



Endocardial Boundary Estimation and Tracking in Echocardiographic Images using Deformable Templates and Markov Random Fields

Max Mignotte^{1,2}, Jean Meunier¹ and Jean-Claude Tardif³

¹DIRO, Département d'Informatique et de Recherche Opérationnelle, Montréal, Québec, Canada; ²INRIA, Institut National de Recherche en Informatique et Automatique, France; ³ICM, Institut de Cardiologie de Montréal, Montréal, Québec, Canada

Abstract: We present a new approach to shape-based segmentation and tracking of deformable anatomical structures in medical images, and validate this approach by detecting and tracking the endocardial contour in an echocardiographic image sequence. To this end, some global prior shape knowledge of the endocardial boundary is captured by a prototype template with a set of predefined global and local deformations to take into account its inherent natural variability over time. In this deformable model-based Bayesian segmentation, the data likelihood model relies on an accurate statistical modelling of the grey level distribution of each class present in the ultrasound image. The parameters of this distribution mixture are given by a preliminary iterative estimation step. This estimation scheme relies on a Markov Random Field prior model, and takes into account the imaging process as well as the distribution shape of each class present in the image. Then the detection and the tracking problem is stated in a Bayesian framework, where it ends up as a cost function minimisation problem for each image of the sequence. In our application, this energy optimisation problem is efficiently solved by a genetic algorithm combined with a steepest ascent procedure. This technique has been successfully applied on synthetic images, and on a real echocardiographic image sequence.

Keywords: Boundary estimation; Deformable templates; Echocardiography; Markov Random Fields; Tracking

1. INTRODUCTION

The segmentation of images is of great importance in medical imaging, and remains a necessary step to obtain qualitative measurements such as the visualisation or the location of objects of interest, the detection of pathological deformations, as well as for quantitative measurements such as area, volume or the analysis of dynamic behaviour of anatomical structures over time.

Among these images, ultrasound images play a crucial role, because they can be produced at video-rate, and therefore allow a dynamic analysis of moving structures. Moreover, the acquisition of these images is non-invasive, rather inexpensive, and does not require radiations compared to

other medical imaging techniques (for example, CT techniques). On the other hand, the automatic segmentation of anatomical structures in ultrasound images is a real challenge due to speckle noise and artifacts which are inherent in these images. In addition to the poor image quality (i.e. low signal-to-noise ratio, low resolution and contrast), observed objects, like the heart, are constantly undergoing motion and non-rigid deformation due to their own deformations, as well as the deformation of adjacent structures. Another common problem are artifacts caused by turbulent blood flow, air in the lungs, ribs, etc., which lead to a loss of signal and temporary occlusion of object boundaries.

Automatic segmentation techniques used in medical applications have been developed for CT, MRI images [1–3] or echocardiographic sequences [4–12] (see also [13] for a complete review and discussion of existing methods for the endocardial border identification in two-dimensional echocardiographic images). Among these segmentation schemes,

two of the more reliable approaches use Bayesian inference, which allows to take into account the available local or global *a priori* knowledge. This available knowledge is then used to constrain the segmentation process. These methods are based on deformable templates or Markov Random Fields (MRF) models. Both have advantages and shortcomings. The increased popularity of deformable models and templates [1–4] can be explained by their ability to express, *via* a global prior model, the geometric information and the shape variability of the object of interest to be detected over time and across individuals. The major difference between deformable models and templates lies in the degree of global *a priori* knowledge which is used for constraining the segmentation process. While the deformable model-based approach imposes only generic smoothness constraints on the extracted shape, the deformable template imposes more specific constraints *via* a predefined prior distribution. In MRF-based segmentation approaches, the proper use of available prior information is expressed by a local prior model. This prior model, described by a Gibbs distribution, expresses the fact that nearby pixels are fairly likely to belong to the same class. This contextual knowledge is captured through the specification of spatially local interactions (called clique potentials) that constraint the segmentation process [14].

The main problem with deformable models is that the likelihood energies are experimentally and heuristically designed, and do not statistically segment the image. As for MRF models, they cannot incorporate any global shape constraints in the segmentation process. Nevertheless, contrary to deformable model-based approaches, the likelihood model of an MRF model-based unsupervised segmentation exploits an accurate statistical modeling of the grey level distribution of each class present in the input image [15]. Since it seems that deformable templates provide an interesting framework for structurally analysing an image and MRF models provide a good mathematical setting for statistically analysing an image, it makes sense to capitalise on the benefits of both methods.

To this end, we propose in this paper a deformable template-based approach for the segmentation and tracking of anatomical deformable structures in medical images. We validate this approach by detecting and tracking the endocardial contour in an echographic image sequence. We first define a prototype template along with a predefined set of admissible transformations in order to capture some available global *a priori* shape knowledge of the endocardial boundary with its inherent natural variability over time. In this Bayesian segmentation, the likelihood model relies on an accurate Markovian statistical modelling of the grey level distribution of each class present in the image. This likelihood is defined as maximal when the deformed template delineates two regions distributed according to the aforementioned distributions. The parameters of this distribution mixture are given by a preliminary statistical estimation method called Iterative Conditional Estimation (ICE) [16]. This estimation scheme relies on a local *a priori* model, and takes into account the distribution shape of each class present in the ultrasound image. Then, the detection problem of the object

of interest is stated in a Bayesian framework as the estimation of the deformation parameters of the template that maximise the posterior Probability Density Function (PDF). To efficiently maximise this function, and contrary to many other approaches which use gradient-based optimisation methods [1,4,17] and rely on a proper initialisation of the template for each image of the sequence, we will show that an alternative approach consists in using a genetic exploration of the parameter space combined with a steepest ascent procedure. The use of a Genetic Algorithm (GA) as part of image segmentation algorithms [18,19] or for deformable model initialisation [20] is not new. Nevertheless, the combination of this global optimisation technique with a local one (a gradient ascent technique) has not been exploited in this detection and segmentation context. The optimisation method we propose is fast, has the capability of avoiding local minima, and no human interaction has to be used to initialise the prototype template on the first frame of the echographic image sequence.

The first contribution of this paper relies on the combination of deformable templates and Markov Random Field (MRF) models for the segmentation and tracking of anatomical structures, such as the endocardial boundary. In this method, we search to model the available local and global *a priori* knowledge (global prior shape and local prior textural knowledges) to more accurately constraint this detection and segmentation problem. The second contribution of this paper consists in tackling the minimisation problem related to this Bayesian segmentation with an unsupervised hybrid optimisation procedure. This leads to a segmentation method that is completely data-driven, and does not require any user interaction.

This paper is organised as follows. Sections 2 and 3 present the deformable template modelling and the proposed Bayesian segmentation approach. In Section 4, we detail the distribution mixture parameters estimation step used in the data likelihood model. The stochastic search method using a GA combined with a steepest ascent procedure is described in Section 5. In Section 6, we report some experimental results obtained on synthetic images and on a real ultrasound image sequence.

2. DEFORMABLE TEMPLATE REPRESENTATION

To model the global shape knowledge of the endocardial contour, we first roughly define it by a set of τ labelled points, equally sampled, which approximate the outline of a circle (see Fig. 1a). A cubic B-spline shape representation involving these control points corresponding to ‘landmarks’ is then defined.

