Automatic 3-D Segmentation of Endocardial Border of the Left Ventricle From Ultrasound Images

Carlos Santiago, Jacinto C. Nascimento, *Member, IEEE*, and Jorge S. Marques

***Abstract*—The segmentation of the left ventricle (LV) is an important task to assess the cardiac function in ultrasound images of the heart. This paper presents a novel methodology for the segmentation of the LV in three-dimensional (3-D) echocardiographic images based on the *probabilistic data association filter* (PDAF). The proposed methodology begins by initializing a 3-D deformable model either semiautomatically, with user input, or automatically, and it comprises the following feature hierarchical approach: 1) edge detection in the vicinity of the surface (low-level features); 2) edge grouping to obtain potential LV surface patches (mid-level features); and 3) patch filtering using a shape-PDAF framework (high-level features). This method provides good performance accuracy in 20 echocardiographic volumes, and compares favorably with the state-of-the-art segmentation methodologies proposed in the recent literature.**

***Index Terms*—3-D echocardiography, deformable models, image segmentation, left ventricle (LV), robust estimation.**

I. Introduction

THE automatic segmentation of the LV endocardium from ultrasound (US) images is a crucial step to characterize heart functionality, e.g., ejection fraction. In a clinical setup, there are several advantages involved in solving this problem in echocardiography. This includes the increase of patient throughput and the reduction of interuser variability in the LV delineation procedure. Although other imaging techniques, such as MRI and CT, provide images with better spatial resolution, echocardiography is still one of the most preferred modalities due to its temporal resolution, versatility, and low cost [1]. However, automatic segmentation in echocardiography is not a simple task since it raises several challenges. First, the US images may contain more than one anatomical structure, meaning that it is necessary to correctly identify the region of interest (i.e., *initialization* procedure). Second, the borders of those structures may not always be present and they are often difficult to locate due to misleading edges and background noise. This problem is particularly accentuated in echocardiographic images, characterized by low signal-to-noise ratio, weak echoes, scattering, echo dropouts, and presence of speckle (i.e., *cluttered background and outliers*).

Deformable models have been one of the most promising approaches to the segmentation in echocardiographic images [2], [3]. However, they still face two major issues. First, the *initialization* has a significant influence in the segmentation. This means that different initializations may provide different final shape estimates, introducing user-dependent bias. An automatic initialization solves this problem by guaranteeing that, for a specific volume, the final segmentations are consistent. Second, the *presence of clutter* hampers the accuracy of the segmentations. Most approaches do not have the ability to identify which image features (e.g., edge points) actually belong to the object of interest and which belong to the background noise or neighboring structures. This may cause the final segmentation to be poorer since erroneous edges have a negative impact in the adaptation of the deformable model. These problems have been the subject of intense research in the past decades and are the main concern of this paper.

This paper proposes a new algorithm based on deformable models, tailored to perform the segmentation of the LV in three-dimensional (3-D) echocardiographic images targeted to account for the difficulties mentioned previously, i.e., *initialization* and *presence of outliers and cluttered background*. To achieve this goal, we propose 1) different initialization procedures along with 2) a probabilistic data association filter (PDAF) allowing for the robustness against the presence of clutter.

II. Related Work

Since the work by Blake *et al.* [4], the sequential state estimation framework has proved useful in a wide range of applications. The main idea is to use the Kalman filter, combining state predictions, and image measurements, to track B-spline contours deformed by linear transformations lying in a shape space (model subspace). The LV segmentation from echocardiographic images is one of such examples that uses the above framework. For instance, in [5], [6], the proposed segmentation is performed in a sequential state-estimation where the extended Kalman filter is used to recursively update global pose and local shape parameters. In the same spirit [7], the state estimation is combined with an active shape model with the predefined deformation modes. This approach is also applied in [8] in the context of two-dimensional (2D) echocardiography. In a similar line of research, state-space tracking is formulated in a information fusion problem [9]. Also, a state-based approach for cardiac deformation analysis has been proposed in [10]. In this paper, we propose a novel algorithm to perform the 3-D segmentation of the LV from echocardiographic volumes. The proposed

Manuscript received October 3, 2013; revised December 20, 2013; accepted February 18, 2014. Date of publication February 26, 2014; date of current version December 30, 2014. This work was supported in part by the FCT Project PEst-OE/EEI/LA0009/2013, by the Project "HEARTRACK"–PTDC/EEA-CRO/103462/2008, and under the FCT Scholarship SFRH/BD/87347/2012.

The authors are with the Institute for Systems and Robotics, Instituto Superior Tecnico, 1049-001 Lisbon, Portugal (e-mail: carlos.santiago@ist.utl.pt; jan@isr.ist.utl.pt; jsm@isr.ist.utl.pt).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JBHI.2014.2308424

methodology relies on the use of state estimation whose update is accomplished via PDAF that has its roots in [11].

There exists several alternatives to perform segmentation/tracking with multiple observations. Although, Kalman filter provides an ideal solution, this is only valid under the Gaussian assumption, which is not the present case since the US is not an unimodal image (e.g., presence of speckle). Nearest neighbor filter, is another approach for localizing the target in clutter environments. In this approach, the closest observation is used to update the state. However, if the selected observation, that is closest to the predicted measurement is wrong, this method does not provide acceptable results. An alternative is to deal with every observations. A possible strategy is to track-and-split individually every possible hypothesis among observations. The disadvantage is that, this can lead to an expensive computational effort if one has relatively large number of observations. Other alternatives are the Bayesian class of approaches, e.g., optimal Bayesian filter that computes the association probabilities for each sequence of observations. Suboptimal version are possible the so-called "N scan back", meaning that the observations of the N latest images are taken into account Again, this approach may have excessive computing in US images. The best compromise is a suboptimal strategy—PDAF, that only accounts for the latest set of observations, i.e., it computes the association probabilities for the current image observations (i.e., "No scan back"). From the aforementioned reasons, PDAF is a suitable choice for the purpose of this paper, since it can provide robust segmentation in the presence of clutter at low computational cost.

Another concern of this paper is the initialization problem. Several strategies can be found in the literature concerning this aspect. For instance, the gradient vector flow technique [12] that has been widely used, including in the LV segmentation problem [13], [14]. Alternatively, in [15], it is proposed to combine a deformable model with an image registration scheme consisting of a mutual information analysis between the echocardiographic volume and a presegmented image (template) of the LV. Another common approach is to statistically determine the most probable location of the LV based on a learning procedure, e.g., [16]–[20]. Other approaches have also been proposed. In [21], the centers of divergence concept is suggested to solve the initialization issue. On the other hand, Tauber *et al.* [22] proposes an automatic initialization algorithm that consists in analyzing the 2-D short axis (SA) planes of the volume. To the best of our knowledge, there have not been other reports of a 3-D (semi) automatic LV segmentation concerning the use of the proposed framework.

III. Proposed Methodology and Contributions

This paper proposes a methodology that is able to deal with both the *initialization* and the *presence of clutter* issues. Fig. 1 illustrates an overview of our approach. It includes three main blocks: 1) the initialization of the surface model; 2) the extraction of the image features, in a multistep process that filters edge points based on their reliability; and 3) the adaptation of the surface towards the LV border. Each of these steps is explained in detail in the following sections.

