

Published in final edited form as:

Inf Process Med Imaging. 2013 ; 23: 656–667.

Bayesian Segmentation of Atrium Wall Using Globally-Optimal Graph Cuts on 3D Meshes

Gopalkrishna Veni, Zhisong Fu, Suyash P. Awate, and Ross T. Whitaker*

Scientific Computing and Imaging (SCI) Institute, University of Utah

Abstract

Efficient segmentation of the left atrium (LA) wall from delayed enhancement MRI is challenging due to inconsistent contrast, combined with noise, and high variation in atrial shape and size. We present a surface-detection method that is capable of extracting the atrial wall by computing an optimal a-posteriori estimate. This estimation is done on a set of nested meshes, constructed from an ensemble of segmented training images, and graph cuts on an associated multi-column, proper-ordered graph. The graph/mesh is a part of a template/model that has an associated set of learned intensity features. When this mesh is overlaid onto a test image, it produces a set of costs which lead to an optimal segmentation. The 3D mesh has an associated weighted, directed multi-column graph with edges that encode smoothness and inter-surface penalties. Unlike previous graph-cut methods that impose hard constraints on the surface properties, the proposed method follows from a Bayesian formulation resulting in soft penalties on spatial variation of the cuts through the mesh. The novelty of this method also lies in the construction of proper-ordered graphs on complex shapes for choosing among distinct classes of base shapes for automatic LA segmentation. We evaluate the proposed segmentation framework on simulated and clinical cardiac MRI.

Keywords

Atrial Fibrillation; Bayesian segmentation; Minimum s - t cut; Mesh Generation; Geometric Graph

1 Introduction

Segmentation of the heart's left atrium (LA) is a highly relevant problem in the clinical domain. In the context of medical imaging, delayed enhancement MRI (DE-MRI) has been shown to produce contrast in myocardium (heart wall) and in regions subjected to fibrosis and scarring [1]. So, these regions are associated with risk factors and treatment of atrial fibrillation (AF). Imaging with DE-MRI is therefore useful for the evaluation of potential effectiveness of radio-ablation therapy and for studying recovery. This AF recovery includes analysis of scarring as well as atrial shape and structural remodeling (SRM) after treatment. Automatic segmentation of the heart wall in this context is quite important; in a single clinic,

The authors would like to acknowledge the Comprehensive Arrhythmia Research and Management (CARMA) Center, and the Center for Integrative Biomedical Computing (CIBC) by NIH Grant P41 GM103545-14, and National Alliance for Medical Image Computing (NAMIC) through NIH Grant U54 EB005149, for providing Utah fibrosis data, and CIBAVision.

© Springer-Verlag Berlin Heidelberg 2013

hundreds of man hours are spent per month in manual segmentation. In DE-MRI images, this becomes a challenging task, because of relatively low and inconsistent contrast, high level of unwanted texture and noise, and high variability of atrial shape. Moreover, this problem gets aggravated by inaccuracies in cardiac gating and the SRM in chronic AF. Thus, this is a difficult image analysis problem, which also represents an ubiquitous challenge in a 3D medical segmentation: segmenting in the presence of relatively poor signal, high noise, and large variations in shape.

Several papers address the problem of segmenting the blood pool in MRI angiography (MRI-A) images [2, 3]. These methods make use of the relatively homogeneous brightness of the blood pool in MRI-A, which is well suited for deformable models or registration-based approaches. However, high-quality properly-aligned blood-pool images are often *not* readily available from DE-MRI protocols. Furthermore, due to thinness of the atrium wall, algorithms based on template registration fail as they often rely on coarse anatomical features. Figure 1 shows examples of DE-MRI images of the LA that depict its varying, low-contrast boundaries, high level of correlated noise, and high shape variability.

