Automatic Cardiac MRI Segmentation Using a Biventricular Deformable Medial Model

Hui Sun[1], Alejandro F. Frangi[2,3], Hongzhi Wang[1], Federico M. Sukno[2,3], Catalina Tobon-Gomez[2,3], and Paul A. Yushkevich[1]

[1] Penn Image Computing and Science Laboratory (PICSL), Department of Radiology, University of Pennsylvania, Philadelphia, USA
[2] Center for Computational Imaging & Simulation Technologies in Biomedicine, Universitat Pompeu Fabra, Barcelona, Spain
[3] Centro de Investigación Biomédica en Red en Bioingeniería, Biomateriales y Nanomedicina (CIBER-BBN), Zaragoza, Spain

Abstract. We present a novel approach for automatic segmentation of the myocardium in short-axis MRI using deformable medial models with an explicit representation of thickness. Segmentation is constrained by a Markov prior on myocardial thickness. Best practices from Active Shape Modeling (global PCA shape prior, statistical appearance model, local search) are adapted to the medial model. Segmentation performance is evaluated by comparing to manual segmentation in a heterogeneous adult MRI dataset. Average boundary displacement error is under 1.4 mm for left and right ventricles, comparing favorably with published work.

1 Introduction

Segmentation of the human myocardium in *in vivo* MRI in a necessary first step for various computational analyses of heart structure and function. Automatic segmentation is a challenge because intensity characteristics of cardiac MRI are complex, and because the myocardium is thin relative to typical MRI voxel size, particularly in the right ventricle (RV). Most existing techniques only segment the left ventricle (LV) (e.g. [11,12,1,6]), or segment either the inner surface (endocardium) or outer surface (epicardium) of the RV [9,17,13]. When characterizing pathology such as myocardial hypertrophy, myocardial infarction, and ventricular arrhythmias, a segmentation of both surfaces in LV and RV is desirable. Examples of dual-surface RV and LV segmentation exist in the literature [10], but there remains substantial room for improvement in accuracy.

When segmenting structures known to be thin *a priori*, it is natural to build thickness constraints into the segmentation algorithm. Standard active shape model (ASM), level set, and registration-based approaches, used widely for cardiac segmentation, do not provide an explicit way of doing so. However, coupled-surface approaches incorporated into the level set framework [16,12] can be used to constrain the distance between endocardial and epicardial surfaces. In this paper, we explore an alternative approach, where a Markovian prior on model thickness is introduced by representing the myocardium using a deformable *medial*

model. The medial model explicitly represents the skeleton of the myocardium as well as myocardial thickness. The boundary of the model (epicardial and endocardial) is derived from the skeleton and thickness data using a simple analytical expression. A parameterization of the volumetric region enclosed by the model is also easily derived. Thus, the model can explicitly represent constraints and prior knowledge regarding myocardial thickness, and there is no possibility of endocardial and epicardial surfaces crossing during model deformation.

Our approach builds on earlier work [14] that demonstrated the feasibility of capturing and analyzing myocardium shape using the *continuous medial representation (cm-rep)* with branching skeletons (Fig. 1). This paper makes the step from shape modeling to segmentation by developing, from a large training set, a shape and thickness prior, which combines standard PCA for shape with a Markov model for thickness; and by incorporating into the cm-rep framework a statistical learning-based appearance model derived from the Active Shape Model (ASM) framework [4,3]. This leads to excellent segmentation accuracy on short-axis cardiac MRI from a heterogeneous patient population.

2   Materials and Methods

2.1   Subjects, Imaging and Manual Segmentation

The primary dataset in this study, used for appearance learning and evaluation, consists of data from 40 subjects in four clinical cohorts: normal controls (n=10), myocardial infarction (MI) (n=10), hypertrophy (n=10) and dilation (n=10). Short-axis cardiac MRI (TR/TE=2.9/1.2 ms, FA=45°, in-slice resolution = $1.56\,\text{mm} \times 1.56\,\text{mm}$, slice thickness= 8 mm, no gap, FOV = $400\,\text{mm} \times 340\,\text{mm}$) were acquired at CETIR Sant Jordi Centre (Barcelona, Spain) using a GE Signa CVi-HDx 1.5T scanner (GE Healthcare, Milwaukee, USA). Author XX manually outlined the endocardial and epicardial surfaces of the LV and RV at end diastole, forming a volumetric two-chamber heart representation that typically spans 8-12 slices from the base to the apex. A secondary dataset of manual segmentations derived from similar short axis MRI of 81 subjects with a variety of common cardiovascular pathologies was used in the construction of the shape and thickness priors.

