

Active Curve Recovery of Region Boundary Patterns

Mohamed Ben Salah, Ismail Ben Ayed, *Member, IEEE*, and
Amar Mitiche, *Member, IEEE Computer Society*

Abstract—This study investigates the recovery of region boundary patterns in an image by a variational level set method which drives an active curve to coincide with boundaries on which a feature distribution matches a reference distribution. We formulate the scheme for both the Kullback-Leibler and the Bhattacharyya similarities, and apply it in two conditions: the simultaneous recovery of all region boundaries consistent with a given outline pattern, and segmentation in the presence of faded boundary segments. The first task uses an image-based geometric feature, and the second a photometric feature. In each case, the corresponding curve evolution equation can be viewed as a geodesic active contour (GAC) flow having a variable stopping function which depends on the feature distribution on the active curve. This affords a potent global representation of the target boundaries, which can effectively drive active curve segmentation in a variety of otherwise adverse conditions. Detailed experimentation shows that the scheme can significantly improve on current region and edge-based formulations.

Index Terms—Image segmentation, boundary patterns, boundary feature distributions, active curves, level sets, similarity measures.

1 INTRODUCTION

IMAGE segmentation is a long standing, extensively researched topic in image processing for its theoretical and methodological challenges, and numerous useful applications. Current major application areas include medical image analysis, remote sensing, robotics, and surveillance [1], [2], [3], [4].

Active contour variational formulations, which define image domain partitions by closed regular plane curves, have been widely used. The corresponding Euler-Lagrange equations are evolution equations which drive the curves to coincide with relevant region boundaries. Implemented via level sets [5], the evolution equations have led to effective, numerically efficient, and stable algorithms in a variety of settings [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]. The objective functional data terms, which measure the conformity of the image to model descriptions, are basically of one of two types, *edge-based*, when they evaluate an image function along the active curve, or *region-based* when they refer to the image within the region enclosed by the curve. Therefore, the corresponding curve evolution velocities are due to the image exclusively along or within the curve.

The Snake model [16] and the geodesic active contour (GAC), which adopted a more effective curve representation

[17], were precursors of a vast literature on edge-based active curve image segmentation [14], [16], [17], [18], [19], [20], [21]. Typically, a decreasing function of the image gradient is integrated along the geodesic contour so that it settles on high-contrast boundaries which are thought to characterize the desired regions. In general, geodesics are seriously challenged when the desired boundaries have segments of low gradient, as is common in many applications. For instance, in magnetic resonance imaging (MRI) and computed tomography (CT) medical images, the organs to segment can have weak, almost nonexistent contrast with neighboring structures. In such cases, the geodesic leaks away from the desired boundary and can vanish.

By referring to the image over regions, the region-based schemes are significantly less sensitive to weak boundary gradient than the geodesic schemes [2], [6], [22], [23]. In general, this is due to a global model description of the image within the extent of each desired region, which penalizes movements of the active curves in or out of the regions they are intended to delineate. Both parametric [6], [24], [25], [26] and nonparametric [8], [10], [11] image descriptions have been used for the purpose. However, and in spite of this accrued robustness, the region-based active curve evolution can be seriously challenged, by definition, when the regions to segment have similarly distributed segments [9]. When these segments occur between regions, the placement of the separating boundary becomes largely ambiguous.

There are methods which combine the advantages of both edge-based and region-based models by using a linear combination of two or more such terms [14], [27], [28]. However, current methods are not applicable when the desired regions are characterized by the distribution of a feature on their boundary, i.e., when region boundaries are considered patterns described by a feature distribution rather than simply the location of the feature as with typical geodesic descriptions. Two examples where this

• M. Ben Salah is with the Department of Computing Science, University of Alberta, 221 Athabasca Hall, Edmonton, Alberta T6G 2E8, Canada. E-mail: bensala@ualberta.ca.

• I. Ben Ayed is with GE Healthcare, 268 Grosvenor, E5-138, London, ON N6A 4V2, Canada. E-mail: ismail.benayed@ge.com.

• A. Mitiche is with the Institut National de la Recherche Scientifique, INRS-EMT, Bureau 6900, 800 De la Gauchetiere West, Montreal (QC), H5A 1K6, Canada. E-mail: mitiche@emt.inrs.ca.

Manuscript received 12 May 2010; revised 30 June 2011; accepted 22 July 2011; published online 8 Oct. 2011.

Recommended for acceptance by D. Cremers.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2010-05-0372.

Digital Object Identifier no. 10.1109/TPAMI.2011.201.




Fig. 1. The segmentation targets (a) elliptical objects and (b) rectangular objects. Only the targeted objects should be segmented. In (c), the segmentation targets the left ventricle in an MRI sequence in spite of weak inner and outer boundary segments with neighboring image objects.

description of boundaries by a feature distribution is befitting are in Fig. 1. One example, Figs. 1a and 1b, is of an image where regions occur with boundaries that are either rectangular or ellipsoidal outline. For each outline pattern there are two regions of varying appearance and the goal of segmentation is to extract, in a single instantiation, the regions of each of the two patterns. Matching the distribution of boundary curvature, measured from the image gradient, against a model distribution, has extracted both rectangular regions in one case (Fig. 1a) and both ellipsoidal regions in the other (Fig. 1b), without prior knowledge of the number of regions. Note that a shape prior constraint will not be able to segment all of the regions of the same figure, unless one such prior is used for each region, and each with an accompanying close by initialization, which supposes an information about the image not available for practical purposes, such as the number of objects as well as the section of the image domain where each occurs [29], [30]. Because a shape prior is an image-independent term added to the segmentation functional so as to bias a detected region to have a given geometric outline modulo a transformation (such as rigid or affine) [7], [31], [32], [33], a shape prior constraint will also require additional optimization over pose transformations or a constraint on the curve deformation with respect to a reference shape [34], [35].

The other example, Fig. 1c, shows an image from a cardiac MRI sequence of the left ventricle (LV). The red and green boundaries are accurate delineations of the inner and outer ventricle boundaries by curves along which the image distribution matches a model distribution. Yet, both boundaries are severely faded in places, and parts of the inner cavity have an image distribution closely resembling that of the ventricle wall. Also, the ventricle wall and its outer surrounding context have parts which look similar. These image particularities are examples of the general segmentation ailments discussed earlier. Note that the outer and inner ventricle boundaries have the same shape, in which case a shape prior will not distinguish them unless the active curve is initialized rather close to the desired boundary, an action which may require manual intervention.

The purpose of this study is to investigate a level set variational segmentation method which drives an active curve to coincide with boundaries on which a feature distribution matches a reference distribution. We have

addressed the problem earlier in [36]. This *TPAMI* version expands on [36] with a broader, more informative discussion of the subject and a more rigorous, wider investigation which includes the use of geometric features along contours. Several new experiments with distributions of curvature computed from the image have been added to enhance the photometric feature experimentation.

Feature distributions are potent global representations of region boundaries [37] which can effectively drive active curve segmentation in a variety of otherwise adverse conditions. We formulate the scheme for both the Kullback-Leibler and the Bhattacharyya similarities, and apply it in two particularly relevant conditions, the simultaneous recovery of all region boundaries consistent with a given outline pattern and segmentation in the presence of faded region boundary segments. The first task uses a geometric feature, rather than a photometric feature as does the second task. Fig. 1 illustrates a case of each of these two tasks. Detailed experimentation (Section 4) shows that the scheme is valid and can improve on region and edge-based methods. Compared to the region-based formulations in [8], [9], and [11], the objectives of the proposed functionals are fundamentally different. For instance, the formulations in [8], [9], and [11] would not distinguish, and it is not their purpose, between the elliptical and rectangular regions in Figs. 1b and 1c because these regions have exactly the same image distributions. The marginal similarity with these studies is in using global measures, but the curve evolution equations we obtained are quite different. Such a difference will be evidenced in the experiments.

Interestingly, each of the evolution equations we obtained can be viewed as a GAC having a variable stopping function. However, the stopping functions have two fundamental differences with the usual GAC stopping function. First, they are functions of both the image and the curve, when the GAC stopping function depends only on the image. Second, they reference global information, namely, the feature distribution, rather than just pixel-wise, as with GAC; such richer information should afford better boundary detection behavior. We will give an interpretation of this behavior.

The remainder of this paper is organized as follows: Section 2 describes the formulation in detail, including the objective function and the similarity measure used. The corresponding Euler-Lagrange curve evolution equations for both the Kullback-Leibler divergence and the Bhattacharyya

measure and the level set equations are derived in Section 3. Section 4 describes the experimental results using geometric and photometric features on various synthetic and real images. Section 5 contains a conclusion.

2 FORMULATION

The formulation in this study seeks region boundaries along which the distribution of a representation feature is closest to that of a model.

2.1 Objective Function

Let $I : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$ be an image function, $\gamma : [0, 1] \rightarrow \Omega$ a simple closed plane parametric curve, and $F : \Omega \subset \mathbb{R}^2 \rightarrow \mathcal{F} \subset \mathbb{R}$ a feature function from the image domain Ω to a feature space \mathcal{F} . Let P_γ be a kernel density estimate of the distribution of F along γ :

$$\forall f \in \mathcal{F} \quad P_\gamma(f) = \frac{\oint_\gamma K(f - F_\gamma) ds}{\mathcal{L}_\gamma}, \quad (1)$$

where F_γ is the restriction of F to γ , \mathcal{L}_γ is the length of γ ,

$$\mathcal{L}_\gamma = \oint_\gamma ds, \quad (2)$$

and K is the estimation kernel. In this work, we consider the Gaussian kernel of width h :

$$K(z) = \frac{1}{\sqrt{2\pi}h^2} \exp\left(-\frac{z^2}{2h^2}\right). \quad (3)$$

Given a model feature distribution \mathcal{M} , let $\mathcal{D}(P_\gamma, \mathcal{M})$ be a similarity measure between P_γ and \mathcal{M} . The purpose is to determine $\tilde{\gamma}$ such that

$$\tilde{\gamma} = \arg \min_\gamma \mathcal{D}(P_\gamma, \mathcal{M}). \quad (4)$$

To apply this formulation, we need to specify the feature function, the model, the similarity, and a scheme to conduct the objective functional minimization in (4).