A variety of conventional segmentation methods have proven to be ineffective. One strategy to address these challenges is to introduce a prior on the segmentation problem, either in the form of probability on specific kinds of shapes or more generally on shape properties, such as smoothness. These priors are combined with image matching terms or simply feature detection to find some ideal compromise between the prior and the data. Level-set methods [4] rely on gradient-descent optimizations, which are sensitive to initializations and local minima. We have found such local optimizations to be particularly ill suited to this problem. Statistical models, such as active shape models [5] have been proven to be effective, but are also limited in their ability to deal with the small and large-scale shape variability. Generally, coarse-to-fine optimization strategies can help avoid local minima, but have proven inadequate for this segmentation problem, mostly because the features of interest (thin, brighter regions and small dark gaps between the atrium and nearby tissues) do not hold up under blurring. While recent developments addressing this problem [6] are promising, they rely on deformable models and/or image registration approaches that tend to also get caught in local minima.

The difficulty of segmentation in this context suggests that this problem would benefit from a global optimization strategy. Recently, Wu and Chen [7] described a scheme by which the problem of finding an optimal function value on a discrete grid (a *surface net* problem) is represented as a minimum- s - t cut on a proper-ordered graph. Optimal solutions to the s - t cut are given by relatively efficient, polynomial-time algorithms. Li *et al.* [8] applied a version of this surface-net formulation to simultaneously segment multiple coupled surfaces in noisy images by including image-based costs and geometric constraints of the underlying graph. That approach has demonstrated some success in several challenging image segmentation problems [8–10]. This surface-net relies on the construction of a properly-ordered graph, which also defines the topology of the resulting segmentation. The construction of such graphs is challenging for complex and irregular anatomical structures, such as LA. Using naive offsets from a base mesh results in “tangling” between columns, and resulting cuts are

not guaranteed to be valid surfaces or regions. Thus, these proper ordered graph-cut methods require a careful construction of the underlying graph.

The contributions of this paper are as follows. We show that a Bayesian formulation with a Markov random field prior can give rise to a certain type of surface-net problem, namely, a *VCE-Net*, which is solvable by the algorithm of Wu and Chen [7]. This formulation gives rise to soft penalties on surface smoothing and surface coupling, which, as we will show, is superior to the hard constraints described by Li *et al.*. The Bayesian formulation also gives rise to a set of *learned* feature detectors, so that the method does not rely on user-defined methods for characterizing edges or regions. We also propose a new method for the graph construction on irregular surfaces that avoids tangling. To address the variability in shapes, we process training examples into clusters to form multiple shape templates, that *compete* in our optimization scheme for the best segmentation. We evaluate the method on a set of synthetic examples and LA DE-MRI images with hand segmentations as the ground truth.

2 Methods

2.1 A Bayesian Formulation of Graph-Cut Segmentation

We treat the problem of segmentation as a maximum a-posteriori estimation. The proposed work differs from many previous Bayesian methods in two important aspects. First, we formulate the segmentation as estimation problem on a graph structure, rather than the image directly. Secondly, we obtain a global optimum to this problem by means of a graph-cut algorithm. The data for this formulation is the image data sampled at locations that are associated with the model. The prior is expressed as a Markov random field (MRF) on the location of the cut in the graph which is related to the formulation introduced by Ishikawa [11]. The graph, which forms a 3D mesh, must approximately adhere the shape to be segmented. It introduces a topological structure on the problem over which the Markov property is introduced.

We begin with a description of the graph structure and associated notation. The graph G is a proper-ordered graph with a set of columns, a neighborhood structure on those columns, and a consistent topological structure as one moves up and down the columns. We define the *base graph* $G^0 = (V^0, E^0)$, as a set of vertices $\{v_i^0 \in V^0\}$, and edges, $\{e_{i,j}^0 = (v_i^0, v_j^0) \in E^0\}$. For a proper-ordered graph, the vertices are arranged logically as a collection of (conceptually) parallel columns that have the same number of vertices. The entire graph G consists of an ordered set of copies of the base graph, and each vertex can be referenced by its column i and the position within that column l , e.g. v_i^l . The collection $G^l = (V^l, E^l)$ of vertices and edges at the same position l across all columns is called a *layer*. For ease of notation, an edge or vertex without a superscript, v_i or e_{ij} , is considered with respect to the base layer, which defines the topology of all columns. We let N be the number of columns and L be the number of vertices in each column (number of layers). The neighboring columns of the i th column are denoted as the set \mathcal{N}_i .