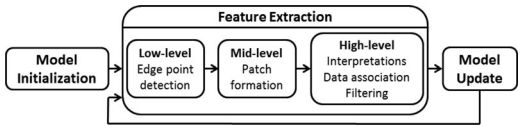

Fig. 1. Block diagram of the segmentation system.

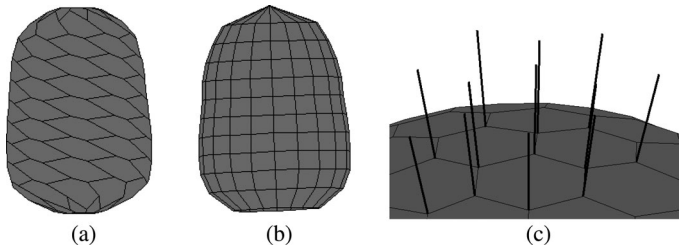

Fig. 2. Surface models used: (a) a simplex mesh (b) a conic-shaped model. Directional search along directions orthogonal to the surface at each vertex of the surface model (c).

Regarding our previous publication [23], this paper provides the following new contributions:

1) Two different initialization approaches are compared: 1) a user-dependent initialization based on the space carving algorithm [24] and 2) an automatic initialization algorithm inspired in the method proposed by [22];
2) In order to mitigate the influence of outliers, the proposed method is inspired in the PDAF framework proposed in [11]. In this paper, we extend the concept of the shape probabilistic data association to 3-D data;
3) We compare the LV segmentations using two different 3-D surface models (conic and simplex mesh models);
4) We extend the experimental evaluation for 20 US volumes of the LV
5) We also compare the quality of the segmentations using several denoising techniques as a preprocessing step.

The paper is organized as follows. Section IV describes the 3-D surface models used to segment the LV. Section V explains the feature extraction steps, from low-level to high-level. Section VII describes the data set used to test the proposed method, as well as the evaluation methodology. Section VIII-B shows the results from all the tests and segmentations made to assess the method. Finally, Section IX concludes the paper with final remarks.

IV. Surface Models and Initialization

In this section, we describe the two 3-D surface models used in the proposed segmentation system. The first model used is a 3-D simplex mesh [25], which consists of a set of N vertices $x = [x_{v_1}, \ldots, x_{v_N}]^T$ [see Fig. 2(a)]. The simplex mesh has a neighborhood system that imposes that each vertex $x_{v_i} \in \mathbb{R}^3$ has exactly three neighboring vertices, which form a triangle that defines a tangent plane to the surface at x_{v_i}. This enables the computation of the normal to the surface and the search line used in the detection of low-level features [see Fig. 2(c)].

As for the second model, we used a conic-shaped deformable surface that consists of a set of N' vertices $x = \left[x_{v_1}, \ldots, x_{v_{N'}}\right]^T$ that form a square grid shape around the LV

and finishes with a triangulation with the apex point $x_{v_{N'}}$ [see Fig. 2(b)]. We designed this model so that each equatorial set of vertices stay in the SA plane (orthogonal to the left ventricular long axis). Consequently, the feature search is performed within the SA plane and the adaptation of the model is strictly orthogonal to the long axis (LA). Therefore, these two surface models constitute two different approaches: a flexible simplex mesh without any node restrictions (apart from surface smoothness, as will be mentioned ahead) and a model with additional constrains since the nodes are organized along parallel planes.

In both models, the update equation follows the one originally defined by Delingette in [26], where the new position of the ith vertex at k iteration $x_{v_i}(k)$ is a sum of two forces plus a momentum term. Formally,

$$x_{v_i}(k) = x_{v_i}(k-1) + c_1 M_{v_i}(k) + c_2 F_{v_i}^{\text{int}}(k) + c_3 F_{v_i}^{\text{ext}}(k) \tag{1}$$

where $M_{v_i}(k) = x_{v_i}(k-1) - x_{v_i}(k-2)$ is the momentum term, used to improve the speed of the adaptation, $F_{v_i}^{\text{int}}(k)$, and $F_{v_i}^{\text{ext}}(k)$ are the internal and external forces, respectively, and c_1, c_2, and c_3 are constants.

The external force $F_{v_i}^{\text{ext}}(k)$, is responsible for pulling each vertex towards the desired position $\widehat{x}_{v_i}(k)$ (which is the high-level feature). This is accomplished by simply computing the desired displacement

$$F_{v_i}^{\text{ext}}(k) = \widehat{x}_{v_i}(k) - x_{v_i}(k-1) \tag{2}$$

where $x_{v_i}(k-1)$ is the vertex position at the previous time step and $\widehat{x}_{v_i}(k)$ is the corresponding estimated location of the LV boundary point. The computation of the high-level features $\widehat{x}(k)$ will be described in Section V-C.

The internal force $F_{v_i}^{\text{int}}(k)$, on the other hand, is responsible for keeping the structure of the surface smooth. The computation of the internal force can be separated into two orthogonal components: a force $F_{v_i}^{n}$ orthogonal to the surface at x_{v_i} and a tangential force $F_{v_i}^{t}$. The way these two components are defined is different and depend on the surface model being used.

Regarding the simplex mesh, the orthogonal force $F_{v_i}^{n}$ adjusts the vertex's height to the tangent plane formed by the three neighbors (related to the curvature of the surface at x_{v_i}). The desired height is defined by the average height of the second degree neighborhood, i.e., of the surrounding nine vertices. This means that the surface tends to have a uniform curvature (see [25] for a detailed description). The tangential force $F_{v_i}^{t}$ attracts each vertex towards the center of the triangle formed by the neighbors, thus ensuring that the vertices are evenly spaced between each other.

In the conic-shaped model, each $F_{v_i}^{n}$ acts orthogonally to the curve formed by the set of vertices located in that SA plane. It adjusts the distance of each vertex x_{v_i} to the center of the curve (radius) based on the average radius of all the vertices in that plane and in the adjacent planes (above and below planes). This means the model will tend to keep a similar radius throughout the SA planes, thus guaranteeing that the surface is smooth. The tangential force, $F_{v_i}^{t}$, on the other hand, keeps x_{v_i} equally spaced to the two neighboring vertices in that plane, in a tangential displacement to the curve. This helps avoid overlapping vertices throughout the adaptation.

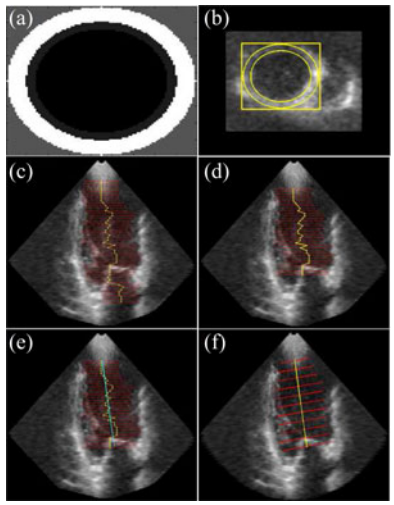

Fig. 3. Automatic initialization procedure. (a) Template used in the detection of the LV center and radius; (b) selection of the best patch location and size; (c) coronal view of the detected LV centers and corresponding radii (size of the red line); (d) selection of the slices that embrace the LV; (e) computation of the LA; (f) correction of the centers and corresponding radii.

