

Comparing Appearance and Edge Information for Myocardial Segmentation of the Left Ventricle in 3D Echocardiography using Graph Cuts

Michael Verhoek^{a,*}, Kashif Rajpoot^a, Andrew Blake^b and J. Alison Noble^a

^aInstitute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford

^bMicrosoft Research Cambridge

Abstract. Fast semi-automatic segmentation of the myocardium in 3D echocardiograms may be useful for diagnosis of heart diseases, by facilitating quantification of wall thickening and local wall motion. Segmentation in ultrasound is challenging, especially 3D myocardial segmentation, due to low signal-to-noise ratio, different tissues having a similar appearance and the large amount of data. In this paper, we employ a semi-automatic graph cuts algorithm to efficiently segment the myocardium. Several types of input information are used: appearance information in the form of intensity and a tissue characterising probability distribution, boundary information in the form of an edge detector, and information on the position of the left ventricle. We apply our algorithm to 11 3D echocardiograms, while varying the importance of region and boundary information, and testing several combinations of input information. The results are compared to expert manual segmentation. We demonstrate that using a tissue characterising distribution and positional information is beneficial to the performance, that edge measures are only beneficial in some cases, and in general that a fast and relatively accurate segmentation can be obtained, with a possible runtime of the order of 30 s and a true positive rate of 80% at a false positive rate of only 6.6%.

1 Introduction

Fast myocardial segmentation in 3D echocardiograms can be useful for assessing cardiac diseases, both visually for the benefit of the clinician, and quantitatively for estimating (local) wall motion and thickening. Ultrasound (US) images can however be challenging to segment, due to attenuation, missing boundaries, low signal-to-noise ratio and different tissues that have a similar appearance. Working in 3D also means a large amount of data has to be processed fast, in order to make optimal use of one of the advantages of US over other modalities, its speed.

Previously, there have been few published papers on myocardial segmentation [1] [2] [3]. As far as we are aware, this paper is the first application of graph cuts to 3D echocardiography. Graph cuts have been applied previously to segmentation of the heart in other imaging modalities [4]. Segmentation of 3D echocardiographic images has been explored using a range of methods as reviewed in [5].

In this paper we use the graph cuts technique to segment the myocardium in 3D echocardiograms. Graph cuts [6] are a promising fast segmentation methodology in computer vision, that is applied increasingly to biomedical images. In our implementation, we employ three types of input: appearance information like intensity and tissue characterisation by making use of the shifted Rayleigh distribution; edge information in the form of a local phase feature asymmetry detector; and positional information. We minimise the amount of user input, restricting this to three mouse clicks to define the position of the left ventricle (LV). The results are referenced against manual expert delineations of the myocardium, to be able to compare the performance of each of the types of input information. In the next section, we first describe the data available for this study and provide a brief explanation of the theory of graph cuts and of the setup of the experiments performed. Then we discuss the results of our experiments.

2 Methods

2.1 Data sets

We obtained 11 3D+time echocardiograms from healthy subjects. These echocardiograms were recorded at the John Radcliffe Hospital in Oxford with a Philips iE33 ultrasound system. In each time sequence, we chose to use the end-systolic frame. These 11 3D images each had a size of $224 \times 208 \times 208$ voxels and a spatial resolution of $0.93 \text{ mm} \times 0.94 \text{ mm} \times 0.85 \text{ mm}$ or $0.82 \text{ mm} \times 0.84 \text{ mm} \times 0.76 \text{ mm}$, depending on the settings of the ultrasound system. A manual segmentation was made for each echocardiogram, where an endocardial and an epicardial contour for the left and right ventricles was delineated by an expert. This was done on every fifth short-axis slice, using an interactive pen

*Corresponding author: Michael Verhoek, mverhoek@robots.ox.ac.uk

display (Wacom, USA). Papillary muscles were annotated as blood pool. Where myocardium tissue appeared to be missing because of shadow or attenuation, the expert used their knowledge of the shape of the myocardium to make the delineation. The expert was able to view orthogonal long-axis slices while segmenting to assess the delineations in those views.