2.2 Features

There are two fundamental types of boundary representation features. One type is *photometric*, where the feature is a function of the image. Examples of such features are the image, $F = I$, its gradient norm $F = \|\nabla I\|$, and, more generally, scalar image filter outputs. The other type of feature function is *geometric*, which pertains to the boundary form, irrespective of the image function. The curvature is in this category. This is a singular feature because it can be estimated from the image under the assumption that the region boundary normals coincide with the isophote normals:

$$F = \kappa_I = \text{div} \left(\frac{\nabla I}{\|\nabla I\|} \right). \quad (5)$$

This is quite convenient and important from an implementation point of view because curvature estimation over the image domain is done once and at the onset. However, it remains intrinsically tied to the boundary geometry and not to the image function. Studies have shown that curvature histograms, which can be viewed as

empirical marginal distributions of the shape considered a random variable, are useful statistics to describe closed regular plane shapes [37]. Ideally, a geometric description is invariant to the shape position, orientation, and size. It must also be robust to the distortions which normally affect the shape. Curvature, which is the rate of change of the tangent angle along the contour [38], is invariant to translation and rotation but varies with scale. However, this variability is taken into account by an affine transformation of the curvature values so that they always correspond to the same set of bins. For practical means, this normalization makes the histograms unaffected by scale. A curvature histogram alone is not, of course, sufficient to describe shape in general. Although it has served our purpose in this study (Section 4), other features have been necessary for a more general encoding of planar shapes [37]. A geometric feature which, unlike curvature, cannot be estimated from the image function would be useful only in conjunction with a photometric feature because the image is the essential support for boundary pattern detection.

Each of the two fundamental types of features, photometric and geometric, corresponds to a fundamental application of the formulation. Geometric features are necessary when the target object boundary has no specific photometric description, either because the description varies with the picture in which the object appears (e.g., as in Fig. 5, where objects can appear with different colors/textures and/or over different backgrounds) or because there are no photometric features which would distinguish the target from other objects in the image (as in Fig. 2). Photometric features are necessary when photometry, not geometry, is distinctive of the target region boundary. The automatic detection of the left ventricle wall in Fig. 10 is an example of this case.

2.3 Model

A model in our context is an *exemplar* of the shape of interest, to be used to estimate, via a histogram, a model distribution of the representation feature. An exemplar should be able to represent the boundary shape of the objects of interest up to allowable transformations such as position, orientation, scale, and nonlinear class noise. There are several sensible ways to draw a model of a targeted shape. One way is to manually trace (using a graphics manipulation package) the contour of an object in the class of those targeted, and use a histogram of the representation feature along the trace as the model distribution. Alternatively, this typical object contour may be extracted by an active geodesic curve, or a region-based segmentation algorithm such as Chan-Vese's [39], using an initialization that is close to and contains the object. When the representation feature pertains to the object contour geometry (e.g., curvature), one can manually draw a cartoon of the object contour and use it to estimate the model distribution of the geometric feature. In general, the application images determine the choice of a model selection scheme over others. We will show examples of each model learning scheme (Section 4).

2.4 Similarity

The Kullback-Leibler divergence and the negative of the Bhattacharyya coefficient are two common similarity




Fig. 2. Detection of regions whose boundaries are consistent with learned outline patterns. Each row depicts a segmentation of the image corresponding to a different model of curvature learned a priori. For example, for the first row the model of curvature is learned *independently* of the shape of a single rectangle and, for the second row, it is learned from a single ellipse with approximately the same aspect ratio as in the figure. In this example, several regions in the image correspond to the same shape instance (a rectangle or an ellipse). By learning the model distribution of (image-based) curvature from a single rectangle or ellipse, the proposed method can successfully delineate all the regions consistent with the learned pattern without additional optimization over pose transformations and without prior knowledge of the number of regions. Columns: (a) initial curve positions, (b) training images and contours, (c) the final segmentations, and (d) segmentation with the GAC model [17] (upper) and the RL piece-wise constant model [6] (lower). The Kullback-Leibler divergence has been used for this example.

functions between distributions. Several studies have used them for foreground-background image segmentation [39]. An experimental study [40] has given some validity to the Bhattacharyya distance by showing that for a variety of distributions the number of misclassified pixels by maximum likelihood and minimum description length (MDL) increases with increasing Bhattacharyya distance between the foreground and background distributions. Efficient applications of the Bhattacharyya distance have been reported in [8] and [11] in active contour segmentation. It has also been implemented to match the distribution along contours of a local image average to the distribution along a model object boundary [36]. As well, the Kullback-Leibler divergence has been part of effective image segmentation formulations [8], [41], [42], [43]. Studies which mention or use both measures have presented them as alternatives.

We implemented the minimization in (4) for both the Kullback-Leibler divergence and the Bhattacharyya distance as the similarity function \mathcal{D} . The Kullback-Leibler divergence between P_γ and \mathcal{M} is

$$\mathcal{D}(P_\gamma, \mathcal{M}) = KL(P_\gamma, \mathcal{M}) = \int_{\mathcal{F}} \mathcal{M}(f) \log \frac{\mathcal{M}(f)}{P_\gamma(f)} df, \quad (6)$$

and the negative of the Bhattacharyya coefficient is

$$\mathcal{D}(P_\gamma, \mathcal{M}) = -B(P_\gamma, \mathcal{M}) = - \int_{\mathcal{F}} \sqrt{P_\gamma(f) \mathcal{M}(f)} df. \quad (7)$$

Higher values of the Kullback-Leibler divergence indicate smaller overlaps between the distributions. The range of the Bhattacharyya coefficient is $[0, 1]$, 0 corresponding to no overlap between the distributions and 1 to a perfect match. The symmetry of the similarity function with respect to its two distribution arguments (the Kullback-Leibler divergence is not symmetric) is not an issue here because we want to assess how close a variable distribution is to a fixed (model) distribution. However, we will see that the curve evolution equations, derived next, show that the Bhattacharyya flow is more general. Moreover, the $[0, 1]$

range of the Bhattacharyya coefficient affords a conveniently practical appraisal of the similarity.

The simpler analytical expression of the Bhattacharyya distance affords a computational advantage, albeit relatively modest considering the high computational capacities of current common computers. Also, a brief look ahead at the curve evolution equations (13) and (17) shows that the expressions are the same except for the first term which is constant and equal to 1 for the Kullback-Leibler (13) and is the Bhattacharyya distance (the maximum of which is 1) in (17). Therefore, the Bhattacharyya distance offers a more general expression, which potentially affords a better representation of similarity in (4). However, in all the examples where we experimented with both similarities, there were no noticeable differences in the results (Section 4).

Next, we derive the Euler-Lagrange descent equations corresponding to (4) for both the Kullback-Leibler and the Bhattacharyya similarities.

3 MINIMIZATION

The data term in (4) is a measure of similarity between distributions over the boundary representation curve γ . The use of such similarity measures in image segmentation often leads to challenging optimization problems because it involves *nonlinear* functions of integrals, or sums in the discrete case, over the segmentation boundaries or regions. A discretization of this term will not conform to graph cut optimization usually used in image segmentation [44], [45], [46], [47], [48], which requires that the objective function be a sum over the segmentation boundary or region of pixel-dependent penalties. For instance, the data term in [44] is a linear combination of sums over the segmentation regions of minus the log likelihood of pixel data given the histograms of the regions, and the data term in [45] is a sum over the segmentation boundary of the dot product between the normal to the boundary and a fixed vector field. Measures of similarity between distributions have been generally avoided in the context of graph cuts because

they cannot be expressed in such forms. The recent studies in [49] and [50] are notable exceptions which optimized a distribution-similarity measure over the segmentation region via graph cuts. To do so, Ben Ayed et al. [49] and Mukherjee et al. [50] used relaxations via bounds or approximations of the cost function so as to befit graph cut optimization. However, in our case, the problem involves a similarity measure over the boundaries, not regions. Therefore, the approximations in [49] and [50] are not applicable. Instead, we will address the problem in (4) by continuous optimization via the associated Euler-Lagrange γ -evolution descent equations. These equations are derived next.

Another important argument in favor of continuous optimization is the fact that graph-cut approaches are prone to the well-known grid bias (or metrication error) [51]. Reducing metric artifacts can be done by increasing the number of neighboring graph nodes, but this may result in a heavy computation and memory load [45].

Let γ be embedded in a one-parameter family of curves indexed by (algorithmic) time t : $\gamma(s, t) : [0, 1] \times \mathbb{R}^+ \rightarrow \Omega$, and deriving the Euler-Lagrange descent equation:

$$\frac{\partial \gamma}{\partial t} = -\frac{\partial \mathcal{D}}{\partial \gamma}. \quad (8)$$

3.1 Kullback-Leibler Divergence

For $\mathcal{D}(P_\gamma, \mathcal{M}) = KL(P_\gamma, \mathcal{M})$, we have

$$\begin{aligned} \frac{\partial \mathcal{D}}{\partial \gamma} &= \frac{\partial KL}{\partial \gamma} = \int_{\mathcal{F}} \mathcal{M}(f) \frac{\partial}{\partial \gamma} \log \left(\frac{\mathcal{M}(f)}{P_\gamma(f)} \right) df \\ &= - \int_{\mathcal{F}} \mathcal{M}(f) \frac{\partial}{\partial \gamma} \log \left(\frac{\oint_{\gamma} K(f - F_{\gamma(s)}) ds}{\mathcal{L}_\gamma} \right) df \\ &= \frac{1}{\mathcal{L}_\gamma} \frac{\partial \mathcal{L}_\gamma}{\partial \gamma} - \int_{\mathcal{F}} \mathcal{M}(f) \frac{\partial}{\partial \gamma} \log \left(\oint_{\gamma} K(f - F_{\gamma(s)}) ds \right) df, \end{aligned} \quad (9)$$

where, we recall, F_γ is the restriction of F to γ . Assuming feature F is independent of γ , both curve length \mathcal{L}_γ and the second integral in (9) can be written as $\oint_{\gamma} h ds$, where h is independent of γ , and their functional derivative with respect to γ is of the form [17]

$$\frac{\partial \oint_{\gamma} h ds}{\partial \gamma} = (-h\kappa + \nabla h \cdot \vec{n}) \vec{n}, \quad (10)$$

where \vec{n} is the inward unit normal to γ and κ its mean curvature function. Therefore,

$$\begin{aligned} \frac{\partial}{\partial \gamma} \log \left(\oint_{\gamma} K(f - F_{\gamma(s)}) ds \right) &= \frac{\frac{\partial}{\partial \gamma} \oint_{\gamma} K(f - F_{\gamma(s)}) ds}{\oint_{\gamma} K(f - F_{\gamma(s)}) ds} \\ &= \frac{-K(f - F_\gamma)\kappa + \nabla K(f - F_\gamma) \cdot \vec{n}}{\oint_{\gamma} K(f - F_{\gamma(s)}) ds} \vec{n}, \end{aligned} \quad (11)$$

and, using (2), $\frac{\partial \mathcal{L}_\gamma}{\partial \gamma} = -\kappa \vec{n}$.