A. Initialization

Both deformable models will be subjected to two different initialization methods: 1) a manual initialization that requires a user input and 2) an automatic initialization that does not need user intervention. In both cases, the initialization process has three steps: 1) identifying the location of the LV center in each SA plane of the image, 2) computing the left ventricular LA, and 3) creating the surface model using the information obtained in 1) and 2).

In the first method, we use an initialization scheme that makes use of the space carving algorithm [24] (see [23] for details). The second initialization is a contribution of this paper, and works as follows. The algorithm produces an initial LV boundary by searching for the LV cross-section in each axial plane. We assume that the LV has approximately a circular shape in these slices. This assumption has previously been proved to lead to good initializations in [22]. The proposed approach follows three main steps:

1) *Detection of the LV center and radius*. The axial planes (volume slices) are sequentially analyzed from apex to base. For each zth slice of the LV volume, we compute the convolution between its intensity [see Fig. 3(b)] and a template [see Fig. 3(a)] varying its center and radius. For each slice, we keep highest convolution score. In the next $z+1$th slice, a sliding window operation is also performed but only in the vicinity of the previous best center c_z and close to the previous radius r_z. The output of this process is illustrated in Fig. 3(c).

   From Fig. 3(c), we see that slices located above and below the LV are also obtained and should be discarded. To determine which axial slices actually embrace the LV

volume, we define a profile of a volume as the signal with the best scores along the slices. Repeating this process for N volumes, we obtain N profiles. To determine which slices embrace the LV, we average all volume profiles, aligned so that the first value corresponds to the first apex slice. Finally, we search for the set of slices that maximize the correlation between their profile and the template. The output of this process is illustrated in Fig. 3 (d).

2) *3D orientation of the LV surface model*. This step aims to correct the positions of the centers [see the irregular yellow lines in the images shown in Fig. 3(c,d)]. In both the manual and automatic initialization, the selected centers are used to determine the left ventricular LA by performing a 3-D linear regression. Fig. 3(e) shows the centers before (yellow color) and after (cyan color) the regression procedure. With the LA, the new centers are sampled evenly spaced from the base to the apex. Fig. 3(f) shows an example of the final centers and radii of the automatic initialization procedure.
3) *Creation of the surface model*. In both models, several points are sampled in each SA plane (orthogonal do the LA) in a circle with the center and radius computed in the previous step. Regarding the conic-shaped model, the sampled points match the model vertices and the neighborhood system forms quadrilateral polygons. An apex point, located along the LA, is added to the model which forms triangles with the vertices from the slice closest to the apex (see Fig. 2). On the other hand, the vertices of the simplex mesh are obtained from the centers of the triangles formed by the sampled points. In this case, a base and an apex point are added. The number of sampled points per slice in the simplex mesh is half of the conic-shaped model so that the final number of model points is similar. In all the tests performed, the number of sample points in the simplex mesh case was set to 10 (20 in the conic-shaped model), resulting in a total number of $20\times$#slices vertices plus two vertices in the base and apex (one vertex in the case of the other model).

V. Feature Extraction

This section addresses how feature detection can be performed under noisy images. The proposed feature extraction block includes three steps that aim to produce more reliable LV boundary points. The first step consists of the detection of 3-D edge points in the vicinity of the surface model, which we denote as low-level features. The second step introduces the notion of mid-level features, which are surface patches formed by grouping low-level features. These patches correspond to edge surface portions that describe the boundary of some structure in the volume. Finally, the third step consists of filtering the mid-level patches based on how reliable they are. It makes use of the data association [11] concept to produce the final set of points—denoted as high-level features, that indicate the estimated location of the LV boundary. These steps are described in the following subsections.

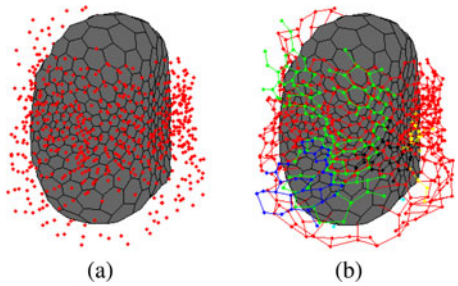

Fig. 4. Mid-level features detection (five patch detected). (a) Shows all the detected low-level features (red dots). (b) Shows the corresponding patches: each color represents a different patch and the edges link neighboring low-level features from the same patch.

A. Edge Point Detection—Low-level Features

Edge point detection is performed by directional search along lines orthogonal to the surface. Edges are detected using the matched filter approach described (we refer to the reader Chap. 5 of [4] for a comprehensive review).

Even though the matched filter is the optimal filter for detecting known shapes in noisy signals [26], in this case, it often leads to the detection of more than one edge point, mainly due to the background noise and to the existence of other cardiac structures in the vicinity of the LV. Fig. 4(a) shows the result of the low-level feature extraction in an echocardiographic volume.

Since it is impossible to know which of the detected features actually belong to the LV border and which are outliers, one should consider all possible combinations of valid and invalid features (interpretations). However, the number of interpretations grows exponentially, which means that this approach is not feasible. Linking neighboring low-level features into patches (surface portions), which we call mid-level features, significantly decreases the number of interpretations.

B. Patch Formation—Mid-level Features

Defining which low-level features belong to a specific patch can be viewed as a labeling problem, similar to the one described in [27]. This means that the label assigned to an edge point defines the patch it belongs to.

This labeling problem is solved using two main assumptions: 1) edge points associated to a specific vertex (i.e., found along the same search line) must belong to different patches, since it is unlikely that a single border has two consecutive edges; 2) two edge points associated to two neighboring vertices are more likely to belong to the same patch if they are close to each other, meaning that patches should have spatial continuity.

Based on the two aforementioned premises, we defined an energy function E associated to the label configuration Λ (complete set of labels assigned to each feature) that is minimum when these two criteria are satisfied. Let $y_{v_i,w}$ be the wth edge point associated to vertex v_i and $y_{v_j,q}$ be the qth edge point associated to vertex v_j. If $y_{v_i,w}$ and $y_{v_j,q}$ have the same label (belong to the same patch) and are detected at neighboring sites (v_i and v_j are neighbors), we define a link l between them (see

Fig. 4). The energy of a label configuration Λ is then

$$E(\Lambda) = \sum_{l} e(l). \tag{3}$$

where $e(l)$ denotes the energy of a single link l and the sum is computed over all links in Λ. The energy function e can be divided into three terms. The first term, e_1, prevents links between features too far apart, therefore promoting patches with spatial continuity:

$$e_1(l) = \begin{cases} \infty & \text{if } \|y_{v_i,w} - y_{v_j,q}\| \geq t_l \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

where t_l is a threshold (maximum allowed distance between neighboring patch features). The second term, e_2, rejects label repetitions along the same search line

$$e_2(l) = \begin{cases} \infty & \text{if the label has already been assigned to an} \\ & \text{edge point associated with } v_i \text{ or } v_j \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

The third term, e_3, promotes assigning neighboring features the same label if they verify the mutual favorite pairing criterion (MFP) [28]

$$e_3(l) = \begin{cases} -1 & \text{if } y_{v_i,w} \text{ and } y_{v_j,q} \text{ verify MFP} \\ 1 & \text{otherwise.} \end{cases} \tag{6}$$

This criterion imposes that the same label should be assigned to $y_{v_i,w}$ and simultaneously $y_{v_j,q}$ if 1) $y_{v_j,q}$ is closer to $y_{v_i,w}$ than any other feature associated with v_j, and 2) $y_{v_i,w}$ is closer to $y_{v_j,q}$ than any other feature associated with v_i.