2.2   Medial Models

The medial model is defined as a pair $(\mathbf{m}, R)$, where $\mathbf{m}$ is the skeleton, i.e., a collection of adjoining co-dimension 1 manifolds in $\mathbb{R}^3$, and R is a positive scalar field defined on $\mathbf{m}$, called the radial field. In practice, a triangular mesh $(\mathbf{m}, R; \mathcal{E})$ is defined by applying Loop [8] subdivision rules (modified to handle junctions between manifolds) to a triangular mesh of control points $(\mathbf{m}_{\text{ctl}}, R_{\text{ctl}}; \mathcal{E}_{\text{ctl}})$, where $\mathcal{E}, \mathcal{E}_{\text{ctl}}$ denote the edges in these meshes. Edges in the medial mesh (either control or subdivided) may belong to one, two or three triangles. The set of edges with three adjacent triangles is called the *medial seam*, and the set of edges adjacent
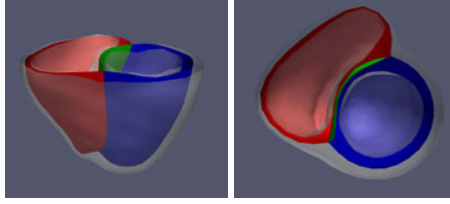
Fig. 1. The medial model of the myocardium that is used as a deformable template in our method. The skeleton of the model is partitioned into red, blue, and green surfaces corresponding to the RV wall, LV wall, and interventricular septum. The translucent gray surface is the boundary of the model, i.e. epicardium and endocardium.

to just one triangle is called the *medial edge*. Below, we treat $(\mathbf{m}, R)$ as smooth manifolds, although in practice they are approximated by discrete meshes.

The *boundary* corresponding to the pair $(\mathbf{m}, R)$ is defined as a set of points $\mathcal{B}$, such that the Blum medial axis transform (MAT) [2] of $\mathcal{B}$ is $(\mathbf{m}, R)$. Generally, $\mathcal{B}$ cannot be described analytically as a function of $(\mathbf{m}, R)$, as MAT may map whole patches of $\mathcal{B}$ to a single point on $\mathbf{m}$. However, if $(\mathbf{m}, R)$ satisfy a set of conditions given in [15], such an analytical description is possible. It is given by $\mathcal{B} = \mathbf{b}^+ \cup \mathbf{b}^-$, where $\mathbf{b}^+$, $\mathbf{b}^-$ are surfaces that lie on the opposite sides of $\mathbf{m}$:

$$\mathbf{b}^{\pm} = \mathbf{m} + R\left[-\nabla_{\mathbf{m}}R \pm \sqrt{1 - \|\nabla_{\mathbf{m}}R\|^2} \cdot \mathbf{N_m}\right] , \qquad (1)$$

where $\nabla_{\mathbf{m}}R$ is the Riemannian gradient of R on $\mathbf{m}$, and $\mathbf{N_m}$ is the unit normal vector to $\mathbf{m}$. The sufficient conditions for $\mathcal{B} = \mathbf{b}^+ \cup \mathbf{b}^-$ given in [15] include equality conditions that must hold along medial edges and seams, as well as inequality conditions that hold everywhere on the skeleton. For brevity, we do not discuss them here in detail. It suffices to say that the constraints ensure that surfaces $\mathbf{b}^+$, $\mathbf{b}^-$ do not self-intersect and that $\mathbf{b}^+$, $\mathbf{b}^-$ join together along medial seams and edges to form a single closed surface $\mathcal{B}$.

The medial model of the myocardium is shown in Fig. 1. Surfaces $\mathbf{b}^+$ and $\mathbf{b}^-$ describe the the epicardium and endocardium. The skeleton $\mathbf{m}$ has three branches, corresponding to the LV wall, the RV wall, and the ventricular septum.