This way of modelling objects has been widely considered in the object recognition literature, and particularly in the active contour approach [21]. Such a scheme captures the global structure of a shape without specifying a parametric form for each class of shapes. Let us note that this original prototype template can also be obtained from a learning population and an *off-line* training procedure. This can be




Fig. 1. Local and global non-affine deformations. (a) Initial template (the grid is not part of the model itself); (b) γ_0 with local random perturbations. Global non-affine deformations on this locally deformed template with: (c) $M = N = 1$ and $\xi_{mn}^x = \xi_{mn}^y = 2 \forall m, n \in \{0, 1\}^2$; (d) $M = N = 2$ and $\xi_{mn}^x = \xi_{mn}^y = 2 \forall m, n \in \{0, 1, 2\}^2$.

done by manually extracting the endocardial outline on different images from an echocardiographic sequence. The extracted shapes have to be normalised in scale and aligned with respect to a common reference frame. A mean shape γ_0 can then be easily computed from this learning population [4,22]. Nevertheless, the circular crude representation of the endocardial boundary remains close to that given by this type of training procedure (cf. Fig. 15b).

The prototype template γ_0 does not describe the possible instances of the object shape to be detected. To take into account the natural variability of the object considered, we introduce a set of admissible affine transformations, ensuring a first crude registration of the shape, and a set of non-affine local and global transformations.

2.1. Description of Local Non-affine Transformations

First, a local deformation process δ applied to the τ control points is introduced. These local deformations applied to the original shape γ_0 allow us to model local random perturbations that can occur for each control point of the template outline [4]. Mathematically, these local deformations are represented by local random translations δ for each control point of γ_0 (cf. Fig. 1b).

2.2. Description of Global Non-affine Transformations

To refine the shape representation, we introduce a global non-affine deformation. The one used in this paper was first proposed by Jain et al. [17]. In this approach, the image is considered to be mapped to the unit square $S = [0, 1]^2$. The deformation is then represented by a displacement function denoted $\mathcal{D}(x, y)$. The mapping $(x, y) \mapsto (x, y) + \mathcal{D}(x, y)$ is thus a deformation of S , a smooth mapping of the unit square onto itself. The space of such displacement functions is spanned by the following orthogonal bases:

$$e_{mn}^x(x, y) = (2 \sin(\pi mx) \cos(\pi my), 0) \quad (1)$$

$$e_{mn}^y(x, y) = (0, 2 \cos(\pi mx) \sin(\pi my)) \quad (2)$$

for $m, n = 1, 2, \dots$. Low values of m and/or n correspond to lower frequency components of the deformation in the x and y directions, respectively. The displacement function is then chosen as follows:

$$\mathcal{D}_\xi(x, y) = \sum_{m=1}^M \sum_{n=1}^N \frac{\xi_{mn}^x e_{mn}^x(x, y) + \xi_{mn}^y e_{mn}^y(x, y)}{\lambda_{mn}} \quad (3)$$

where $\lambda_{mn} = \pi^2 (n^2 + m^2)$ are the normalising constants. The parameter vectors ξ are the projections of the displacement function on the orthogonal basis previously defined. To allow a sufficiently wide range of possible deformations, while keeping the number of parameters reasonable, we use $M = N = 2$. Figure 1 illustrates the series of deformations of the initial template using higher order terms and different values of M and N . We can notice that these deformations can efficiently express the inherent elasticity of the endocardial contour. We also notice that the deformation becomes more complex as higher frequency components are added to the displacement function \mathcal{D} .

2.3. Description of Affine Transformations

Finally, we introduce a set of admissible linear transformations on γ . These deformations involve translation, scaling, rotation and stretching of the template.

Let now $\gamma\theta$ be a deformed version of the original prototype γ_0 according to the aforementioned transformations with parameter vector θ . A global configuration of the deformable template is thus described by five parameters corresponding to affine transformations, four pairs of global non affine deformation parameters (ξ_{mn}^x, ξ_{mn}^y) , and $(\tau - 2)$ pairs of translation vector $(\delta_{x_i}, \delta_{y_i})$, respectively. The circular prototype template along with the set of aforementioned transformations constitute our global prior model describing the different possible configurations of the shape of interest.

3. SEGMENTATION MODEL

3.1. Introduction

As mentioned in the introduction, the automatic segmentation of anatomical structures in ultrasound images is

especially challenging due to artifacts caused by the speckle noise effect. This partially correlated noise, inherent to the imaging process, lead to a poor image quality (with low resolution, contrast and signal-to-noise ratio) on which the sharpness of anatomical structure contours is noticeably reduced.

In a commonly used deformable model-based Bayesian segmentation approach, these contours (or some spatial and/or temporal gradient measures derived from the input image) are generally exploited in the data likelihood model. In a Bayesian approach, this likelihood model aims at expressing a measure of similarity between the deformed template and the object(s) present in the image and constrains the deformable template to be attracted and aligned, *via* an appropriate likelihood energy term, to the salient spatial or temporal edges of the object to be detected [4,17,23–25].

Due to the speckle noise, our likelihood model cannot rely efficiently on this type of gradient measures. Instead, we propose here a region oriented statistical approach. More precisely, we propose to model and use the grey level statistical distribution of each homogeneous region existing in an echocardiogram, and to exploit an *a priori* information about the location of the endocardium with respect to each other anatomical structures present in the ultrasound image. This *a priori* anatomical information expresses the fact that the endocardial contour delineates two homogeneous regions (\mathcal{R}_0 , \mathcal{R}_1). The first one, called the *blood* zone (\mathcal{R}_0), arises from the low acoustic wave reverberation in the ventricular cavity (filled with blood). The second one, called the *muscle* region (\mathcal{R}_1), is due to the acoustic signal reverberation on the cardiac muscle. Each aforementioned region has completely different grey level intensity distribution, and each pixel of the input image has a membership likelihood to each class $\{e_0 = \textit{blood}, e_1 = \textit{muscle}\}$ that we can define by the following Probability Density Functions (PDFs):

$$P_{Y_s/X_s}(y_s/x_s = e_i), \quad i \in \{0, 1\} \quad (4)$$

where y_s designates the grey level of the pixel at site s and x_s is its associated class label. Y_s and X_s represent the random variables associated to the realisations x_s and y_s , respectively. Assuming that R distinct homogeneous regions are present in the input image ($R = 2$ in our application), the maximisation over θ of the following expression:

$$P_{Y|\theta}(y/\theta) = \prod_{j=0}^{R-1} \sum_{s:x_s=e_j} P_{Y_s/X_s}(y_s/x_s) \quad (5)$$

allows us to find the optimal template (or the optimal deformation parameter vector θ of the original prototype template γ_0) which best matches the image. Equivalently, we can minimise $-\ln P_{Y|\theta}(y/\theta)$ and use this previous expression as our energy function. This energy function is not heuristic and statistically segments the image in a Maximum Likelihood sense. Let us also note that this term defines the likelihood distribution used in a classical Markovian segmentation.

