

Accurate Segmentation of the Left Ventricle in Computed Tomography Images for Local Wall Thickness Assessment

J. Peters^{1,*}, J. Lessick², R. Kneser¹, I. Wächter², M. Vembar³,
O. Ecabert¹, and J. Weese¹

¹ Philips Research Europe – Aachen, 52066 Aachen, Germany

jochen.peters@philips.com

² Philips Healthcare – Advanced Technology Centre, Haifa 31004, Israel

³ Philips Healthcare – CT Clinical Science, Cleveland, OH 44143, USA

Abstract. In recent years, the fully automatic segmentation of the whole heart from three-dimensional (3D) CT or MR images has become feasible with mean surface accuracies in the order of 1mm. The assessment of local myocardial motion and wall thickness for different heart phases requires highly consistent delineation of the involved surfaces. Papillary muscles and misleading pericardial structures lead to challenges that are not easily resolved. This paper presents a framework to train boundary detection functions to explicitly avoid unwanted structures. A two-pass deformable adaptation process allows to reduce false boundary detections in the first pass while detecting most wanted boundaries in a second pass refinement. Cross-validation tests were performed for 67 cardiac datasets from 33 patients. Mean surface accuracies for the left ventricular endo- and epicardium are 0.76mm and 0.68mm, respectively. The percentage of local outliers with segmentation errors $> 2\text{mm}$ is reduced by a factor of 3 as compared to a previously published approach. Wall thickness measurements in full 3D demonstrate that artifacts due to irregular endo- and epicardial contours are drastically reduced.

Keywords: cardiac wall thickness, papillary muscles, boundary detection, image segmentation, multi-slice computed tomography.

1 Introduction

Fully automatic segmentation of three-dimensional (3D) cardiac images enables efficient assessment of global and local functional cardiac parameters such as ejection fraction or wall thickness. Model-based segmentation approaches [1,2,3] enable full-heart segmentation with surface-to-surface accuracies in the order of 1mm in both computed tomography (CT) and magnetic resonance (MR) image volumes (see [4]). Shape-constrained deformable models for whole heart segmentation have been described in detail in [1]. There, a triangulated mesh

* Corresponding author.

models the endocardial surfaces of the four cardiac chambers as well as the left ventricular (LV) epicardium and short trunks of the major vasculature. The mesh is adapted towards corresponding image boundaries. External image forces are balanced by internal forces that roughly maintain the model shape. This approach reaches an average accuracy of 0.82mm on 28 CT datasets. Clinically, global parameters such as cardiac chamber volumes and myocardial mass are sufficiently accurate and automation results in appreciable time savings and excellent reproducibility [5,6,7,8].

To characterize irregularities of the heart motion, a local assessment of wall motion and wall thickness for different cardiac phases is desirable. For this task, a highly consistent segmentation of the involved surfaces is needed. For the endocardium, papillary muscles and trabeculations lead to an ambiguity that must be resolved. To measure blood volume, the segmentation should consistently follow the bloodpool border. Wall thickness measurements should be based on a convex hull around the bloodpool that excludes the papillaries from the myocardium. For the epicardium, it is important to avoid local confusions with the nearby pericardium or lung transition.

In [9], manual contours from one cardiac phase are propagated to other phases. Thereby, user-defined deviations between the wanted contours and the visible bloodpool-papillary transitions are propagated. In [10], a convex hull is established around classified LV bloodpool voxels. Similarly, in [11], papillaries and trabeculations at the endocardium are addressed by post-processing of a model-based segmentation: If contrast suffices to classify the LV bloodpool voxels near the initial segmentation, a convex mesh is established around the classified voxels. [12] presents a model-based segmentation with locally disabled image forces near the papillary muscles without contour post-processing.

This paper presents a method to change the boundary detection for model-based segmentation such that (1) papillaries and trabeculations are excluded from the myocardium and (2) the epicardial segmentation is no longer misled by false boundaries. Apart from quantitative improvements, we demonstrate that obvious errors for local wall thickness estimation are eliminated.