2.3   Statistical Shape Model

Training data are used to construct statistical models for myocardial shape and thickness, as follows. The myocardial model (Fig. 1) is deformed to binary segmentations of the myocardium in training data. The likelihood and prior terms used in this deformation are specified in [14]. In brief, the likelihood measures volume overlap between the binary segmentation and the model's interior, and the prior ensures model validity by penalizing violation of constraints in [15]. Rudimentary correspondence between training data is achieved by penalizing the distortion in area element of $\mathbf{m}$ during deformation (this is roughly equivalent to equal arc length parameterization for 2D models). The Generalized

Procrustes algorithm is applied to points in the skeleton mesh to remove pose variability from fitted medial models.

Variability in myocardial shape is modeled using a multivariate Gaussian distribution: principal component analysis (PCA) is applied to the x, y, z coordinates of vertices on the skeleton $\mathbf{m}$. Thickness variability is modeled by a more localized Gibbs distribution (i.e., thickness is treated as a Markov random field):

$$-\log\ p(R_1, \ldots, R_N; \mu_\bullet, \sigma_\bullet)\ \sim$$

$$\sim\ \lambda_1 \sum_{i=1}^{N} \frac{(R_i - \mu_i)^2}{\sigma_i^2} + \lambda_2 \sum_{(j,k) \in \mathcal{E}} \frac{(|R_j - R_k| - \mu_{jk})^2}{\sigma_{jk}^2}, \quad (2)$$

where $i = 1 \ldots N$ indexes vertices in the skeleton, and $\{\mu_i, \sigma_i, \mu_{jk}, \sigma_{jk}\}$ are parameters estimated from the training data. The reason we use different priors for $\mathbf{m}$ and R stems from the fact that thickness varies much more smoothly than the x, y, z coordinates of the skeleton. Thus, the MRF is an appropriate model for thickness, but it is too restrictive for shape. We performed experiments that show that using PCA for both $\mathbf{m}$ and R leads to worse segmentation performance than using PCA for $\mathbf{m}$ and MRF for R.

2.4   Appearance Model

At each vertex in the medial model, we build a model of local appearance, which is subsequently used to drive image segmentation. Recall that (1) associates each vertex $\mathbf{m}_i$ on the skeleton with two boundary vertices (BV), $\mathbf{b_i^+}$ and $\mathbf{b_i^-}$. At each BV, we train a classifier to discriminate between a "well-placed" BV and a "misplaced" BV. A well-placed BV lies on the corresponding anatomical boundary in the training image, and a misplaced BV lies some distance away from the anatomical boundary. Well-placed BVs are obtained by fitting models to manual segmentations of the myocardium in the training data. Misplaced BVs are obtained by applying displacements to the well-placed BVs along the direction from $\mathbf{b_i^+}$ to $\mathbf{b_i^-}$ (called the *chord direction* because $\mathbf{b_i^+}\mathbf{b_i^-}$ is a chord of the sphere $(\mathbf{m}_i, R_i)$). So training exemplars for each classifier include well-placed and misplaced versions of a given BV across all subjects included in the training subset. To further increase the number of training exemplars and make classifiers less sensitive to location, we include, as training exemplars for each classifier, misplaced and well-placed versions of the BVs in the two-ring neighborhood of the BV associated with the classifier.

A rich set of features is used to build these classifiers. First, for each MRI slice, we compute a set of rotation-invariant texture descriptors at different scales [7]. Using Einstein notation, with L_α denoting image derivative in direction $\alpha \in \{x, y\}$, these descriptors are given by L, $L_\alpha L_\alpha$, $L_{\alpha\alpha}$, $L_\alpha L_{\alpha\beta} L_\beta$, $L_{\alpha\beta} L_{\beta\alpha}$. These texture descriptors are sampled around each BV using a cylindrical sampling grid oriented along the chord direction. Linear interpolation is used to sample texture descriptors between slices. Thousands of features are obtained for each BV. For each feature, a simple threshold-based *weak classifier* is constructed. AdaBoost [5] is used to combine these weak classifiers into a single *strong classifier*.
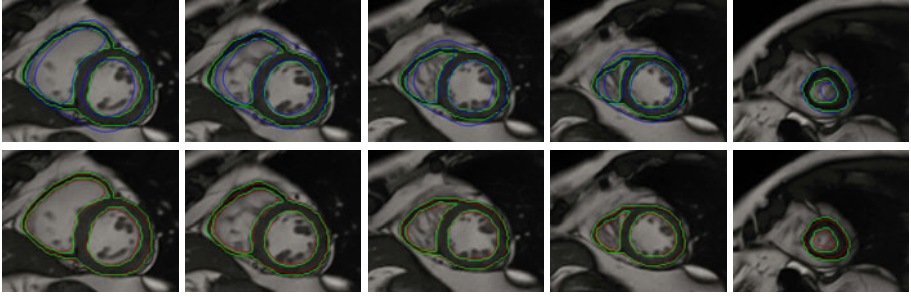
2.5   Automatic Segmentation