2.2 Graph Cuts

Our discussion of the theory of the graph cuts technique will be brief and the reader is referred to [6] for a more detailed discussion. The basic idea is to represent an image by a graph $G = (V, E)$, where the nodes are $V = \{s, t\} \cup P$. P is the set of nodes that represent the voxels, and s and t are auxiliary nodes. Each node $p \in P$ is connected to both s and t via the so-called t-links and each node p is also connected to neighbours using a neighbourhood system N , these links are n-links. Both types of links make up E . All links are assigned a weight according to an energy function, and the graph is cut in two, so that each node p is still connected to either s or t . The algorithm finds the cut with the minimum cost (the sum of link weights cut), which corresponds to the global energy minimum of the segmentation, assuming the submodularity condition is satisfied, as explained in [7]. For segmentation purposes, the energy function is often given as

$$E(A) = \lambda R(A) + B(A), \quad (1)$$

where A is a binary labelling, λ is the ratio of importance between both terms, $R(A)$ is a region or data term and $B(A)$ is a boundary or regulariser term. R penalises when the label of a voxel does not correspond to the label suggested by the prior model of the data, and B penalises when neighbouring voxels have a different label. Link weights in the graph are directly related to these terms: weights on t-links are determined by R , weights on n-links are determined by B .

In our implementation, we employ the following forms for these energy terms. R is a sum of log-likelihood functions, based on the probability of a voxel belonging to one of the labels. In our case, this is based on the normalised intensity histogram of voxels in foreground (\mathcal{F}) and background (\mathcal{B}), which are compiled before segmentation.

$$R(A) = \sum_{p \in P} R_p(A_p), \quad (2)$$

$$R_p(A_p = \text{foreground}) = -\ln P(I_p | \mathcal{F}), \quad (3)$$

$$R_p(A_p = \text{background}) = -\ln P(I_p | \mathcal{B}), \quad (4)$$

where A_p is the label of voxel p and I_p is the intensity value of voxel p ; this can be the image intensity value or some other appearance measure. B is a modified Ising prior [8]:

$$B(A) = \sum_{\{p,q\} \in N} B_{p,q} |A_p - A_q|, \quad (5)$$

$$B_{p,q} = \alpha + \exp\left(-\frac{(I_p - I_q)^2}{\rho}\right) \frac{\beta}{\text{dist}(p, q)}, \quad (6)$$

where α and β are factors that give the function a more useful shape with respect to R , ρ is a factor describing the scanner noise and dist is a function describing the distance between neighbouring voxels p and q . To solve the minimum energy graph cut, a C++ code written by Boykov and Jolly was used that finds the Max Flow using their version of the Ford-Fulkerson algorithm [9]. Values used throughout our experiments were $\alpha = 150$, $\beta = 10000$, $\rho = 10$ (note that varying α and β together changes the relative importance of the R and B terms, we chose to vary this through λ); neighbourhood system N : 14 neighbours (6 direct neighbours and 8 farthest neighbours, out of the 3^3 nearest neighbourhood).

2.3 Implementation and Experiments

In our experiments, we segmented the myocardium in 11 3D echocardiograms using the graph cuts technique, while varying the ratio λ in eq. (1). Since there was significant variation in brightness between echocardiograms, a histogram equalisation step was performed to reduce this variation. The histogram of each echocardiogram was approximately matched to the histogram of one of the echocardiograms. This pre-processing facilitated the use of the input histograms in the graph cuts algorithm, enabling use of histograms from other images as an input for each segmentation. We used several types of information as input for the algorithm, as described below.

Image intensity. This is the simplest form of information about a region. Using the manual segmentations, for each image we collected image intensity values in voxels labelled as myocardium and as non-myocardium. From these,

normalised foreground and background 14-bin histograms were generated. When segmenting each image, we used the histograms of the 10 other images to make up \mathcal{F} and \mathcal{B} .

Shifted Rayleigh distribution. Several probability distributions to characterise different tissues in US images have been tested [10]. In [1], the use of a shifted Rayleigh distribution was proposed. For each voxel in the image, we estimate the σ^2 parameter, in a neighbourhood of 5^3 voxels around it, using the fast integral image method [11]. From these 11 parametric images we generated a parametric histogram in the same way as the image intensity histograms were generated.