This gives

$$\begin{aligned} \frac{\partial \gamma}{\partial t} &= \frac{\kappa}{\mathcal{L}_\gamma} \vec{n} - \left[\frac{\kappa}{\mathcal{L}_\gamma} \int_{\mathcal{F}} \frac{\mathcal{M}(f)}{P_\gamma(f)} K(f - F_\gamma) df \right] \vec{n} \\ &\quad + \frac{\vec{n}}{\mathcal{L}_\gamma} \nabla \left(\int_{\mathcal{F}} \frac{\mathcal{M}(f)}{P_\gamma(f)} K(f - F_\gamma) df \right) \cdot \vec{n} \\ &= \frac{1}{\mathcal{L}_\gamma} \left[\underbrace{\left(1 - \int_{\mathcal{F}} \frac{\mathcal{M}(f)}{P_\gamma(f)} K(f - F_\gamma) df \right) \kappa}_{\text{Stopping force}} \right. \\ &\quad \left. + \underbrace{\nabla \left(\int_{\mathcal{F}} \frac{\mathcal{M}(f)}{P_\gamma(f)} K(f - F_\gamma) df \right) \cdot \vec{n}}_{\text{Refinement force}} \right] \vec{n}, \end{aligned} \quad (12)$$

which can be written as

$$\frac{\partial \gamma}{\partial t} = \left(\underbrace{\mathcal{G}_{KL}(P_\gamma, \mathcal{M}, F_\gamma) \kappa}_{\text{Stopping}} - \underbrace{\nabla \mathcal{G}_{KL}(P_\gamma, \mathcal{M}, F_\gamma) \cdot \vec{n}}_{\text{refinement}} \right) \vec{n}. \quad (13)$$

The evolution (13) is of an ordinary geodesic active contour form [17] except that the function of space and time \mathcal{G}_{KL} multiplying the curvature can, at some times during curve evolution, evaluate to negative at some points, i.e., it does not necessarily evaluate to positive everywhere at all times. However, the evolution is not to be assimilated to an inverse heat flow [52], [53] because, first, \mathcal{G}_{KL} does not necessarily evaluate to negative everywhere and at all times, as with inverse heat flow, and, second, the evolution speed is also modulated by the gradient of this function projected on the curve normal, $\nabla \mathcal{G}_{KL} \cdot \vec{n}$. The equation behavior can be examined according to two cases.

Case 1. The curve is close to the desired boundary. When in the vicinity of the target boundary, close to adhering, the curve has a feature density close to the reference density, i.e., $P(F_{\gamma(p)}) \approx M(F_{\gamma(p)})$, which implies that $\mathcal{G}_{KL} \approx 0$. Consequently, the curve behavior is predominantly modulated by the gradient term, which drives it to adhere to the desired boundary because it constrains it to move so as to coincide with local highs in the model and curve distributions similarity, just as the common GAC gradient term guides the curve toward local highs in image contrast [17].

Case 2. The curve is distant from the desired boundary. Away from the target boundary, the curve and its model are dissimilar and their feature distributions have little overlap. Therefore, for most points p on the curve, and recalling that (13) references points on the curve, not on the model, we have $P(F_{\gamma(p)}) > M(F_{\gamma(p)})$ and, consequently, $\mathcal{G}_{KL} > 0$.

The preceding argument points to a stable behavior of the evolution equation in general. In the event \mathcal{G}_{KL} at some point evaluates to negative at some time during curve evolution, the gradient term $\nabla \mathcal{G}_{KL} \cdot \vec{n}$ acts as a stabilizer of the curvature term because, as previously mentioned, it constrains the curve to move along its normal to fit highs in the similarity between its feature distribution and the model distribution. There has been no evidence to the contrary in all our experiments. All tests were in agreement with a stable descent algorithm to maximize the similarity between

the evolving curve feature distribution and the reference distribution (Section 4).

A final pointer to an expected good evolution behavior is the presence of global information in the descent equation, namely, the feature distribution (histogram), in addition to pixel-wise information. In general, global information affords added strength to local descriptions. In practice, the descent equation is discretized following [5] and involves choosing an appropriate time step according to a data dependent recipe.

3.2 Bhattacharyya Measure

We have

$$\frac{\partial \mathcal{D}}{\partial \gamma} = -\frac{\partial B}{\partial \gamma} = -\frac{1}{2} \int_{\mathcal{F}} \sqrt{\frac{\mathcal{M}(f)}{P_{\gamma}(f)}} \frac{\partial P_{\gamma}(f)}{\partial \gamma} df. \quad (14)$$

To compute this functional derivative, we need to compute $\frac{\partial P_{\gamma}(f)}{\partial \gamma}$:

$$\frac{\partial P_{\gamma}(f)}{\partial \gamma} = \frac{\mathcal{L}_{\gamma} \frac{\partial}{\partial \gamma} \oint_{\gamma} K(f - F_{\gamma(s)}) ds - \frac{\partial \mathcal{L}_{\gamma}}{\partial \gamma} \oint_{\gamma} K(f - F_{\gamma(s)}) ds}{\mathcal{L}_{\gamma}^2}. \quad (15)$$

Similarly to the previous section (see (11)), the application of (10) to $\oint_{\gamma} K(f - F_{\gamma(s)}) ds$ and to \mathcal{L}_{γ} in the numerator in (15) gives, after some algebraic manipulations,

$$\forall f \in \mathcal{F}, \quad \frac{\partial P_{\gamma}(f)}{\partial \gamma} = \frac{1}{\mathcal{L}_{\gamma}} (-K(f - F_{\gamma}) \kappa + \nabla K(f - F_{\gamma}) \cdot \vec{n} + P_{\gamma}(f) \kappa) \vec{n}. \quad (16)$$

Embedding this result in (14) leads to, after algebraic manipulations,

$$\begin{aligned} \frac{\partial \gamma}{\partial t} = \frac{1}{2\mathcal{L}_{\gamma}} & \left[\underbrace{\left(B(P_{\gamma}, \mathcal{M}) - \int_{\mathcal{F}} \sqrt{\frac{\mathcal{M}(f)}{P_{\gamma}(f)}} K(f - F_{\gamma}) df \right) \kappa}_{\text{Stopping force}} \right. \\ & \left. + \underbrace{\nabla \left(\int_{\mathcal{F}} \sqrt{\frac{\mathcal{M}(f)}{P_{\gamma}(f)}} K(f - F_{\gamma}) df \right) \cdot \vec{n}}_{\text{Refinement force}} \right] \vec{n}. \end{aligned} \quad (17)$$

Here also, the evolution equation can be viewed as a geodesic evolution with a variable stopping function \mathcal{G}_{Bh} . Since $B(P_{\gamma}, \mathcal{M})$ is explicitly independent of image coordinates, $\nabla B(P_{\gamma}, \mathcal{M}) = 0$. Therefore, we can write evolution (17) in GAC-like form:

$$\frac{\partial \gamma}{\partial t} = \left(\underbrace{\mathcal{G}_{Bh}(P_{\gamma}, \mathcal{M}, F_{\gamma}) \kappa}_{\text{Stopping}} - \underbrace{\nabla \mathcal{G}_{Bh}(P_{\gamma}, \mathcal{M}, F_{\gamma}) \cdot \vec{n}}_{\text{refinement}} \right) \vec{n}, \quad (18)$$

where

$$\mathcal{G}_{Bh}(P_{\gamma}, \mathcal{M}, F_{\gamma}) = B(P_{\gamma}, \mathcal{M}) - \int_{\mathcal{F}} \sqrt{\frac{\mathcal{M}(f)}{P_{\gamma}(f)}} K(f - F_{\gamma}) df. \quad (19)$$

The Bhattacharyya evolution equation can be interpreted in a way similar to the Kullback-Leibler flow. Away from the desired boundary, we can assume that the curve and the model are dissimilar and, therefore, that the feature and the reference distributions have little overlap, which implies for most pixels that $P(F_{\gamma(p)}) > M(F_{\gamma(p)})$ and, consequently, $\mathcal{G}_{Bh} > 0$. When the curve is close to adhering to the target boundary, $P(F_{\gamma(p)}) \approx M(F_{\gamma(p)})$, which implies that $\mathcal{G}_{KL} \approx 0$ and that the curve behavior would be predominantly driven by the gradient term to adhere to the desired boundary. In the event \mathcal{G}_{Bh} in the stopping term evaluates to negative sometimes at some some points, the refinement term acts as a stabilizer to produce a stable descent equation.

