

Automated Segmentation of the Left Ventricle in Cardiac MRI

Michael R. Kaus¹, Jens von Berg¹, Wiro Niessen², and Vladimir Pekar¹

¹ Philips Research Laboratories, Sector Technical Systems,
Röntgenstr. 24-26, D-22335 Hamburg, Germany

² Image Sciences Institute, University Medical Center,
Heidelberglaan 100, 3584 CX Utrecht, The Netherlands
michael.kaus@philips.com

Abstract. We present a fully automated deformable model technique for myocardium segmentation in 3D MRI. Loss of signal due to blood flow, partial volume effects and significant variation of surface grey value appearance make this a difficult problem. We integrate various sources of prior knowledge learned from annotated image data into a deformable model. Inter-individual shape variation is represented by a statistical point distribution model, and the spatial relationship of the epi- and endocardium is modeled by adapting two coupled triangular surface meshes. To robustly accommodate variation of grey value appearance around the myocardiac surface, a prior parametric spatially varying feature model is established by classification of grey value surface profiles. Quantitative validation of 60 end-diastolic 3D MRI datasets demonstrates accuracy and robustness, with 1.28 ± 0.81 mm mean deviation from manual segmentation. We investigate the extension to 4D by incorporating a constraint on the allowed deformation based on a learned example and show illustrative results for 4D MRI.

Keywords: cardiac MRI, image segmentation, deformable models

1 Introduction

Cardiac MRI provides important information for diagnosis and treatment of cardiovascular diseases by enabling quantitative assessment of functional parameters [4]. However, the lack of computational tools still implies that not all the information contained in the data is currently used. In order to fully exploit the information in the MRI data in clinical practice, automated extraction of the myocardium is essential. Considering the amount of data involved (10–20 frames by 8–12 image slices), clinically feasible delineation should be automated. However, poor object features resulting from signal loss due to blood flow and partial volume effects, and significant variation of grey value appearance make finding a robust and accurate solution a difficult problem.

Recently proposed approaches to automated segmentation incorporate prior knowledge to improve robustness [5,8,9,10]. Lorenzo-Valdes et al. achieved convincing results by volumetric atlas matching using B-spline registration [9]. However, this approach requires several minutes for a single time frame. The fast and robust framework of active

shape and appearance models by Cootes et al. [1] was extended and applied to cardiac 3D and 2D time series [8], achieving promising results. However, statistical shape models cannot capture variability outside of the learning set, which is likely to occur in the case of pathology, and appearance modeling may fail in the presence of large grey value variability across subjects and time.

Earlier work by this group proposed shape constrained deformable models, enabling local deviation from a statistical shape model by embedding it into a 3D triangular deformable mesh [6,11]. This approach used deterministic feature search and was successfully applied to segment bones in CT. However, like many other deformable model approaches [7], multiple objects as the endo- and epicardium are not considered, nor objects of significant shape change over time and across subjects. In addition, methods with deterministic feature functions assume constant grey value appearance of the object's surface, which is not true for the myocardiatic surface due to different surrounding tissue as lung parenchyma, fat, and the right ventricle.

We propose to integrate several prior information sources into shape constrained deformable models including a deterministic, parametric model of the variation of surface features, spatial relationships to handle multiple objects, and object deformation in time. In the following we present the modeling and segmentation methods with application to cardiac 3D MRI time series.

2 Methods

The shape model is represented by a triangular mesh. After initial positioning, the mesh is adapted to the image by iterating the two-step procedure consisting of i) surface detection in the image for each triangle, and ii) reconfiguration of the vertex coordinates by minimizing $E = E_{\text{ext}} + \alpha E_{\text{int}}$. The parameter α weighs the relative influence of an external energy E_{ext} , which drives the mesh towards detected surface points, and an internal energy E_{int} , which maintains the vertex configuration of an statistical shape model. In the following we outline the deformable model framework [11] and then present how to represent, learn and integrate prior knowledge into the method.

2.1 Deformable Shape Models

Surface Detection is carried out for each triangle center \mathbf{x}_i . We seek the point $\tilde{\mathbf{x}}_i$ along the triangle normal \mathbf{n}_i which maximizes the cost function of a feature function F (see below) and distance $j\delta$ to the triangle center according to

$$\tilde{\mathbf{x}}_i = \mathbf{x}_i + \delta \mathbf{n}_i \arg \max_{j=-l, \dots, l} \{F(\mathbf{x}_i + j\delta \mathbf{n}_i) - Dj^2\delta^2\}, \quad (1)$$

where $2l + 1$ is the number of points investigated, δ specifies the distance between two points on the profile, and D controls the tradeoff between feature strength and distance.

