Pattern Recognition Letters 32 (2011) 1240–1249

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Using a geometric formulation of annular-like shape priors for constraining variational level-sets

M. Alessandrini [a], T. Dietenbeck [b], O. Basset [b], D. Friboulet [b], O. Bernard [b,*]

[a] ARCES, Università di Bologna, Bologna, Italy
[b] CREATIS, CNRS UMR5220, INSERM U630, Université de Lyon, Insa-Lyon, Villeurbanne, France

ARTICLE INFO

Article history:
Received 8 March 2010

Communicated by W. Zhao

Keywords:
Level-sets
Segmentation
Annular shapes
Least square fitting

ABSTRACT

In this paper we address the segmentation of images exhibiting annular like shapes which may be approximated by two elliptical contours. Such patterns are indeed recurrent in many image processing applications. In this context, we develop a level-set framework specifically dedicated to the detection of annular shapes. Thanks to a fast solution to the least-squares fitting problem of similar patterns, our model handles the segmentation task efficiently with a single level-set function. The behavior of this approach is illustrated on images from various fields. An evaluation is then performed for the myocardium detection in MRI and ultrasound cardiac images.

1. Introduction

Segmentation is an important topic in image processing. Many approaches have been proposed for this problem, among which level-set based methods have received an increasing attention (Caselles et al., 1997; Chan and Vese, 2001; Bernard et al., 2009). In this context, many applications using only image information (i.e. image gray levels, image gradient, image statistics) have been developed (Osher and Fedkiw, 2002). It is now well established that such approaches fail to segment meaningful objects from images when they are occluded, when some parts of them are in low contrast regions or even missing. In these cases, higher-level prior knowledge about the expected shape can be exploited in order to successfully segment the desired target.

The introduction of general shape priors has been widely studied in level-set literature (Leventon et al., 2000; Rousson and Paragios, 2002; Chen et al., 2002; Pardo et al., 2004; Chan and Zhu, 2005; Cremers et al., 2007). Level-set-based methods correspond to a class of deformable models where the object to be segmented is captured by propagating an interface represented by the zero level-set of a smooth function. The evolution of this interface is derived through a variational formulation: the segmentation problem is expressed as the minimization of an energy functional that reflects the properties of the objects to be recovered. In this context, the shape information is generally taken into account by considering a linear combination of two terms in the energy functional. The first one, referred as *data attachment term*, drives the evolving interface to particular image features. The second, referred as *shape prior term*, preserves the similarity between the segmenting contour and a reference shape model.

The choice of the prior model is a fundamental issue: typically it is derived from the statistical analysis of a set of training examples (Leventon et al., 2000; Chen et al., 2002; Chan and Zhu, 2005). Since orientation and scale are a priori unknown, a transformation is performed at each iteration to align the segmented contour with the prior shape for similarity measurement purposes. This step is mainly addressed in a steepest gradient descent framework (Chen et al., 2002; Chan and Zhu, 2005).

In this paper[1] we focus on the segmentation of shapes which may be approximated by two ellipses (an inner and an outer one), using a variational level-set framework. For brevity sake, we will refer in the sequel to the shape comprised between these two elliptical contours as annular. Indeed, many natural scenes involve such pattern, as shown in Fig. 1. While the use of one ellipse has been often exploited in the field of image processing as a shape constraint (Pang et al., 2005; Guerrero et al., 2007), the expression of an annular constraint into a level-set framework has not been investigated yet. Using this shape information may be a powerful tool to recognize these patterns, particularly in the case of difficult images with occlusions or low SNR. Starting from the framework developed in (Chen et al., 2002), the approach presented in this paper brings the following contributions:

[1] A preliminary version of this work appeared in (Alessandrini et al., 2010). The present paper describes in details the formal aspects and the behavior of the proposed method, evaluates the performances from a data set including 115 medical images and compares it to the shape prior-based algorithm of Chen et al. (2002).

* Corresponding author. Fax: +33 4 72 43 85 26.
E-mail address: olivier.bernard@creatis.insa-lyon.fr (O. Bernard).

0167-8655/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2011.03.018

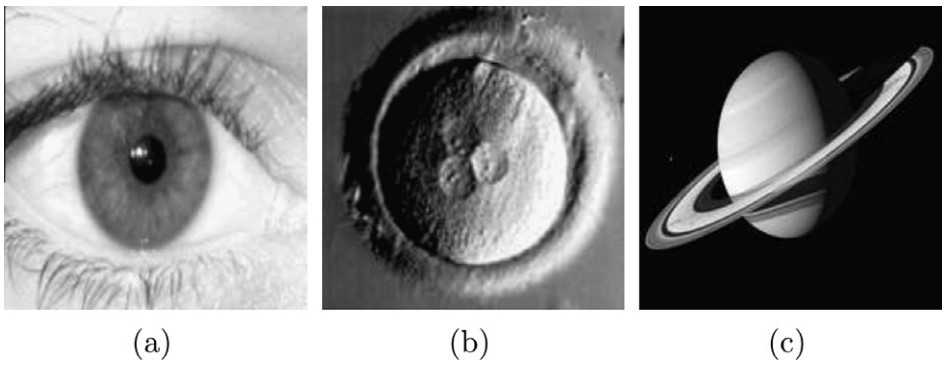

Fig. 1. Natural images in which annular patterns occur: (a) iris and pupil recognition (Jang et al., 2008), (b) segmentation of zona pellucida (Karlsson et al., 2004), (c) detection of the rings of Saturn.

- a parametric implicit function defining the distance from two ellipses is introduced;
- an efficient solution to the least-squares fitting problem of annular shapes is developed;
- thanks to its parametric representation, the shape prior may be adjusted to the evolving contour through an efficient least square strategy. Correspondingly, the computational cost associated to the introduction of the prior shape information is made essentially negligible.

This framework provides two additional benefits. First, since a parametric representation is adopted, we avoid a training phase for model selection, for which the availability of training samples and their number are difficult issues. Secondly, the detection of both contours is addressed with a single level-set function, which is desirable in terms of memory consumption and computation. The resulting algorithm is derived on the base of the traditional variational level-set paradigm (Chen et al., 2002; Chan and Vese, 2001). While the proposed method is applied to elliptic annular shape in this paper, it should be noted that the underlying framework may be applied to build a prior from any shape that can be described as a pair of similar implicit parametric contours and for which an efficient least squares fitting can be derived.

The paper is organized as follows. In Section 2 we recall the general form of the level-set framework, describe our annular shape prior and derive the minimization of the corresponding energy functional. In Section 3, implementation issues are addressed. We explain in particular how the level-set is initialized for each experiment. In Section 4 we give insights on possible applications by presenting results on natural images from several fields. We then provide an evaluation of the method in the field of cardiac imaging from MRI and ultrasound (US) data. A comparison with the shape prior algorithm of Chen et al. (2002) is presented as well. The main conclusions of this work are given in Section 5.

2. Proposed method

2.1. Level-set framework

Let Ω be a bounded open subset of $\mathbb{R}^d$ and let $I : \Omega \to \mathbb{R}^m$ be a given d-dimensional image. In the level-set formalism, the evolving interface $\Gamma \subset \mathbb{R}^d$ is represented as the zero level-set of a Lipschitz-continuous function ϕ of dimension $d+1$ that satisfies

$$\begin{cases} \phi(\boldsymbol{x}) < 0, & \forall \boldsymbol{x} \in \Omega_{\text{in}}, \\ \phi(\boldsymbol{x}) > 0, & \forall \boldsymbol{x} \in \Omega_{\text{out}}, \\ \phi(\boldsymbol{x}) = 0, & \forall \boldsymbol{x} \in \Gamma, \end{cases} \tag{1}$$

where Ω_{in} is a region in Ω bounded by $\Gamma = \partial\Omega_{\text{in}}$. The region Ω_{out} is defined as $\Omega \backslash \Omega_{\text{in}}$.