3.3 Level Set Implementation

Active curve $\gamma(s, t) : [0, 1] \times \mathbb{R}^+ \rightarrow \Omega$ is implicitly represented by the zero level set of a function $\phi(\mathbf{x}, t) : \Omega \times \mathbb{R}^+ \rightarrow \mathbb{R}$, i.e., $\gamma = \{\mathbf{x} \in \Omega \mid \phi(\mathbf{x}, t) = 0\}$. Recall [5] that when γ evolves according to

$$\frac{\partial \gamma(s, t)}{\partial t} = V(s, t) \vec{n}(s, t), \quad (20)$$

then ϕ evolves according to

$$\forall \mathbf{x} \in \Omega, \quad \frac{\partial \phi(\mathbf{x}, t)}{\partial t} = V(\mathbf{x}, t) \|\nabla \phi(\mathbf{x}, t)\|, \quad (21)$$

with the convention that $\phi > 0$ inside the zero level set and \vec{n} is oriented inward. Therefore, the level set evolution equations corresponding to the flows (13) and (18) are given by

$$\begin{aligned} V(\mathbf{x}, t) = \mathcal{G}(P_{\phi}, \mathcal{M}, F(\mathbf{x})) \kappa_{\phi}(\mathbf{x}, t) \\ - \nabla \mathcal{G}(P_{\phi}, \mathcal{M}, F(\mathbf{x})) \cdot \frac{\nabla \phi(\mathbf{x}, t)}{\|\nabla \phi(\mathbf{x}, t)\|}, \end{aligned} \quad (22)$$

where \mathcal{G} is \mathcal{G}_{KL} for the Kullback-Leibler flow and \mathcal{G}_{Bh} for the Bhattacharyya. These stopping functions are variable of the curve and therefore must be updated during evolution using the sample feature distribution within a narrow band δ around the zero level set of ϕ [5]:

$$P_{\phi}(f) = \frac{\int_{-\delta < \phi(\mathbf{x}) < \delta} K(f - F(\mathbf{x})) d\mathbf{x}}{\int_{-\delta < \phi(\mathbf{x}) < \delta} d\mathbf{x}}. \quad (23)$$

κ_{ϕ} is the mean curvature function of ϕ :

$$\kappa_{\phi}(\mathbf{x}, t) = \text{div} \left(\frac{\nabla \phi(\mathbf{x}, t)}{\|\nabla \phi(\mathbf{x}, t)\|} \right), \quad \forall \mathbf{x} \in \Omega. \quad (24)$$

Geodesic evolution is often quickened by an additional constant speed c along the curve normal [17], resulting in the level set motion [13]:

$$\frac{\partial \phi(\mathbf{x}, t)}{\partial t} = (V(\mathbf{x}, t) + c) \|\nabla \phi(\mathbf{x}, t)\|. \quad (25)$$

The algorithm can be summarized as follows:

Algorithm 1. Feature distribution matching method

Input: Image I and feature distribution model \mathcal{M} .

Output: $\gamma(s, \infty)$ delineating the object of interest.

```

begin
  Initialize the level set function  $\phi$ ;
  repeat
    • Update feature distribution  $P_\phi$  according to (23);
    • Update the evolution speed  $V(x, t)$  according to (22) with  $\mathcal{G}$  given by
       $\mathcal{G}_{KL}$  in (13) or by  $\mathcal{G}_{Bh}$  in (19);
    • Update the level set function according to (25);
  until Convergence ;
end

```

4 EXPERIMENTS

We will describe two sets of experiments and comparisons illustrating various applications of the proposed formulation; one uses a geometric feature, namely, curvature, and the other photometric features, namely, the image gradient and a neighborhood image average. In general, the choice of a feature will depend on the application. The purpose of the first set of experiments (Section 4.1) is to demonstrate that the method can effectively extract, in a single instantiation, all of the regions in an image whose (geometric) feature boundary distribution follows a learned outline pattern. The examples include evaluations over color images from the ETHZ database [54], [55] and over medical images. The purpose of the second set of experiments (Section 4.2) is to show that the formulation can efficiently segment images in the presence of regions which have fading contrast at some of their boundary segments. The examples include a task of anatomical tracking.

In the experiments with the photometric features, we compared this study contour distribution matching functional (abbreviated **CDM** hereafter) to the following functionals:

RDM. The region-based distribution matching functional in [8]. Optimization of this functional seeks a region so that the image distribution within the region most closely matches a learned model.

RL. The region-based likelihood functional commonly used in image segmentation [6], [10], [14], [48]. Optimization of this functional seeks a two-region partition maximizing the conditional probability of pixel data given the learned models within the segmentation regions.

ROP. Concatenation of RDM and the region-based overlap functional in [9], which embeds information about the overlap between the distribution of the image data within the segmentation regions.

GAC. The classical geodesic active contour functional [17] commonly used in image segmentation as an edge-based constraint, which biases the segmentation boundaries toward high gradients of image data.

With the geometric feature, it was relevant to compare to **RL**, **GAC**, and **GAC-SP** which is GAC with a shape prior term.

In all the experiments, the feature distribution is estimated using the kernel width $h = 1$ and the narrow band parameter $\delta = 1$.

4.1 Image-Based Geometric Feature: Extraction of Region Boundaries of a Given Outline Pattern

The purpose here is to recover region boundaries consistent with an outline pattern without prior knowledge of the number of regions. Intensity-based methods would not allow doing this because, as illustrated in the simple synthetic image of Fig. 2, the targeted regions may have the same intensity distribution as unwanted differently shaped regions. Instead, we will use a geometric feature, namely curvature. As explained in Section 2.2, curvature is estimated from the image under the assumption that the region boundary normals coincide with the isophote normals. Using curvature affords a scheme which handles differences in the pose and number of targeted regions. This is in sharp contrast with shape prior methods which require the knowledge of the number of regions and inclusion of pose parameters in the optimization. We begin by showing synthetic examples which we used to verify the relevance of curvature as a feature for the scheme to locate all of the instances of a targeted shape in an image. We follow with real images, including a tracking example. Since we have no knowledge about the target object's position, the initial contour is systematically placed wide out in the image domain to ensure that it encloses all objects. In spite of narrow banding, we noted that this can cause the progression of the CDM curve to be slow in reaching object boundaries. To accelerate the evolution toward region boundaries, we add a nonweighted GAC term to the functional (any term which would guide the curve evolution toward high-contrast boundaries would be acceptable). However, when region boundaries are reached, the GAC term weakens considerably and the CDM becomes preponderant because it causes the curve to adhere to the desired boundaries but move away from all others and vanish. The addition of a GAC term is an implementation aid to faster execution which does not affect the meaning of the CDM functional. It also has the beneficial side effect of reinforcing the positivity of the term multiplying curvature in the evolution equation.

4.1.1 Synthetic Examples

The purpose is to verify in a very simple example that the formulation can detect all of the shapes of the given class, irrespective of scale (and position). Fig. 2 is a synthetic image of a row of three ellipses of appearances, but with approximately the same aspect ratio and a row of different rectangles. The purpose is to segment either all of the rectangular regions, in one instantiation, or all of the ellipsoidal regions but not both. The model curvature (5) distributions are learned from one rectangle and one ellipse independently from those appearing in Fig. 2. Fig. 2c show two final segmentations of the squares (upper) and the ellipses (lower). Fig. 2b show the training images with the contours, and Fig. 2a show the initial positions.

The contour evolves first toward high-contrast boundaries. Once the region boundaries are reached, the proposed CDM flow causes the active curve to remain in coincidence with the desired boundaries but leave and vanish from the others. For example, along the boundary of the rectangles (second row), the feature distribution does not match the model distribution of curvature along an ellipse. Therefore, the contour continues to evolve inside the rectangles, and




Fig. 3. Segmentation of spine bones in a CT image. The model distribution is learned beforehand from a manual segmentation of a single rectangular bone (the middle one). (a) Final contour position. (b) Segmentation with the GAC model. The proposed functional succeeded in guiding the active curve toward the spine bones while omitting the other neighboring parts which have a similar intensity profile. The Bhattacharyya coefficient has been adopted for this example.

delineates only the elliptical regions at convergence. Evidently, neither region-based segmentation (via RL, for instance) nor edge-based segmentation (via GAC, for instance) will be able to distinguish one type of regions from the other (last column of Fig. 2).

4.1.2 Real Image Examples

The purpose of this experiment is to show the advantage over standard algorithms such as GAC in the case of multiple occurrences of the desired object in a real image. Fig. 3 depicts an example of segmentation of vertebrae in a CT scan of the human spine. The distribution of curvature has been approximated by a histogram along the outline of an exemplar vertebra. Although the results are not totally accurate, CDM (Fig. 3a) has been able to correctly segment two vertebrae, do a decent outlining of the others (the poor segmentation of the upper vertebra is due to border effects as no initialization could include this vertebra), and ignore the bone structures on the right. Of course, region-based methods cannot handle this example because the image

profile within the vertebrae is very similar to that of other surrounding structures. Edge-based functionals such as GAC will bias the active contour to all high image gradients which do not correspond necessarily to the edges of the vertebrae. GAC (Fig. 3b) has not been able to do as well on the vertebrae and has included the bone structures on the right.