Above is the topological structure of the graph; here we describe its geometry. Each node in the graph has an associated position in the 3D volume/image, which we denote as

$x_i^l = x(v_i^l) \in \mathbb{R}^3$. Associated with each x_i^l , there is a *set* of image coordinates, which form an image patch for that vertex, which we call \mathcal{P}_i^l . Associated with each patch is a probabilistic model of the intensity patterns one would find in the image at those locations, which is like the formulation of [5].

We now model a set of image measurements associated with a segmentation on the graph. We introduce a probabilistic model with respect to a single segmentation and extend that to coupled surfaces subsequently. We define the surface segmentation as a subset of nodes in the graph $\mathcal{S} \subset V$. Because we restrict the optimal cut to have only one vertex per column, we can parametrize the cut with respect to the base mesh, thus \mathcal{S} can be represented as the function $S : V^0 \mapsto [0, \dots, L-1]$. Furthermore, $S(i)$, combined with the topology introduced by the base mesh and the 3D coordinates of the vertices describes a surface in 3D. Thus, we are describing a *surface estimation* problem.

For any given vertex in the graph, v_i^l , we can sample the image I as prescribed by the patch \mathcal{P}_i^l . We call the set of image patches for all vertices in the graph as I^V and the set of patches associated with segmentation to be I^S . For a particular segmentation, there is an associated patch I_i^S for each column i .

Now we introduce the probabilistic model, the posterior probability of a segmentation conditioned on image data as follows. Using Bayes rule and considering only terms in the optimization we have:

$$P(\mathcal{S} | I^V) \propto P(I^V | \mathcal{S}) P(\mathcal{S}) \quad (1)$$

Next we introduce specific models. For the image intensity model we assume independence of image patches and use an isotropic Gaussian, with a mean for each column that is learned from a set of training examples. That is,

$$P(I^V | \mathcal{S}) \propto \prod_{i=1}^N P(I_i^S) = \prod_{i=1}^N \exp\left(-\frac{1}{2\sigma^2} \|I_i^S - \mu_i\|^2\right) \quad (2)$$

where μ_i is an average patch template learned for surface with physical locations of column i in training examples, and σ is a standard-deviation parameter associated with this data.

For the surface prior, we use a MRF on the function $S(i)$. Let $\mathcal{C} \subset V^0 \times V^0$ be the set of cliques in the base graph, defined by the neighborhood structure, and $C(S(j), S(k))$ is the pairwise clique potential. We use a Gibbs potential on these cliques for the MRF prior, which gives:

$$P(\mathcal{S}) = \exp\left(-\sum_{(v_j, v_k) \in \mathcal{C}} C(S(j), S(k))\right), \quad (3)$$

where the clique potential $C(\cdot)$ typically takes the form $f(|S(j) - S(k)|)$. Here f is monotonic and convex (for optimization to be feasible). In this paper, we use $f(d) = \alpha d^{1+\gamma}; \gamma > 0$.

We minimize the negative log posterior to get the optimal segmentation as:

$$\operatorname{argmin}_{\mathcal{S}} \left[\lambda \sum_{i=1}^N \|I_i^S - \mu_i\|^2 + \sum_{(v_j, v_k) \in \mathcal{C}} C(S(j), S(k)) \right], \text{ where } \lambda = 1/(2\sigma^2). \quad (4)$$

Segmentation of LA wall requires extraction of epicardial and endocardial surfaces. So, we extend the model to two surfaces/segmentations, $\mathcal{S}^1, \mathcal{S}^2$:

$$P(\mathcal{S}^1, \mathcal{S}^2 | I^V) \propto P(I^V | \mathcal{S}^1, \mathcal{S}^2) P(\mathcal{S}^1, \mathcal{S}^2) \quad (5)$$

We use the same independence assumption with different mean patches for the different surfaces. As we use the MRF for intra-surface smoothness, we propose an inter-surface probability to model interactions between surfaces.