The **external energy** term drives the mesh towards the detected surface points:

$$E_{\text{ext}}(\mathbf{x}) = \sum_{i=1}^T w_i (\tilde{\mathbf{x}}_i - \mathbf{x}_i)^2, \quad w_i = \max \{0, F(\tilde{\mathbf{x}}_i) - Dj^2\delta^2\}, \quad (2)$$

T being the number of triangles. The weights w_i give the most promising surface points $\tilde{\mathbf{x}}_i$ the largest influence during mesh reconfiguration.

The **internal energy** maintains the distribution of the mesh vertex coordinates \mathbf{v}_j w.r.t. the edges of a given initial mesh $\tilde{\mathbf{v}}_{jk} = \tilde{\mathbf{v}}_j - \tilde{\mathbf{v}}_k$

$$E_{int} = \sum_{j=1}^V \sum_{k \in N(j)} (\mathbf{v}_j - \mathbf{v}_k - s\mathbf{R}\tilde{\mathbf{v}}_{jk})^2, \quad (3)$$

where $N(j)$ is the set of neighbors of vertex j , and V is the number of vertex coordinates [11]. The rotation \mathbf{R} and the scaling s of the mesh are estimated each iteration using a fast closed-form point-based registration method based on singular value decomposition. Since the energies in (2), (3) are quadratic, energy minimization results in the efficient solution of a sparse linear system using the conjugate gradient method.

2.2 Integration of Prior Knowledge

Statistical Shape Model. A point distribution model (PDM) is a parametric representation of a set of learning shapes [1], where an arbitrary shape $\mathbf{s} = [\mathbf{v}_1^T \dots \mathbf{v}_V^T]^T$ can be approximated by the sum of the mean shape \mathbf{m} and $L - 1$ weighted eigenmodes \mathbf{e}_l

$$\mathbf{s} \approx \mathbf{m} + \sum_{l=1}^{L-1} p_l \mathbf{e}_l, \quad p_l = (\mathbf{s} - \mathbf{m})^T \mathbf{e}_l. \quad (4)$$

The model is automatically generated from a set of L segmented images [6]. First, a shape template is generated by triangulation of an arbitrarily selected image, which is adapted to each of the remaining images by rigid and non-rigid adaptation (see above). Finally, \mathbf{m} and \mathbf{e}_l are calculated using principal component analysis on the landmark vectors, which consist of the vertex coordinates of the adapted shape template.

The shape model can be adapted to image data by finding the optimal weights p_i [1], but shapes outside the variation of the training set cannot be described. In contrast, we allow local deviations but encourage the distribution of the mesh vertices to stay close to the PDM by embedding it into the internal energy [11], extending (3) to

$$E_{int} = \sum_{j=1}^V \sum_{k \in N(j)} (\mathbf{v}_j - \mathbf{v}_k - s\mathbf{R}(\mathbf{m}_j - \mathbf{m}_k + \sum_{l=1}^{\tilde{L}} p_l (\mathbf{e}_{lj} - \mathbf{e}_{lk})))^2, \quad (5)$$

where \tilde{L} is the number of eigenmodes used, \mathbf{m}_j , \mathbf{e}_{lj} , are the vertices of the PDM, and the p_l are calculated each iteration according to (4).

Simultaneous Adaptation of Spatially Related Surfaces. Our previous work considered segmenting single objects [11]. However, the epi- and endocardium are two structures which are not connected but spatially related. This relationship may be violated during mesh adaption, yielding e.g. too close or even intersecting meshes. This may result e.g. in the attraction of the endocardium to epicardiac image features (Fig. 3(a),(c)).