The problem of segmenting one object from the background is then handled by the evolution of one level-set driven by the minimization of a specific energy criterion; its steady state partitions the image into two regions that delimit the boundaries of the object to be segmented. Since we are addressing the segmentation of annular shapes, which are bounded by two unconnected contours, we propose to adopt a level-set function ϕ of the kind depicted in Fig. 2.

2.2. Energy formulation

As mentioned in the introduction, we set our approach in a variational level-set framework. Many energy functionals have been proposed in literature to incorporate shape priors into level-set formulation. As in (Leventon et al., 2000; Chan and Zhu, 2005; Cremers et al., 2007) we adopt the following general expression for the energy functional:

$$E = E_{\text{data}} + \alpha \cdot E_{\text{shape}}, \tag{2}$$

where E_{data} represents the chosen data attachment term and E_{shape} embeds the shape prior. The weight α corresponds to a positive hyper-parameter that balances the influence between the two terms. The expression for the two energy terms adopted for the segmentation task is derived in the following of this section, along with the variational framework to address the optimization problem.

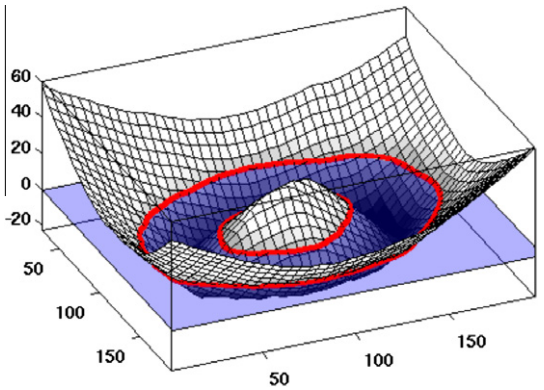

Fig. 2. Example of level-set function used in this work. The zero level is represented as a red bold line and provides the 2 contours corresponding to the boundary of the annular target. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2.3. Data attachment term

At this stage, E_{data} could be any of the data attachment terms described in literature. In this work we adopt as data attachment term the localized version of the Chan-Vese model, recently proposed in (Lankton and Tannenbaum, 2008). This formulation handles local statistics and is therefore well suited for segmenting objects whose boundaries are discontinuous or heterogeneous, where standard region based methods that use global statistics fail. Using their general framework, the minimization of E_{data} with respect to ϕ leads to the following level-set equation:

$$\frac{\partial \phi}{\partial \tau}(\boldsymbol{x}) = f(\boldsymbol{x}) \cdot \delta(\phi(\boldsymbol{x})), \tag{3}$$

where $\delta(\cdot)$ is the Dirac delta function and:

$$f(\boldsymbol{x}) = \int_{\Omega} B(\boldsymbol{x},\boldsymbol{y})\delta(\phi(\boldsymbol{y})) \cdot ((I(\boldsymbol{y}) - u_{\boldsymbol{x}})^2 - (I(\boldsymbol{y}) - v_{\boldsymbol{x}})^2) d\boldsymbol{y}. \tag{4}$$

The function $B(\cdot)$ is a binary mask defined as:

$$B(\boldsymbol{x},\boldsymbol{y}) = \begin{cases} 1, & \text{if } \boldsymbol{y} \in N(\boldsymbol{x}), \\ 0, & \text{otherwise}, \end{cases} \tag{5}$$

where $\boldsymbol{y}$ is a spatial variable and $N(\boldsymbol{x})$ corresponds to a user-defined neighborhood at point $\boldsymbol{x}$. The quantities $u_{\boldsymbol{x}}$ and $v_{\boldsymbol{x}}$ correspond to the localized version of the inside and outside average intensity values measured in the window $N(\boldsymbol{x})$. This data attachment term is independent of the shape prior to be described in the next section, so the user may in practice adopt any criterion which suits the image features of his particular applications.

2.4. Main contribution: annular shape prior term

Inspired by the framework proposed in (Chen et al., 2002), we introduce an annular shape constraint into our level-set framework by minimizing the following energy criterion:

$$E_{\text{shape}}(\phi, \boldsymbol{\lambda}) = \int_{\Omega} \Psi^2(\boldsymbol{x}, \boldsymbol{\lambda}) \|\nabla \phi(\boldsymbol{x})\| \delta(\phi(\boldsymbol{x})) d\boldsymbol{x}, \tag{6}$$

where $\Psi(\boldsymbol{x},\boldsymbol{\lambda})$ is the implicit function representing the distance of a point $\boldsymbol{x}$ from the annular shape defined by the parameters $\boldsymbol{\lambda}$. Clearly (6) reads as a measure of the distance between the active contour and the shape prior, and therefore imposes a similarity between the segmentation result and the prior itself. We propose here to adopt the following parametric expression for Ψ:

$$\Psi(\boldsymbol{x}, \boldsymbol{\lambda}) = \max\{\mathcal{E}(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{out}}), -\mathcal{E}(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{in}})\}, \tag{7}$$

where $\boldsymbol{\lambda} = [\boldsymbol{\lambda}_{\text{in}}, \boldsymbol{\lambda}_{\text{out}}]$ and $\boldsymbol{\lambda}_{\text{in}}$ and $\boldsymbol{\lambda}_{\text{out}}$ represent the parameters of the inner and outer ellipses. The function $\mathcal{E}(\cdot)$ corresponds to the algebraic distance of a point $\boldsymbol{x} = (x, y)$ to an ellipse, represented by the standard quadratic equation for conic sections, and defined as:

$$\mathcal{E}(\boldsymbol{x}, \boldsymbol{\lambda}_i) = \lambda_{i1}x^2 + \lambda_{i2}xy + \lambda_{i3}y^2 + \lambda_{i4}x + \lambda_{i5}y + \lambda_{i6}, \text{ with } \lambda_{i2}^2 < 4\lambda_{i1}\lambda_{i3}. \tag{8}$$

Fig. 3 gives an illustration of the parametric implicit function Ψ adopted in the proposed framework. We note here that another, more trivial, choice for Ψ could be adopted, consisting in the pointwise product between the two ellipse distance functions. However, that choice would no longer represent a distance function (Berg, 1998) and, further, it would not allow an optimized solution to the least-squares fitting problem, as the one we are presenting in the sequel. It is also to be noted that (8) corresponds to a representation of the ellipse through the algebraic distance. While the Euclidean distance could also be used, the algebraic distance has the advantages of yielding an analytical evaluation of (7) and allowing the use of the fast elliptic fitting algorithm of Fitzgibbon et al. (1999), as described in the next section. On the opposite, the Euclidean distance implies numerically evaluation of (7) and thus heavier iterative techniques (Faber and Fisher, 2001; Ahn et al., 2001).

The minimization of energy (6) leads to finding a geodesic zero level-set (in our case two contours) in a Riemannian space derived from the shape prior content. As compared to the initial work of Caselles et al. (1997) and derived approaches (Chen et al., 2002), the arc length function of (6) is no longer weighted by an image based information but only by our shape prior term. From the observation that the minimum of this expression is reached when the zero level of ϕ perfectly fits the zero level of the parametric implicit function Ψ, one can anticipate that the minimization of this energy criterion will make the level-set evolve toward an annular shape. The numerical minimization of (6) is addressed using a two phase scheme (Chan and Vese, 2001; Chen et al., 2002). Specifically, keeping $\boldsymbol{\lambda}$ fixed, the minimization of E_{shape} with respect to ϕ leads to the following equation:

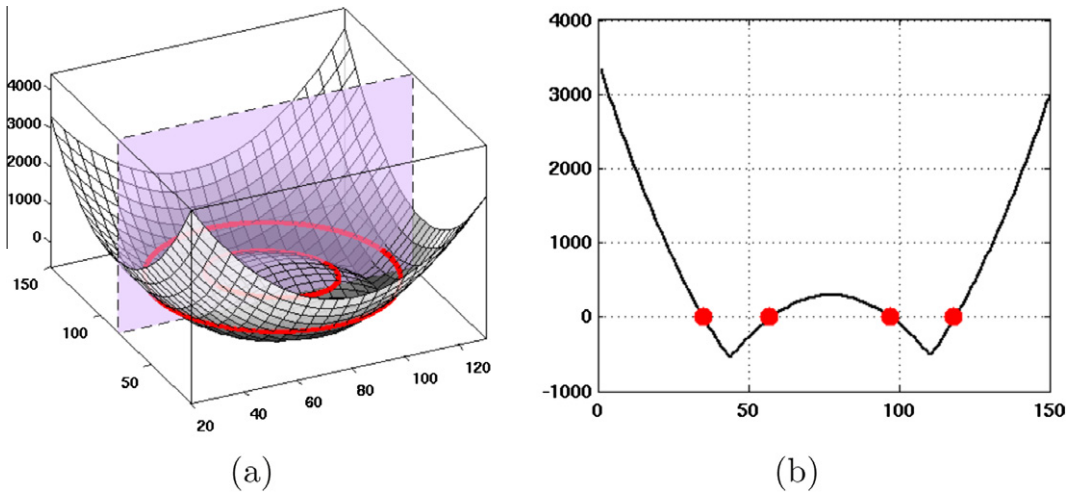

Fig. 3. Illustration of the parametric implicit function Ψ. The zero level is represented in red. (a) 3D representation. (b) Visualization of a 2D slice corresponding to the color plan given in (a). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$\frac{\partial \phi}{\partial \tau}(\boldsymbol{x}) = g(\boldsymbol{x}, \boldsymbol{\lambda}) \cdot \delta(\phi(\boldsymbol{x})), \tag{9}$$

where

$$g(\boldsymbol{x}, \boldsymbol{\lambda}) = 2\Psi(\boldsymbol{x}, \boldsymbol{\lambda}) \frac{\nabla \Psi(\boldsymbol{x}, \boldsymbol{\lambda}) \cdot \nabla \phi(\boldsymbol{x})}{\|\nabla \phi(\boldsymbol{x})\|} + \Psi^2(\boldsymbol{x}, \boldsymbol{\lambda}) \mathrm{div}\left(\frac{\nabla \phi(\boldsymbol{x})}{\|\nabla \phi(\boldsymbol{x})\|}\right). \tag{10}$$

Then, keeping ϕ fixed, the minimization of E_{shape} with respect to $\boldsymbol{\lambda}$ is obtained as follows. By noting that $\|\nabla \phi\| = 1$ because of the signed distance property (Osher and Fedkiw, 2002) and that $\delta(\phi)$ selects the points that belong to the contour Γ, the optimal solution can be written as:

$$\hat{\boldsymbol{\lambda}} = \mathrm{argmin}_{\boldsymbol{\lambda}} \oint_{\Gamma} \Psi^2(s, \boldsymbol{\lambda}) ds = \mathrm{argmin}_{\boldsymbol{\lambda}} \sum_{\boldsymbol{x}_i \in \Gamma} \Psi^2(\boldsymbol{x}_i, \boldsymbol{\lambda}). \tag{11}$$

The rightmost of (11) is justified by the fact that the image space is in practice discrete. The parameters vector $\boldsymbol{\lambda}$ is thus updated at each iteration according to the direct least-squares fitting problem defined by (11). To the best of author's knowledge, no method for the fitting of annular shapes has been yet proposed. We derive such a solution in the following.

2.4.1. Least-squares fitting of annular shapes

Considering (7), we can rewrite the sum in (11) as:

$$J(\boldsymbol{x}, \boldsymbol{\lambda}) = \sum_{\boldsymbol{x} \in \Gamma_{in}} \mathcal{E}^2(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{in}}) + \sum_{\boldsymbol{x} \in \Gamma_{out}} \mathcal{E}^2(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{out}}), \tag{12}$$

where the partition $\Gamma = \{\Gamma_{in}, \Gamma_{out}\}$ has been introduced:

$$\Gamma_{in}(\boldsymbol{\lambda}_{\text{in}}, \boldsymbol{\lambda}_{\text{out}}) = \{\boldsymbol{x} \in \Gamma | \mathcal{E}(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{out}}) < -\mathcal{E}(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{in}})\}, \tag{13}$$

$$\Gamma_{out}(\boldsymbol{\lambda}_{\text{in}}, \boldsymbol{\lambda}_{\text{out}}) = \{\boldsymbol{x} \in \Gamma | \mathcal{E}(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{out}}) \geqslant -\mathcal{E}(\boldsymbol{x}, \boldsymbol{\lambda}_{\text{in}})\}. \tag{14}$$

From this formulation, we observe that (12) can be minimized by fitting two separate ellipses on Γ_{in} and Γ_{out}, for which fast direct solvers exist (Fitzgibbon et al., 1999). Consequently we propose to minimize J by alternatively fitting the two ellipses and updating Γ_{in} and Γ_{out} according to (13) and (14). The resulting algorithm is summarized in Algorithm 1. By doing so, the energy J is ensured to decrease at each step. In Algorithm 1 we call fitLS the function performing the direct least-squares ellipse fitting described in (Fitzgibbon et al., 1999).

Algorithm 1. Annular least square fitting algorithm

Input data

$$\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \quad \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}, \quad tol = 0.01, \quad k = 1$$

Initialization

$$E^{(0)} = J\left(\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}\right)$$

$$\Gamma_{in}^{(0)} = \Gamma_{in}(\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}), \quad \Gamma_{out}^{(0)} = \Gamma_{out}\left(\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}\right)$$

while $\epsilon > tol$ **do**

$$\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)} = \mathrm{fitLS}(\Gamma_{in}^{(k-1)}); \quad \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)} = \mathrm{fitLS}\left(\Gamma_{out}^{(k-1)}\right)$$

$$\Gamma_{in}^{(k)} = \Gamma_{in}(\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)}); \quad \Gamma_{out}^{(k)} = \Gamma_{out}\left(\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)}\right)$$

$$E^{(k)} = J\left(\hat{\boldsymbol{\lambda}}_{\text{in}}^{(k)}, \hat{\boldsymbol{\lambda}}_{\text{out}}^{(k)}\right)$$

$$\epsilon = \|E^{(k)} - E^{(k-1)}\| / \|E^{(k-1)}\|; \quad k = k + 1$$

end while

In Fig. 4, two examples are reported. In both cases, the two initial ellipses defined by $\hat{\boldsymbol{\lambda}}_{\text{in}}^{(0)}$ and $\hat{\boldsymbol{\lambda}}_{\text{out}}^{(0)}$ are coincident. We adopt at first three different initializations from the same dataset (blue contours in Fig. 4(a)). In each case, the algorithm converges to the same results after only *3* iterations, as shown in Fig. 4(b). This shows the robustness of our method with respect to the initialization. We then tested our fitting algorithm on a more challenging case, where some of the points that describe both the inner and outer contours are superimposed (Fig. 4(c)). Fig. 4(d) shows the good behavior of our algorithm in such situation. In this case, convergence has been reached after only 4 iterations.

The computational complexity of Algorithm 1 is dominated by the few single ellipse least-squares fitting operations involved, indicated by the function fitLS. As shown in (Fitzgibbon et al., 1999), such a problem is equivalent to an eigenvalue problem for a 6×6 matrix, which can be efficiently solved with $\mathcal{O}(n^3)$ floating point operations (Golub and Van Loan, 1996). As an example, the solution of Fig. 4 is obtained in 13 ms in a MATLAB (R2010b, The Math Works) implementation of the algorithm, executed on a 2.27 GHz Intel Core i5 laptop equipped with 4 GB of RAM and running Windows 7 64-bit.

As detailed in the sequel, the segmentation method consists in a two stage procedure, i.e. alternating the fit of the annular shape to the level-set described in Algorithm 1 and the constrained evolution of the level-set through (9). It is then to be noted that the fitting stage can be made very fast by initializing Algorithm 1 with the fitting result computed in the previous iteration. In practice, the experiments showed that this strategy allows performing the fitting stage in one single iteration.

As a further remark, it should be noted that, although the framework has been specialized for annular shapes, the proposed formalism is general. In particular it can be extended to geometric primitives for which an efficient solution to the least squares fitting problem can be individuated, to be used in place of FitLS in Algorithm 1.