The automatic segmentation is initialized manually by identifying five landmarks in the most basal slice of the MRI image and one landmark at the apex. The template medial model is affinely aligned to the six landmarks. Subsequently, the model is deformed by iteratively applying two steps: *local search* and *Bayesian deformation*. Both of these steps are adaptations of the segmentation algorithm used in active shape models [4,3], and we only summarize them here. During local search, we displace each BV along the chord direction, so as to maximize the probability of it being a "well-placed" BV; this probability is estimated using the AdaBoost classifier trained for that BV. Given two displaced BVs $\hat{\mathbf{b}}_{\mathbf{i}}^{+}$ and $\hat{\mathbf{b}}_{\mathbf{i}}^{-}$, we compute a new position $\hat{\mathbf{m}}_i$ such that the triangles $\hat{\mathbf{b}}_i^{+}\,\hat{\mathbf{m}}_i\,\hat{\mathbf{b}}_i^{-}$ and $\mathbf{b}_i^{+}\,\mathbf{m}_i\,\mathbf{b}^{-}{}_i$ lie in the same plane and have the same shape; the new radius value $\hat{R}_i$ is equal to the radius of the circumscribed circle of $\hat{\mathbf{b}}_i^{+}\,\hat{\mathbf{m}}_i\,\hat{\mathbf{b}}_i^{-}$. The mesh $(\hat{\mathbf{m}},\hat{R},\mathcal{E})$ resulting from the local search step is not guaranteed to satisfy any of the constraints for a "valid" medial model; nor does it respect statistical shape models computed from training data. During the Bayesian deformation step, we deform the template myocardial model to optimize a log posterior energy that is decomposed into log likelihood and log prior terms. The log likelihood simply measures similarity between $(\mathbf{m}, R)$ in the model and $(\hat{\mathbf{m}},\hat{R})$ computed by the local search:

$$\log p(\hat{\mathbf{m}},\hat{R}\,|\,\mathbf{m}, R) = \sum_{i=1}^{N}[\|\mathbf{m}_i - \hat{\mathbf{m}}_i\|^2 + (R_i - \hat{R}_i)^2]\,.$$

This likelihood treats $\mathbf{m}$ and R as having the same units, which has not been a problem in practice. A likelihood term based on Mahalanobis distance could also be used as an alternative. Prior terms include (a) validity priors that ensure constraints on the medial model are satisfied; (b) PCA-based statistical shape prior on $\mathbf{m}$; (c) Markovian prior on R in (2); (d) a penalty on distortion in area element, used to maintain rudimentary correspondence between subjects.

3   Results

Segmentation performance is evaluated using a cross-validation strategy. We perform ten experiments in which the primary dataset (n=40) is divided into a training subset (n=24, 6 images from each cohort) and test subset (n=16, 4 images from each cohort). In each experiment, we build shape and thickness priors and train AdaBoost classifiers for appearance modeling using the training subset. We then apply the segmentation method to the test subset. We measure the difference between automatic segmentation results and manual segmentations and report average errors over the ten cross-validation experiments. Errors are reported in terms of the widely used point-to-mesh distance metric (e.g.,[1]). For each point on the mesh, the closest point (not necessarily a vertex) on the other mesh is located and the Euclidean distance between these two points is