3.2. Joint Model

In that prospect, we let the template define two regions (cf. Fig. 2): the first is the set of pixels *inside* the region delimited by the prototype template γ_0 , and is denoted γ_0^\bullet . The second is the outside region, defined by the set of pixels located on the contour of the deformed template with a slightly superior scale. This last one is denoted γ_0° . From this, we propose the following joint model through the Gibbs distribution:

$$P_{\theta,y}(\theta, y) = \frac{1}{Z} \exp\{-\epsilon(\theta, y)\} \quad (6)$$

where Z is a normalising constant and ϵ is an objective function measuring how well a given instance of deformed template γ_θ fits the content of the input image y . The energy function $\epsilon(\theta, y)$ is composed of two terms, as explained below:

- *Likelihood energy term*: this first energy term measures the likelihood of the image, and can be compared to the log-likelihood distribution defined in Eq. (5):

$$\epsilon_l(\theta, y) = -\frac{1}{N_{\gamma_0^\bullet}} \sum_{s \in \gamma_0^\bullet} \ln P_{Y_s/X_s}(y_s/e_0) - \frac{1}{N_{\gamma_0^\circ}} \sum_{s \in \gamma_0^\circ} \ln P_{Y_s/X_s}(y_s/e_1) \quad (7)$$

where the summation of the first and second term of $\epsilon_l(\theta, y)$ is over all the $N_{\gamma_0^\bullet}$ pixels inside the region defined by the deformed template γ_θ and over all the $N_{\gamma_0^\circ}$ pixels delimited by γ_θ , respectively. Note that the two weighting factors are necessary to get a scale-invariant likelihood measure. This function attains its minimum value when the deformed template delimits exactly two homogeneous regions with grey level distribution corresponding to *blood* and *muscle* class for the region γ_0^\bullet and γ_0° , respectively.

- *Prior energy term*: the second term penalises the deviation of the deformed template γ_θ from the original prototype




Fig. 2. The two different regions delimited by the template at a specific location and transformation. γ_0^\bullet is the set of pixels inside the region delimited by the deformed prototype template γ_θ . γ_0° is the set of pixels located on the contour of γ_θ with a slightly superior scale.

γ_0 . This function does not penalise affine transformations¹ and is close to that proposed by Jain et al. [17], except for the second term:

$$\epsilon_p(\theta, \gamma) = \left(\sum_{\xi_i \in \xi} [(\xi_i^x)^2 + (\xi_i^y)^2] + \sum_{\delta_i \in \delta} (\delta_i)^2 \right) \quad (8)$$

where (ξ_i^x, ξ_i^y) and δ_i correspond to displacement function parameters of the global non-affine deformation, and the translation parameter vector of the non-affine local deformations, respectively (see Section 2).

3.3. MAP Detection

Using these two energy terms, the joint distribution $P_{\Theta, Y}(\theta, \gamma)$ can be written as

$$P_{\Theta, Y}(\theta, \gamma) = \frac{1}{Z} \exp - \underbrace{\{\epsilon_l(\theta, \gamma) + \rho \epsilon_p(\theta, \gamma)\}}_{\epsilon(\theta, \gamma)} \quad (9)$$

where the factor ρ provides a relative weighting between the two penalty terms, and allows us to control the ‘rigidity’ of the template. The posterior distribution deduced from Eq. (9) is given by

$$P_{\Theta|Y}(\theta/\gamma) = \frac{1}{Z_\gamma} \exp \{-\epsilon(\theta, \gamma)\} \quad (10)$$

where Z_γ is a normalising constant depending on γ only. We formulate now the detection problem as the search of the *Maximum A Posteriori* (MAP) estimation of θ

$$\hat{\theta}_{MAP} \in \arg \max_{\theta} \{P_{\Theta|Y}(\theta/\gamma)\} \quad (11)$$

$$\in \arg \min_{\theta} \epsilon(\theta, \gamma) \quad (12)$$

The search of $\hat{\theta}_{MAP}$ is therefore equivalent to the minimisation of the non concave following objective function:

$$\begin{aligned} \epsilon(\theta, \gamma) = & -\frac{1}{N_{\gamma_0}} \sum_{s \in \gamma_0} \ln P_{Y_s|X_s}(y_s/e_0) - \frac{1}{N_{\gamma_0}} \sum_{s \in \gamma_0} \ln P_{Y_s|X_s} \\ & (y_s/e_1) + \rho \left(\sum_{\xi_i \in \xi} [(\xi_i^x)^2 + (\xi_i^y)^2] + \sum_{\delta_i \in \delta} (\delta_i)^2 \right) \end{aligned} \quad (13)$$

This function is minimal when the deformed template delineates two homogeneous regions (*blood* and *muscle* for our application) distributed according to the grey-level distribution corresponding to each region, and when the deformed template is not too different from the original prototype (in the non-affine deformations sense). To get an unsupervised deformable-based segmentation scheme, we have now to estimate the distribution mixture parameter vector Φ , i.e. the distribution parameters associated to each region, *blood* and *muscle* of the ultrasound image. We consider this in the next section.

4. Distribution mixture parameter estimation

We consider a couple of random fields $Z = (X, Y)$ (called the ‘complete data’), where $Y = \{Y_s, s \in S\}$ represents the field of observations located on a lattice S of ν sites s (associated to the ν pixels of the input image), and $X = \{X_s, s \in S\}$ the label field (related to ν *blood* or *muscle* class labels of the segmented image). Each Y_s takes its value in $\{0, \dots, 255\}$, and each X_s in $\{e_0 = \textit{blood}, e_1 = \textit{muscle}\}$. The distribution of (X, Y) is herein defined, firstly, by prior distribution $P_X(x)$, supposed to be stationary and Markovian, and secondly, by site-wise likelihoods $P_{Y_s|X_s}(y_s/x_s)$ whose parameter vector $\Phi_{(x_s)}$ depends on class label x_s . To take into account the speckle noise phenomenon [26] in the reverberation areas, we model each $P_{Y_s|X_s}(y_s/x_s)$ by a shifted Rayleigh law with parameter vector $\Phi_{(x_s)} = (\min, \alpha)$, namely; $\mathcal{R}_Y(y; \Phi) = [(y_s - \min)/\alpha^2] \exp[-(y_s - \min)^2/2\alpha^2]$ with $y > \min$ and $\alpha > 0$.

To determine $\Phi = (\Phi_{(e_1)}, \Phi_{(e_2)})$, we resort to the ICE algorithm [27]. This procedure relies on an estimator $\hat{\Phi}(X, Y)$ for completely observed data case. This iterative method starts from an initial parameter vector $\Phi^{[0]}$ (not too far from the optimal one), and generates a sequence of parameter vectors leading to the optimal parameters (in the least squares sense) with the following iterative scheme:

$$\Phi^{[p+1]} = \frac{1}{\varphi} [\hat{\Phi}(x_{(1)}, \gamma) + \dots + \hat{\Phi}(x_{(n)}, \gamma)] \quad (14)$$

where $x_{(i)}$, $i = 1, \dots, \varphi$ are realisations of X drawn according to the posterior distribution $P_{X|Y, \Phi}(x/\gamma, \Phi^{[p]})$. To decrease the computational load, we can take $\varphi = 1$ without altering the quality of the estimation [16]. Finally, we can use the Gibbs sampler algorithm [28] to simulate realisations of X according to the posterior distribution. For the local *a priori* model of the Gibbs sampler, we adopt a standard isotropic Potts model with the 8-connexity spatial neighbourhood [14]. In this model, there are four parameters, called ‘the clique parameters’, associated to the horizontal, vertical, right and left diagonal binary cliques, respectively (see Fig. 3). Given this *a priori* model, the prior distribution $P_X(x)$ can be written as

$$P_X(x) = \exp \left(- \sum_{\langle s, t \rangle} \beta_{st} (1 - \delta(x_s, x_t)) \right) \quad (15)$$

where summation is taken over all pairs of spatial neighbouring sites and δ is the Kronecker delta function. To favour homogeneous regions with no privileged orientation in the




Fig. 3. Second order neighbourhood and associated two-site cliques.