$$P(\mathcal{S}^1, \mathcal{S}^2) = \exp \left(- \sum_{(v_j, v_k) \in \mathcal{C}} C(S_j^1, S_k^1) \right) \exp \left(- \sum_{(v_j, v_k) \in \mathcal{C}} C(S_j^2, S_k^2) \right) \exp \left(- \sum_{j=1}^N g(S_j^1 - S_j^2 - \Delta_j) \right), \quad (6)$$

where Δ_j is the ideal inter-surface distance, which may vary with column and learned from training examples, and $g(S_j^1 - S_j^2 - \Delta_j)$ must meet the same conditions of $f()$ in the clique penalty, but must also enforce $S_j^1 < S_j^2$. For this work we use

$$g(d) = \begin{cases} \alpha' d^{1+\gamma'} & d > -\Delta_j \\ \infty & d \leq -\Delta_j \end{cases} \quad (7)$$

The optimization problem for coupled surfaces is therefore:

$$\operatorname{argmin}_{\mathcal{S}^1, \mathcal{S}^2} \left[\lambda \sum_{i=1}^N \left(\|I_i^{S^1} - \mu_i^1\|^2 + \|I_i^{S^2} - \mu_i^2\|^2 \right) + \sum_{(v_j, v_k) \in \mathcal{C}} (C(S^1(j), S^1(k)) + C(S^2(j), S^2(k))) + \sum_{j=1}^N g(S_j^1 - S_j^2 - \Delta_j) \right] \quad (8)$$

2.2 Graph Cut Formulation

From the objective functions in the previous section, we now construct a revised graph and define an optimal graph cut that is equivalent to the above optimization. The construction of the derived graph follows, generally, the method proposed by [7] for converting this optimization into an s - t cut. Wu *et al.* [7] detail general strategies for solving *surface-net* problems of the type described by Eq. 8. They describe both the Vnet problem, which imposes hard constraints on inter-column behavior and the VCEnet problem, which allows for soft penalties. Previous work including [9, 10] shows the use of the Vnet solution for image segmentation. The Bayesian formulation in the previous section leads to a VCEnet problem, which we also extend to coupled surfaces.

We now briefly review the conversion to the graph-cut problem. The weights on vertices and edges on the extended graph are denoted by $w(v)$ and $c(e)$, respectively. Every vertex in the base layer is connected by a directed edge with a cost $+\infty$ to every other base vertex in its adjacent (neighboring) columns. This makes the base layer strongly connected. For each

vertex in layer $l \in [1, L - 1]$, a weight of $w_i^l = c_i^l - c_i^{l-1}$ is assigned. A directed edge $e_{i,j}^{l,l-1}$ with a cost $+\infty$ is let from that vertex to the one below it.

The MRF property is incorporated as follows. For every pair of adjacent columns in G , a sequence of directed edges, $e_i^{l,l-d}$, $d = \{l, \dots, 0\}$ go from a vertex v_i^l in i -column to vertices v_i^{l-d} for all $j \in \mathcal{N}_i$, as shown in Figure 2(a). For notational convenience we first define an intermediate function to edge weights

$$q(e_{i,j}^{l,l-d}) = f(d), d=0, \dots, l, \quad (9)$$

where $f(d)$ is the penalty, which derives from the clique potential, on the difference in the “height” of adjacent cuts. The weights on these edges are defined through a finite-difference scheme for second derivatives (along columns) of q :

$$w(e_{ij}^{l,0}) = q(e_{ij}^{l,1}) - q(e_{ij}^{l,0}) \quad (10)$$

$$w(e_{ij}^{l,m}) = q(e_{ij}^{l,m+1}) + q(e_{ij}^{l,m-1}) - 2q(e_{ij}^{l,m}), m=1, \dots, l-1. \quad (11)$$

For the penalty on *inter-surface* distance, we extend the method of [8] to the VCENet construction. We construct two identical disjoint subgraphs, using the procedure above, one for each surface. In addition, a set of directed arcs are added between a pair of subgraphs such that the consistency is maintained between a pair of mutually interacting surfaces. To achieve this interaction, we include a set of arcs between corresponding columns of two subgraphs which are penalized by soft constraints. The formulation resembles the one above; however, all edges are between corresponding columns in the two subgraphs. For ease of notation, all references to vertices associated with the second/inner surface will have a hat (i.e., $\hat{\cdot}$). So, v_i^l and \hat{v}_i^l are corresponding vertices on the two subgraphs. We denote edges between the two surface graphs with a \sim .