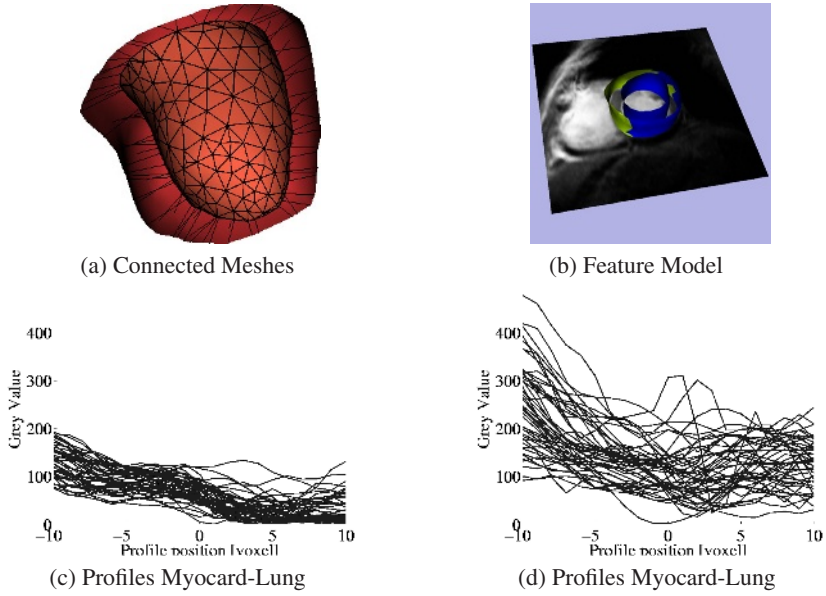


Fig. 1. Myocardial mesh with connections between endo- and epicardium (both cut in half for visualization purposes, (a)), and the spatially varying feature model (color-coded gradient direction, (b)). Depending on the triangle location, profile statistics may exhibit small (c) and large (d) variations, drastically reducing the specificity (d) of profile matching.

We propose to integrate this information to improve the segmentation robustness and accuracy inspired by [3,12]. Additional vertex connections are introduced that connect neighboring vertices of the existing endo- and epicardiac surface model (Fig. 1(a)). Both surface meshes are adapted simultaneously while constraining their initial configuration and the vertex connections by adding a connection energy to the total energy

$$E = E_{\text{ext}} + \alpha E_{\text{int}} + \beta E_{\text{conn}}, \quad E_{\text{conn}} = \sum_{j \in C} (\mathbf{v}_j^{\text{epi}} - \mathbf{v}_{k(j)}^{\text{endo}} - s \mathbf{R} \tilde{\mathbf{v}}_{jk})^2, \quad (6)$$

$\tilde{\mathbf{v}}_{jk} = \tilde{\mathbf{v}}_j^{\text{epi}} - \tilde{\mathbf{v}}_k^{\text{endo}}$ refers to a connected pair of vertices of the epi- and endocardiac mean surfaces, C is the set of connected vertices, β weighs the connection energy, and $E_{\text{int,ext}}$ are calculated as in (2), (3) for both surfaces.

The connections are established once as part of the model construction phase. For each vertex in the epicardiac surface the closest vertex in the epicardium is found in the sense of the Euclidean distance. Since the shortest distance is not symmetric, the same is done for the epicardium's vertices. The resulting set of vertex pairs is then reduced so that each vertex has at most one connection, favoring shorter connections. This results in connections approximately orthogonal to the surface meshes in most cases (Fig. 1(a)).

Spatially Varying Feature Model. Deterministic, parameterized feature models have been successfully used, e.g. combining image gradient strength and direction w.r.t. the triangle normal, $F = \mathbf{n}^T \mathbf{g}$, (1), for vertebrae in CT [11]. However, the gradient direction varies across the myocardiac surface, and the gradient strength alone is not a robust

feature (Section 3). Statistical profile matching based on training data [1,8] models variable grey value appearance, but is computationally expensive in 3D and achieves low specificity (and robustness) if the inter-subject profile variance is high (Fig. 1(c,d)).

We propose a hybrid approach, where the parameters of a deterministic feature model are learned once from a training sample as part of the model construction. Specifically, for each triangle we learn a grey value range $[t1_j; t2_j]$, a gradient range $[g1_j; g2_j]$, and a gradient direction $s_j = \pm 1$ (inwards, outwards w.r.t. the mesh normal, Fig. 1(b)). Feature search is made more robust to outliers by extending it to

$$\tilde{F}(\mathbf{x}) = s_j F(\mathbf{x}) \mathbf{1}_{I(\mathbf{x}) \in [t1_j; t2_j], g_j \in [g1_j; g2_j]} \cdot \quad (7)$$

The s_j are automatically determined using the triangular meshes generated from the above learning examples during PDM construction. First, grey value profiles are sampled in all L images at all T triangle centers along the triangle normal. All triangles are assigned a particular tissue interface class $c_{j,l}$, $j = 1 \dots T$, $l = 1 \dots L$, using square-error (k-means) classification [2] on the entire profiles. We argue that this classification is more robust against outliers than local profile statistics since the sample size is significantly larger (approximately $L \cdot T/C$ vs. L , C : number of classes). We chose the three classes myocardium (MC)–fat, MC–right ventricle, and MC–lung parenchyma, where each class was assigned the proper gradient direction. After classification, each triangle j of the shape model is assigned the class label C_j and gradient direction s_j following $C_j = \arg \max_{c=0,1,2} \sum_{l=1}^L \mathbf{1}_{c_{j,l}=c}$. This is based on the property that the triangles j, l have approximately anatomically corresponding locations on the learning examples [6]. Similarly, the grey and gradient value ranges for a triangle are the mean \pm (std deviation) of grey values and gradients in the corresponding triangle centers in the learning set.