We conclude this section by pointing out that the penalty term in (6) is not new in literature: Chen et al. introduced an analogous one (Chen et al., 2002), in order to embed a shape constraint in the geodesic active contour flow (Caselles et al., 1997). Nevertheless, the difference between the presented framework and the one in (Chen et al., 2002) is substantial: in (Chen et al., 2002) Ψ represents the distance, computed numerically, of a point to an arbitrary shape, and consequently a similarity transformation between the segmenting contour and the prior shape has to be searched in a steepest descent scheme. In our work, we develop an analytic representation for Ψ, expressed by (7), allowing thereby to express the parameters update step as a direct least squares fitting. As a consequence, this stage of the algorithm avoids slow steepest descent iterative procedures. These aspects will be further stressed in the results section, where a comparison of the proposed algorithm with the one by Chen is presented, both in terms of segmentation accuracy and execution time.

3. Implementation issues

We implemented our level-set evolution equation using standard finite difference scheme (Osher and Fedkiw, 2002), where the implicit function is represented by a signed distance function ϕ. In order to improve efficiency, we only compute values of ϕ in a narrow band around the zero level set. Consequently, we re-initialize ϕ every iteration using a fast marching scheme (Osher and Fedkiw, 2002). In order to perform segmentation with an annular shape prior, the level-set evolves according to the following equation:

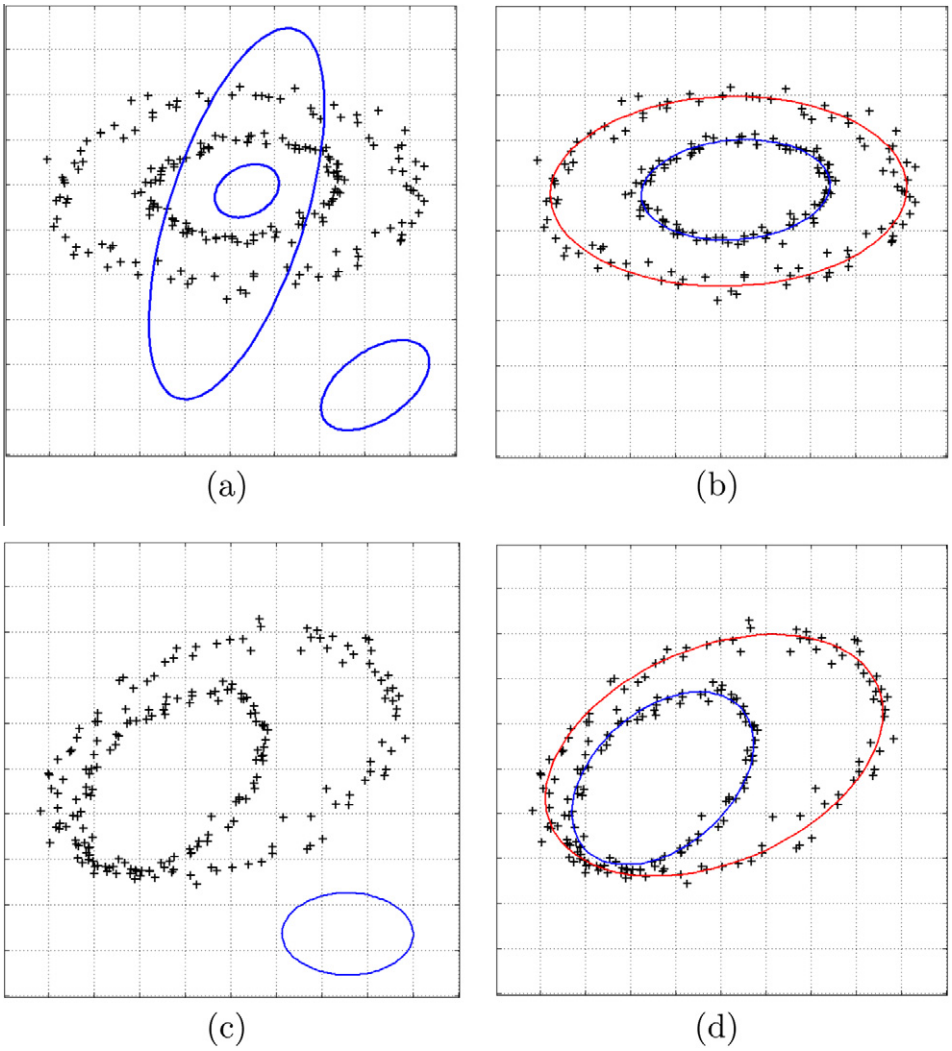

Fig. 4. Illustration of the use of annular least square fitting. In both cases, the two initial ellipses are coincident. First row: illustration on a simple case. (a) Initialization: the three ellipses corresponds to three different initializations. (b) Results obtained at convergence: the final annular contours are identical. Second row: illustration on a more challenging case. (c) Initialization. (d) Results obtained at convergence after only 4 iterations.

$$\frac{\partial \phi}{\partial \tau}(\boldsymbol{x}) = \delta(\phi(\boldsymbol{x}))[f(\boldsymbol{x}) + \alpha \cdot g(\boldsymbol{x}, \boldsymbol{\lambda})], \tag{15}$$

where $f(\cdot)$ is the data attachment term given in (4) and $g(\cdot)$ is the shape prior term given in (10). The choice of the weight term α is dependent on the specific application. For instance, when the quality of the image is poor, it is recommended to use higher value of α in order to put more importance on the shape prior term. When presenting results in the experiments section, we will adopt the α value which produces the best segmentation result. The experimental α values obtained are in the range $[0.5; 1.5]$.

Concerning the initialization, let us note that the localized framework of Lankton and Tannenbaum (2008) requires that it must be made close to the desired solution. Indeed while such a local data attachment term allows making the algorithm more robust to the variation of image properties, it makes it more sensitive to initialization. This is thus an inherent trade-off of this type of data attachment term. For applications where a fully automatic procedure is needed, note however that this limitation maybe tackled by using a preprocessing step providing a rough detection of the object (see e.g. Dias and Leitao, 1996 or Paragios et al., 2005). Such a strategy is however inherently application-dependent and is thus out of the scope of this paper. In the experiments described in the next sections, the following simple procedure is thus considered. The user is asked to position six points, as this is the minimum number of points needed to define 2 concentric ellipses. The five first points are used to set an ellipse (one ellipse being uniquely defined by 5 points) and the last point is used to obtain the second concentric ellipse (see Fig. 5). For each experiment, the initial points used to achieve the given results are displayed in green.

4. Results

4.1. Application to natural images from several fields

Annular-like patterns are recurrent in images from several fields: medicine, biology and astronomy are some examples. In all these cases the proposed methodology may be exploited in order to detect structures of interest. Fig. 5 illustrates the application of the proposed method for the segmentation of the pupil and iris in human eye, of the zona pellucida in human embryo, and the rings of Saturn. We used for these experiments the same α value equal to 0.8. We observe that in all those situations gray scale information alone is not sufficient for obtaining correct results. In the eye image (Fig. 5(a)), the eyelash hinders a correct iris detection, moreover the variations in the iris color make pupil