For this example, we plotted in Fig. 4a the model distribution and the distribution of curvature on the active curve at convergence. The proposed curve evolution method recovered accurately the model. The Bhattacharyya measure obtained at convergence is equal to 0.98, which means a very close match. In Fig. 4b, we plotted the evolution of the optimized Bhattacharyya measure as a function of the number of iterations. An inspection of the graph shows that there is a minor transient decrease from a maximum at iteration 175 before leveling off after iteration 230. This is also happening at about iteration 500 in the graph of Fig. 4, although the transient decrease is even slighter. This behavior is probably an artifact of the

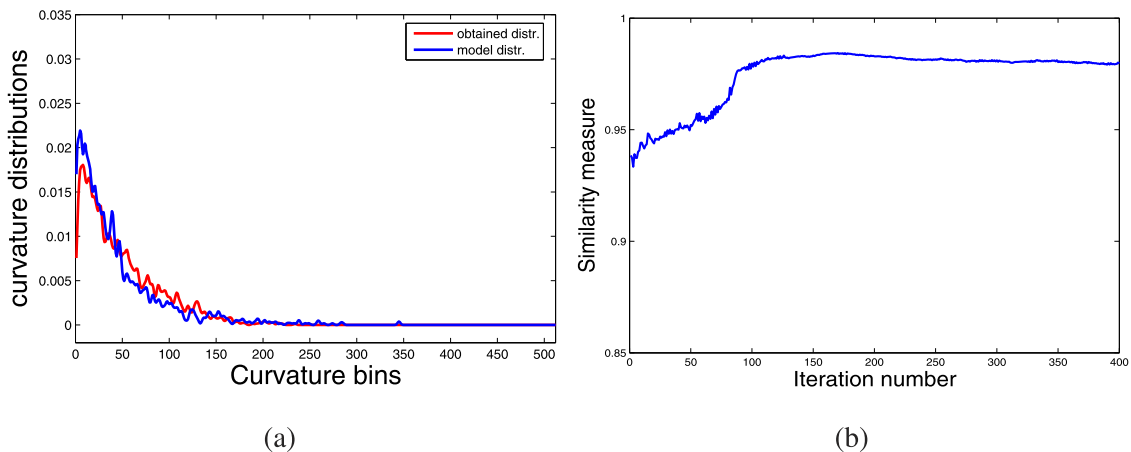


Fig. 4. Segmentation of spine bones in a CT image. (a) Distribution of the curvature along the final curve and the model distribution. (b) Evolution of the functional as a function of the number of iterations.

combined effect of digital curve processing and equivocal object boundaries causing the curve to drift incidentally about a local minimum of the functional in the immediate vicinity of the desired boundary. The principal digital approximations in this application occur with 1) the discretization of the active curve via the zero level set, 2) the estimation of the derivatives in the calculation of curvature, and 3) the binning of curvature values (for histogram calculations).

4.1.3 ETHZ Database

In this section, we use images from the ETHZ database [54], [55]. This data set contains about 200 images of objects of one of five types of shapes, such as bottles and Apple logos. It has been used to test algorithms which detect and then recognize objects in images. The objects appear in various sizes, positions, colors, and there are within-class variations in shape. The ground truth object shape is provided with each image. Also provided is an edge map obtained by the Berkeley natural boundary detector [56], [57]. Although many of the database images are not of use to validate our algorithm, others afford a good testbed. To be of use in our application, an image should, of course, contain several distinct objects to be able to show the detection of the relevant object while ignoring the others. Moreover, the target object should be present in other images, modulo interesting shape variations, to be able to learn a model independently of the test image. Also of use are the images which contain several instances of the relevant object, ideally each modulo a variation in the shape, to show that the algorithm can detect all of the instances.

For each class of objects, we learn the model histogram on the ground truth of an image different from the test images. Given the model histogram, it is then expedient to effect the algorithm on the edge map rather than the image. In all the examples, this resulted in an efficient fast implementation of the algorithm.

Fig. 5 depicts a sample of the results obtained with the ETHZ data set. The first row contains the test images. The initial curves, which are placed so as to enclose all the objects, are also shown in these images. The second row contains the edge maps of the images in the first row. The model curvature histograms are evaluated on the model shapes in the images of the third row. A single model shape is used for a class of objects and it comes from an image different from the test images. The last four rows show positions of the evolving curve until convergence depicted in the last row of Fig. 5. When an object boundary is reached, the distribution matching flow causes the curve to coincide with the desired boundaries but close away from the undesired ones and vanish.

As done in all the tests described above, the initial contour must be placed so as to enclose all the desired objects in the image, for instance by being wide enough to enclose all objects. When the initial contour does not contain a desired object, or contains it only partially, the detection will fail. This is shown in Fig. 6. In Fig. 6a in the first row, there are three different initial contours all enclosing the desired object. As a result, they all converge to coincide

with the target object (the bottle) as shown in images Fig. 6b, Fig. 6c, and Fig. 6d of the first row, while ignoring the other objects. In the second row, Fig. 6a shows an initialization which does not enclose the desired object (bottle). As a result, the algorithm fails to detect the desired boundary (see Fig. 6b). Fig. 6c has been constructed by copying a rectangular portion from the image background and pasting it so as to occlude partially the bottle. The resulting bottle image no longer resembles the model and should be ignored by the active curve because the formulation has no explicit provision to handle occlusion. Fig. 6d shows that this is the case as the contour passes through the desired boundary. A shape prior with pose parameters started close to the object would have been able to recover the desired contour.

We tested the dependence of the results on the training by segmenting one image using three distribution models different from the actual image. Fig. 7a shows the input image with the initial contour for all the three experiments. Columns (b), (c), and (d) of Fig. 7 show the results corresponding to each distribution model: The model shapes are shown in the first row and the final curve positions in the second row. They all converge to coincide with the target object (the bottle) with some minor differences.

When a desired object occurs more than once in the same image, a shape prior will not be able to segment all of the regions in the image. Moreover, shape priors which include pose parameters require an initialization that is close to the target object, which is often impractical. In Fig. 8, we show the results obtained on two different test images using the GAC model with a template matching shape prior term as in [58]. For both images, the template used actually corresponds to one of the objects in the image, the smallest ellipse in the first image and the black cup to the right of the second image. Using one of the desired objects as the template simplifies the problem because this forgoes the need to optimize with respect to the pose parameters. As expected, the contours evolved toward the objects corresponding to the templates but missed all the other desired object instances.

4.1.4 A Tracking Example

Fig. 9 depicts tracking of both the left ventricle cavity (first row) and the right ventricle (second row) using the curvature as feature. For each frame, the model distributions were learned from the result of the previous frame. The first frame of the sequence was segmented manually. Based on the learned outline pattern, the proposed method succeeds to distinguish between the left and the right ventricles in the considered sequence.

4.2 Photometric Feature: Segmentation in the Presence of Fading Contrast along Boundary Segments

In this set of experiments, we show application of the proposed curve evolution to difficult situations where parts of the target boundary correspond to weak transitions of image data. This frequently occurs in medical images. We ran several experiments of segmentation and tracking of the left ventricle inner and outer boundaries in cardiac