Edge-indicator measure. Edge information was also used as an input. The edges are derived using a local-phase inspired 3D feature detector, called feature asymmetry (FA) measure [12]. The FA measure is designed to detect low-contrast step-like edges (i.e., endocardial or epicardial edges) from echocardiographic images. The FA measure for 3D feature detection uses the monogenic signal [13], which is an isotropic extension of the analytic signal, and is given by:

$$FA(x) = \sum_{sc} \frac{[\text{odd}_{MG}^{sc}(x) - \text{even}_{MG}^{sc}(x)] - T_{sc}}{\sqrt{(\text{odd}_{MG}^{sc}(x))^2 + (\text{even}_{MG}^{sc}(x))^2 + \epsilon}}, \quad (7)$$

where sc represents the scale variable as the FA measure is computed over 3-scales, T_{sc} is a scale-specific threshold parameter, ϵ is a small constant to avoid division by zero, and $[\dots]$ zeros the negative values. Here, $\text{odd}_{MG}^{sc}(x)$ and $\text{even}_{MG}^{sc}(x)$ are the odd- and even-symmetric filter responses at scale sc using the monogenic signal, respectively. This information was implemented into the graph cuts framework, by defining

$$\hat{B}_{p,q} = \begin{cases} B_{p,q} \text{ from eq. (6)} & \text{if } FA(p) \leq \tau \\ 0 & \text{if } FA(p) > \tau, \end{cases} \quad (8)$$

where τ is a threshold and using this $\hat{B}_{p,q}$ rather than $B_{p,q}$ (as defined in eq. (6)). The edge image was on a scale of 0 to 255, but thresholded at $\tau = 130$ to ignore most noise.

Positional information. In [3], training of a classifier using the estimated position of the LV as a feature proved beneficial for the results. In graph cuts, since there is no training step, we decided to employ a simple user input: the user provides an approximation estimate of the location of the centre of the LV and its size, by clicking three times in a 3D image (three orthogonal views were visible). This is translated into a capped cylinder in which the left and right ventricles are to be found. This provides a weighting of the region term in eq. (3): we modify this to be $\hat{R}_p(\text{fg}) = R_p(\text{fg}) \exp(-0.075 \cdot \text{Pos}(x_p))$, where x_p is the position of voxel p within the image. If x_p is inside the cylinder, $\text{Pos}(x_p) = 0$; if x_p is outside it, $\text{Pos}(x_p)$ is the Euclidean distance between the cylinder and x_p .

Experiments. We ran the segmentation algorithm on all 11 images, for 25 values of λ , varying from 1 to 2100. We did this for the normal intensity images, and for the shifted Rayleigh parametric images, without edge or positional information. This was repeated, but now using edge information. Then these four experiments were repeated, now making use of the positional information. The resulting segmentations were collected and compared against the manual segmentations, calculating the number of true positives, false positives, false negatives and true negatives summed over the 11 images, for each value of λ , in each of the experiments.

3 Results and Discussion

In Fig. 1(a), several slices of 3D echocardiograms are shown, presenting the information available (intensity image, manual segmentation, shifted Rayleigh parametric image, position image and edge image) and in Fig. 1(b-d) examples of the results of the segmentation algorithm using the various types of input information. In Fig. 2, we present quantitative results in the form of ROC curves generated from the calculated true and false positives and negatives.

From Fig. 2 it can be seen that the addition of position information improves the segmentation result considerably. Figure 1(b2-3) illustrates that this is caused by the fact that the atrium walls are segmented as well when position information is not available. Also, bright near-field clutter can be seen as myocardium by the segmentation algorithm, which can be prevented to some extent when using position information. From the ROC curves it is also clear that using shifted Rayleigh parametric images rather than normal intensity images improves segmentation. This suggests that in a tissue characterisation measure image the myocardium can be found better than in intensity images. It would be worthwhile to empirically test which distributions would perform best in tissue characterisation, as has been done in [10], but for 3D echocardiograms. Fig. 1(c2-3) also shows that segmentations based on the parametric images yield smoother results than intensity images, this is caused by the smaller amount of noise in the parametric images.