Part of our design for this segmentation problem is that one surface should always lie inside the other surface (or “below”, if we imagine all columns standing vertically). To achieve this, we include a directed edge between graphs, $\tilde{e}_{i,i}^{l,l+1} = (v_i^l, \hat{v}_i^{l+1})$ with weight $w(e_{i,i}^{l,l+1}) = +\infty$. Similarly, we construct a set of weighted edges that capture the second derivative of the inter-surface penalty when the inner/outer constraint is met as shown in Figure 2b.

$$\tilde{e}_{i,i}^{l,l+\Delta_i+d} = (v_i^l, \hat{v}_i^{l+\Delta_i+d}), d = -\Delta_i+1, \dots \quad (12)$$

$$\text{and } w(\tilde{e}_{i,i}^{l,l+\Delta_i+d}) = g(d+1) - 2g(d) + g(d-1). \quad (13)$$

Subsequently, we obtain optimal segmentation of coupled surfaces by finding a minimum s -excess set in the derived graph, as described in Wu *et al.* [7]. This minimum s -excess set is computed by applying a minimum s - t cut in the transformed graph, G_{st} .

2.3 Building a Valid 3D Mesh

In the previous section, we described topology of the underlying graph based on a triangle structure per each layer. Here we describe the assignment of 3D positions to mesh vertices and triangulation of each layer so that these layers form a nested set of watertight meshes in 3D. This complete collection including a set of vertices, their 3D positions, and the prismatic topology of the nested meshes form a proper-ordered (PO) mesh.

For constructing the PO-mesh, we use an extension of the dynamic-particle-system method proposed by Meyer *et al.* [12]. This method computes thin-layers of triangular prisms that conform to shapes. A mesh is built using a template shape (described in the next section), which approximates the LA that we intend to segment. This template shape is represented as the zero level-set of a signed distance transform in the volume. So the following paragraph describes how to generate layers of high-quality meshes on top of this template.

The meshing strategy uses a cluster of points called particles. These particles are distributed on an implicit surface by interactively minimizing a potential function. The potential function based on pairwise distances defines a repulsive interaction between particles as,

$U_{i,j}^{l,l} = \Phi(|\mathbf{x}_i^l - \mathbf{x}_j^l|)$. We denote the sum of this collection of repulsive potentials within each layer as \mathcal{R} . These particle systems have been shown to form consistent, nearly regular packings on complex surface [12]. Once points have been distributed on an implicit surface (with sufficient density), a Delaunay tetrahedralization scheme can be used to build a watertight triangle mesh of the surface [13].

To build a nested set of surface meshes, we require a collection of offset surfaces, both inside and out, that not only inherit the topology of the base surface, but also represent valid, watertight 3D triangle meshes. This is crucial, because the cuts, which pass through vertices from different layers, must also form watertight triangle meshes. Thus, it results to bend the columns in order to avoid tangling of columns/triangles as the layers extend outward from the mean shape. For this, we introduce a collection of particle systems, one for each layer in the graph/mesh, and we couple these particles by an attractive force (Hooks law) between

layers. Thus, there is an additional set of potentials of the form $U_{i,i}^{l,l+1} = |\mathbf{x}_i^l - \mathbf{x}_i^{l+1}|^2$, and we denote the sum of the attractive forces of neighboring particles between layers as \mathcal{A} .

To optimize an ensemble of particle systems for L layers, we perform gradient descent, using asynchronous updates, as in [12], on the total potential $\mathcal{R} + \beta\mathcal{A}$. Figure 3a illustrates a nested 3-layered mesh for one of the LA templates. The parameter β controls the relationship between attraction across layers and repulsion within layers and is tuned to prevent tangling. For this paper, we have used $\beta = 10$. The optimization requires an initial collection of particles. So, we place a particle at each point where the adjacent voxels have values on either side of the level set. This gives an average density of approximately one particle per unit surface area (in voxel units). The physical distance between layers must be

inversely proportional to the particle density within layers. This is a compromise between the tangling that results from large offsets and the extra computation associated with many thin layers. Since a good mesh constraints the topology and the set of possible segmentations, we try different meshes based on the assumption that all good segmentations can be represented as spatially varying offsets of a mean. This corresponds to around 14,000 particles per each mesh layer for heart images and 2000 particles for simulated images. We have used a total of 30 layers, spaced at 0.5 pixels each, which gives each template a capture range of approximately 15 pixels.