¹Due to the fact that this function does not penalise affine transformations like the stretching transformation, the ‘mean shape’ γ_0 that is used as prior information is also an ellipse. This prior information seems to be, anatomically speaking, natural and true for *short* or *long* axis views of the endocardial boundary.

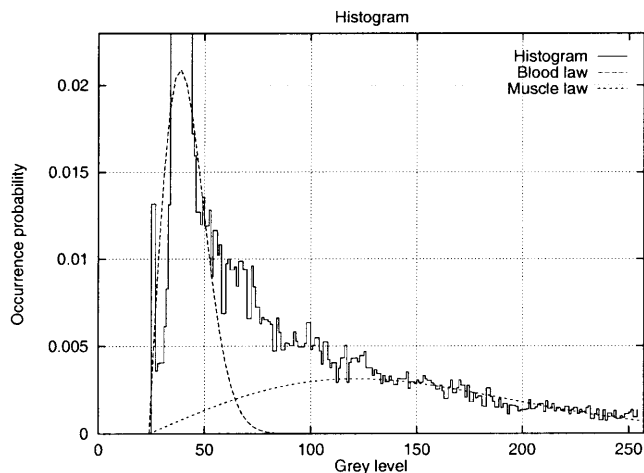


Fig. 4. Image histogram of the picture reported in Fig. 7a (solid curve) and estimated probability density mixture obtained with the ICE procedure (dotted and dashed curves).

Gibbs sampler simulation process, we choose $\beta_{st} = 1$. Finally, $\Phi^{[p+1]}$ is computed from $\Phi^{[p]}$ in the following way:

- *Stochastic step*: using the Gibbs sampler, one realisation x is simulated according to the posterior distribution $P_{X|Y}(x/y)$, with parameter vector $\Phi^{[p]}$.
- *Estimation step*: the parameter vector $\Phi^{[p+1]}$ is estimated with the Maximum Likelihood (ML) estimator of the ‘complete data’ [29]. If $N_0 = \#\{s \in S: x_s = e_0\}$ is the number of pixel of the *blood* area, we have for $\Phi_{(e_0)}$

$$\hat{\min} = \min_{s \in S: x_s = e_0} (y_s) - 1, \quad \hat{\alpha}^2 = \frac{1}{2N_0} \sum_{s \in S: x_s = e_0} (y_s - \hat{\min})^2 \quad (16)$$

- Repeat until convergence is achieved, i.e. if $\Phi^{[p+1]} \neq \Phi^{[p]}$, we return to stochastic step.

Figure 4 represents the estimated distribution mixture on the echogram shown in Fig. 7a, and the histogram of this ultrasound image: the two site-wise likelihoods $P_{Y_s|X_s, \Phi}(y_s/e_i, \Phi)$, $i = 0, 1$, weighted by the estimated proportion π_i of each class e_i are superimposed to the image histogram. Corresponding estimates obtained by the ICE procedure, requiring about 12 iterations, are given in Table 1. The quality of the estimations is difficult to appreciate in the absence of ground truth values. We can roughly

Table 1. Estimated parameters for the picture reported in Fig. 7a. π stands for the proportion of the two classes within the ultrasound image. \min and α are the Rayleigh law parameters

	ICE procedure		
$\Phi_{(blood)}^{final}$	0.48 _(π)	24 _(\min)	207 _(α^2)
$\Phi_{(muscle)}^{final}$	0.52 _(π)	24 _(\min)	9436 _(α^2)

perform such an evaluation by comparing the image histogram (solid curve) with the probability density mixture corresponding parameters (dotted and dashed curves). Estimation results on synthetic images with ground truth values will be given in Section 6, in order to appreciate the robustness of this procedure and the quality of the estimations. Experiments have shown that the distribution mixture parameters keep constant for all the images extracted of the same echographic sequence. For this reason, and to decrease the computational load, the mixture parameters are computed once and for all on the first frame of the echographic sequence. Now that we are able to estimate the parameters of the mixture-based data model, we can turn our attention to the optimisation problem of ϵ .

5. OPTIMISATION PROBLEM

The objective function to be minimised in Eq. (12) is a complex function with several local extrema over the deformation parameter space. A global search is usually impossible due to the size of the configuration space. Instead, we have implemented a Genetic Algorithm-based optimisation technique.

Genetic Algorithms (GA) are a class of robust stochastic search and global optimisation procedures which mimic the evolution of natural systems [30]. The algorithm acts in an iterative way by allowing parallel evolution in a population of \mathcal{N} individuals. Each individual represents a point of the search space, and is a candidate solution to the optimization problem. It is represented by a string or ‘chromosome’, which is composed of a list of L features (corresponding to the L searched parameters). The parameters have to be encoded in an appropriate manner. The most common approach is to quantise the parameter values and to binary code them. The fitness of the various individuals (the tentative ‘solution’) to the environment is expressed by a fitness function, which, after the characteristics contained in a chromosome have been decoded, gives a ‘performance’ value to the string. Genetic search is carried out in a sequence of ‘generations’. In each generation, a new population of \mathcal{N} chromosomes is created with genetic operators. These operators mimic the biological phenomena of *selection*, *crossover* and *mutation*. The choice of the solution upon which they are used is dictated by the evolutionary principle of the ‘survival of the fittest’. The algorithm begins with an initial population of \mathcal{N} chromosomes randomly chosen, and terminates when either a specified number of iterations has been performed or a maximally fitted individual has emerged.

In our application, let us recall that we have to optimise an L dimensional function. Each of the L parameters θ_l is binary quantified on q bits in order to take any value within the predefined range $[\theta_{L\min}, \theta_{L\max}]$. Therefore, the i th chromosome, denoted $[\theta]_i$, is a string of qL bits length, i.e.

$$[\theta]_i = \underbrace{(c_{i11}^1, c_{i12}^1, \dots, c_{i1q}^1)}_{\theta_1} \underbrace{c_{i21}^1, c_{i22}^1, \dots, c_{i2q}^1}_{\theta_2}, \dots, \underbrace{c_{iL1}^1, c_{iL2}^1, \dots, c_{iLq}^1}_{\theta_L} \quad (17)$$

where c_{kl}^i designates the l th bits associated to the k th para-

meter, and θ is the parameter vector associated to the chromosome $[\theta]_i$ after decoding.

5.1. Fitness Measure

We can easily derive a fitness measure \mathcal{F} (to be maximised) directly from Eq. (12) (the energy function ϵ to be minimised). To turn $\epsilon(\theta, y)$ into a fitness measure for use in a genetic algorithm, we can choose,

$$\mathcal{F}([\theta]_i) = \exp \{-\epsilon(\theta, y)\} \quad (18)$$

The following is the detail of the *selection*, *crossover* and *mutation* operators. The associated parameters used in our application will be given in Section 6.