To achieve the minimum energy configuration, we propose a labeling algorithm that uses a label propagation scheme (see [23] for detailed description of the labeling process).

Fig. 4(b) shows the output of the labeling process based on the detected edge points, shown in Fig. 4(a).

C. Data Association—High-level Features

The final step of the feature extraction block consists in estimating the LV boundary based on the detected patches. Since undesired patches (outliers and not belonging to the LV) may be detected, we use a robust shape estimation approach that is able to deal with outliers. This method is inspired in data association filtering [11], [29].

In each iteration k of the model adaptation, we define a set of possible patch interpretations $I_k = \{I_{0\,k}, \ldots, I_{m_k\,k}\}$, where a patch interpretation $I_{i\,k}$ is a binary sequence of patch labels, classifying each patch as valid (label "1") or invalid (label "0"). Then, each interpretation receives a confidence degree, the association probability $\alpha_{i\,k}$, which depends on the size and position of the patches with respect to the surface model. The computation of the association probabilities is explained further.

1) State Model: The model assumes that the LV boundary is described by a set of 3-D points $x_k = [x_{v_1\,k}, \ldots, x_{v_N\,k}]^T$ of dimension $N \times 3$. Since this is a static problem, (we only address the 3-D segmentation of the LV, although the extension to 3D + T is straightforward), the motion model is described by

$$x_k = x_{k-1} + w_k \tag{7}$$

where x_k is the state vector (mesh vertices) at the kth iteration of the estimation process and $w_k \sim \mathcal{N}(0, Q)$ is white noise with Gaussian distribution.

The observation model, on the other hand, depends on the patch interpretation, i.e., it depends on which observations are considered as valid. For an interpretation $I_{i\,k}$, the observation model is defined as

$$y_{i\,k} = C_{i\,k} x_k + v_{i\,k} \tag{8}$$

where $C_{i\,k}$ is the observation matrix associated with the ith interpretation $I_{i\,k}$ and $v_{i\,k} \sim \mathcal{N}(0, R_i)$ is white noise with Gaussian distribution. $C_{i\,k}$ is an $M_i \times N$ matrix obtained from the identity matrix by removing the rows associated to invalid features. Thus, the observation matrix $C_{i\,k}$ assigns the set of points from x_k that originated the set of observations $y_{i\,k} = [y_{1\,k}, \ldots, y_{M_i\,k}]^T$, with dimension $M_i \times 3$, that are considered valid according to $I_{i\,k}$.

2) Data Association: The estimation of the state vector $\widehat{x}_{k|k}$ involves the computation of the expected value of the posterior probability density function $p(x_k \mid L^k)$, where $L^k = \{L_1, \ldots, L_k\}$ is the cumulative set of detected patches from iteration 1 to k. The PDAF formulation performs the decomposition of the estimation with respect to the latest set of measurements (i.e., *no scan back*). This means that, the probability of the state vector given the past observations L^{k-1} can be approximately defined (see [11]) by a Gaussian distribution

$$P(x_k \mid L^{k-1}) = \mathcal{N}\left[x_k;\, \widehat{x}_{k|k-1}, P_{k|k-1}\right] \tag{9}$$

where $\widehat{x}_{k|k-1}$ and $P_{k|k-1}$ are the predicted mean vector and covariance matrix, respectively. Under this assumption, the conditional mean of the state at iteration k is

$$\begin{aligned} \widehat{x}_{k|k} &= E[x_k \mid L^k] = \sum_{i=0}^{m_k} p\{I_{i\,k}|L^k\} E[x_k|I_{i\,k}, L^k] \\ &= \sum_{i=0}^{m_k} \alpha_{i\,k|k} \widehat{x}_{i\,k|k} \end{aligned} \tag{10}$$

where $\alpha_{i\,k|k} = p\{I_{i\,k}|L^k\}$ is the probability of the interpretation $I_i(k)$, also called association probability, and $\widehat{x}_{i\,k|k}$ is the state vector estimate conditioned on $I_i(k)$. Each $\widehat{x}_{i\,k|k}$ is updated by the Kalman filtering

$$\widehat{x}_{i\,k|k} = \widehat{x}_{k|k-1} + K_{i\,k}\,\nu_{i\,k},\ i = 1, \ldots, m_k \tag{11}$$

where $K_{i\,k}$ and $\nu_{i\,k}$ are the Kalman gain and the innovation, respectively, associated with the interpretation $I_{i\,k}$. Replacing (11) in (10), we obtain

$$\widehat{x}_{k|k} = \widehat{x}_{k|k-1} + \sum_{i=1}^{m_k} \alpha_{i\,k} K_{i\,k} \nu_{i\,k}. \tag{12}$$

A similar analysis can be performed to obtain the covariance estimate (see [11], [30] for details).

Equation (12) means that the state estimate $\widehat{x}_{k|k}$ is a weighted-sum of the estimates according to each interpretation. The

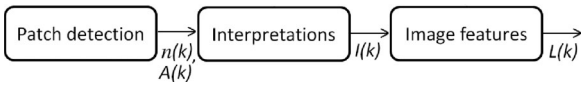

Fig. 5. Data generation model.

weight each interpretation receives depends on the corresponding association probabilities $\alpha_{i\,k}$. The following subsection describes the computation of these weights.

3) Association Probabilities: The association probabilities depend on the data model that generated the detected patches. We assume that the data generation model can be described as shown in Fig. 5, and, consequently, that it depends on the variables listed below.

As we will often encounter terms involving the dependence of k, we will omit this dependence so that the equations can be more compactly and clearly described. The variables are as follows:

1) n—number of patches;
2) $L = \{L_1, \ldots, L_n\}$—vector of the detected patches, where $L_j = \{y_1^j, \ldots, y_{A_j}^j\}$ includes all the low-level features that belong to that patch;
3) $A = \{A_1, \ldots, A_n\}$—A_j being the area of the jth patch, i.e., the number of low-level features that belong to the jth patch;
4) $I = \{I_1, \ldots, I_m\}$—set of interpretations, where $I_i = \{I_i^1, \ldots, I_i^n\}$ is a binary sequence that considers the jth patch valid if $I_i^j = 1$ and invalid if $I_i^j = 0$, and $m = 2^n$ is the number of data interpretations.[1]

The generative model of the data is characterized by the joint distribution $P(L, I_i, A, n \mid L^{k-1})$. This probability can be factorized as follows

$$P(L, I_i, A, n \mid L^{k-1}) = P(I_i \mid L, A, n, L^{k-1}) P(L, A, n \mid L^{k-1}) \tag{13}$$

where the unknown term is the association probability $\alpha_i = P(I_i \mid L, A, n, L^{k-1})$ that must be computed. Using the Bayes' rule, this can be written as

$$\alpha_i \propto P(L|I_i, A, n, L^{k-1})\, P(I_i|A, n, L^{k-1}) \tag{14}$$

the term $P(L \mid I_i, A, n, L^{k-1})$ is the likelihood of the set of patches L according to I_i, and $P(I_i \mid A, n, L^{k-1})$ is the prior probability of the interpretation I_i. In the aforementioned data model, it will be assumed that high association probabilities are assigned to interpretations containing valid patches that are: 1) large and 2) located close to the surface model. The two terms in (14) are defined as follows:

Likelihood: The likelihood term determines the probability of the patches being correct based on their distance to the surface model. We assume that if a patch is close to the surface, it is more likely that it belongs to the LV boundary, whereas if it is farther away from the surface, it probably belongs to another anatomical structure.