Figure 1. Qualitative results. Example slices illustrating the input images and segmentation results. a1-a5: Sample slices of the input information. Short-axis slices of subject 2; a1. intensity image; a2. manual segmentation (white contour on top of the intensity image); a3. shifted Rayleigh parametric image; a4. edge indicator image; a5. position information (brightness indicates distance from cylinder). b1-b4: Sample segmentation results (white contour on top of the intensity image), long-axis slices of subject 2, $\lambda = 500$; b1. manual segmentation; b2. segmentation result using intensity image; b3. segmentation result using intensity image and position information; b4. segmentation result using intensity image, position and edge information. c1-c4: Sample segmentation results, short-axis slices of subject 3, $\lambda = 2100$; c1. shifted Rayleigh parametric image; c2. segmentation result using intensity image and position information; c3. segmentation result using Rayleigh image and position information; c4. segmentation result using Rayleigh image, position and edge information. d1-d2: Sample slices showing edge information is not used to the fullest extent, short-axis slices of subject 2, $\lambda = 2100$; d1. segmentation result using Rayleigh image and location and edge information; d2. edge indicator image FA.

The use of edge information gives mixed results, as can be seen in Fig. 2: when applied to intensity images, a clear improvement is seen when adding edge information (illustrated in Fig. 1(b3-4)). When applied to shifted Rayleigh parametric images however, no clear improvement can be observed, and the two ROC curves almost coincide (Fig.2, and illustrated in Fig. 1(c3-4)). We are led to believe that the edge information runs into the restrictions that the GC algorithm itself imposes on this type of information. This can be made clear when looking at eq. (5): the boundary penalty only comes into play when neighbouring voxels are labelled *differently* in the current segmentation. In this setup, the algorithm does not look at the edge measure as long as neighbouring voxels have the same label, i.e. it only takes edge measures into account when the intensity image or parametric image gives it reason to do so. However, we want a boundary penalty when neighbouring voxels are labelled *the same* and the edge measure indicates there they are supposed to be labelled differently. One way to improve the handling of an edge measure, would be the introduction of a ‘reversed’ version of eq. (5), i.e. $B(A) = \sum B_{p,q} \delta(A_p, A_q)$, using the Kronecker delta. In future work, we intend to implement this proposal. In Fig. 1(d1) it can be seen that no boundaries are found, even though they are clearly present in the edge image Fig. 1(d2); cases like these would benefit from this reversal of the boundary penalty.

The graph cut algorithm is relatively fast, segmenting a $224 \times 208 \times 208$ voxel image in about 30-60 s on a normal desktop computer (Intel Xeon 3.2 GHz CPU, 2GB RAM). Our code has not been optimised for speed, if we were to do so, run times would probably be halved or smaller. Implementation on a GPU would possibly enable run time in the order of seconds or less. The best possible segmentation results (averaged over all images) are a TPR (true positive rate) of 80% at an FPR (false positive rate) of only 6.6%. This is a good accuracy, especially since myocardial segmentation is more challenging than endocardial segmentation.

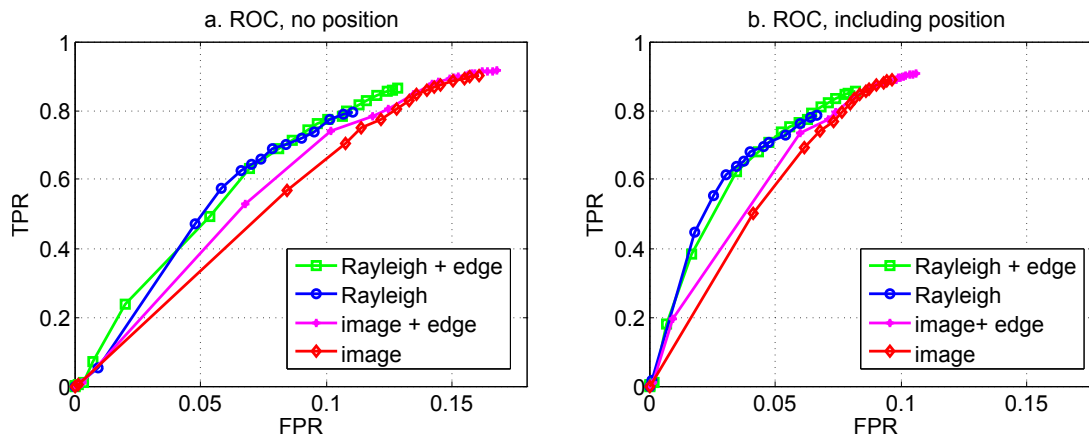


Figure 2. Quantitative results. ROC curves demonstrate an increase in segmentation performance for Rayleigh over intensity image, for using position information, and for using edge information in the intensity image. a. No position information including, b. position information included.