5.2. Selection

Individuals with higher fitness survive and individuals with lower fitness die. Let us assume that at iteration k , the population of the GA is the set of \mathcal{N} chromosomes, i.e.

$$\text{POP}^k = \{[\theta]_1^k, \dots, [\theta]_{\mathcal{N}}^k\} \quad (19)$$

Generation of the next population is based on the evaluation of \mathcal{F} for all individuals of POP^k . More precisely, we probabilistically select each chromosome for ‘reproducing’ in the next generation, using their relative fitness

$$p([\theta]_i^k) = \frac{\mathcal{F}([\theta]_i^k)}{\sum_{j=1}^{\mathcal{N}} \mathcal{F}([\theta]_j^k)} \quad (20)$$

5.3. Crossover

A pair of chromosomes is picked up at random, and the single-point crossover operator is applied according to a fixed crossover probability. For this operation, a random number in the range of 0 to the length L_q of the string is generated. This is called the crossover point. The portions of two strings lying to the right of the crossover point are interchanged to yield two new strings, as shown in Fig. 5.




Fig. 5. An example of crossover with the partial exchange of information.

5.4. Mutation

Mutation consists in considering in turn each bit of a given chromosome and changing its value with a predefined low probability called the mutation rate.

To speed up the convergence rate, we have developed two strategies and we have combined them:

1. The first one is an elite-preservation strategy [30]: the individual with the highest fitness survives to be an individual of the next generation.
2. The second strategy (called hybrid GA [30]) consists in associating the genetic search with a local optimisation technique. In each generation, a percentage of the best individuals are used to initialise a gradient ascent technique. Therefore, these best individuals explore local neighbourhoods in the parameter space to find a point of higher fitness.

In our application, these two strategies are used together in an efficient way to create a global optimisation technique called a ‘hybrid genetic algorithm with an elitist strategy’.

6. EXPERIMENTAL RESULTS

In this section, the real echocardiograms (short axis view) of a normal subject were provided by the Montreal Heart Institute. The ultrasound imaging system was a Hewlett Packard SONOS 1500 with a transducer frequency of 2.5 MHz. The images were Time Gain Compensated (TGC) for attenuation.

6.1. Distribution Mixture Estimation Step

Based on the distribution mixture parameters given by the ICE procedure (see Table 1), we can compute the probability map for the *blood* and *muscle* classes. This map exhibits sharp boundaries between high and low probability regions for each class. As a result, the probability maps have a high




Fig. 6. Smoothed probability map for the ultrasound image reported in Fig. 7a. Each probability map is estimated by computing the membership likelihood of each pixel to each class ($e_0 = \text{blood}$, $e_1 = \text{muscle}$) defined by Eq. (4) and based on distribution mixture parameters (see Table 2 and Fig. 4) given by the ICE procedure (described in Section 4). (a) Probability map associated to the *blood* class; (b) probability map associated to the *muscle* class.




Fig. 7. Successive generations of the genetic search on an ultrasound image showing the endocardium (first frame of the sequence). Best 5% deformed templates before the gradient ascent procedure (on the left) and the best resulting template after the gradient ascent procedure (on the right). Hybrid GA optimisation after. (a) Two generations ($\mathcal{F}(\epsilon(\theta, y)) = 0.091$); (b) three generations ($\mathcal{F}(\epsilon(\theta, y)) = 0.119$); (c) four generations ($\mathcal{F}(\epsilon(\theta, y)) = 0.173$); (d) eight generations ($\mathcal{F}(\epsilon(\theta, y)) = 0.186$). An unsupervised and reliable detection of the endocardial border for the first frame of the sequence is obtained.




Fig. 8. Successive generations of the genetic search on an ultrasound image showing the endocardium. Best 5% deformed templates before the gradient ascent procedure (on the left) and the best resulting template after the gradient ascent procedure (on the right). Hybrid GA optimisation after. (a) 0 generation (initial population of the GA) ($\mathcal{F}(\epsilon(\theta, y)) = 0.100$); (b) two generations ($\mathcal{F}(\epsilon(\theta, y)) = 0.123$); (c) four generations ($\mathcal{F}(\epsilon(\theta, y)) = 0.155$); (d) eight generations ($\mathcal{F}(\epsilon(\theta, y)) = 0.157$). An unsupervised and reliable detection of the endocardial border is obtained.




Fig. 9. Successive generations of the genetic search on synthetic images with two different synthetic Rayleigh noises. (a) $\alpha_5^2 = 700$, $\alpha_1^2 = 2000$, (b) $\alpha_5^2 = 1000$, $\alpha_1^2 = 1500$. Best 5% deformed templates before the gradient ascent procedure (on the left) and the best resulting template after the gradient ascent procedure (on the right). From top to bottom, hybrid GA optimisation after. (a) Two generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.149$), six generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.167$), 10 generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.174$); (b) two generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.116$), 10 generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.119$), 24 generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.120$). In both cases, an accurate and reliable detection and segmentation is obtained, and can be compared to the ground truth segmentation (at the bottom). Rates of correct classification are (a) 97.20%, (b) 93.46%.

gradient at the boundary of the different regions and low gradient everywhere else. To avoid that the gradient-based local optimisation procedure gets trapped in a local minima, and to speed up the convergence rate, we decide to convolve each probability map five times with the following 3×3 2D Gaussian mask: [1, 2, 1; 2, 4, 2; 1, 2, 1]. Besides, due to the speckle noise effect, the probability maps, associated to each class are very noisy. The smoothing process counteracts this problem as well. Figure 6 shows the smoothed probability map for the ultrasound image reported in Fig. 7a.

6.2. Detection Step

This detection step has been carried out with the hybrid genetic algorithm using the elitist strategy described in Section 5, and after the distribution mixture estimation step proposed in Section 4. To reduce the size of the parameter search space, the stochastic exploration based on the GA is used to only estimate the affine transformation parameters (five parameters to be estimated). This affine deformation parameter estimation ensures a first crude registration of the

shape to be detected. The local exploration technique of the best individuals selected by the GA, which is based on the steepest ascent procedure, is then used to estimate the whole parameter vector (including the global and local non-affine deformations parameters).

Tests have shown that this optimisation procedure is not very sensitive to the control parameters.² In our application, these parameters are commonly used [30], and are the following: population size = 100, crossover rate = 0.8, mutation rate = 0.008, maximum number of generations = 30. Parameters value are quantified on $q = 8$ bits. At each generation, we select 5% of the best individuals (i.e. five best parameter vectors) for the hybridation with the local optimisation technique. The prototype template is a connected 12-points model ($\tau = 12$). The weighting factor penalising the prior term with respect to the likelihood term is set to 0.1 and the size of the image is 256×256 pixels. All these parameters keep constant for the detection and the tracking step.

Therefore, to summarise, 100 individuals (or 100 ellipses: circle + affine transforms) are initially randomly generated in the image (first generation or first step of the genetic search). Each individual represents a candidate solution (a candidate contour) for the optimisation problem. The genetic procedure then acts in a iterative way by creating a new population with genetic operators. At each generation, we select 5% of the best individuals (i.e. five best parameter vectors) for the initialisation of the gradient ascent technique.

Using this procedure, a reliable detection of the endocardial contour is obtained for all images of the echographic sequence. Our GA takes about 5–20 generations to converge to the true solution. In fact, the convergence rate can vary depending on the complexity of the objective function $\epsilon(\theta, y)$ to be minimised (or the complexity of the input image and/or the shape to be extracted). Tests have shown that, after a learning step (i.e. typically 3–6 generations or iterations of the algorithm), the genetic exploration of the search space parameter provides good initialisations for the local exploration technique. This optimisation technique allows us to reach a minima nearby the global minima without requiring the high computational load required by other global optimisation algorithm such as the simulated annealing algorithm proposed in Kervrann and Heitz [4]. In our application, the optimisation procedure for the detection process takes about three minutes (average CPU time) on a standard Sun/Sparc2 workstation.