Generation and Adaptation of a 4D Deformation Model. Capturing significant LV shape variation with deformable models is difficult because they are local methods with limited capture range, and may adapt to false features due to poor image quality and other nearby structures. One approach is to initialize the surface mesh with a close approximation to the object in the image [8,9,10]. However, constraining a mesh adaptation to the previous segmentation result may fail if the deformation between two time frames is too large, and a prior 4D model of the LV for independent initialization of each time frame may not describe a particular subject well.




Fig. 2. 4D motion model constructed from a single subject (bottom row, every 5th frame, endocardium cut open for visualization purposes) and adaption to a different subject (top row). The model $M(i)$ constrains the mesh adaptation to image $I(i)$, initialized with the previous adaptation result $S(i-1)$.

We propose to combine a prior 4D shape model with the subject-specific adaptation result of the previous image (Fig. 2). The 4D model consists of meshes $M(t)$ with T triangles, V time-dependent vertex coordinates $v_j(t)$, and T feature attributes s_j . Here we assume a constant feature model, enabled by the fixed mesh topology during adaptation, i.e. constant number of triangles T and vertices V .

For each time frame $I(t)$, the mesh $M(t)$ represents the shape model, and the mesh $S(t-1)$ adapted to the previous image is used as an initialization. Each time frame is segmented independently using the internal energy of (3), where $\tilde{\mathbf{v}}_{jk}$ correspond to the vertex coordinates of $M(t)$, and $\mathbf{v}_j, \mathbf{v}_k$ are initialized with the vertex coordinates of $S(t_{n-1})$ prior to the first adaptation. In our experiments, we had as many shape models as time frames. If the numbers of shape models and time frames doesn't match, interpolation between shape model vertex coordinates is possible [10].

Due to the limited availability of segmented image data we derive a 4D LV deformation model from a single subjects manually segmented 4D MRI. A triangular mesh is constructed from each time frame by adapting the myocardiac mesh of the end-diastole generated earlier to the segmented image. In contrast to grey value images, this strategy is possible with segmented images because they provide perfect surface features.

3 Experimental Results

For 3D (4D) adaptation the shape model is automatically positioned in the center of the (first) image volume. In case of unfavorable FOV, more sophisticated positioning is possible [8]. The parameter settings used are summarized in Table 4. An adaptation to a single time frame took approximately 8 seconds on a 1.7 GHz Pentium 4.

Image Data. We used cardiac MRI datasets (ECG-triggered FFE-EPI, Philips ACS-NT 1.5 T, 20 time frames, 256×256 images with 7–10 slices, $1.25 \times 1.25 \times 10$ mm) from 60 subjects with undetermined pathology and compared the automated results with manual segmentations by calculating the mean and maximum closest distance between mesh vertices and binary image. For each test case, a model was built from the remaining 59 subjects (leave-one-out test). The 4D deformation model was generated from this one manually segmented time series, and adapted to time frames of other subjects. Due to the lack of ground truth, the assessment was only done visually.

3D Adaptation. Fig. 3 illustrates the effect of the various sources of prior knowledge on the model adaptation. Separate adaptation of the two surface meshes with a global feature function $F = \mathbf{n}^T \mathbf{g}$ (Fig. 3(a)) results in severe adaptation errors since it models a bright object on a dark background. Simultaneous adaptation of connected meshes





Fig. 3. Effect of prior knowledge on model adaptation. F is the feature function.



Parameter	Value
Energy weight α	0.25
Distance weight D	100
Profile step size δ	1 mm
Profile length l	10
Num eigenmodes \tilde{L}	10
Num mesh reconfigurations	20

Error [mm]:	Mean	Std Dev	Max
Initialization	2.34	1.46	20.35
Simple Gradient	3.24	1.65	21.25
Combination	1.28	0.82	11.18

Fig. 4: Adaptation results (end-diastolic phase) after mesh initialization in the volume’s center (data points sorted and connected with lines for visualization purposes), parameter settings.

approximately maintains the distances between the meshes (Fig. 3(b)), but does not prevent attraction of the epicardiac mesh to the LV. Separate adaptation with a spatially varying feature model (Fig. 3(c)) correctly adapts the epicardiac mesh in places where the myocardiac tissue interfaces with the right ventricle. Nevertheless, the partial differences between the feature model and this subject result in false attraction to the epicardiac fat. Constraining the spatial relationship of the surface meshes (Fig. 3(d)) compensates the partially wrong feature model. For wall thickness analysis the papillary muscles must be inside the endocardiac surface, which is achieved by maintaining the shape of the model and the spatial relationship between the meshes.