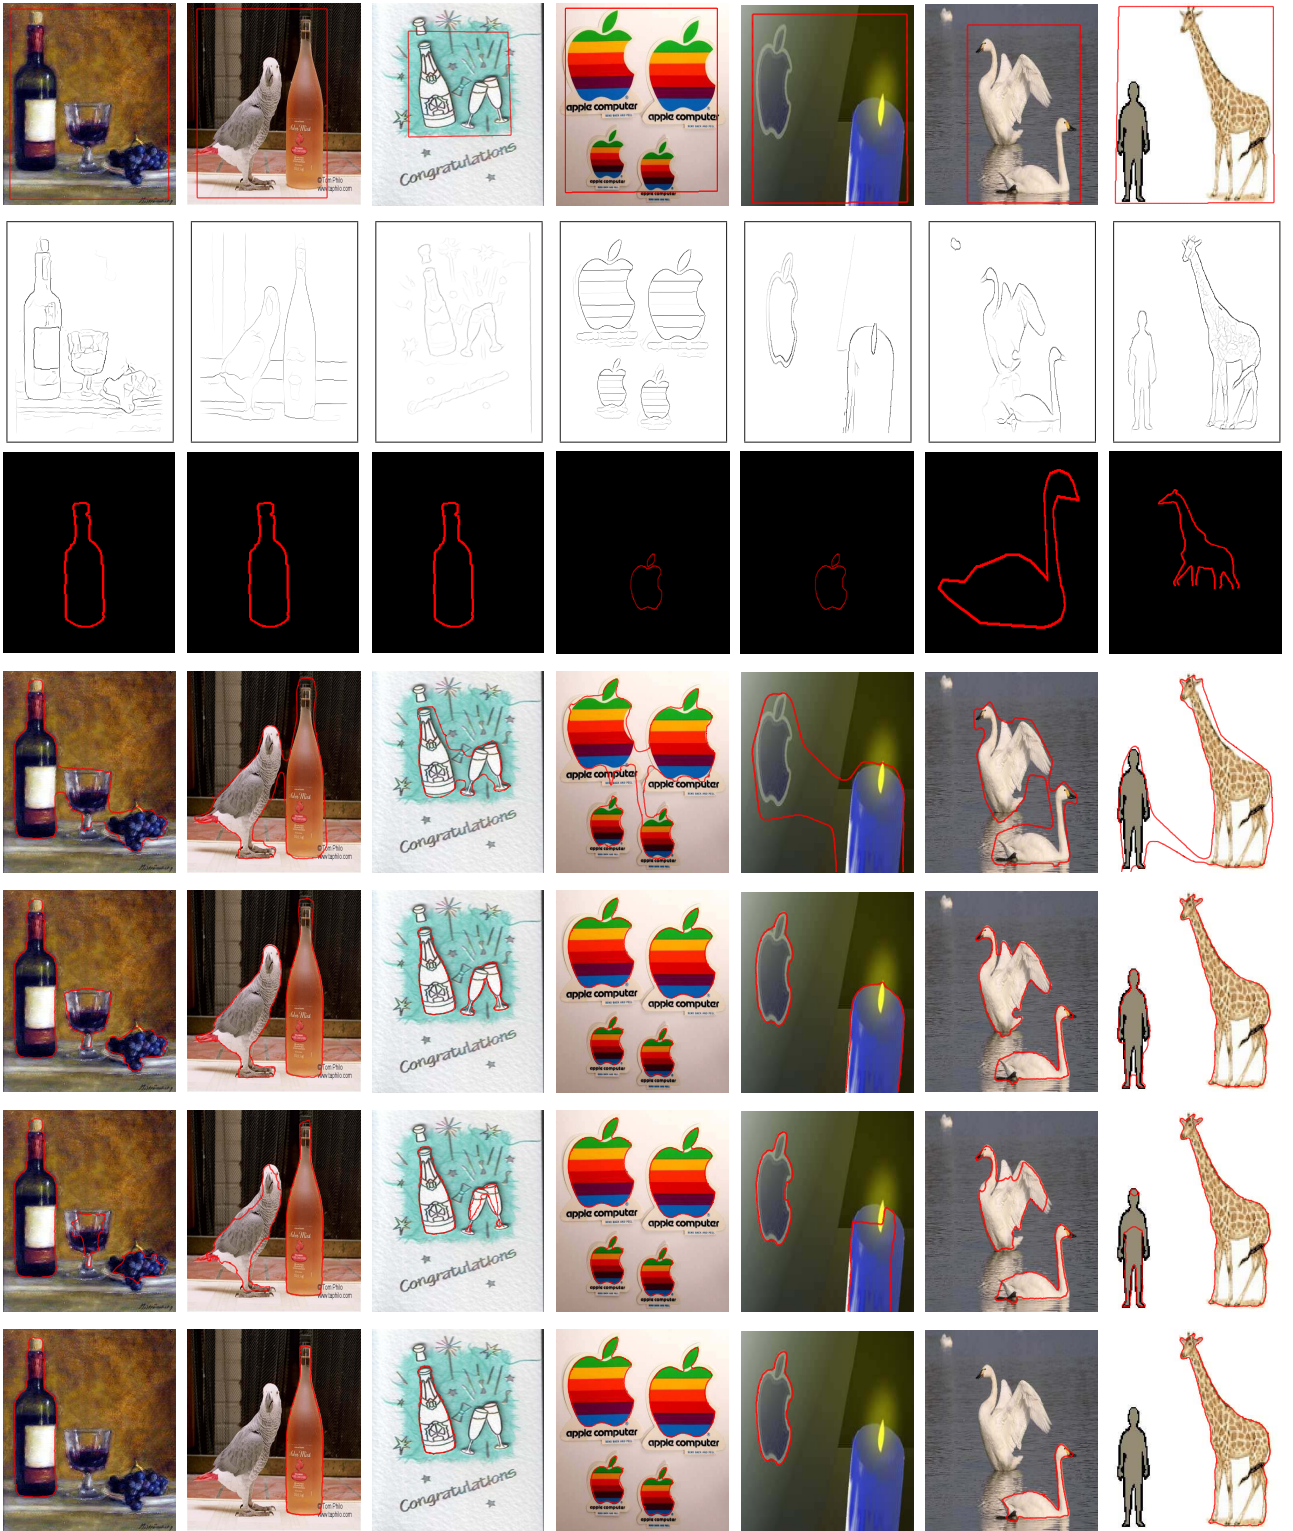


Fig. 5. A sample of the results on the ETHZ data set. Row 1: Initial curve. Row 2: Edge contours. Row 3: Shape model. Rows 4-6: Intermediate curve positions. Row 7: Final curve position. The Kullback-Leibler divergence has been adopted for these examples.

magnetic resonance image sequences. To explicitly demonstrate the positive effect of the proposed contour-based functional, we first give two typical segmentation examples. In these examples, we run tests that show how the proposed functional can lead to improvements in accuracy over other region-based and edge-based functionals. Then, we give a representative sample of the tracking results supported by

quantitative performance evaluations by comparisons with manual delineations.

4.2.1 Segmentation of the Inner Boundary of the LV

Accurate LV segmentation is acknowledged as a difficult problem, and is essential in automating the diagnosis of cardiovascular diseases [28]. A typical example is shown in

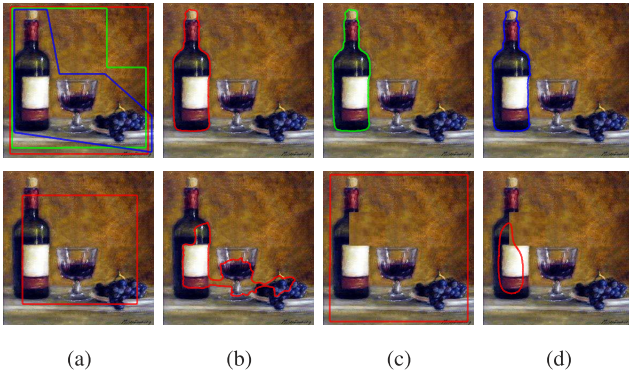


Fig. 6. Effect of initialization and occlusion. Top row: (a) Different initializations, (b)-(d) corresponding final curve positions. Bottom row: (a) Initialization with a curve which does not contain the desired object (the bottle) and (b) position of the curve showing it has missed the object; (c) image with occlusion; the object contour no longer resembles the bottle model and should be ignored by the moving curve and (d) position of the curve corresponding to (c) after it has missed the object boundary.

Fig. 10 where the purpose is to find the boundary between the heart cavity and the background. The manual segmentation by an expert is depicted by the green curve in Fig. 10b. This example is difficult because the papillary muscles within the cavity and the background are connected and have the same intensity profile. Therefore, a part of the target boundary corresponds to very weak image transitions, i.e., the norm of the image gradient is null or nearly null.

To illustrate how the proposed functional CDM refines the segmentation in such cases, we report comparisons with the curve evolution methods based on RDM, RL, ROP, and GAC. For all the functionals, we used the same initialization depicted by the red curve in Fig. 10c, and learned model distributions from the ground truth. We assessed the similarities between the ground truth and the segmentations obtained with CDM, GAC, RL, RDM, and ROP using two measures: the *Dice Metric* (*DM*) [28] and the *Root Mean Squared Error* (*RMSE*) with symmetric nearest neighbor correspondences [59]. The Dice Metric is region-based and is given by¹ $DM = \frac{2A_{am}}{A_a + A_m}$, with A_a , A_m , and A_{am} denoting the areas of the automatically detected region (region inside the curve), the corresponding ground-truth region, and the intersection between them, respectively. The *RMSE* is contour-based, and measures the distance between manual and automatic boundaries over N points as follows:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{x}_i - \tilde{x}_i)^2 + (\hat{y}_i - \tilde{y}_i)^2}, \quad (26)$$

where (\hat{x}_i, \hat{y}_i) is a point on the automatic boundary and $(\tilde{x}_i, \tilde{y}_i)$ the corresponding point on the manual boundary [59].

For comparative purposes, we report the *DM* and *RMSE* corresponding to all the functionals in Table 1, and show the curve at convergence (red curve) with the ground truth curve (green curve) in Fig. 10.

The geodesic active contour, biases the curve toward high gradients of the image (Fig. 10e), thereby yielding the




Fig. 7. Segmentation results dependence on varying model distributions. (a) Initial curve. (b)-(d) Three different models (top) and the corresponding final curve positions (bottom).

lowest conformity to the ground truth, i.e., the highest *RMSE* and the lowest *DM* (refer to Table 1). Furthermore, region-based functionals fail to accurately recover the target boundary due to the lack of relevant image information on the curve (refer to Figs. 10f, 10g, and 10h). The CDM, contrarily to the other methods, accurately refined the results. As depicted by the red curve in Fig. 10d, it stopped the curve at convergence in a position very similar to the ground truth, yielding the lowest *RMSE* and highest *DM*, which correspond to the best conformity to the ground truth (refer to Table 1). This example explicitly illustrates the usefulness of the proposed curve evolution in the application at hand. To quantitatively demonstrate the positive effect of the proposed functional, we report in Table 2 the following statistics of the segmentations obtained with RDM, ROP, and CDM:

- *Contour-based similarity measure.* The Bhattacharyya measure of similarity between the distribution of image feature within a narrow band around the curve and a model, i.e., CDM at convergence.
- *Region-based matching measure.* The Bhattacharyya measure of similarity between the distribution of image feature inside the curve at convergence and a model, i.e., RDM at convergence.
- *Region-based overlap measure.* The Bhattacharyya measure of similarity between the distribution of image feature inside and outside the curve at convergence.




Fig. 8. Results using the GAC model with a shape prior term. (a) and (b) Final curve positions.

1. Note that *DM* is always in $[0, 1]$, where *DM* equal to 1 indicates a perfect match between manual and automatic segmentations.




Fig. 9. Tracking of both the left ventricle cavity (first row) and the right ventricle (second row) in an MR sequence containing 25 frames (f_x depicts frame x). Feature: The curvature. For each frame, the model distributions were learned from the result of the previous frame. The first frame of the sequence was segmented manually.

For all the considered methods, we obtained almost the same region-based measures (refer to the first two rows in Table 2), although the segmentation results at convergence are different as shown in Figs. 10d, 10g, and 10h. On the contrary, curve evolution with CDM increased the contour-based similarity, leading to a measure very different from the ones obtained with RDM and ROP (refer to the last rows in Table 2).

These statistics show that the region-based functionals in [9] and [8] fail in this example, whereas the distribution of an image feature on the contour can limit the space of possible solutions to a contour very close to the ground truth.

For this example, we plotted the distribution of the image feature along the final contour and the model in Fig. 11a. The proposed curve evolution method recovered accurately the learned model. We also plotted in Fig. 11b the evolution




Fig. 10. Detection of the inner boundary of the LV (endocardium) in a Magnetic Resonance image. The purpose is to find the boundary between the heart cavity (foreground) and the background, as shown by the manual segmentation provided by an expert and depicted by the green curve in (b). (d) The curve at convergence with the proposed functional CDM (red curve) superimposed on the ground-truth curve (green curve). (e)-(h) The results with, respectively, GAC [17], RL [10], RDM [8], and ROP [9]. For all the functionals, we used the same initialization depicted by the red curve in (c), and model distributions were learned from the ground truth. Image feature for CDM: Average intensity of pixel neighborhood (rectangular neighborhood of size 5×5 centered on the pixel). $c = 0$ (no constant velocity). The curve evolution is limited to the yellow box containing the region of interest in (a).

TABLE 1
Evaluation of DM and $RMSE$ (in Pixels)
Obtained by Curve Evolution
with the Proposed Functional (CDM) and Other Functionals

Method	GAC	RL	RDM	ROP	CDM
DM	0.71	0.77	0.91	0.91	0.97
$RMSE$	4.78	4.44	1.73	1.76	0.79

CDM yielded the lowest $RMSE$ and highest DM and, therefore, the best conformity to the ground truth.

of the optimized functional, the Bhattacharyya measure in this case, as a function of the iteration number. The Bhattacharyya measure converged approximately to its maximum possible value ($B(P_\gamma, \mathcal{M}) = 0.99 \approx 1$).