[1]Recall that, for a large number of patches, the cardinality $|I|$ is also large, and thus, m grows exponentially with the number of patches. However, this does not happen in practice, since interpretations containing overlapped valid patches (labeled "1") are discarded, so that only a smaller subset of interpretations $m \ll 2^n$ needs to be explicitly examined.

Assuming that each patch L_j is independently generated, the likelihood can be expressed as the product of each individual probability

$$P(L \mid I_i, A, n, L^{k-1}) = \prod_{j=1}^{n} p(L_j \mid I_i, A, n, L^{k-1}). \tag{15}$$

We define the probability $p(L_j \mid I_i, A, n, L^{k-1})$ for two different cases, depending on whether L_j is considered valid or invalid according to I_i. Following [11], we assign a Gaussian distribution, $\mathcal{N}(d; 0, \sigma^2)$ (d defined below), if L_j is considered as valid; an uniform distribution (with the same length as the search line V), otherwise. Formally,

$$p(L_j \mid I_i, A, n, L^{k-1}) = \begin{cases} \beta \mathcal{N}(d; 0, \sigma^2) & \text{if } L_j(k) \text{ is valid} \\ V^{-1} & \text{otherwise} \end{cases} \tag{16}$$

where β is a normalization constant. The parameter d is a measure of the average distance between the low-level features of the jth patch, $\{y_1^j, \ldots, y_{A_j}^j\}$, and the corresponding vertices of the surface model, $\{x_1^j, \ldots, x_{A_j}^j\}$ and defined as$d = \|\bar{e}\|$, where

$$\bar{e} = \frac{1}{A_j} \sum_{q=1}^{A_j} \left(x_q^j - y_q^j\right). \tag{17}$$

Prior: The prior probability of an interpretation I_i determines the probability of I_i being correct based on the size of its valid patches. The higher the area of valid patches, the higher its prior probability. Regarding invalid patches, these should have the opposite effect: if patches with high areas are labeled as invalid in I_i, then it should receive a small prior probability.

Similarly to the likelihood case, we assume that the prior factorizes as follows

$$P(I_i \mid A, n, L^{k-1}) = \prod_{j=1}^{n} p(I_i^j \mid A, n, L^{k-1}) \tag{18}$$

where $p(I_i^j \mid A, n, L^{k-1})$ is the probability of the label assigned to patch j in interpretation i. The following definition satisfies previous statements

$$\begin{aligned} p(I_i^j = 1 \mid A, n, L^{k-1}) &= a \log(A_j + 1) + b = P_j \\ p(I_i^j = 0 \mid A, n, L^{k-1}) &= 1 - P_j \end{aligned} \tag{19}$$

where:

$$a = \frac{P_B - P_A}{\log(A_{\max} + 1)}, b = P_A \tag{20}$$

where P_A, P_B are constants, $A_{\max}$ is the total number of vertices in the surface model, and A_j is the area of the patch L_j (number of features it comprises).

Combining (19) with (18) yields

$$p(I_i \mid A, n, L^{k-1}) = \prod_{j:I_i^j=1} P_j \prod_{j:I_i^j=0} 1 - P_j. \tag{21}$$

In practice, small patches (such that $A_j < 10\% A_{\max}$) are discarded to further reduce the number of possible interpretations and the computational burden of the algorithm.

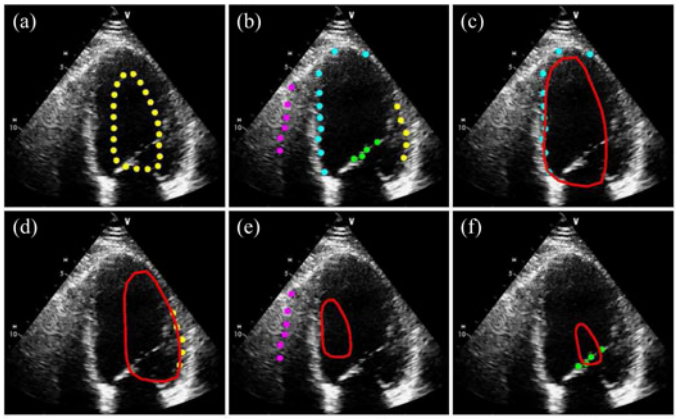

Interpretations	*S1*	S2	S3	S4	α
I_0	0	0	0	0	0.0
I_1	0	0	0	1	0.0
I_2	0	0	1	0	0.0
I_3	0	0	1	1	0.0
I_4	0	1	0	0	0.0003
I_6	0	1	1	0	0.0001
I_8	1	0	0	0	00451
I_9	1	0	0	1	0.1096
I_{12}	1	1	0	0	0.8447

Fig. 6. 2-D shape segmentation example. (a) initial contour guess; (b) four detected segments; shape estimate provided by (c) cyan (valid) segment; (d) yellow (valid) segment; (e) magenta (invalid) segment; (f) green (invalid) segment; (g) final PDAF segmentation; (h) data interpretations and association probabilities.

VI. Illustration of the Proposal

Before testifying the usefulness of the approach in the 3-D experiments, this section illustrates, through a 2-D example,[2] the main characteristics of the PDAF presented in Section V-C. This aims to highlight why probabilistic data association, as a bottom-up approach, is tailored and suited to provide the final boundary segmentation.

Suppose we have an initial guess [see yellow dots in Fig. 6(a)], from which we intend to segment the endocardium. So, a contour expansion is expected. From the orthogonal lines (although not shown) radiating from these samples points, several transitions are detected in the image, and organized in contour-segments [equivalent to the patches in 3D—see Fig. 6(b)]. In this case, four segments ($S = 4$) are detected: two reliable or true segments (cyan and yellow colors), that belong to the LV boundary; and two outlier segments (magenta and green colors) that do not belong to the LV boundary. The goal is to perform a correct segmentation in the presence of unreliable structures that may jeopardize the location of the LV.

The PDAF framework starts by considering all the combinations of the segments, i.e., the data interpretations. This is accomplished by giving two possible labels for each segment (i.e., label "1" and "0" if the segment is considered as true or false, respectively) and considering all possible binary label combinations within the segments. In this case, from 16 (2^S) possible data interpretations, only nine are considered for data association.[3] The table in Fig. 6 shows which interpretations are considered.