2 Improved Boundary Detection Training

A crucial point for robust and accurate segmentation is the detection of corresponding image boundaries for each model surface element. Discriminative boundary detectors can be trained using the “Simulated Search” described in detail in [4]. This method evaluates the performance of a set of boundary detection functions on a set of representative images with corresponding reference segmentations. The set of tested boundary detectors is generated from some templates with parameters estimated from the training examples. The evaluation metric to rank the proposed detectors simulates the geometric segmentation errors without shape constraints: For each triangle, boundary detection is performed from systematically modified triangle poses to simulate the mesh adaptation. Each detected boundary point is compared to the reference segmentation

and the geometric distance to the reference surface is averaged over all simulated mesh states and all training images. Finally, the boundary detector with minimal simulated root-mean-square (RMS) error is assigned to the triangle.

For the task of accurate LV endo- and epicardium segmentation, two problems motivate algorithmic extensions of the above optimization metric. First, image boundaries are *missing* for some surface regions. Bloodpool boundaries at the papillaries shall *not* attract the mesh if a convex hull around the bloodpool is required. Furthermore, the epicardium often touches the diaphragm and no epicardial edges are visible there. Second, the epicardium is often parallel to nearby boundaries such as the pericardium or the lung transition (see Fig. 1). In all these cases, the image contains boundaries close to the wanted contours. These, however, shall *not* be detected.

We can summarize our training requirements: (1) At papillary muscles and trabeculations, bloodpool contours *inside* the reference endocardium shall *not* be detected. (2) For the epicardium, nearby competing boundaries shall be rejected. The simulated geometric error alone does not fully address these goals. E.g., a detector that often locates the weak epicardial transition but sometimes detects false points or rejects all edges may have a larger error than another detector that reliably locates the pronounced lung transition if both boundaries are adjacent. At the papillaries, learning to detect bloodpool contours is evidently good if geometric proximity is optimized and if no better edges can be found. Rejecting all edges produces simulated errors equal to the simulated mesh modifications, which are typically larger than the papillaries.

To include the above explicit goals (1) and (2) into the training, the optimization criterion of the Simulated Search was extended. First, regions containing *unwanted* boundaries were explicitly annotated in the image volume (see Fig. 2). All detected boundaries in these “forbidden” regions that would pull the mesh towards unwanted boundaries are counted. The resulting percentage p of “forbidden” detections is linearly combined with the geometric RMS error d from the standard training procedure to arrive at a new optimization criterion:

$$\text{Cost}(F, i) = (1 - \alpha) \cdot d(F, i) + \alpha \cdot D \cdot p(F, i) \quad (1)$$

Here, F is the tested boundary detector for the given triangle i . d and p include all simulated detections in all images. D converts the percentage p into a distance




Fig. 1. Closeup of endo-, epi-, and pericardium. Black outlined arrows indicate the wanted epicardial boundary. White arrows indicate competing edges (pericardium, lung, diaphragm). Endo- and epicardial reference segmentations are shown on the right.




Fig. 2. Masks describing “forbidden” regions for the endocardium (blue, including the papillaries) and for the epicardium (yellow, covering various false edges; the green mask covers trabeculations in the right ventricle). Note that the masks do *not* include the immediate neighborhood of the reference segmentations.

comparable to typical simulated values of d . $\alpha = 0$ restores the standard training, while $\alpha = 1$ would focus only on avoiding forbidden edges and ignore the geometric accuracy of all detected boundaries.

3 Model for Accurate LV Segmentation

3.1 Reference Segmentations and Shape Model

Segmentations of 67 image volumes from 33 patients at different cardiac phases were carefully inspected and corrected in short and long axis views as well as in axial, coronal, and sagittal views. For the endocardial contour, a convex hull tightly enclosing the visible bloodpool was created. In some images, however, faint contrast and image artifacts made the proper definition of this contour difficult. Epicardial contours needed corrections where the pericardium or a nearby transition to the lung or beyond the diaphragm had been segmented. A smoothly interpolating mesh was used in regions of missing visible edges at the diaphragm. From the resulting 67 meshes, a new shape model was trained that better represents the convex shape of the endocardium.

