
15. U. Grenander, A. Srivastava, and M. Miller. Asymptotic performance analysis of bayesian target recognition. *IEEE Trans. Inform. Theory*, 46:1658–1665, 2000.
16. I. Jermyn. Invariant bayesian estimation on manifolds. *Ann. Statistics*, 33(2):583–605, 2005.
17. M. Kass, A. Witkin, and D. Terzopolous. Snakes: Active contour models. *Intl. Journal of Computer Vision*, 1(4):321–331, 1988.
18. S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *Proc. Intl. Conf. Computer Vision*, pages 810–815, 1995.
19. E. Klassen, A. Srivastava, W. Mio, and S. H. Joshi. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 26(3):372–383, 2004.
20. M. Leventon, W. Eric L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Proc. IEEE Conf. Comp. Vision and Pattern Recognition*, pages 316–323, 2000.
21. C. Li, C. J. Liu, and M. D. Fox. Segmentation of external force field for automatic initialization and splitting of snakes. *Pattern Recognition*, 38(11):1947–1960, 2005.
22. P. W. Michor and D. Mumford. Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.*, 8:1–48, 2006.
23. M. I. Miller and L. Younes. Group actions, homeomorphisms, and matching: A general framework. *Intl. Journal of Computer Vision*, 41(1/2):61–84, 2001.
24. W. Mio and A. Srivastava. Elastic-string models for representation and analysis of planar shapes. In *Proc. IEEE Conf. Comp. Vision and Pattern Recognition*, pages 10–15, 2004.
25. V. Patrangenaru and R. Bhattacharya. Large sample theory of intrinsic and extrinsic sample means on manifolds. *Ann. Statistics*, 31(1):1–29, 2003.
26. J. Shah. H^0 -type Riemannian metrics on spaces of planar curves. *Preprint*, 2005. <http://arxiv.org/abs/math.DG/0510192>.
27. C. G. Small. *The Statistical Theory of Shape*. Springer, 1996.
28. A. Srivastava, A. Jain, S. Joshi, and D. Kaziska. Statistical shape models using elastic-string representations. In *Proceedings of Asian Conference on Computer Vision (ACCV), LNCS 3851, P. J. Narayanan (Eds.)*, pages 612–621. Springer-Verlag, 2005.
29. A. Srivastava, S. H. Joshi, W. Mio, and X. Liu. Statistical shape analysis: Clustering, learning and testing. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 27(4):590–602, 2005.
30. L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(11):1061–1075, 1992.
31. H. L. Van Trees. *Detection, Estimation, and Modulation Theory, vol. I*. John Wiley, N.Y., 1971.
32. M. Vaillant, M.I. Miller, L. Younes, and A. Troune. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23:161–169, 2004.
33. B. Vemuri and Y. Chen. Joint image registration and segmentation. In S. Osher and N. Paragios, editors, *Geometric Level Set Methods in Imaging, Vision, and Graphics*, pages 251–269. Springer-Verlag, 2003.
34. D.C. Wilson, M. Geiser, and Larocca J. Automated analysis of echocardiographic apical 4-chamber images. In *Proc. Intl. Society for Optical Eng. in Math. Modeling, Estimation, and Imaging*, pages 128–139, 2000.
35. C. Xu and J. Prince. Generalized gradient vector flow external force for active contours. *Signal Processing*, 71(2):131–139, 1998.
36. C. Xu and J. Prince. Snakes, shapes, and gradient vector flow. *IEEE Trans. Image Processing*, 7(3):359–369, 1998.
37. A. J. Yezzi and A. Mennucci. Conformal metrics and true "gradient flows" for curves. In *Proc. IEEE Intl. Conf. Computer Vision*, pages 913–919, 2005.