To better illustrate how important is the weighting (achieved through the computation of the association probabilities) of the segments, we start by illustrating the influence of each segment individually, by considering it as being reliable (i.e., unit association probability). Fig. 6(c) and (d) illustrate the contour segmentation (in red) when the segments S_1 and S_2 are considered individually. We can see that the segmentation is not well represented in the opposite side of the segment, since the model is attracted by the only available data. Now, observing the singly influence of the invalid segments (S_3 and S_4), we see again the "miopia" of the model, since it collapses, scaling and translating towards the available segment.[4] Fig. 6(g) shows the final segmentation using PDAF. It shows that associating the segments with different weights for the data interpretations is the key, since the contour model model is "looking" for the best coherence among segments to form the LV contour. The table in Fig. 6 shows the weights of the nine interpretations considered, and it is possible to conclude that I_{12} ($\alpha \approx 0.84$) is the interpretation that influences the most the final shape estimates, correctly neglecting the presence of the invalid segments S_3 and S_4.

VII. Data Analysis

The experimental data used to evaluate the proposed algorithm consists of 20 echocardiographic volumes with $240 \times 208 \times 256$ voxels.[5] The data was acquired with a resolution in x, y, and z of 0.834, 0.822, and 0.726 mm, respectively.

The annotations for the volumes were given by an expert and consist of 2-D manual segmentations of the LV in four equally spaced slices for each volume. The expert performed the manual segmentations twice in different days, enabling us to determine the intrauser variability. The expert contours were also used to evaluate the accuracy of the computer segmentation by comparing both, as will be explained in the next subsection.

Furthermore, several preprocessing techniques were tested to improve the segmentation results, including:

1) A 2-D mask median filter (2DMF)
2) *Total variation* (TV) proposed in [31],
3) *Squeeze Box filter* (SBF) proposed in [32],
4) A 3-D mask median filter (3DMF)

The filters 1), 2), and 3) are 2-D and were applied to each SA slice. Section VIII-B shows the performance of the algorithm using these techniques.

To perform a quantitative evaluation, we used two similarity metrics that are common in the literature for contour comparison. More specifically, we used the Dice similarity coefficient

[2]A 2-D example has been selected for the clearness of the theoretical exposition.

[3]Recall the assumption of the PDAF [11] is that only one measurement is considered as being target originated. This prevents overlapping segments (or patches) to be considered simultaneously as valid. Thus, interpretations containing overlapped valid segments are discarded, as already discussed in Section V-C3.

[4]In this experiment, the LV contour "lives" in the Euclidean similarities space, i.e., translation, rotation, and isotropic scale.

[5]All the LV volumes presented in this paper were acquired at Hospital Universitario Río Hortega (Valladolid).

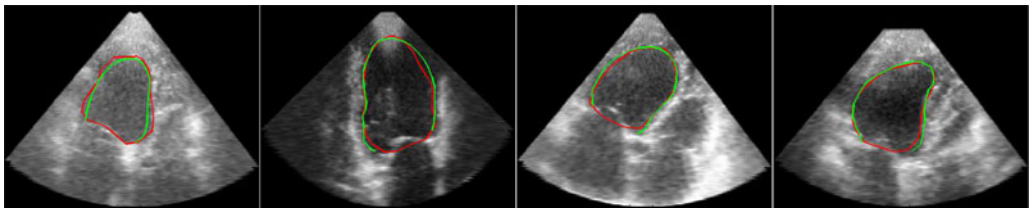

Fig. 7. Slices extracted from four different echocardiographic volumes showing: (in red) the cross-section of the segmentation using the simplex mesh and the TV filter; and (in green) the gold-standard.

(DSC) [33] and average distance (AV) between the estimated contour and the ground truth.

VIII. Experimental Evaluation

In this section, we first discuss the influence of the various algorithm parameters. Then, we present the results and discuss over the qualitative and quantitative evaluation of the obtained segmentations. We compare our results against other state of the art methods and against the intrauser variability. We also compare the segmentations obtained with and without the S-PDAF step. Finally, we assess the robustness of the algorithm to the initialization.

A. Parameter Selection

The set of parameters included in the proposed algorithm were all tuned by performing several tests and selecting the values that led to the best results. The set of parameters tuned were the following: c_1, c_2 and c_3 for (1), and t_l for (4). Different values were tested for each parameter, including: $c_1, c_2, c_3 \in [0, 1]$ and $t_l \in [5, 17]$ mm.

Regarding the parameters in (1), our tests showed that higher values of c_1, c_2 and c_3 occasionally led to instability issues in the surface model. More specifically, as $c_1 \to 1$ (associated with the momentum term), some vertices experienced oscillations issues before converging to the final position, particularly if c_2 was also high. However, this behavior only occurred (noticeably) for $c_1 > 0.5$. The improvement in convergence speed was also negligible, so higher values were discarded. As for c_3 (associated with the external force), high values also caused some occasional instabilities in the surface model, mainly due to the fact that abrupt deformations of the surface can cause it to lose its structure. Consequently, the normal to the surface at some vertices would be inconsistent, eventually leading to incorrect segmentations. Apart from these exceptions, all values led to good segmentations. The best results were obtained with $c_1 = 0.1, c_2 = 0.3, c_3 = 0.2$ for the simplex mesh and $c_1 = 0.1, c_2 = 0.2, c_3 = 0.3$ for the conic model. These were the values used for the remainder tests.

As for t_l, this parameter depends on the number of vertices used in the surface model (since it determines the maximum distance between neighbor low-level features of the same patch). For the case of 20 vertices per slice, experiments showed that $t_l = 10$ mm was the value that led to better overall results.

Summarizing, the set of parameters selected in the following tests were: $c_1 = 0.1, c_2 = 0.3, c_3 = 0.2$ for the simplex mesh; $c_1 = 0.1, c_2 = 0.2, c_3 = 0.3$ for the conic model; and $t_l = 10$ in both cases. The iterative process is stopped when no significant changes are made to the model.

Finally, the parameters P_A and P_B, used in the S-PDAF in (20), also need to be set. P_A defines the minimum prior probability of a valid patch j (when $A_j = 0$), and $1 - P_B$ defines the minimum prior probability of an invalid patch (when $A_j = A_{\max}$). The main purpose of these parameters is to avoid assigning a prior probability of 0 to a patch labeled as valid (or invalid), because this would lead to some interpretations having an association probability of 0. They also define the area for which a patch j has the same prior probability when considered valid and invalid. We chose $P_A = 1 - P_B$ so that the equality happens half way through $A_{\max}$ (in logarithmic scale). We chose $P_A = 0.05$ and $P_B = 0.95$ to have a wider range in prior probability values.

TABLE I

Best Performance for: (a) Each Surface Model, $\mu(\sigma)$, Across all Filters and Initialization Types; Each (b) Each Initialization Type, $\mu(\sigma)$, Across all Models and Filters; and (c) Each Filter, $\mu(\sigma)$, Across all Models and Initialization Types

	Simplex Mesh	Conic
d_{AV}	**2.26 (0.30)**	2.49 (0.93)
d_{DSC}	**0.92 (0.02)**	0.92 (0.04)

(a)

	Automatic	Manual
d_{AV}	**2.26 (0.30)**	2.48 (0.47)
d_{DSC}	**0.92 (0.02)**	0.91 (0.02)

(b)

	None	2DMF	TV	SBF	3DMF
d_{AV}	2.64 (0.74)	2.44 (0.66)	**2.26 (0.30)**	2.34 (0.51)	2.52 (0.69)
d_{DSC}	0.90 (0.05)	0.91 (0.04)	**0.92 (0.02)**	**0.92 (0.03)**	0.91 (0.03)

(c)

B. Segmentation Results

We applied the segmentation algorithm to the 20 volumes described in Section VII. Some examples can be seen in Fig. 7.