3.2 Segmentation Chain with Two-Pass Deformable Adaptation

We follow the segmentation framework from [1]: The heart is first roughly localized by a Generalized Hough Transform [13]. The mesh is then adapted towards corresponding image boundaries. First, a global similarity transformation followed by a piecewise affine transformation initializes the mesh close to the desired boundaries. Finally, the mesh is deformed elastically but image forces are balanced by shape-preserving forces. For this deformable adaptation, we adopt a two-pass scheme from [12] that allows to reduce false positive boundary detections in the first pass while improving true positive detections in the second pass. In the first pass, the whole heart model is adapted with a search range of 10mm. Here, boundary detectors trained with large α in (1) can be used to minimize detection in forbidden regions. The shape model will smoothly interpolate the

surface across missing boundaries. In the second pass, only LV endo- and epicardial mesh parts are further adapted using a short search range of 2mm. The mesh can thus adapt to nearby boundaries missed in the first pass. Many forbidden edges are beyond the short search range and boundary detectors optimized for this range with small α can focus on local accuracy.

3.3 Boundary Detection Functions

Our boundary detectors combine edge detection with additional rejection criteria [1,4]: We either project the image gradient onto the triangle normal (to suppress edges with false orientation) as in [1,4] or we calculate the difference between inside and outside gray values (averaged over several points). To reject false boundaries, one or several local image features are evaluated and compared to trained acceptance intervals. If any feature value violates its interval, the edge is rejected. Image features offered to the training covered the gray values on both sides of the triangle (possibly averaged over several points) and Taylor coefficients of gray values along the triangle normal direction as in [1,4]. Here, variances of gray values parallel and normal to the mesh have been added to the feature set.

3.4 Detector Training and Selection

Using the extended training framework of Sec. 2, we first studied the trade-off between simulated geometric errors and forbidden boundary detections. Optimizing all $\{F_i\}$ for increasing α in (1), d increases and p decreases as expected. Plotting the mesh average of both quantities versus α , p shrinks most rapidly for small α , while d increases most for large α . Increasing d and shrinking p may indicate that no edge was detected for some triangles. (This contributes to d with the simulated displacements.) Under the influence of shape constraints, local errors may remain small if some neighboring edges are still found. After model adaptation, we typically find minima at medium values of α for both the geometric errors and the percentages of triangles ending up in forbidden regions.

To select the best combination of first and second pass boundary detectors, we decided to train and test endo- and epicardial detectors *separately*. This has two advantages: First, it allows to efficiently explore a multitude of training setups with different image features for boundary rejection since the “problems” at endo- and epicardium are of different nature. Second, it allows the use of different values of α for both mesh regions. To select the final boundary detectors from the different tested training setups, separate shape-constrained endo- and epicardial models were adapted to all training images in the two-pass scheme. Amongst several detector combinations with close-to-minimal surface-to-surface errors the percentage of triangles ending up in forbidden regions was minimized. This resulted in one detector combination for both the endocardium and the epicardium. These detectors were finally intergrated into the full heart model.

4 Results

4.1 Surface-to-Surface Errors

To evaluate the accuracy of the segmentation, we measure constrained surface-to-surface distances: For each triangle, the distance to the closest surface point in a corresponding patch of the compared mesh is determined. Corresponding patches around a triangle are constrained by anatomical labels and limited to a geodesic radius of 10mm over the mesh surface. The distances are measured from the automatic mesh to the reference mesh and vice versa since they are not symmetric. From each distance pair, the *maximum* is retained. These distances over all test images and all mesh triangles per anatomical region are finally evaluated statistically.