Figures 7 and 8 show two real echocardiograms extracted from an echocardiographic sequence and illustrate the best deformed template and the best 5% set of templates before the gradient ascent technique, and thus before the estimation of parameters associated to global and local non-affine transformations, for successive iterations or generations

of the genetic search. In spite of the speckle noise and a random initial population for the GA, an accurate and reliable detection of the endocardial border is obtained. Figure 9 shows two synthetic images with different white Rayleigh noise levels and the resulting detection/segmentation obtained by the proposed deformable template-based unsupervised segmentation method, along with their associated rates of correct classification. Rates of correct classification are obtained by dividing the number of correctly classified pixels with the number of pixels within the synthetic shape (ground truth). Real parameters of the distribution mixture and resulting segmentations can be compared to estimated parameters by the ICE procedure (see Table 2) and to the ground truth segmentations (cf. bottom of Fig. 9). Figures 10 and 11 show synthetic ultrasound images with the texture model developed in Meunier and Bertrand [31] (with correlated Rayleigh noise) and the corresponding segmentation results. We notice that faithful noise model estimations and accurate segmentations are obtained. These experiments demonstrate that the proposed detection and segmentation scheme is reliable, robust and avoid a manual initialisation of the model.

6.3. Tracking Step

The tracking strategy used in our application is the following: the final estimate of the previous time frame is used as an initialisation for the steepest ascent procedure for the current time frame. If the resulting value of energy ϵ at convergence is higher than a given threshold, we can assume that the inter-frame motion is small, and this strategy allows to provide a proper initialisation of the template for the local exploration technique used on the next frame. In this case, the result of the gradient ascent technique gives the final result for the current time frame. Otherwise (i.e. if ϵ

Table 2. Estimated parameters for the picture reported in Fig. 9. π stands for the proportion of the two classes within the synthetic image. min and α are the Rayleigh law parameters. From top to bottom (a) Initialisation of the ICE procedure; (b) Estimated parameters by the ICE procedure; (c) Real parameters

Initialisation			
$\Phi_{(e_0)}^{[0]}$	$0.50_{(\pi)}$	$00_{(min)}$	$150_{(\alpha^2)}$
$\Phi_{(e_1)}^{[0]}$	$0.50_{(\pi)}$	$00_{(min)}$	$10000_{(\alpha^2)}$
ICE procedure			
$\Phi_{(e_0)}^{final}$	$0.09_{(\pi)}$	$19_{(min)}$	$672_{(\alpha^2)}$
$\Phi_{(e_1)}^{final}$	$0.91_{(\pi)}$	$19_{(min)}$	$2037_{(\alpha^2)}$
Real parameters			
$\Phi_{(e_0)}^{real}$	$0.09_{(\pi)}$	$19_{(min)}$	$700_{(\alpha^2)}$
$\Phi_{(e_1)}^{real}$	$0.91_{(\pi)}$	$19_{(min)}$	$2000_{(\alpha^2)}$

² Let us recall that this GA-based optimisation technique is only used to provide good initialisations for the gradient-based local exploration technique. This explains why the proposed hybrid optimisation procedure is relatively ‘insensitive’ to these genetic control parameters (if these are contained within the commonly used range of values proposed by Goldberg).




Fig. 10. Successive generations of the genetic search on synthetic images with the ultrasound texture model developed in Meunier and Bertrand [31]. Best 5% deformed templates before the gradient ascent procedure (on the left) and the best resulting template after the gradient ascent procedure (on the right). From top to bottom, hybrid GA optimisation after. (a) Two generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.289$), four generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.409$), 18 generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.439$); four generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.269$), eight generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.299$), 12 generations ($\mathcal{F}(\epsilon(\theta, \gamma)) = 0.406$). In both cases, an accurate and reliable detection and segmentation is obtained, and can be compared to the ground truth segmentation (at the bottom). Rates of correct classification (a) 97.20%, (a) 95.08%.

is lower than a given threshold), we decide to use the genetic algorithm combined with the steepest ascent procedure to re-estimate the whole parameter vector θ . To this end, each of the L ($L = 35$) parameters θ_i is quantified on 8 bits in order to take any value within the range $[\theta_i - \delta\theta_i, \theta_i + \delta\theta_i]$, with θ_i the estimated parameter value given at the previous time. Tests have shown that the tracking procedure often provides good initialisations from one frame to the next, avoiding having to resort to GA optimisation. Nevertheless, this GA-based stochastic minimisation procedure remains necessary to recover the good shape if the inter-frame motion is important. In this tracking process, the gradient ascent technique takes about ten seconds and

GA-based optimisation takes about one minute on a standard Sun/Sparc2 workstation. Figure 12 summarises, step-by-step, the overall estimation procedure of the endocardial boundary.

Figure 13 shows the tracking of the endocardial contour in an echocardiographic image sequence (each frame size is 256×256) at different time frames during the cardiac cycle. Figure 14 shows two other detection/segmentation results obtained with our procedure on real ultrasound images. The best resulting template is drawn on a rectangular grid to visualise the global non-affine deformation estimated by the hybrid GA. The global prior model seems to be flexible enough for representing accurately the inherent




Fig. 11. Top: synthetic ultrasound images with the ultrasound texture model developed in Meunier and Bertrand [31] and for different ratios of noise level (the backscattering ratio between the two regions is respectively: 10, 7, 6 and 5). Bottom: segmentations obtained with our method. (a) Rate of correct classification: 95.98%; (b) rate of correct classification: 95.12%; (c) rate of correct classification: 95.12% (d) rate of correct classification: 93.65%.

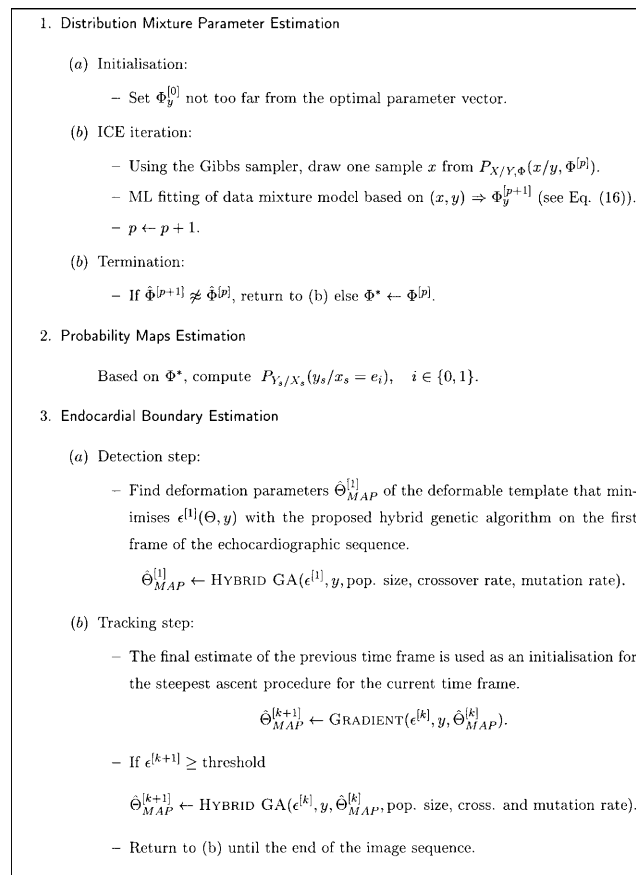


Fig. 12. Unsupervised endocardial boundary estimation algorithm.