For the quantitative results, three comparisons were interesting to evaluate: 1) which deformable model had the best overall performance; 2) which initialization performed better; and 3) which filter led to the best results and how much the accuracy of the segmentation improved. To help the comparison between each of them, Table Isummarizes the best mean performance for each comparison.

Table I (a) shows that the best mean performance of the simplex mesh led to more accurate results, achieving 9% less error than the conic model (with the d_{AV} metric). Comparing the initialization procedures, Table I (b) shows that the automatic approach led to overall best performances. The difference between the two was also of approximately 9%. Note that the standard deviation of d_{AV} also decreased in the automatic approach. As for the different preprocessing techniques, all of them

TABLE II
SUMMARY OF STATE OF THE ART REPORTED RESULTS FOR THE AVERAGE POINT-TO-MESH DISTANCE BETWEEN SEGMENTATION AND EXPERT CONTOURS (IN MILLIMETERS)

	Orderud [36]	Hansegård [20]	Hansegård [21]	Leung [18]	Dikici [35]	Yang [17]	Juang [1]	Proposed
Mean	2.7	3.4	2.2	1.19	2.76	1.28	2.39	2.26
Variance	-	2.3	0.56	0.47	1.08	1.11	3.21	0.40
#Volumes	-	-	-	-	-	1143	4	20

In [36] we averaged the systolic and diastolic results. Dashed entries mean unavailable information.

improved the accuracy of the segmentation. Comparing the best mean values of d_{AV} in Table I (c), we see that the filters 2DMF, TV, SBF, 3DMF, led to an improvement of approximately: 8%, 14%, 11%, and 5%, respectively. The TV filter also achieved a smaller variance in the segmentation accuracy.

The best overall result was obtained in an automatic approach, with the simplex mesh and the TV filter, achieving the values $d_{\mathrm{AV}} = 2.26 \pm 0.30$ mm and $d_{\mathrm{DSC}} = 0.92 \pm 0.02$.

Table II shows a comparison between this result and other state of the art segmentation methods. The proposed approach achieved similar results to the other methods. Besides the methods in Table II, our approach also compares favorably with other recent works. For instance, Zheng *et al.* [34] report an average Dice coefficient of 0.81 (0.16), and an average point-to-mesh distance of 1.25 (1.34) and 2.44 (2.32) millimeters on both of their methods. Lynch *et al.* [35], on the other hand, achieve an accuracy of 1.13 (0.55). However, a fair comparison is not possible because they use different image modalities, namely CT and MRI, respectively.

As a final remark, we should stress that the main reason for the observed performance (as shown in Table II) is that, opposing to other related works, our algorithm partitions the contour into part-of-segments. This leads to a competing (partial) segmentations through the weights (i.e., association probabilities) given to the data interpretation (i.e., combinations of valid/invalid part-of-contours).[6] Each data interpretation embraces a specific part of the contour. Say that, we have a linear combination of Kalman filters, each one specialized in a given interpretation, and acting on given part of the contour. The resulting state update is a linear combination of this bank of Kalman filters. This is the point that makes PDAF robust for the segmentation.

IX. CONCLUSION

This paper presents a 3-D segmentation system that is able to segment the LV in echocardiographic volumes. The methodology is based on the use of a probabilistic data association filter, originally proposed in the context of control theory, to improve the detection of the ventricle boundary. The proposed segmentation system was coupled with an automatic initialization method to produce a fully automatic segmentation procedure that achieved competitive results when compared to a semiautomatic initialization based on space-carving.

The results show that the proposed system performs well in a varied set of echocardiographic volumes. It also shows that the accuracy of the segmentations is competitive with other state-of-the-art works.

[6]For instance in the example of Section VI, the winner/dominant interpretation is I_{12}. However, the other interpretations have also its contribution for the final shape estimates.

ACKNOWLEDGMENT

We would like to thank Hospital Universitario Río Hortega (Valladolid) for providing the echocardiography data used in this study.

REFERENCES

[1] R. Juang, E. McVeigh, B. Hoffmann, D. Yuh, and P. Burlina, "Automatic segmentation of the left-ventricular cavity and atrium in 3D ultrasound using graph cuts and the radial symmetry transform," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2011, pp. 606–609.

[2] T. McInerney and D. Terzopoulos, "Deformable models in medical image analysis," in *Proc. IEEE Math. Meth. Biomed. Image Anal., Workshop*, 1996, pp. 171–180.

[3] A. Hammoude, "Endocardial border identification in two-dimensional echocardiographic images: Review of methods," *Comput. Med. Imag. Graphics*, vol. 22, no. 3, pp. 181–193, 1998.

[4] A. Blake and M. Isard, *Active Contour*. New York, NY, USA: Springer, 1998.

[5] F. Orderud, J. Hansengård, and S. Rabben, "Real-time tracking of the left ventricle in 3D chocardiography using a state estimation approach," in *Proc. Mid. Image Comput. Comput.-Assisted Intervention Conf.*, 2007, vol. 4791, pp. 858–865.

[6] F. Orderud and S. Rabben, "Real-time 3D segmentation of the left ventricle using deformable subdivision surfaces," in *Proc. Comput. Vis. Pattern Recog.*, 2008, pp. 1–8.

[7] J. Hansengård, F. Orderud, and S. Rabben, "Real-time active shape models for segmentation of 3D cardiac ultrasound," *Comp. Anal. Images Patterns*, pp. 157–164, 2007.

[8] G. Jacob, J. Noble, C. Behrenbruch, A. Kelion, and A. banning, "A shape-space-based approach to tracking myocardial borders and quantifying regional left-ventricular function applied in echocardiography," *IEEE Trans. Med. Imag.*, vol. 21, no. 3, pp. 226–238, Mar. 2002.

[9] D. Comaniciu, X. Zhou, and S. Krishnan, "Robust real-time myocardial border tracking for echocardiography: An information fusion approach," *IEEE Trans. Med. Imag.*, vol. 23, no. 7, pp. 849–860, Jul. 2004.

[10] H. Liu and P. Shi, "State-space analysis of cardiac motion with biomechanical constraints," *IEEE Trans. Imag. Proc.*, vol. 16, no. 4, pp. 901–917, Apr. 2007.

[11] Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*. New York, NY, USA: Academic, 1988.

[12] C. Xu and J. L. Prince, "Snakes, shapes and gradient vector flow," *IEEE Trans. Imag. Proc.*, vol. 7, no. 3, pp. 359–369, Mar. 1998.

[13] X. Hang, N. Greenberg, and J. Thomas, "Left ventricle quantification in 3D echocardiography using a geometric deformable model," *Comput. Cardiol.*, pp. 649–652, 2005.

[14] Y. Honggang, M. Pattichis, and M. Goens, "Robust segmentation and volumetric registration in a multi-view 3D freehand ultrasound reconstruction system," in *Proc. Asilomar Conf. Signals Syst. Comput.*, 2006, pp. 1058–6393.