4.2.2 Segmentation of the Outer Boundary of the LV

In the example depicted in Fig. 12, the purpose is to detect the outer boundary of the LV (epicardium). This problem is also acknowledged as difficult because the image gradients between the bottom part of the LV and the background are very small (refer to the manual delineation in Fig. 12a). The model distribution is learned from a previous frame and the initial curve is depicted with the red square in Fig. 12b. The curve evolution is shown by several intermediate steps in Fig. 12c. The proposed functional allowed successfully stopping the curve at the bottom part of the LV, where the background and heart myocardium are connected and have approximately the same intensity profile. In comparison to

TABLE 2
Left Ventricle Inner Boundary Example—Region-Based and
Contour-Based Bhattacharyya Similarity Measures Obtained
with RDM [8], ROP [9], and the Proposed Energy (CDM)

Method	RDM	ROP	CDM
Region-based matching measure	0.99	0.99	0.99
Region-based overlap measure	0.63	0.64	0.64
Contour-based matching measure	0.88	0.85	0.99

the expected delineation, the proposed method yielded a $RMSE$ equal to 1.66 pixels and a DM equal to 0.95. At convergence, the Bhattacharyya measure is equal to 0.91, whereas its initial value was 0.68.

4.2.3 Tracking Examples

Tracking the LV inner and outer boundaries in cardiac MR sequences is an essential yet challenging task in cardiac image analysis [28], [60]. We applied the proposed method to a set of sequences in order to track LV inner and outer boundaries. For each frame, the model distribution and the initial curve are obtained from the result of the previous frame. The first frame of each sequence is segmented manually. The tracking performance appraisal is carried out by comparison with independent manual segmentations over five sequences. Each sequence contains 10 frames, which amounts to segmenting 45 images automatically. We give a representative sample of the results obtained in Fig. 13

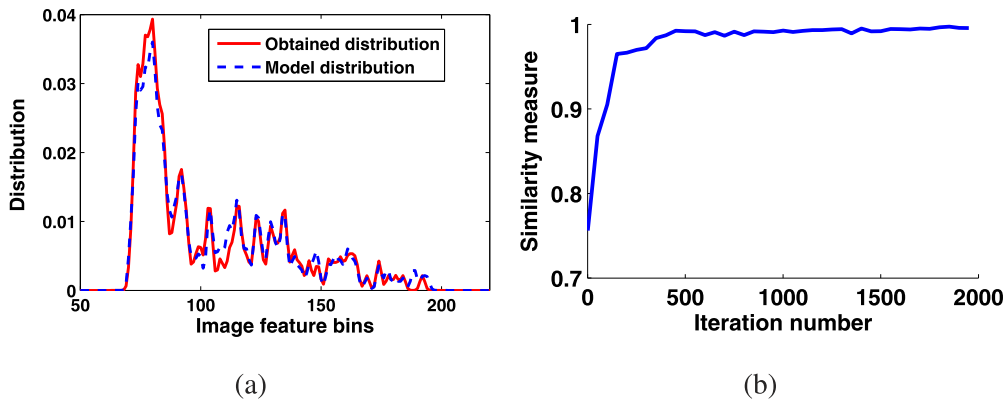


Fig. 11. Detection of the inner boundary of the LV. (a) Distribution of the image feature on the final curve and the model. (b) Evolution of the functional as a function of the iteration number.




Fig. 12. Detection of the outer boundary of the LV (epicardium) in an MR image. Model distribution \mathcal{M} is learned from a previous frame. $c = 0.1$. The image feature is the average intensity of pixel neighborhood (rectangular neighborhood of size 5×5 centered on the pixel). (b)-(d) The initial curve, several intermediate evolution steps, and the final curve (red curve), respectively, superimposed to the expected delineation (green curve).




Fig. 13. A representative sample of the results of tracking the inner (endocardium) and outer (epicardium) boundary of the LV in five MR sequences, each containing 10 frames (systole phase of the cardiac cycle). Image feature: Image gradient. $c = 0$. $sxfy$ depicts frame y in sequence x . A video of sequence 2 is uploaded with the submission to illustrate the tracking.

for visual inspection. The method succeeded in stopping the curve at relevant positions where the image transitions are very small. We report in Table 3 the statistics of the performance measures (DM and $RMSE$). We obtained an average DM equal to 0.89 for the LV inner boundaries and to 0.93 for the outer ones. Note that an average DM higher than 0.80 indicates an excellent agreement with manual segmentations [61], and an average DM higher than 0.90 is, generally, difficult to obtain [60]. For instance, the study in [60] reports an average DM equal to 0.81. The performance measures obtained demonstrate that the proposed functional leads to competitive results.

TABLE 3
Tracking Performance Evaluation of the Proposed Curve Evolution over Five Cardiac Sequences by Comparisons with Independent Manual Segmentations—Means and Standards Deviations of $RMSE$ and DM

Measure (<i>mean</i> \pm <i>std</i>)	$RMSE$	DM
<i>endocardium (inner)</i>	1.96 ± 0.59	0.89 ± 0.01
<i>epicardium (outer)</i>	2.21 ± 0.22	0.93 ± 0.01

5 CONCLUSION

This study addressed the problem of recovering region boundary patterns in an image by the minimization of an active curve functional, which measures the similarity between a feature distribution on the curve and a learned model distribution. This distribution matching drives the active curve until it settles on the boundaries of interest, i.e., boundaries on which the feature follows the model distribution. The method was formulated for the Kullback-Leibler divergence and the Bhattacharyya measure. It was applied in two challenging circumstances, specifically the extraction of boundaries fitting a learned outline pattern and segmentation in the presence of boundaries with weakly contrasted segments. The scheme used the distribution of a geometric feature for the first task and an photometric feature for the second task. The formulation is fundamentally different from region-based schemes, which cannot distinguish between regions having the same image distributions. The evolution equations we obtained can be viewed as GACs having variable stopping functions, which have two major differences from the usual GAC stopping function. First, they are functions of both the image and the curve, rather than just the image, as with the usual GAC. Second, they use global information, namely, the feature distribution on the curve, rather than just pixelwise, as with GAC. Several experiments confirmed that the proposed method outperforms region and edge-based formulations in adverse conditions.

REFERENCES

- [1] M. Holtzman-Gazit, R. Kimmel, N. Peled, and D. Goldsher, "Segmentation of Thin Structures in Volumetric Medical Images," *IEEE Trans. Image Processing*, vol. 15, no. 2, pp. 354-363, Feb. 2006.
- [2] I. Ben Ayed, A. Mitiche, and Z. Belhadj, "Multiregion Level Set Partitioning on Synthetic Aperture Radar Images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 5, pp. 793-800, May 2005.
- [3] F.N. Mortensen, *Progress in Autonomous Robot Research*. Nova Science Publishers, 2008.
- [4] G. Monteiro, J. Marcos, M. Ribeiro, and J. Batista, "Robust Segmentation for Outdoor Traffic Surveillance," *Proc. 15th IEEE Int'l Conf. Image Processing*, Oct. 2008.
- [5] J.A. Sethian, *Level Set Methods and Fast Marching Methods*. Cambridge Univ. Press, 1999.
- [6] T. Chan and L. Vese, "Active Contours without Edges," *IEEE Trans. Image Processing*, vol. 10, no. 2, pp. 266-277, Feb. 2001.
- [7] D. Cremers, M. Rousson, and R. Deriche, "A Review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape," *Int'l J. Computer Vision*, vol. 72, no. 2, pp. 195-215, 2007.
- [8] D. Freedman and T. Zhang, "Active Contours for Tracking Distributions," *IEEE Trans. Image Processing*, vol. 13, no. 4, pp. 518-526, Apr. 2004.
- [9] I. Ben Ayed, S. Li, and I. Ross, "A Statistical Overlap Prior for Variational Image Segmentation," *Int'l J. Computer Vision*, vol. 85, no. 1, pp. 115-132, 2009.
- [10] M. Rousson and D. Cremers, "Efficient Kernel Density Estimation of Shape and Intensity Priors for Level Set Segmentation," *Proc. Medical Image Computing and Computer-Assisted Intervention*, pp. 757-764, 2005.
- [11] O.V. Michailovich, Y. Rathi, and A. Tannenbaum, "Image Segmentation Using Active Contours Driven by the Bhattacharyya Gradient Flow," *IEEE Trans. Image Processing*, vol. 16, no. 11, pp. 2787-2801, Nov. 2007.
- [12] X. Xie and M. Mirmehdi, "MAC: Magnetostatic Active Contour Model," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 30, no. 4, pp. 632-646, Apr. 2008.
- [13] C. Li, C. Xu, C. Gui, and M.D. Fox, "Level Set Evolution without Re-Initialization: A New Variational Formulation," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2005.
- [14] N. Paragios and R. Deriche, "Geodesic Active Regions and Level Set Methods for Supervised Texture Segmentation," *Int'l J. Computer Vision*, vol. 46, no. 3, pp. 223-247, 2002.
- [15] R. Kimmel and A.M. Bruckstein, "Regularized Laplacian Zero Crossings as Optimal Edge Integrators," *Int'l J. Computer Vision*, vol. 53, no. 3, pp. 225-243, 2003.
- [16] M. Kass, A.P. Witkin, and D. Terzopoulos, "Snakes: Active Contour Models," *Int'l J. Computer Vision*, vol. 1, no. 4, pp. 321-331, 1988.
- [17] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic Active Contours," *Int'l J. Computer Vision*, vol. 22, no. 1, pp. 61-79, 1997.
- [18] N. Paragios, O. Mellina-Gottardo, and V. Ramesh, "Gradient Vector Flow Fast Geometric Active Contours," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 26, no. 3, pp. 402-407, Mar. 2004.
- [19] S. Kichenassamy, A. Kumar, P.J. Olver, A. Tannenbaum, and A.J. Yezzi, "Gradient Flows and Geometric Active Contour Models," *Proc. Fifth IEEE Int'l Conf. Computer Vision*, pp. 810-815, 1995.
- [20] X. Bresson, S. Esedoglu, P. Vanderghynst, J. Thiran, and S. Osher, "Fast Global Minimization of the Active Contour/Snake Model," *J. Math. Imaging and Vision*, vol. 28, no. 2, pp. 151-167, 2007.
- [21] A. Vasilievskiy and K. Siddiqi, "Flux Maximizing Geometric Flows," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 12, pp. 1565-1578, Dec. 2002.
- [22] S.C. Zhu and A. Yuille, "Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for Multiband Image Segmentation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 118, no. 9, pp. 884-900, Sept. 1996.
- [23] C. Samson, L. Blanc-Feraud, G. Aubert, and J. Zerubia, "A Level Set Model for Image Classification," *Int'l J. Computer Vision*, vol. 40, no. 3, pp. 187-197, 2000.
- [24] A. Mansouri, A. Mitiche, and C. Vazquez, "Multiregion Competition: A Level Set Extension of Region Competition to Multiple Region Partitioning," *Computer Vision and Image Understanding*, vol. 101, no. 3, pp. 137-150, 2006.
- [25] I. Ben Ayed, A. Mitiche, and Z. Belhadj, "Polarimetric Image Segmentation via Maximum Likelihood Approximation and Efficient Multiphase Level Sets," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 28, no. 9, pp. 1493-1500, Sept. 2006.
- [26] M. Ben Salah, A. Mitiche, and I. Ben Ayed, "Effective Level Set Image Segmentation with a Kernel Induced Data Term," *IEEE Trans. Image Processing*, vol. 19, no. 1, pp. 220-232, Jan. 2010.
- [27] C. Vazquez, A. Mitiche, and R. Laganieri, "Joint Segmentation and Parametric Estimation of Image Motion by Curve Evolution and Level Sets," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 28, no. 5, pp. 782-793, May 2006.
- [28] I. Ben Ayed, S. Li, and I. Ross, "Embedding Overlap Priors in Variational Left Ventricle Tracking," *IEEE Trans. Medical Imaging*, vol. 28, no. 12, pp. 1902-1913, Dec. 2009.
- [29] D. Cremers, N. Sochen, and C. Schnorr, "Towards Recognition-Based Variational Segmentation Using Shape Priors and Dynamic Labeling," *Proc. Int'l Conf. Scale Space Theories in Computer Vision*, pp. 388-400, 2003.
- [30] T. Chan and W. Zhu, "Level Set Based Shape Prior Segmentation," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, vol. 2, pp. 1164-1170, June 2005.
- [31] M.E. Leventon, W.E. Grimson, and O. Faugeras, "Statistical Shape Influence in Geodesic Active Contours," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 316-323, June 2000.
- [32] D. Cremers, S. Osher, and S. Soatto, "Kernel Density Estimation and Intrinsic Alignment for Shape Priors in Level Set Segmentation," *Int'l J. Computer Vision*, vol. 69, no. 3, pp. 335-351, 2006.
- [33] M. Rousson and N. Paragios, "Shape Priors for Level Set Representations," *Proc. European Conf. Computer Vision*, 2002.
- [34] A. Foulonneau, P. Charbonnier, and F. Heitz, "Affine-Invariant Geometric Shape Priors for Region-Based Active Contours," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 28, no. 8, pp. 1352-1357, Aug. 2006.
- [35] A. Foulonneau, P. Charbonnier, and F. Heitz, "Multi-Reference Shape Priors for Active Contours," *Int'l J. Computer Vision*, vol. 81, no. 1, pp. 68-81, 2009.
- [36] I. Ben Ayed, A. Mitiche, M. Ben Salah, and S. Li, "Finding Image Distributions on Active Curves," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 3225-3232, 2010.
- [37] S.C. Zhu, "Statistical Modeling and Conceptualization of Visual Patterns," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 6, pp. 691-712, June 2003.