Fig. 13. Tracking of the endocardial contour in the echographic sequence at different time frames during the cardiac cycle. From top left to bottom right: frame 1, 4, 6, 9, 12, 13, 18, 20, 27, 34, 35, 40, 41, 43, 46, 48.

natural variability of the endocardial contour during the cardiac cycle. Figure 15a shows the best resulting templates obtained on the 20 first frames of the image sequence reported in Fig. 13. The mean shape is computed from this set of deformed templates (see Fig. 15b). We can see that this mean shape is not too different from the proposed circular original template γ_0 .

6.4. Validation and Discussion

To confirm the effectiveness of this automatic segmentation method, we have provided a comparison of accuracy with manual tracings without inclusion of the papillary muscles³

³ We recall that our segmentation model of the endocardial contour is supposed to *exclude* the cardiac muscles such as the papillary muscles.

(see Fig. 16). These contours were traced by two experts from the Montreal Heart Institute on each frame of the echographic sequence (50 frames).

To give quantitative measurements, we have computed

- The mean and standard deviation for intra-observer variability, expressed by the rate of identically classified pixels between two manual tracings given by the same expert.
- The extra-observer variability, expressed by the rate of identically classified pixels between two manual tracings given by two different experts.
- The rate of correct classification⁴ between the ground-

⁴ We recall that the rate of identical (or correct) classification is expressed by the number of identically (or correctly) classified pixels divided by the number of pixels within the endocardial contour.




Fig. 14. Top: detection/segmentation obtained with our procedure on real ultrasound images. Bottom: the best resulting template is drawn on a rectangular grid to visualise the global non-affine transformation estimated by the GA. (a) Frame 33; (b) frame 45.




Fig. 15. Set of deformed templates and mean shape. (a) Best resulting templates obtained on the 20 first frames of the image sequence reported in Fig. 13. (b) Mean shape computed from this set of deformed templates.

truths and 50 estimated curves by our segmentation method.

Each expert has thus provided two sets of 50 manual tracings (namely, *contour1-1* and *contour1-2* for the first expert and *contour2-1* and *contour2-2* for the second expert). These manual tracings allow us to estimate two rates of intra-observer correct classification, four rates of inter-observer correct classification and four rates of correct classification for our automatic segmentation method (by assuming that the ground-truths given by an expert correspond to the true classification). These percentages of correct classification are given in Table 3.

The mean rates of intra-observer correct classification are,

respectively, 92.4% (std. deviation $\sigma = 2.4\%$) for the first expert and 84.2% ($\sigma = 3.2\%$) for the second expert. The four mean rates of inter-observer correct classification are, respectively, 85.8% ($\sigma = 3.9\%$), 83.3% ($\sigma = 4.4\%$), 83.8% ($\sigma = 4.0\%$) and 85.3% ($\sigma = 4.2\%$).

The variability between the ground-truths given by an expert and the 50 estimated curves by our algorithm, expressed in terms of the rate of correct classification is, on average, 79%, (see Table 3). This rate of correct classification remains a bit lower than that given by the inter-observer rates (on average, 84.5%). This can be explained by the fact that the papillary muscles are not always clearly visible, particularly near end-diastole (see Fig. 13, frame 1, for example). In this situation, the algorithm includes a part of the papillary muscle. The mean rate of correct classification between the ground-truths given by an expert including the papillary muscles and the computed curves during this period is, as expected, much higher: 88.4%. As for the other frames acquired around the end-systole (see Fig. 13, frame 35, for example), the papillary muscles are clearly visible on most images, and the algorithm obtains an acceptable 84.0% of correct classification during this period. Therefore, the algorithm performs well with errors comparable to the inter-observer errors, but will include the papillary muscle when those structures are barely visible.

The inclusion of papillary muscles may appear as a limitation of the proposed method if there is a need to exclude them, even if they are barely visible. To overcome this limitation, a solution consists in manually adjusting the weighting factor (ρ in Eq. (9)) controlling the ‘rigidity’ of the template. For instance, a lower value when the papillary muscles are not well defined can largely solve this problem. Nevertheless, a lower value for this parameter cannot be used for all the frames of the echographic sequence, because a sufficient value is sometimes necessary to constrain the segmentation procedure.

7. CONCLUSION

In this paper, we have developed a robust algorithm to detect and track anatomical deformable structures like the endocardial border in an ultrasound image. We have stated the segmentation issue in the Bayesian framework, and take into account some available *a priori* knowledge of this problem. First, the proposed global prior model, integrating an original prototype template along with predefined global and local deformations, has shown itself to be very flexible and well suited to model the high variability of the endocardial contour over long image sequence. Another important source of *a priori* knowledge is information about the intensity and texture of the tissue of different anatomical structures. These factors are highly dependent on the imaging process and are taken into consideration by modelling the speckle distribution of each class of the ultrasound image by a Rayleigh law. Parameters of each PDF are given by a preliminary Markovian estimation step, and are then exploited in the data likelihood model in order to statistically segment the image. The proposed method can be easily




Fig. 16. Examples of manual and computed endocardial contours in successive frames of the echographic sequence. Top: segmentation results. Bottom: manual tracing given by an expert.

Table 3. Rates of identical (or correct) classification. Respectively, intra-observer measures, extra-observer measures and rate of correct classification obtained by our segmentation method

Intra-observer measures			
contour1-1	contour1-2	92.4%	($\sigma = 2.4\%$)
contour2-1	contour2-2	84.2%	($\sigma = 3.2\%$)
Inter-observer measures			
contour1-1	contour2-1	85.8%	($\sigma = 3.9\%$)
contour1-2	contour2-2	83.3%	($\sigma = 4.4\%$)
contour1-1	contour2-2	83.8%	($\sigma = 4.0\%$)
contour1-2	contour2-1	85.3%	($\sigma = 4.2\%$)
Algorithm measures			
Algorithm	contour 1-1	77.6%	($\sigma = 7.8\%$)
Algorithm	contour 1-2	78.7%	($\sigma = 8.9\%$)
Algorithm	contour 2-1	80.2%	($\sigma = 10\%$)
Algorithm	contour 2-2	78.9%	($\sigma = 11\%$)

generalised with multidimensional pixel value and/or regions parameterised with a distribution mixture, and presents several attractive features compared to other related approaches. In particular, it seems to be well suited to handle ultrasound images with strong speckle noise on which edge information or some gradient measures cannot be exploited. Finally, we have shown that this problem can be handled as an equivalent energy minimisation problem for each image of the echographic sequence. To this end, the optimisation problem considered is tackled using a genetic exploration combined

with a steepest ascent procedure. This combined local and global optimisation procedure is fast, robust, simple and well suited for our application compared to other optimisation techniques, such as the gradient-based method or the simulated annealing algorithm [28]. Besides, this method does not require initialisation of the template close to the desired solution. Initialisation may be defined at random, leading to segmentation procedures that are completely data driven.

Acknowledgments

The authors thank INRIA (Institut National de la Recherche en Informatique et Automatique, France) for financial support of this work (postdoctoral grant).