4 Conclusion

We investigated various input measures in a novel graph cuts myocardial segmentation algorithm, applied to 3D echocardiograms. We compared against manual segmentations and found that the results were fast and relatively accurate, with a TPR of 80% at an FPR of only 6.6%. Of the possible input measures, we found that using a Rayleigh parametric image performs better than a normal intensity image, that using a position measure is very beneficial and that adding edge information only improves the result for intensity images; for Rayleigh images, the segmentation is already at its peak performance, but we suggest that the results from using edge measures in general can be improved by changing the way boundary penalties are given within the algorithm. In future work, we intend to compare our results against a delineation using Random Forests, an advanced machine learning approach [3]. The implementation in [3] runs at a similar speed as the method described in this paper.

Acknowledgements

We would like to thank Dr. Cezary Szmigielski and Prof. Harald Becher for providing us with the 3D echocardiograms and their advice on cardiac segmentation. We would also like to acknowledge Dr. Vladimir Kolmogorov and Dr. Yuri Boykov for their MinCut/MaxFlow graph cut solving code. This work was supported by Microsoft Research through its PhD Scholarship Programme.

References

1. Y. Zhu, X. Papademetris, A. Sinusas et al. "Segmentation of myocardial volumes from real-time 3d echocardiography using an incompressibility constraint." In *MICCAI*, pp. 44–51. 2007.
2. M. Lynch, O. Ghita & P. Whelan. "Left-ventricle myocardium segmentation using a coupled level-set with a priori knowledge." *Computerized Medical Imaging and Graphics* **30(4)**, pp. 255–262, June 2006.
3. V. Lempitsky, M. Verhoek, J. A. Noble et al. "Random forest classification for automatic delineation of myocardium in real-time 3d echocardiography." In *Functional Imaging and Modelling of the Heart 2009*. To appear June 2009.
4. Y. Boykov & M.-P. Jolly. "Interactive organ segmentation using graph cuts." In *MICCAI*, pp. 147–175. 2000.
5. J. Noble & D. Boukerroui. "Ultrasound image segmentation: a survey." *IEEE Trans. Med. Imaging* **25(8)**, pp. 987–1010, 2006.
6. Y. Boykov & O. Veksler. *Handbook of Mathematical Models in Computer Vision*, chapter 5 (Graph Cuts in Vision and Graphics: Theories and Applications). Springer, 2006.
7. V. Kolmogorov & R. Zabih. "What energy functions can be minimized via graph cuts?" *IEEE PAMI* **26(2)**, pp. 147–59, 2004.
8. Y. Boykov & M.-P. Jolly. "Interactive graph cuts for optimal boundary & region segmentation of objects in n-d images." In *International Conference on Computer Vision (ICCV), Vancouver, Canada, Vol. 1*, pp. 105–12. Jul 2001.
9. L. Ford & D. Fulkerson. "Maximal flow through a network." *Canad. J. Math.* **8**, pp. 399–404, 1956.
10. Z. Tao, H. Tagare & J. Beaty. "Evaluation of four probability distribution models for speckle in clinical cardiac ultrasound images." *IEEE Trans. Med. Imaging* **25(11)**, pp. 1483–1491, 2006.
11. P. Viola & M. Jones. "Robust real-time face detection." In *International Conference on Computer Vision*. 2001.
12. K. Rajpoot, V. Grau & J. A. Noble. "Local-phase based 3d boundary detection using monogenic signal and its application to real-time 3-d echocardiography images." In *International Symposium on Biomedical Imaging*. 2009.
13. M. Felsberg & G. Sommer. "A new extension of linear signal processing for estimating local properties and detecting features." In *DAGM-Symposium*, pp. 195–202. Springer-Verlag, 2000.