2.4 Learning Template Meshes and Feature Detectors

Here we describe the construction of template shapes and the mechanism for computing costs on nodes from input images. The shapes of LA in the context of AF are highly variable. To address this, we rely on a *training set* of presegmented images. For this paper, the training set consisted of 32 segmented DE-MRI images of the LA. The work in this paper represents a prototype, and we anticipate a production-scale system that relies on hundreds of training images. These training images enable two things. First, training images give us a way of constructing a collection of PO-graphs, so that new images can be segmented as cuts through one of these graphs. Second, training images give us examples of patch profiles for the features that define epi- and endocardial surfaces, which leads to the costs at each node in the PO-graph.

We begin by clustering the examples based on their shapes. For this, we compute distance transforms of each endocardial surface. Training images are aligned via translation to ensure common center of mass for the blood pool (region bounded by the endocardium). This demands careful manual initialization of a template which will be handled in our future work by inducing other transformations. We then compute clusters using k-means using mean-squared distance metric between volumes. Based on the cluster residual curve, 5 clusters are chosen. However, one of the clusters has been removed from the test, because it contained only two (high distorted) examples. Surface meshes associated with the distance-transform means of these four clusters are shown in Figure 3b.

The cost associated with each vertex reflects the degree to which that vertex is a good candidate for a boundary, which will be found via a graph cut. At each vertex, the training data is used to derive a patch profile along a line segment, or *stick* perpendicular to the surface. We sample the stick at a spacing of one voxel. In our case, a patch size of 11 is considered along the normal direction of the surface. The intensity along each stick on each vertex of each template is computed by a weighted average of intensities of sticks for each feature point in each training image. Thus, for a particular vertex in a particular cluster, the intensities along a stick would correspond to an average of several hundreds of neighboring sticks from different images (that share the same blood-pool center). Thus the average stick at a vertex would be an isotropic Gaussian weighted average of all the nearby sticks (within the cluster) with standard deviation of 2 pixels. Figure 3c shows a diagram of the stick configuration and several stick intensity profiles for parts of a particular template.

3 Experiments and Results

For validation, we apply Bayesian framework based graph cut method on 100 simulated images of size $64 \times 64 \times 96$ voxels, and 30 DE-MRI images of the left atrium of size $400 \times 400 \times 107$ voxels. In all of our experiments, 30 mesh layers were generated, spaced at 0.5 voxels each, which gives each template a capture range of approximately 15 voxels. The scaling and exponential parameters, α and γ , for the convex function f of the graph's smoothness penalties are fixed as 300 and 2, respectively. The corresponding values for the function g of the graph's inter-surface penalties are set to 2 and 2, respectively. The values of these parameters reflect the complexity of surfaces and the inter-surface coupling between them.

To segment a given test image, we depend on the user input to position the template. The algorithm is robust to this position, as long as the nested mesh, corresponding to the template, does not lie outside or inside the desired surface (e.g. ± 5 voxels). We sample the input image along all of the sticks at all nodes. Then, we compute a posterior probability on each test stick with the corresponding template stick. This results in the assignment of costs, weights, edge capacities, and then an optimal cut. Likewise, we employ all of the learned templates to the input image, choosing the segmentation that produces the best average probability with the local intensity models for the optimal cut. A pair of optimal mesh surfaces are then recovered from the computed minimum s - t cut. Based on the extracted topological mesh structure, defined by the cut, it is scan converted to reproduce segmented volume(s).

In case of the simulated data, 30 training datasets and 100 test sets were considered for analysis. All these images include two oblong non-crossing surfaces with the inner surface translated randomly (Gaussian distribution) in 3D to mimic variations in heart-wall thickness; each image was corrupted with Rician noise ($\sigma = 30$ for the underlying Gaussian model) and a smoothly-varying bias field. Figure 4 illustrates the effectiveness of the proposed method in extracting smoother boundaries for outer and inner surfaces as compared to hard penalties.