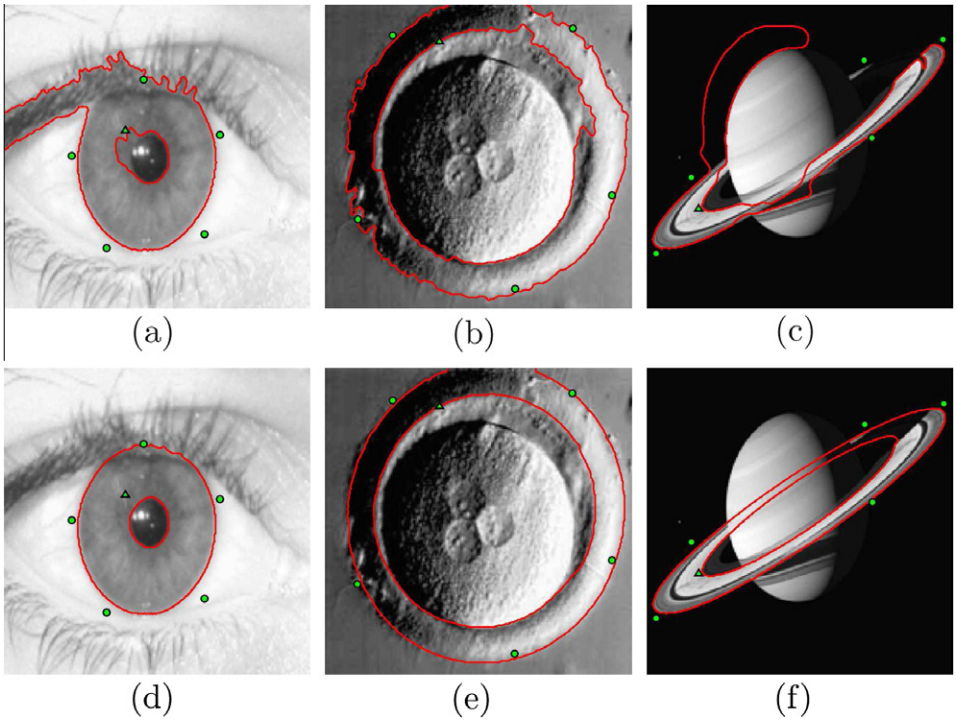

Fig. 5. Segmentation results obtained on natural scene images. The green points correspond to the initialization. Images (a)–(c) show the result obtained when the proposed approach is applied without the shape constraint. Images (d)–(f) show the result obtained when the shape prior is applied. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

segmentation imprecise. A similar effect can be observed in the embryo image (Fig. 5(b)) where the segmentation is inaccurate due to low contrast profiles and to proximity of the cell to the internal boundary of the zona pellucida. In the Saturn image, the ring portion hidden by the planet body cannot be followed without a shape prior approach (Fig. 5(c)).

The results obtained using the shape prior term are given in the second row. In the 3 cases, they are in good agreement with the expected shapes. They illustrate the usefulness of the proposed annular shape prior for such images. Regarding the computational cost, Table 1 provides the cpu times corresponding to the segmentations performed with our approach and shown in Fig. 5(d) to Fig. 5(f). Note that the 3 images share the same size of 200×200 pixels. In order to give an insight on the impact of the shape prior term on the overall computational burden, Table 1 also reports the cpu time corresponding to the fitting step of the algorithm, as performed through Algorithm 1. These results confirm that the improvement in segmentation accuracy ascribable to the shape prior term comes at a substantially negligible price in terms of computational complexity.

4.2. Evaluation on cardiac images

A remarkable case in which annular-like patterns are present is cardiology imaging. Indeed both endo- and epicardial borders of the left ventricle (the innermost and the outer layers of tissue surrounding the cardiac muscle), when visualized in a short axis view can be well approximated by ellipses. The joint segmentation of endo- and epicardium has a high clinical relevance, since it allows some critical indices of heart functionality to be evaluated. Nevertheless, especially in the case of ultrasound, due to its intrinsic complexity, this issue has been addressed in a very limited number of works (Noble and Boukerroui, 2006).

In order to evaluate the performance of our algorithm in accomplishing this task, a set of 56 cardiac MRI images and 59 ultrasound images were considered. The size of MRI images was 141×141 pixels, and the size of the ultrasound images was 249×168 pixels. Pixel size was equal to 1×1 mm^2 for MRI and 0.27×0.27 mm^2 for US. Segmentation result was assessed by measuring the correspondence between the detected contours and the manual reference given by an expert cardiologist. The algorithm was initialized as described in Section 3 by a non expert but experienced user. Three error metrics were adopted, i.e. the Dice coefficient D, the Hausdorff distance HD and the Root Mean Square Distance $RMSD$, all expressed in pixels and defined in the next section.

Table 1
Cpu times corresponding to the segmentation with the proposed approach on natural images.

	Eye image (Fig. 5(d))	Embryo image (Fig. 5(e))	Saturn image (Fig. 5(f))
Overall cpu time(s)	8.31	17.93	13.86
Cpu time for the shape Prior fitting(s)	0.12	0.36	0.33

4.2.1. Definition of the error metrics

Let $\mathcal{R}$ be the reference contour and $\mathcal{S}$ the detected one. Denoting as $\boldsymbol{r}$ and $\boldsymbol{s}$ a point belonging to $\mathcal{R}$ and $\mathcal{S}$ respectively, the Hausdorff distance and Root Mean Square Distance are then defined as:

$$HD(\mathcal{S},\mathcal{R}) = \max\{\sup_{\boldsymbol{s}} d(\boldsymbol{s},\mathcal{R}); \sup_{\boldsymbol{r}} d(\boldsymbol{r},\mathcal{S})\} \tag{16}$$

and

$$RMSD(\mathcal{S},\mathcal{R}) = \frac{1}{2}\left\{\sqrt{\frac{1}{L(\mathcal{S})}\sum_{\mathcal{S}} d^2(\boldsymbol{s},\mathcal{R})} + \sqrt{\frac{1}{L(\mathcal{R})}\sum_{\mathcal{R}} d^2(\boldsymbol{r},\mathcal{S})}\right\}, \tag{17}$$

where $d(\boldsymbol{a},\mathcal{A})$ represents the distance from the point $\boldsymbol{a}$ to the contour $\mathcal{A}$ and $L(\mathcal{A})$ is the contour length. Finally, defining Ω_S and Ω_R

as the sets of pixels contained inside the segmented and the reference region, the Dice coefficient writes as:

$$D(\Omega_S, \Omega_R) = \frac{2Area(\Omega_S \bigcap \Omega_R)}{Area(\Omega_S) + Area(\Omega_R)}. \quad (18)$$

This coefficient thus measures the correspondence between the two regions, and varies from 0) to 1: it is 1 when the two regions are coincident and 0 when they have null intersection.

Table 2
Performance metrics value (MEAN ± STANDARD DEVIATION) for the proposed algorithm on 56 cardiac MRI images (a) and 59 US images (b). The percentage between brackets shows the relative difference between the mean values of the corresponding metric.

	DICE	HD (PIXELS)	RMSD (PIXELS)
(a) MRI Images			
WITH SHAPE PRIOR (α = 0.8)	**0.82 ± 0.05**	**3.63 ± 0.73**	**1.46 ± 0.21**
WITHOUT SHAPE PRIOR	0.80 ± 0.05	4.62 ± 0.82 (+27%)	1.72 ± 0.27 (+18%)
(b) US Images			
WITH SHAPE PRIOR (α = 1.2)	**0.89 ± 0.02**	**5.57 ± 1.18**	**2.22 ± 0.46**
WITHOUT SHAPE PRIOR	0.87 ± 0.03	7.78 ± 2.12 (+39%)	2.55 ± 0.49 (+15%)

The RMSD measures the average distance between the two contours, and is thus suited for evaluating the global performance of the algorithm. On the opposite, the Hausdorff metric represents the maximum distance between the contours, and is therefore a local index of performance (Huttenlocher et al., 1993; Chalana and Kim, 1997). More specifically, the HD keeps trace of the localized discrepancies which are largely hidden in the average operation involved in the RMSD computation. While RMSD and HD depend on the image size, the Dice coefficient returns instead an intrinsic index of performance, since it measures the superposition in percentage between the two enclosed areas. These three metrics are adopted to provide an exhaustive evaluation of the behavior of the proposed algorithm.