References

1. Zhang Y, Jain AK, Dubuisson-Jolly MP. Object tracking using deformable templates. Proc International Conference on Computer Vision 1998; 440-445
2. Rueckert D, Burger P. Geometrically deformed templates for shape-based segmentation and tracking in cardiac MR images. Proc International Workshop EMMCVPR'97: Energy Minimisation Methods in Computer Vision and Pattern Recognition, LNCS 1223, Venice, Italy, May 1997; 83-98. Springer Verlag
3. Fenster SD, Kender JR. Sectorised snakes: evaluating learned-energy segmentations. Proc International Conference on Computer Vision 1998; 420-426
4. Kervrann C, Heitz F. A hierarchical Markov modeling approach for the segmentation and tracking of deformable shapes. Graphical Models and Image Processing 1996; 60(3):884-900
5. Dias JMB, Leitao JMN. Wall position and thickness estimation from sequences of echocardiographic images. IEEE Trans Medical Imaging 1996; 15(1):25-38
6. Chakraborty A, Staib LH, Duncan JS. Deformable boundary

- finding in medical images by integrating gradient and region information. *IEEE Trans Medical Imaging* 1996; 15:859–870
7. Song-Chun Z, Yulle A. Region competition: unifying snakes, region growing, and bayes/MDL for multiband image segmentation. *IEEE Trans Pattern Analysis and Machine Intelligence* 1996; 18(9):884–900
 8. Echocardiographical Cardiac Function Assessment and Wall Motion Visualisation Using Fuzzy Logic and the Wavelet Transform. PhD thesis, Signal Processing Division, University of Strathclyde, Glasgow, 1998
 9. Haynor DR, Chalana V, Linker DT, Kim Y. A multiple active contour model for cardiac boundary detection on echocardiographic sequences. *IEEE Trans Medical Imaging* 1996; 15(3):290–298
 10. Chiou GI, Hwang JN. A neural network based stochastic active contour model (nns-snake) for contour finding of distinct features. *IEEE Trans Image Processing* 1985; 4(10):1407–1416
 11. Krucinski S, Mikic I, Thomas JD. Segmentation and tracking in echocardiographic sequences: Active contours guided by optical flow estimates. *IEEE Trans Medical Imaging* 1998; 17(2):274–285
 12. Kamaleddin Setarehdan S, Soraghan JJ. Automatic cardiac lv boundary detection and tracking using hybrid fuzzy temporal and fuzzy multiscale edge detection. *IEEE Trans Biomedical Eng* 1999; 46(11):1364–1379
 13. Hammoude A. Endocardial border identification in two-dimensional echocardiographic images: review of methods. *Computerized Medical Imaging and Graphics* 1998; 22(3):181–193
 14. Besag J. On the statistical analysis of dirty pictures. *J Roy Statistical Soc* 1986; B-48:259–302
 15. Mignotte M, Collet C, Pérez P, Bouthemy P. Unsupervised Markovian segmentation of sonar images. *Proc International Conference on Acoustics, Speech, and Signal Processing* 1997; 4:2781–2785
 16. Braathen B, Masson P, Picczynski W. Global and local methods of unsupervised Bayesian segmentation of images. *Graphics and Vision* 1993; (1):39–52
 17. Jain AK, Zhong Y, Lakshmanan S. Object matching using deformable templates. *IEEE Trans Pattern Analysis and Machine Intelligence* 1996; 18(3):267–278
 18. Mignotte M, Collet C, Pérez P, Bouthemy P. Hybrid genetic optimization and statistical model-based approach for the classification of shadow shapes in sonar imagery. *IEEE Trans Pattern Analysis and Machine Intelligence* 2000; 22(2):129–141
 19. Yuen SY, Ma CH. Genetic algorithm with competitive image labelling and least square. *Pattern Recognition* 2000; 33(12):1949–1966
 20. Cootes TF, Hill A, Taylor CJ. A generic system for image interpretation using flexible template. *Proc British Machine Vision Conference* 1992; 276–285
 21. Nakazawa Y, Komatsu T, Saito T. A robust object-specified active contour model for tracking smoothly deformable line-features and its application to outdoor moving image processing. *Proc International Conference on Image Processing* 1996; 689–692
 22. Cootes TF, Taylor CJ, Cooper DH, Graham J. Active shape models—their training and application. *CVGIP: Image Understanding* 1994; 61:38–59
 23. Lakshmanan S, Grimmer D. A deformable template approach to detecting straight edges in radar images. *IEEE Trans Pattern Analysis and Machine Intelligence* 1996; 18(4):438–443
 24. Dubuisson Jolly MP, Lakshmanan S, Jain AK. Vehicle segmentation and classification using deformable templates. *IEEE Trans Pattern Analysis and Machine Intelligence* 1996; 18(3):293–308
 25. Kluge K, Lakshmanan S. Lane boundary detection using deformable templates: effects of image subsampling on detected lane edges. *Proc Asian Conference on Computer Vision* 1995; 141–145
 26. Goodman JW. Some fundamental properties of speckle. *J Optical Soc America* 1976; 66(11):1145–1150
 27. Salzenstein F, Pieczynski W. Unsupervised Bayesian segmentation using hidden markovian fields. *Proc International Conference on Acoustics, Speech, and Signal Processing* 1995; 2411–2414
 28. Geman S, Geman D. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans Pattern Analysis and Machine Intelligence* 1984; 6(6):721–741
 29. Schmitt F, Mignotte M, Collet C, Thourel P. Estimation of noise parameters on sonar images. *Signal and Image Processing, SPIE* 2823, Denver, CO, 1996; 1–12
 30. Goldberg DE. *Genetic Algorithm*. Addison Wesley, 1989
 31. Meunier J, Bertrand M. Ultrasonic texture motion analysis: theory and simulation. *IEEE Trans Medical Imaging* 1995; 14(2):274–285
-
- Max Mignotte** received the Master of Sciences (Electronics and Telecommunications) (1992) and the DEA (Postgraduate degree) in Digital Signal, Image and Speech processing from the INPG University, France (Grenoble) in 1993, and the PhD degree in electronics and computer engineering from the University of Bretagne Occidentale (UBO) and the digital signal laboratory (GTS) of the French Naval academy, France, in 1998. He was an INRIA post-doctoral fellow at University of Montreal (DIRO), Canada (Québec), from 1998 to 1999. He was also a NSERC post-doctoral fellow, a teaching assistant and a lecturer at University of Montreal from September 1999 to July 2000. He is currently with DIRO at the Computer Vision & Geometric Modeling Lab as an Assistant Professor at University of Montreal. His current research interests include statistical methods and Bayesian inference for image segmentation (with hierarchical Markovian, statistical templates or active contour models), parameters estimation, tracking, classification, deconvolution and restoration issues in medical or sonar imagery.
-
- Jean Meunier** received the BSc degree in physics from the University of Montreal in 1981, the MScA degree in applied mathematics in 1983, and PhD in biomedical engineering in 1989 from the Ecole Polytechnique de Montréal. In 1989, after postdoctoral studies at the Montreal Heart Institute, he joined the Department of Computer Science at the University of Montreal, where he is currently full professor. He is also a regular member of the Biomedical Engineering Institute at the same institution. His research interests are in computer vision and its applications to medical imaging. His current research focuses on motion assessment and analysis in biomedical images.
-
- Correspondence and offprint requests to:* Professor M. Mignotte, Department d'Informatique et de Recherche Operationnelle, C.P. 6128, Succ. Centre-ville, 2920 Chemin de la Tour, H3C 3J7 Montreal, Canada. Email: mignotte@iro.umontreal.ca.