We evaluated the segmentation accuracy for LA based on leave-one out strategy for a test dataset, against templates from the training data. We compared the segmented boundaries of epicardial and endocardial surfaces using our method to that of hard constraints. Since the geometric constraints and soft penalties in the proposed graph cut formulation are analogous to the energy based formulation in deformable models [4], we compared our results with level set based methods. Figure 5a presents segmentation boundaries for epicardial and endocardial surfaces obtained by the proposed algorithm along with others. The cost function image, derived from a-posterior probability, creates a platform on which graph cuts work. Figure 5b illustrates our segmentation result on cost function image corresponding to the epicardial surface. The rationale behind presenting this result is to show how the algorithm is able to extract smoother and accurate boundaries in some areas of the image where even the costs, which are derived from the sophisticated feature detector, could not be defined properly.

The qualitative comparison between the proposed method with others clearly indicates that our method surmounts other techniques in not only extracting correct surfaces, but also in maintaining smoothness along the surfaces and consistency in between them. The irregularity in the surfaces that we notice due to the hard constraints were greatly eliminated.

To evaluate the segmentation accuracy quantitatively, we used distance metric. The distance metric is based on the aggregate of pairwise distances between corresponding points on the ground truth and our segmentation. For each point on our segmented surface, we measure the distance to the nearest point on the ground truth; and vice-versa. For a perfect delineation of the boundary, all these distances would be zero. In the case of simulated examples, this distance metric which was computed over all the images came out to be 0.1879 voxels for the outer surface and 0.2639 voxels for the inner surface. For LA data, we obtained this metric value of 2.5068 voxels for epicardium and 2.6321 voxels for the endocardium. This indicates that the segmentations acquired by the proposed method lie very close to the ground truth.

For quantitative comparison, we studied Dice measures on heart wall using soft against hard constraints. The Dice metric provides the percent overlap between the ground truth and segmented regions. Figure 6 shows the histogram of Dice measures. In both simulated as well as LA cases, the metric values by inducing soft penalties on geometric constraints overpowered hard penalties. For synthetic data, the Dice values indicate excellent matches. However, in the case of myocardium, the dice values are little lower due to its varying thickness (2–6 mm) and undefined ground truth. The ground truth is a single hand segmentation from an expert. Therefore, much of the observed error is near the veins, which are subject to inter-rater variability, as the cutoff between atrium and vessel is not well defined. Also the ground truths for the wall do not form a complete boundary around the blood pool (even ignoring the vessels). Furthermore, we expect the improvement in results by increasing the number of training images so that more templates are formed in order to better match a given input image.

References

1. McGann CJ, Kholmovski EG, Oakes RS, Blauer JJ, Daccarett M, Segerson N, Airey KJ, Akoum N, Fish E, Badger TJ, DiBella EV, Parker D, MacLeod RS, Marrouche NF. New magnetic resonance imagingbased method for defining the extent of left atrial wall injury after the ablation of atrial fibrillation. *J. American College of Cardiology*. 2008; 52(15):1263–1271.
2. John, M.; Rahn, N. Automatic left atrium segmentation by cutting the blood pool at narrowings. In: Duncan, JS.; Gerig, G., editors. MICCAI 2005. LNCS. Vol. 3750. Heidelberg: Springer; 2005. p. 798-805.
3. Karim R, Mohiaddin R, Rueckert D. Left atrium segmentation for atrial fibrillation ablation. *SPIE Conference Series*. 2008; 6918
4. Malladi R, Sethian J, Vemuri B. Shape modeling with front propagation: a level set approach. *IEEE Transactions on PAMI*. 1995; 17(2):158–175.
5. Cootes T, Taylor C, Cooper D, Graham J. Active shape models-their training and application. *Computer Vision and Image Understanding*. 1995; 61(1):38–59.
6. Gao Y, Gholami B, MacLeod R, Blauer J, Haddad W, Tannenbaum A. Segmentation of the endocardial wall of the left atrium using local region-based active contours and statistical shape learning. *SPIE*. 2010; 7623