4.2.2. Evaluation of shape prior contribution

In order to evaluate the effect of the shape prior information we ran our algorithm with and without using the shape prior term. When the shape information was taken into account a value α = 0.8 was used for MRI images and α = 1.2 for ultrasound images. This difference is due to the fact that ultrasound images quality is intrinsically low and high values of α must be used in order to deal

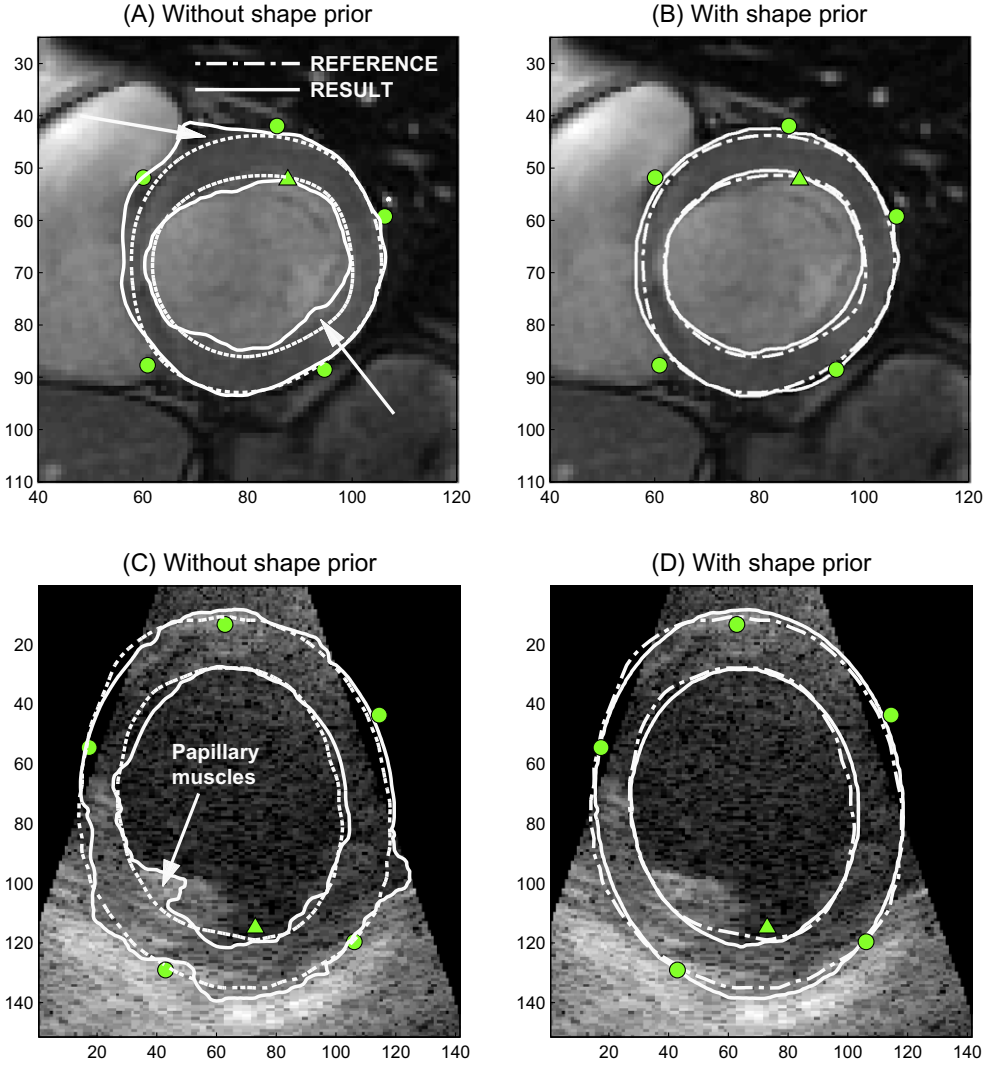

Fig. 6. Example results on cardiac MRI ((a)–(b)) and US ((c)–(d)). In the MRI case the shape constraint prevents the contour from adhering to the right ventricle (upper arrow in (a)) and allows to cope with the low contrast of the endocardium (lower arrow in (a)). In the US case, it avoids irregularities in the endocardial contour (white arrow in (c)), principally ascribable to the papillary muscles. On all images, much regular results are obtained with the shape prior approach, with an improved correspondence with the cardiologist reference. Error metrics for those images are: (a) HD = 5 pixels, RMSD = 1.52 pixels, D = 0.83; (b) HD = 2.23 pixels, RMSD = 0.99 pixel, D = 0.89; (c) HD = 10.7 pixels, RMSD = 2.72 pixels, D = 0.87; (d) HD = 3.61 pixels, RMSD = 1.62 pixels, D = 0.91.

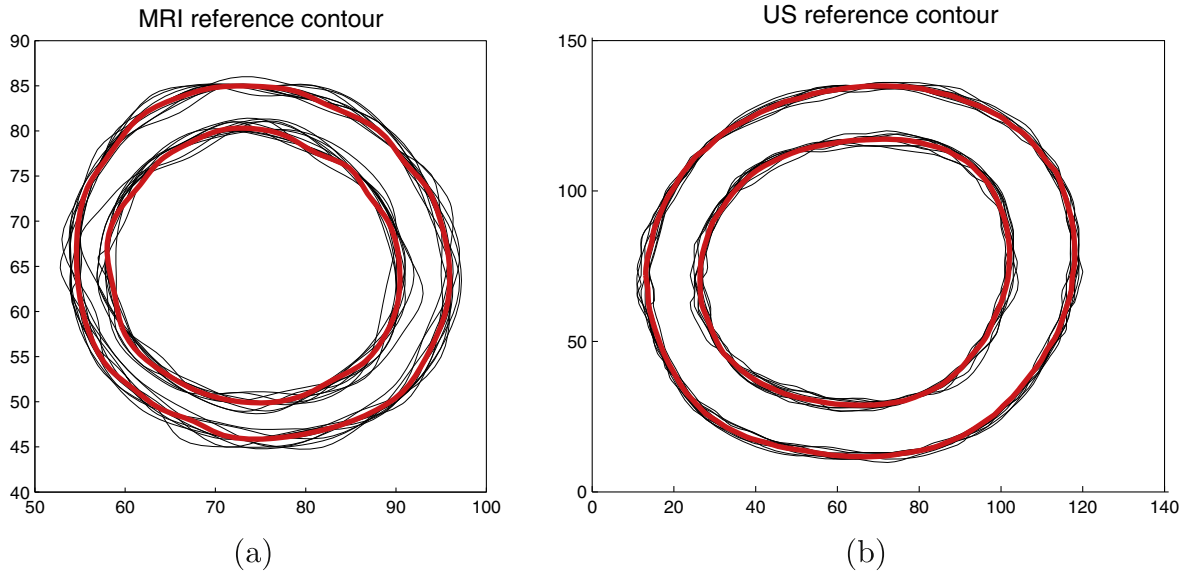

Fig. 7. Prior contour computation for (a) MRI and (b) US images. The cluster of reference curves is represented as thin contours and the corresponding average shape as a bold red contour. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

with this feature. The complete set of results for MRI and ultrasound images is reported in Table 2.[2] These results clearly show that the introduction of the shape prior yields an improvement of the segmentation, yielding a slightly higher Dice and smaller HD and RMSD.

Fig. 6 provides segmentation examples allowing a finer interpretation of these results. It may be observed that the shape prior allows to avoid the local irregularities, linked to the attraction of the active contours to small scale noisy structures. This explains in particular the large decrease of HD associated with the shape information.

4.2.3. Comparison with Chen algorithm

We compare in this section the performance of our algorithm with the approach described by Chen et al. (2002). The basic equations for Chen algorithm are reported in Appendix A. The use of Chen's approach implies the learning of the shape prior through a training phase performed on a set of manually segmented contours. In this experiment, the prior contour was obtained for each modality from a training image set, built by randomly selecting reference contours. Specifically, 10 images out of 56 were used for the MRI case and 10 out of 59 for the US case. For each of these two training sets, the prior contour was obtained by averaging the reference contours drawn by the physician (see Chalana and Kim (1997)). Due to potential misalignment among the references, a registration step must precede the averaging operation, as detailed in Chen et al. (2002). The so obtained prior contours are illustrated in Fig. 7. From this figure it can also be clearly seen that ellipses represent a reliable approximation of myocardial contours.