[15] V. Zagrodsky, V. Walimbe, C. Castro-Pareja, J. X. Qin, J.-M. Song, and R. Shekhar, "Registration-assisted segmentation of real-time 3D echocardiographic data using deformable models," *IEEE Trans. Med. Imag.*, vol. 24, no. 9, pp. 1089–1099, Sep. 2005.

[16] L. Yang, B. Georgescu, Y. Zheng, P. Meer, and D. Comaniciu, "3D ultrasound tracking of the left ventricle using one-step forward prediction and data fusion of collaborative trackers," in *Proc. Comput. Vis. Pattern Recog.*, 2008, pp. 1–8.

[17] L. Yang, B. Georgescu, Y. Zheng, Y. Wang, P. Meer, and D. Comaniciu, "Prediction-based collaborative trackers (pct): A robust and accurate approach toward 3d medical object tracking," *IEEE Trans. Med. Imag.*, vol. 30, no. 11, pp. 1921–1932, Nov. 2011.
[18] K. Y. E. Leung, M. G. Danilouchkine, M. van Stralen, N. Jong, A. F. W. van der Steen, and J. Bosch, "Probabilistic framework for tracking in artifact-prone 3D echocardiograms," *Med. Imag. Anal.*, vol. 14, no. 6, pp. 750–758, 2010.
[19] K. Y. E. Leung, M. van Stralen, G. van Burken, N. de Jong, and J. G. Bosch, "Automatic active appearance model segmentation of 3D echocardiograms," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2010, pp. 320–323.
[20] J. Hansegård, S. Urheim, K. Lunde, and S. Rabben, "Constrained active appearance models for segmentation of triplane echocardiograms," *IEEE Trans. Med. Imag.*, vol. 26, no. 10, pp. 1391–1400, Oct. 2007.
[21] C. Tauber, H. Batatia, and A. Ayache, "Quasi-automatic initialization for parametric active contours," *Pattern Recog. Lett.*, vol. 31, pp. 83–90, 2010.
[22] D. Barbosa, T. Dietenbeck, B. Heyde, H. Houle, D. Friboulet, J. D'hooge, and O. Bernard, "Fast and fully automatic 3D echocardiographic segmentation using B-spline explicit active surfaces," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2012, pp. 1088–1091.
[23] C. Santiago, J. S. Marques, and J. C. Nascimento, "A robust deformable model for 3D segmentation of the left ventricle from ultrasound data," in *Proc. Math. Statist. Math. Methodologies Pattern Recog. Mach. Learning*, 2013, vol. 30, pp. 163–178.
[24] K. N. Kutulakos and S. M. Seitz, "A theory of shape by space carving," *Int. J. Comput. Vis.*, no. 38, pp. 199–218, 2000.
[25] H. Delingette, "General object reconstruction based on simplex mesh," *Int. J. Comput. Vis.*, no. 32, pp. 111–142, 1999.
[26] H. V. Trees, *Detection, Estimation and Modulation Theory*. New York, NY, USA: Wiley, 2001.
[27] J. S. Marques and M. A. T. Figueiredo, "Image super-segmentation: Segmentation with multiple labels from shuffled observations," in *Proc. IEEE Int. Conf. Image Process.*, 2011, pp. 2849–2852.
[28] D. P. Huttenlocher and S. Ullman, "Recognizing solid objects by alignment with an image," *Int. J. Comput. Visi.*, vol. 5, no. 2, pp. 195–212, 1990.
[29] J. Nascimento and J. Marques, "Robust shape tracking in the presence of cluttered background," *IEEE Trans. Multimedia*, vol. 6, no. 6, pp. 852–861, Dec. 2004.
[30] J. Nascimento and J. Marques, "Robust shape tracking with multiple models in ultrasound images," *IEEE Trans. Imag. Proc.*, vol. 17, no. 3, pp. 392–406, Mar. 2008.
[31] J. M. Sanches, J. C. Nascimento, and J. S. Marques, "Medical image noise reduction using the Sylvester–Lyapunov equation," *IEEE Trans. Imag. Proc.*, vol. 17, no. 9, pp. 1522–1539, Sep. 2008.
[32] P. C. Tay, S. T. Acton, and J. A. Hossack, "Ultrasound despeckling using an adaptive window stochastic approach," in *Proc. IEEE Int. Conf. Image Process.*, 2006, pp. 2549–2552.
[33] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.
[34] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu, "Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using Marginal Space Learning and steerable features," *IEEE Trans. Med. Imag*, vol. 27, no. 11, pp. 1668–1681, Nov. 2008.
[35] M. Lynch, O. Ghita, and P. F. Whelan, "Segmentation of the left ventricle of the heart in 3D+t MRI data using an optimized nonrigid temporal model," *IEEE Trans. Med. Imag*, vol. 27, no. 2, pp. 195–203, Feb. 2008.
[36] E. Dikici and F. Orderud, "Graph-cut based edge detection for Kalman filter based left ventricle tracking in 3D+T echocardiography," *Comput. Cardiol.*, pp. 205–208, 2010.

Carlos Santiago received the Graduation degree in biomedical engineering and M.Sc. degrees from the Technical University of Lisbon, Lisbon, Portugal, in 2009 and 2011, respectively. He is currently working toward the Ph.D. degree at Instituto Superior Técnico, Lisbon.

He is currently affiliated with the Institute for Systems and Robotics, Lisbon. His research interests include statistical image processing, shape analysis, and medical image analysis.

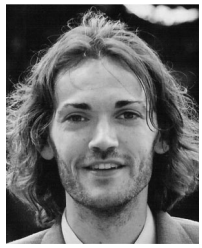

Jacinto C. Nascimento (S'00–M'06) received the E.E. degree from Instituto Superior de Engenharia de Lisboa, Lisbon, Portugal, in 1995, the M.Sc. and Ph.D. degrees from Instituto Superior Técnico, Technical University of Lisbon, Lisbon, in 1998 and 2003, respectively.

Currently, he is an Assistant Professor with the Department of Informatics and Computer Engineering, Instituto Superior Técnico, Lisbon, and a Researcher at the Institute for Systems and Robotics, Lisbon. He has published 30 publications in international journals and 80 in conference proceedings (many of which of the IEEE), has served on program committees of many international conferences, and has been a Reviewer for several international journals. His research interests include statistical image processing, pattern recognition, machine learning, medical imaging analysis, video surveillance, and general visual object classification.

Jorge S. Marques received the E.E. and Ph.D. degrees, and the Aggregation Title from the Technical University of Lisbon, Portugal, in 1981, 1990, and 2002, respectively.

Currently, he is an Associate Professor with the Department of Electrical and Computer Engineering, Instituto Superior Técnico, Lisbon, and a Researcher at the Institute for Systems and Robotics, Lisbon. He has published over 150 papers in international journals and conferences and he is the author of the book "Pattern Recognition: Statistical and Neural Methods" (IST Press, 2005, 2nd ed., in Portuguese). He was the Co-Chairman of the IAPR Conference IbPRIA 2005, President of the Portuguese Association for Pattern Recognition from 2001 to 2003, and Associate Editor of the *Statistics and Computing Journal*, Springer. His research interests are in the areas of statistical image processing, shape analysis, and pattern recognition.