7. Wu, X.; Chen, DZ. Optimal net surface problems with applications. In: Widmayer, P.; Triguero, F.; Morales, R.; Hennessy, M.; Eidenbenz, S.; Conejo, R., editors. ICALP 2002. LNCS. Vol. 2380. Heidelberg: Springer; 2002. p. 1029-1042.
8. Li K, Wu X, Chen D, Sonka M. Optimal surface segmentation in volumetric images-a graph-theoretic approach. PAMI. 2006; 28(1):119–134.
9. Dou, X.; Wu, X.; Wahle, A.; Sonka, M. CVPR. IEEE; 2008. Globally optimal surface segmentation using regional properties of segmented objects; p. 1-8.
10. Song, Q.; Wu, X.; Liu, Y.; Smith, M.; Buatti, J.; Sonka, M. Optimal graph search segmentation using arc-weighted graph for simultaneous surface detection of bladder and prostate. In: Yang, G-Z.; Hawkes, D.; Rueckert, D.; Noble, A.; Taylor, C., editors. MICCAI 2009, Part II. LNCS. Vol. 5762. Heidelberg: Springer; 2009. p. 827-835.
11. Ishikawa H. Exact optimization for markov random fields with convex priors. PAMI. 2003; 25(10):1333–1336.
12. Meyer M, Kirby R, Whitaker R. Topology, accuracy, and quality of isosurface meshes using dynamic particles. IEEE TVCG. 2007; 12(5):1704–1711.
13. Amenta N, Bern M, Eppstein D. The crust and the beta-skeleton: Combinatorial curve reconstruction. Graphic Models Image Proc. 1998; 60(2):125–135.





Fig. 1.
Slices of left atrium DE-MRI images showing the challenges in segmentation

**Fig. 2.**

(a) Inter-column arcs. (b) Inter-surface arcs. Blue arrows from column i_1 to i_2 represent arcs subjected to ideal inter-surface distance, i .




Fig. 3.

(a) An example of several layers of PO-meshes for the LA. (b) Examples of average shapes, derived from k -means clustering on distance transforms of training images, around which the PO-meshes are constructed. (c) A mock up of a simplified PO-mesh in 2D with examples of feature detectors learned from the training data— actually PO-meshes for the LA have over 400,000 vertices.




Fig. 4.

Segmentation boundaries for outer and inner surfaces on synthetic data corresponding to ground truth, graph cuts with hard constraints and the proposed algorithm




Fig. 5.

(a) Surface boundaries of epi and endocardial surfaces corresponding to ground truth, level sets, graph cuts with hard constraints and proposed method. (b) Segmentation result of epicardial surface using the proposed algorithm (red) and ground truth boundary (green) overlaid on corresponding cost function image.

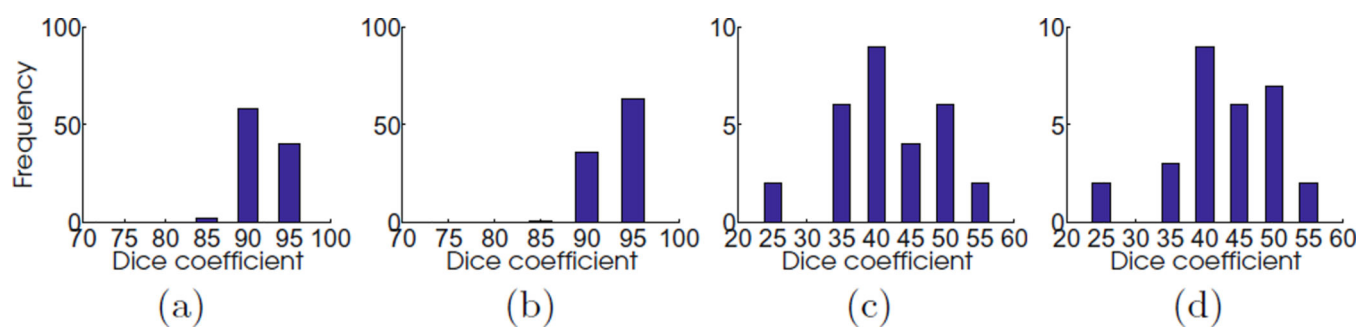


Fig. 6. Histogram of dice coefficients for the (a) middle region, graph cuts with hard constraints, (b) middle region, soft constraints, (c) heart wall, graph cuts with hard constraints, (d) heart wall, soft constraints