The complete set of results for MRI and US images is reported in Table 3. These results show that the proposed algorithm outperforms Chen's approach for the present application. Let us note that the difference in performance is particularly pronounced for the Hausdorff distance. This behavior is linked to the fact that in Chen's approach, the only non-rigid degree of freedom applied to the shape prior in the course of the segmentation is a uniform scaling. On the opposite, our approach performs a complete fit of the prior at each iteration, thus yielding a better adjustment to the myocardial shape variations. This aspect is illustrated in Fig. 8.

[2] Note that *HD* and *RMSD* have intrinsically larger values when computed on ultrasound images due to the bigger size of the latter.

Table 3
Performance metrics value (MEAN ± STANDARD DEVIATION) for the proposed algorithm and Chen's one on MRI (a) and US images (b). Training images are eliminated from the data set before evaluation. The total number of testing images is 46 for MRI and 49 for US. The percentage between brackets shows the relative difference between the mean values of the corresponding metric.

	DICE	HD (PIXELS)	RMSD (PIXELS)
(a) MRI Images			
PROPOSED ALGORITHM	**0.82 ± 0.05**	**3.59 ± 0.74**	**1.45 ± 0.26**
CHEN'S ALGORITHM	0.76 ± 0.18	4.89 ± 1.06 (+36%)	1.74 ± 0.45 (+20%)
(b) US Images			
WITH SHAPE PRIOR	**0.89 ± 0.02**	**5.63 ± 1.23**	**2.23 ± 0.47**
CHEN'S ALGORITHM	0.86 ± 0.03	7.05 ± 1.55 (+25%)	2.50 ± 0.72 (+12%)

Another important issue concerns the computational speed. The average cpu time required for segmenting one MRI image was 2.1 s for the proposed algorithm versus 11.1 s for Chen's one. When considering ultrasound images, these values were 13.5 s for our method and 36.7 s for Chen's one. This speed-up is a consequence of the availability of a parametric representation for the prior, combined with the existence of a fast fitting procedure for such a geometry, as detailed in Section 2.4.1. In particular, the most costly step in Chen's algorithm is the initial alignment of the prior shape onto the initial contour (cf. Appendix A). For the considered images, this step can take up to several thousands of iterations. On the opposite, in our approach the initial registration corresponds to fitting the initial contour with an annular shape, which is typically done in few iterations as detailed in Section 2.4.1.

4.3. Dependence on the initialization

As previously mentioned, the adoption of the localized framework of Lankton and Tannenbaum (2008) imposes a good initialization in order to guarantee adequate results. In practice, this implies that the distance between the initial contour and the object to be detected should be on the order of the extent of the mask $B(\cdot)$ in (5). In this section, we thus propose to study the influence of

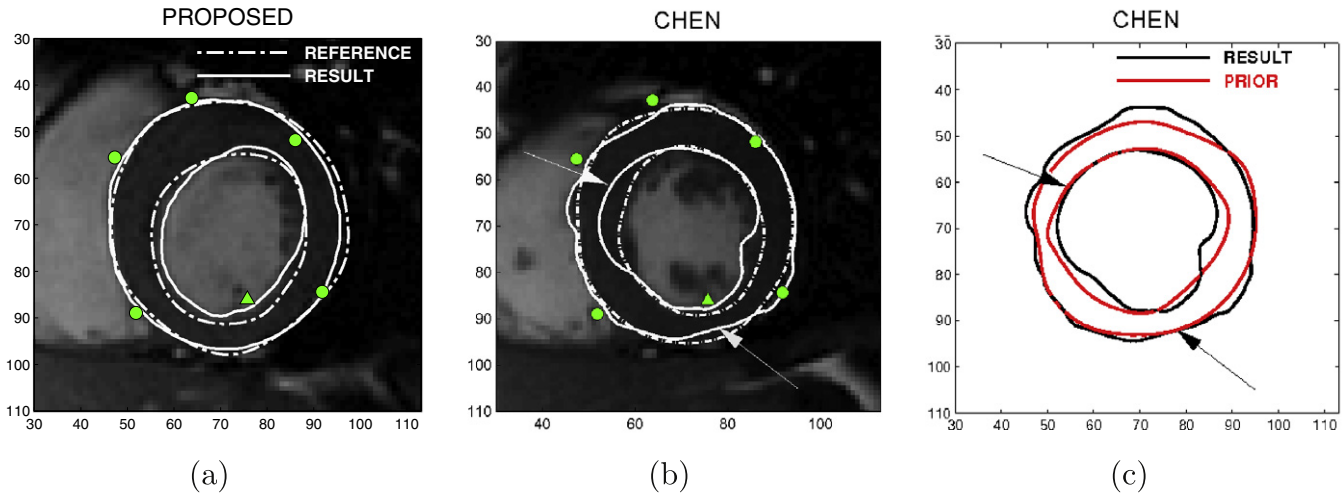

Fig. 8. Comparison between the proposed (a) and Chen (b) algorithm applied to an MRI image. The prior shape for Chen algorithm is not appropriate for this particular frame and compensates for the lack of data information in a way which is not consistent with the real image content, as shown by the arrows in (b) and (c). Differently, the adopted parametric formulation easily allows to second myocardial shape. Error measures are: (a) D = 0.87, HD = 3.60 pixels and RMSD = 1.42 pixels; (b) D = 0.76, HD = 5.88 pixels and RMSD = 2.09 pixels.

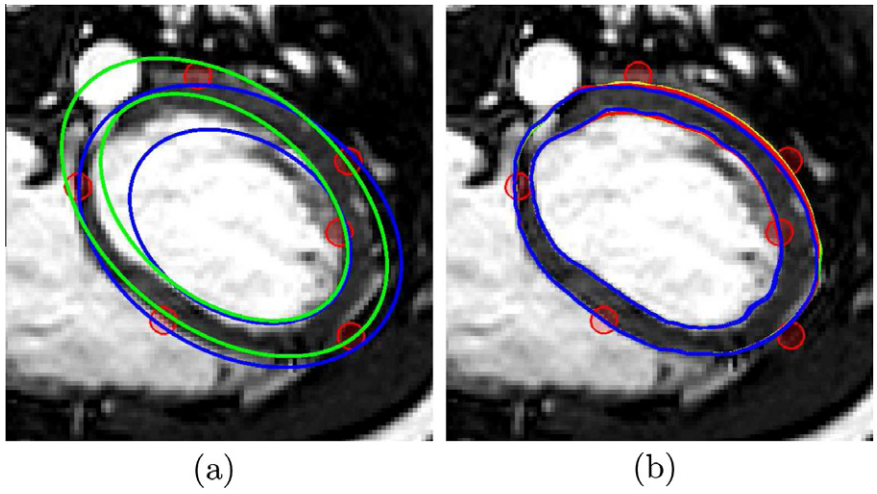

Fig. 9. Evaluation of the influence of the initialization points. (a) Regions where the initial points are randomly positioned with two particular derived contours. (b) Segmentation results obtained from 10 different initializations.

initialization by applying a random variation to the location of the 6 points used in our initialization procedure.

This experiment is illustrated in Fig. 9 and is done as follows. We delineate from a specific MRI image six regions in which initial points are randomly positioned (red disk in Fig. 9(a)). Each region corresponds to a disk with a radius of 3 pixels. This dimension has been chosen to simulate the uncertainty in the positioning of the points by the user. In order to show the influence of such random scheme in the initialization, we display in Fig. 9(a) two particular initial contours obtained from the proposed random initializations and Fig. 9(b) shows the final contours obtained from 10 different initializations. In order to evaluate the variability of the resulting segmentations, we compute the similarity measures by considering all the possible combinations of two final contours. We obtained mean values of D = 0.98, HD = 1.1 pixels and RMSD = 0.3 pixels. These values illustrate the robustness of the proposed method with respect to the initialization phase.

As a conclusion, although the local data attachment term implies an initialization close to the desired contour, within this limit, this initialization thus needs not to be very accurate.