- [38] M.P. Do Carmo, *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.
- [39] A. Mitiche and I. Ben Ayed, *Variational and Level Set Methods in Image Segmentation*, first ed. Springer, Sept. 2010.
- [40] F. Goudail, P. Refregier, and G. Delyon, "Bhattacharyya Distance as a Contrast Parameter for Statistical Processing of Noisy Optical Images," *J. Optical Soc. Am. A*, vol. 21, no. 7, pp. 1231-1240, 2004.
- [41] A. Mansouri and A. Mitiche, "Region Tracking via Local Statistics and Level Set PDEs," *Proc. IEEE Int'l Conf. Image Processing*, pp. 605-608, 2002.
- [42] A. Myronenko and X.B. Song, "Global Active Contour-Based Image Segmentation via Probability Alignment," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 2798-2804, 2009.
- [43] F. Lecellier, S. Jehan-Besson, J. Fadili, G. Aubert, and M. Revenu, "Optimization of Divergences within the Exponential Family for Image Segmentation," *Proc. Second Int'l Conf. Scale Space and Variational Methods in Computer Vision*, pp. 137-149, 2009.
- [44] S. Vicente, V. Kolmogorov, and C. Rother, "Joint Optimization of Segmentation and Appearance Models," *Proc. 12th IEEE Int'l Conf. Computer Vision*, 2009.
- [45] V. Kolmogorov and Y. Boykov, "What Metrics Can be Approximated by Geo-Cuts, or Global Optimization of Length/Area and Flux," *Proc. 10th IEEE Int'l Conf. Computer Vision*, 2005.
- [46] Y. Boykov and M.-P. Jolly, "Interactive Graph Cuts for Optimal Boundary and Region Segmentation of Objects in N-D Images," *Proc. Eighth IEEE Int'l Conf. Computer Vision*, pp. 105-112, 2001.
- [47] Y. Boykov, O. Veksler, and R. Zabih, "Fast Approximate Energy Minimization via Graph Cuts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1222-1239, Nov. 2001.
- [48] Y. Boykov and G. Funka Lea, "Graph Cuts and Efficient N-D Image Segmentation," *Int'l J. Computer Vision*, vol. 70, no. 2, pp. 109-131, 2006.
- [49] I. Ben Ayed, H.-M. Chen, K. Punithakumar, I. Ross, and S. Li, "Graph Cut Segmentation with a Global Constraint: Recovering Region Distribution via a Bound of the Bhattacharyya Measure," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2010.
- [50] L. Mukherjee, V. Singh, and C.R. Dyer, "Half-Integrality Based Algorithms for Cosegmentation of Images," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
- [51] T. Pock, D.C.A. Chambolle, and H. Bischof, "A Convex Relaxation Approach for Computing Minimal Partitions," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
- [52] F. Guichard and J.M. Morel, "Image Analysis and P.D.E.s," IPAM GBM Tutorial, 2001.
- [53] A. Steiner, R. Kimmel, and A.M. Bruckstein, "Planar Shape Enhancement and Exaggeration," *Graphical Models and Image Processing*, vol. 60, no. 2, pp. 112-124, 1998.
- [54] V. Ferrari, T. Tuytelaars, and L.V. Gool, "Object Detection by Contour Segment Networks," *Proc. European Conf. Computer Vision*, May 2006.
- [55] V. Ferrari, F. Jurie, and C. Schmid, "From Images to Shape Models for Object Detection," *Int'l J. Computer Vision*, 2009.
- [56] D. Martin, C. Fowlkes, and J. Malik, "Learning to Detect Natural Image Boundaries Using Local Brightness, Color, and Texture Cues," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 26, no. 5, pp. 530-549, May 2004.
- [57] A. Berg, T. Berg, and J. Malik, "Shape Matching and Object Recognition Using Low Distortion Correspondance," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, 2005.
- [58] N. Paragios, M. Rousson, and V. Ramesh, "Matching Distance Functions: A Shape to Area Variational Approach for Global to Local Registration," *Proc. European Conf. Computer Vision*, pp. 775-790, 2002.
- [59] X. Papademetris, A. Sinusas, D. Dione, R. Constable, and J. Duncan, "Estimation of 3-D Left Ventricular Deformation from Medical Images Using Biomechanical Models," *IEEE Trans. Medical Imaging*, vol. 21, no. 7, pp. 786-800, July 2002.
- [60] M. Lynch, O. Ghita, and P.F. Whelan, "Segmentation of the Left Ventricle of the Heart in 3-D+t MRI Data Using an Optimized Nonrigid Temporal Model," *IEEE Trans. Medical Imaging*, vol. 27, no. 2, pp. 195-203, Feb. 2008.
- [61] C. Pluempitwiriwawej, J.M.F. Moura, Y.-J. LinWu, and C. Ho, "STACS: New Active Contour Scheme for Cardiac MR Image Segmentation," *IEEE Trans. Medical Imaging*, vol. 24, no. 5, pp. 593-603, May 2005.



Mohamed Ben Salah received the PhD degree in computer science from the National Institute of Scientific Research (INRS-EMT), Montreal, Quebec, Canada, in January 2011. He joined the Department of Computing Science, University of Alberta in February 2011 as a postdoctoral fellow. His research interests include image and motion segmentation with focus on level set and graph cut methods, and shape priors.



Ismail Ben Ayed received the PhD degree (with the highest honor) in computer science from the National Institute of Scientific Research (INRS-EMT), University of Quebec, Montreal, Canada, in May 2007. Since then, he has been a research scientist with GE Healthcare, London, Ontario, Canada. His research interests include computer vision, image processing, machine learning, optimization, and their applications in medical image analysis. He coauthored a book, more than 40 papers in reputable journals and conferences, and three patents. He received a GE innovation award and two NSERC fellowships. He is a member of the IEEE.



Amar Mitiche received the Licence Ès Sciences degree in mathematics from the University of Algiers and the PhD degree in computer science from the University of Texas at Austin. Currently, he is working as a professor at the Institut National de Recherche Scientifique (INRS), Department of Telecommunications (INRS-EMT), in Montreal, Quebec, Canada. His research is in computer vision. His current interests include image segmentation, motion analysis in monocular and stereoscopic image sequences (detection, estimation, segmentation, tracking, 3D interpretation) with a focus on level set and graph cut methods, and written text recognition with a focus on neural networks methods. He is a member of the IEEE Computer Society.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.