Fig. 2. Example of automatic segmentation in a single subject. The top row shows the manual segmentation in green and the model initialized by landmarks in blue. The bottom row shows the segmentation result in red, with manual segmentation in green. From left to right, slices progress from most basal slice to the apex.
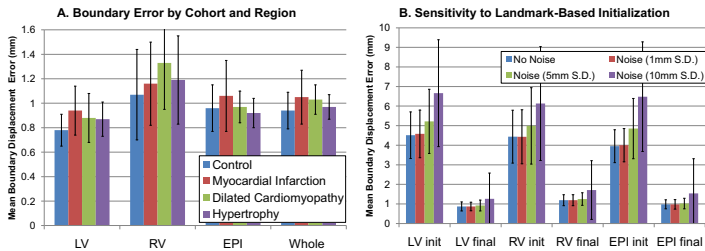
Fig. 3. Accuracy of the proposed method, expressed in terms of average mesh-to-point boundary displacement error. **A.** Average error in four cohorts. Error is computed separately for endocardial LV and RV surfaces, epicardium, and whole myocardium. **B.** Sensitivity to initialization. Random displacement with standard deviation of 1, 5 or 10mm is applied to the landmarks used during initialization. Boundary displacement error before (Init) and after segmentation (Final) is plotted for endocardial LV and RV surfaces and the epicardium. Error bars indicate standard deviation in both plots.

calculated. This distance is computed for each point on the mesh and the weighted average (according to the area) defines the point-to-mesh distance to the other mesh. The distance is calculated from model-based segmentation mesh to the ground-truth and vice versa to make the measurement symmetric.

Fig 2 shows an example of automatic segmentation, compared with initial landmark-based placement of the model and manual segmentation. The model correctly finds boundaries of the myocardium, separating it even from other structures with the same intensity, such as papillary muscles. Fig. 3A plots the average segmentation error in different regions of the myocardium for each clinical cohort. Overall, the error is close to 1 mm, slightly less in the LV and slightly higher in the RV. Errors are slightly higher in presence of pathology than in controls. Sensitivity to initialization is illustrated in Fig. 3B, showing that the method tolerates small errors in initial model placement.

Table 1. Errors ($\pm$ s.d. if known) reported in prior cardiac MR segmentation work

Method	Resolution (mm^3)	LV (mm)	RV (mm)	EPI (mm)
Mitchell '02 [11]^{p,c}	$1.56 \times 1.56 \times 9$	$2.75 \pm 0.86^*$	-	$2.63 \pm 0.76^\dagger$
Lotjonen '04 [10]^c	$1.0 \times 1.0 \times$ variable**	2.01 ± 0.31	2.37 ± 0.50	$2.77 \pm 0.49^\ddagger$
van Assen '08 [1] ^{c}	$1.5 \times 1.5 \times 10$	1.72	-	$1.55^\dagger$
Zhuang '08 [17]^c	$2 \times 2 \times 2$	2.4 ± 1.1	2.6 ± 1.5	$1.3 \pm 0.21^\dagger$
Jolly '09 [6]	$1.25 \times 1.25 \times 8$	2.26	-	$1.97^\dagger$
Peters '10 [13]^p	$0.6 \times 0.6 \times 0.8$	0.69^{***}	0.74	$0.83^\dagger$
Our approach ^{p,c}	$1.56 \times 1.56 \times 8$	0.87 ± 0.23	1.19 ± 0.28	$0.98 \pm 0.23^\ddagger$

^{p} Evaluation in patient data. ^{c} Evaluation in control data. *Distance measured in 2D. **Short and long axis slices. ***Error measured as surface to surface distance. †LV epicardium only. ‡LV+RV epicardium.

4   Discussion and Conclusions

Our results show that the automatic segmentation method is rather accurate relative to manual segmentation, with errors on the order of 1 mm. Table 1 shows that the accuracy of our method compares favorably with recently published boundary-based segmentation cardiac techniques on similar data, although one should use caution when comparing techniques that use different MRI sequences, different segmentation protocols and different error criteria. Even when the landmarks used for initialization are randomly displaced, our method achieves very good accuracy. Additional experiments not detailed here show that the MRF thickness prior is largely responsible for the strong performance of our method, as without this prior, the accuracy degrades substantially. Such a prior is made possible by the adoption of medial modeling.