5. Conclusion

In this paper we addressed the segmentation of shapes which may be approximated by two elliptical contours by introducing an annular shape prior. By defining the distance from two ellipses as a parametric implicit function, we have shown that an efficient solution to the least-squares fitting problem of annular shapes can be developed, avoiding iterative strategies like steepest gradient descent and yielding an efficient algorithm. The application of the method has been illustrated on several images from different fields and an evaluation on MRI and ultrasound cardiac data has been performed.

Appendix A. Chen algorithm

We briefly recall here the basics of Chen's algorithm. The interested reader is refered to the original paper (Chen et al., 2002) for a detailed description. The segmentation process is defined by:

$$\min_{\phi,\mu,R,T} \int_{\Omega} \delta(\phi(\boldsymbol{x}))\left\{e(|\nabla I|) + \frac{\lambda}{2}\Psi^2(\mu R(\theta)\boldsymbol{x} + T)\right\}|\nabla\phi(\boldsymbol{x})|d\boldsymbol{x}, \tag{A.1}$$

where Ψ is the distance from a reference shape, obtained from a training phase, $e(\cdot)$ is an edge indicator function, while μ, θ and T represent scale, rotation and translation for the optimal registration of the prior shape onto the active contour. R is the rotation matrix associated to the angle θ. Ψ can be computed with numerical methods as fast marching. The optimization is handled using the gradient descent method:

$$\frac{\partial\phi}{\partial\tau} = \delta(\phi)div\left\{\left(e(|\nabla I|) + \frac{\lambda}{2}\Psi^2\right)\frac{\nabla\phi}{|\nabla\phi|}\right\}, \tag{A.2}$$

$$\frac{\partial\mu}{\partial\tau} = -\lambda\int_{\Omega}\delta(\phi)\Psi\nabla\Psi(R\boldsymbol{x})|\nabla\phi|d\boldsymbol{x}, \tag{A.3}$$

$$\frac{\partial\theta}{\partial\tau} = -\lambda\int_{\Omega}\delta(\phi)\mu\Psi\nabla\Psi\left(\frac{dR}{d\theta}\boldsymbol{x}\right)|\nabla\phi|d\boldsymbol{x}, \tag{A.4}$$

$$\frac{\partial T}{\partial\tau} = -\lambda\int_{\Omega}\delta(\phi)\Psi\nabla\Psi|\nabla\phi|d\boldsymbol{x}, \tag{A.5}$$

where Ψ must be evaluated at $\mu R\boldsymbol{x} + T$.

Since at time zero the reference shape and the initial contour may be substantially misaligned, a registration step must precede the segmentation flow. This is implemented by iterating Eqs. (A.3) to (A.5) until Ψ^2 reaches a minimum bound.

We observe that image gradient, exploited for edges location, is not a reliable indicator in low SNR situations, as in ultrasound images. For this reason, and for comparison purpose, we use in all the experiments the localized framework described in Section 2.3 as data attachment term. This implies that the left hand side of (A.2) is in practice replaced by Eq. (4).

References

Ahn, S.J., Rauh, W., Warnecke, H.J., 2001. Least-squares orthogonal distances fitting of circle, sphere, ellipse, hyperbola, and parabola. Pattern Recognit. 34, 2283–2303.

Alessandrini, M., Dietenbeck, T., Basset, O., Friboulet, D., Bernard, O. 2010. Using a geometric formulation of annular-like shape priors for constraining variational level-sets. In: Proc. IEEE Internat. Conf. Image Process. (ICIP2010), Hong Kong, pp. 669–672.

Berg, J.M., 1998. On parameter estimation using level sets. SIAM J. Control Optim. 37, 1372–1393.

Bernard, O., Friboulet, D., Thevenaz, P., Unser, M., 2009. Variational *b*-spline level-set: A linear filtering approach for fast deformable model evolution. IEEE Trans. Image Imag. 18, 1179–1191.

Caselles, V., Kimmel, R., Sapiro, G., 1997. Geodesic active contours. Int. J. Comput. Vision 22, 61–79.

Chalana, V., Kim, Y., 1997. A methodology for evaluation of boundary detection algorithms on medical images. IEEE Trans. Med. Imag. 16, 642–652.

Chan, T., Vese, L., 2001. Active contours without edges. IEEE Trans. Image Imag. 10, 266–277.

Chan, T., Zhu, W. 2005. Level set based shape prior segmentation. In: IEEE Comput. Soc. Conf. Comput. Vision Pattern Recognit., CVPR 2005, vol. 2, pp. 1164–1170.

Chen, Y., Tagare, H.D., Thiruvenkadam, S., Huang, F., Wilson, D., Gopinath, K.S., Briggs, R.W., Geiser, E.A., 2002. Using prior shapes in geometric active contours in a variational framework. Int. J. Comput. Vision 50, 315–328.

Cremers, D., Rousson, M., Deriche, R., 2007. A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape. Int. J. Comput. Vision 72, 195–215.

Dias, J., Leitao, J., 1996. Wall position and thickness estimation from sequences of echocardiographic images. IEEE Trans. Med. Imag. 15, 25–38.

Faber, P., Fisher, R., 2001. Euclidean fitting revisited. In: Arcelli, C., Cordella, L., di Baja, G. (Eds.), Visual form 2001, Lect. Notes Comput. Sci., vol. 2059. Springer, Berlin/Heidelberg, pp. 165–175.

Fitzgibbon, A., Pilu, M., Fisher, R., 1999. Direct least square fitting of ellipses. IEEE Trans. Pattern Anal. Mach. Intell. 21, 476–480.

Golub, G.H., Van Loan, C.F., 1996. In: Matrix computations, 3rd ed. Johns Hopkins University Press, Baltimore, MD, USA.

Guerrero, J., Salcudean, S., McEwen, J., Masri, B., Nicolaou, S., 2007. Real-time vessel segmentation and tracking for ultrasound imaging applications. IEEE Trans. Med. Imag. 26, 1079–1090.

Huttenlocher, D., Klanderman, G., Rucklidge, W., 1993. Comparing images using the hausdorff distance. IEEE Trans. Pattern Anal. Mach. Intell. 15, 850–863.

Jang, Y.K., Kang, B.J., Park, K.R., 2008. A study on eyelid localization considering image focus for iris recognition. Pattern Recognit. Lett. 29, 1698–1704.

Karlsson, A., Overgaard, N.C., Heyden, A., 2004. Automatic segmentation of zona pellucida in HMC images of human embryos. In: Internat. Conf. Pattern Recognit., vol. 3, pp. 518–521.

Lankton, S., Tannenbaum, A., 2008. Localizing region-based active contours. IEEE Trans. Image Imag. 17, 2029–2039.

Leventon, M.E., Eric, W., Grimson, L., Faugeras, O. 2000. Statistical shape influence in geodesic active contours. In: IEEE Comput. Soc. Conf. Comput. Vision Pattern Recognit., CVPR 2000, p. 1316.

Noble, A., Boukerroui, D., 2006. Ultrasound image segmentation: A survey. IEEE Trans. Med. Imag. 25, 987–1010.

Osher, S.J., Fedkiw, R.P., 2002. Level Set Methods and Dynamic Implicit Surfaces, 1st ed. Springer.

Pang, B., Zhang, D., Wang, K., 2005. The bi-elliptical deformable contour and its application to automated tongue segmentation in chinese medicine. IEEE Trans. Med. Imag. 24, 946–956.

Paragios, N., Jolly, M.P., Taron, M., Ramaraj, R. 2005. Active shape models and segmentation of the left ventricle in echocardiography. In: Internat. Conf. Scale Space Theories PDEs Methods Comput. Vision, pp. 131–142.

Pardo, X.M., Leborán, V., Dosil, R., 2004. Integrating prior shape models into level-set approaches. Pattern Recognit. Lett. 25, 631–639.

Rousson, M., Paragios, N., 2002. Shape priors for level set representations. Proc. 7th Eur. Conf. Comput. Vision-Part II. Springer-Verlag, London, UK, pp. 78–92.