Experiments were conducted on 67 retrospectively ECG-gated cardiac multi-slice CT image volumes from 33 patients (1–3 cardiac phases per patient). These had been obtained using 16-, 40-, and 64-channel CT scanners (Brilliance CT, Philips Healthcare, Cleveland, Ohio, USA). Bloodpool contrast and image quality varied considerably. The mean in-plane voxel resolution was 0.45mm (range 0.30–0.49mm) using a 512×512 matrix. The mean slice thickness was 0.72mm (range 0.33–2.0mm). Our baseline model had been generated and trained as described in [1] on 28 datasets from 13 patients using old reference segmentations that had *not* been corrected at the LV endo- and epicardium. For the improved LV segmentation, we performed a 3-fold cross-validation on all 67 datasets. Here, a new shape model as well as the LV endo- and epicardial boundary detectors were trained on 22 patients and the datasets of the left-out 11 patients were segmented with the resulting cross-validation model. The boundary detectors of all other structures were taken from the baseline model.

Table 1 summarizes the resulting error statistics for the LV surfaces. Mean errors for epi- and endocardium are reduced by 30% and 40%, resp. In view of our carefully corrected reference segmentations we consider this reduction as a clear improvement. More important, however, is the distribution of errors over the mesh and across images. Local outliers with large errors exceeding 2mm that are beyond typical inter-observer uncertainties are reduced by a factor of 3.

Table 1. Mean surface-to-surface errors [in mm] with standard deviations covering variations both across mesh triangles and across images. The error distribution describes the percentages [in %] of triangles within the listed error bins [in mm] in all images (NumTriangles × NumImages data points). The LV epicardium includes the septal boundary towards the right ventricle.

	Baseline model					New model				
	Mean ± SDev	[0, 1)	[1, 2)	[2, 3)	≥ 3	Mean ± SDev	[0, 1)	[1, 2)	[2, 3)	≥ 3
LV endo	1.26 ± 1.35	58.5	22.3	8.9	10.2	0.76 ± 0.88	77.0	16.6	3.6	2.9
LV epi	0.96 ± 1.02	67.1	21.6	6.8	4.4	0.68 ± 0.62	79.1	16.9	3.1	1.0

4.2 Qualitative Results and Wall Thickness Measurements

Fig. 3 shows some segmentation examples illustrating the achieved improvements. The consistent segmentation of the endocardium and the improvements for the epicardium enable us to measure the LV wall thickness. To account for the bending 3D geometry of the myocardium towards the apex, we do *not* measure wall thickness in short-axial slices. Rather, the normal of the endocardial wall is estimated for endocardial triangles from a regression plane fitted through the triangle vertices and their neighbors. We use 1st and 2nd order vertex neighbors for robust direction estimation. A ray is then cast from the triangle center outwards, and the distance to the epicardial wall is taken as wall thickness. We restrict our measurements to the myocardial region assigned to the 17 AHA segments proposed in [14].




Fig. 3. Segmentations obtained with the baseline model (rows 1 and 3) and with the cross-validation models of Sec. 4.1 (rows 2 and 4). The first two examples stem from the same patient. The first column is at mid-diastole, all other examples are end-systolic.