As any complex Bayesian segmentation approach, our method requires many parameters to be set, including the relative weights of different prior terms, numbers and types of features used to train the appearance model, range of the local search, etc. In future work, it will be essential to measure sensitivity of the method to these parameters. Furthermore, it is critical to estimate how well the method generalizes to cardiac MRI from different MRI scanners and cohorts.

In conclusion, we have presented and evaluated a novel technique for biventricular myocardial segmentation. To our knowledge, it is the first such technique to use medial models and to leverage thickness priors that such models enable. The thickness prior, combined with established techniques from the boundary-based cardiac segmentation literature result in excellent accuracy relative to manual segmentation.

References

1. van Assen, H.C., Danilouchkine, M.G., Dirksen, M.S., Reiber, J.C., Lelieveldt, B.F.: A 3-D active shape model driven by fuzzy inference: Application to cardiac CT and MR. IEEE T. Inf. Technol. B 12(5), 595–605 (2008)
2. Blum, H., Nagel, R.: Shape description using weighted symmetric axis features. Pattern Recognit. 10(3), 167–180 (1978)
3. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models - their training and application. Comput. Vis. Image Underst. 1(61), 38–59 (1994)
4. Cristinacce, D., Cootes, T.: Facial feature detection using adaboost with shape constraints. In: British Machine Vision Conference, vol. 1, pp. 231–240 (2003)
5. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. J. Comput. Syst. Sci. 55(1), 119–139 (1997)
6. Jolly, M.P., Xue, H., Grady, L., Guehring, J.: Combining registration and minimum surfaces for the segmentation of the left ventricle in cardiac cine MR images. Med. Image Comput. Comput. Assist Interv. 12(Pt 2), 910–918 (2009)
7. Koenderink, J.J., van Doorn, A.J.: Representation of local geometry in the visual system. Biol. Cybern. 55(6), 367–375 (1987)
8. Loop, C., DeRose, T.: Generalized b-spline surfaces of arbitrary topology. In: Computer Graphics (ACM SIGGRAPH Proceedings), pp. 347–356 (1990)
9. Lorenzo-Valdés, M., Sanchez-Ortiz, G.I., Elkington, A.G., Mohiaddin, R.H., Rueckert, D.: Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm. Med. Image Anal. 8(3), 255–265 (2004)
10. Lotjonen, J., Kivisto, S., Koikkalainen, J., Smutek, D., Lauerma, K.: Statistical shape model of atria, ventricles and epicardium from short- and long-axis MR images. Med. Image Anal. 8(3), 371–386 (2004)
11. Mitchell, S.C., Bosch, J.G., Lelieveldt, B.P.F., van der Geest, R.J., Reiber, J.H.C., Sonka, M.: 3-D active appearance models: segmentation of cardiac MR and ultrasound images. IEEE T. Med. Imaging 21(9), 1167–1178 (2002)
12. Paragios, N.: A Variational Approach for the Segmentation of the Left Ventricle in Cardiac Image Analysis. Int. J. Comput Vision 50(3), 345–362 (2002)
13. Peters, J., Ecabert, O., Meyer, C., Kneser, R., Weese, J.: Optimizing boundary detection via simulated search with applications to multi-modal heart segmentation. Med. Image Anal. 14(1), 70–84 (2010)
14. Sun, H., Avants, B.B., Frangi, A.F., Sukno, F., Gee, J.C., Yushkevich, P.A.: Cardiac medial modeling and time-course heart wall thickness analysis. Med. Image Comput. Comput. Assist Interv. 11(Pt. 2), 766–773 (2008)
15. Yushkevich, P.A., Zhang, H., Gee, J.: Continuous medial representation for anatomical structures. IEEE T. Med. Imaging 25(2), 1547–1564 (2006)
16. Zeng, X., Staib, L., Schultz, R., Duncan, J.: Volumetric layer segmentation using coupled surfaces propagation. In: Proc. CVPR IEEE, pp. 708–715 (1998)
17. Zhuang, X., Rhode, K., Arridge, S., Razavi, R., Hill, D., Hawkes, D., Ourselin, S.: An atlas-based segmentation propagation framework locally affine registration–application to automatic whole heart segmentation. Med. Image Comput. Comput. Assist Interv. 11(Pt 2), 425–433 (2008)