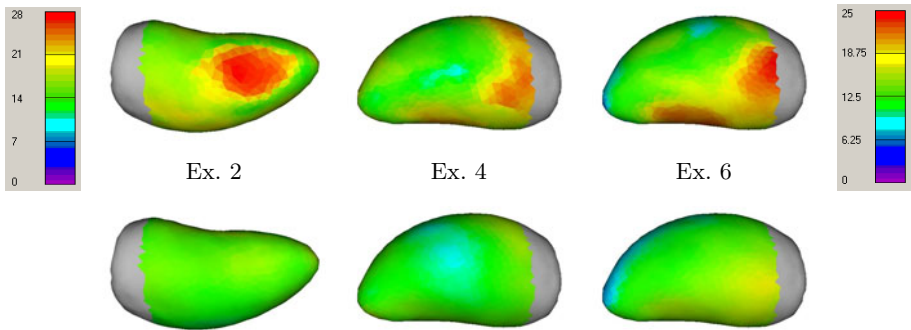


Fig. 4. Wall thickness measurements for examples 2, 4, and 6 from Fig. 3. Example 2: inferior view (color bar 0–28mm). Examples 4 and 6: lateral view (color bar 0–25mm, all measurements projected onto the mean mesh). Local maxima due to irregular endocardial and epicardial surfaces with the baseline model (top row) are eliminated with the cross-validation models (bottom).

This region excludes the LV outflow tract towards the aorta and the transition towards left atrium. Results are shown in Fig. 4.

5 Discussion and Conclusion

The quantitative and qualitative results demonstrate clear improvements of the LV segmentation. Using the extended optimization criterion of the Simulated Search and carefully corrected reference segmentations, boundary detectors could be trained that effectively avoid unwanted edges. Locally visible structures are well segmented in the two-pass adaptation and outliers with large errors are drastically reduced. The consistent exclusion of papillary muscles from the myocardium and the accurate segmentation of the epicardial surface are the basis for local wall thickness measurements.

References

1. Ecabert, O., et al.: Automatic model-based segmentation of the heart in CT images. *IEEE Trans. Medical Imaging* 27(9), 1189–1201 (2008)
2. van Assen, H.C., et al.: A 3-D active shape model driven by fuzzy inference: Application to cardiac CT and MR. *IEEE Trans. Information Tech. Biomed.* 12(5), 595–506 (2008)
3. Zheng, Y., et al.: Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features. *IEEE Trans. Medical Imaging* 27(11), 1668–1681 (2008)
4. Peters, J., et al.: Optimizing boundary detection via simulated search with applications to multi-modal heart segmentation. *Medical Image Analysis* 14(1), 70–84 (2010)

5. Plumhans, C., et al.: Comparison of manual, semi- and fully automated heart segmentation for assessing global left ventricular function in multidetector computed tomography. *Investigative Rad.* 44(8), 476–482 (2009)
6. Ghersin, E., et al.: Clinical evaluation of a fully automated model-based algorithm to calculate left ventricular volumes and ejection fraction using multidetector computed tomography. *Acute Cardiac Care* 11, 47–55 (2009)
7. Coche, E., et al.: Quantitative right and left ventricular functional analysis during gated whole-chest MDCT: A feasibility study comparing automatic segmentation to semi-manual contouring. *European Journal of Radiology* (in Press), <http://dx.doi.org/10.1016/j.ejrad.2009.05.037>
8. Abadi, S., et al.: Feasibility of automatic assessment of four-chamber cardiac function with MDCT: Initial clinical application and validation. *European Journal of Radiology* (in Press), <http://dx.doi.org/10.1016/j.ejrad.2009.01.035>
9. Hautvast, G., et al.: Automatic contour propagation in cine cardiac magnetic resonance images. *IEEE Trans. Medical Imaging* 25(11), 1472–1482 (2006)
10. Cocosco, C.A., et al.: Automatic image-driven segmentation of the ventricles in cardiac cine MRI. *J. Magn. Reson. Imaging* 28(2), 366–374 (2008)
11. Zheng, Y., et al.: Left ventricle endocardium segmentation for cardiac CT volumes using an optimal smooth surface. In: *Proc. SPIE Medical Imaging*, vol. 7259, 72593V–1–11 (2009)
12. Lehmann, H., et al.: Integrating viability information into a cardiac model for interventional guidance. In: Ayache, N., Delingette, H., Sermesant, M. (eds.) *FIMH 2009*. LNCS, vol. 5528, pp. 312–320. Springer, Heidelberg (2009)
13. Ballard, D.H.: Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recogn.* 13(2), 111–122 (1981)
14. Cerqueira, M.D., et al.: Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. *Circulation* 105, 539–542 (2002)