These findings are consolidated in our leave-one-out experiments. Accurate adaptation is not possible with a simple gradient feature alone due to disregard of prior information (Table 4). Mean and maximum errors after automated initialization and mesh adaptation are 1.28 ± 0.82 mm or 1-2 voxels (Fig. and Table 4), where the maximum errors occur in slice direction (10 mm slice thickness). The experiments showed no significant difference between the adaptation accuracy of the endo- and epicardium. In two cases (23, 54, Fig. 4) the initial mesh was a better approximation to the subject’s LV than the adaptation result. In these cases the spatial distribution of features in the model did not describe the subject’s grey value appearance well and could not be compensated by the shape constraints (i.e. surface and connections), resulting in partial attraction of the endocardiac mesh to epicardiac fat.

4D Adaptation. The 4D model (Fig. 2, bottom row) derived from segmented images of a subject was adapted to time series of a different subject (top row). Visual inspection of the adaptation of the model suggests feasibility of the method, which was not possible using the 3D end-diastolic model alone due to the large deformation during the heart cycle. As in the 3D case, similar few false attraction occurred, where the i.e. endocardiac mesh adapted to epicardiac fat.

4 Conclusion

We presented a novel method for modeling and fully automated segmentation of LV in 3D cardiac MRI time-series with a mean error of 1.28 ± 0.82 mm for the end-diastolic

phase. While visual inspection suggests feasibility of the method in 4D, quantitative validation for 4D segmentation is required to assess the quality of the segmentation and quantitative measurements of functional parameters.

We demonstrated that the integration of several sources of prior knowledge into a deformable model significantly increases robustness and accuracy of segmentation. Information about shape and spatial configuration of objects by simultaneous adaptation of coupled meshes helps compensating the lack of prominent features. Learning of parameterized, deterministic local feature functions enables robust surface detection for objects with spatially varying grey value appearance on the surface. The integration of a temporal deformation model enabled to segment objects of large shape variation.

Future work will investigate the advantages of a true 4D approach by e.g. simultaneous adaptation of connected 3D meshes, and the extension of the 4D deformation model to represent the deformation of a population of subjects.

Acknowledgments. Cardiac MRI have been kindly provided by Warren J. Manning, MD, Co-Director of the Cardiac MR Center of the Beth Israel Deaconess Medical Center, Boston, MA. We thank our colleagues T. Netsch, K. Meetz, Philips Research Hamburg, and S. Lobregt, R. Truyen, F. Gerritsen, Philips Medical Systems Best, for useful discussions.

References

1. T.F. Cootes, G.J. Edwards, and C.J. Taylor. Active appearance models. In *Proc. of ECCV'98*, volume 2, p. 484–498. Springer Verlag, 1998.
2. R.O. Duda and P.E. Hart. *Pattern Classification and Scene Analysis*. Wiley and Sons, 1973.
3. P.T. Fletcher, S.M. Pizer, A.G. Gash, and S. Joshi. Deformable m-rep segmentation of object complexes. In *Proc. of ISBI*, p. 26–29. IEEE Press, 2002.
4. A. Frangi, W. Niessen, and M.A. Viergever. “Three-Dimensional Modeling for Functional Analysis of Cardiac Images: A Review”. *IEEE Trans. Med. Imag.*, 20(1): 2–25, 2001.
5. M.-P. Jolly, N. Duta, and G. Funka-Lea. Segmentation of the left ventricle in cardiac MR images. In *Proc. of ICCV*, p. 501–508. IEEE Computer Society, 2001.
6. M.R. Kaus, V. Pekar, C. Lorenz, R. Truyen, S. Lobregt, and J. Weese. Automated 3D PDM construction from segmented images using deformable models. *IEEE Trans. Med. Imag.*, 22(8): 1005–1013, 2003.
7. T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: A survey. *Med. Img. Anal.*, 1(2): 91–108, 1996.
8. S.C. Mitchell, B.P.F. Lelieveldt, R.J. van der Geest, H.G. Bosch, J.H.C. Reiber, and M. Sonka. Multistage hybrid active appearance model matching: Segmentation of left and right ventricles in cardiac MR images. *IEEE Trans. Med. Imag.*, 20(5): 415–423, 2001.
9. M.Lorenzo-Valdes, G.I. Sanchez-Ortiz, R. Mohiaddin, and D. Rückert. Atlas-based segmentation and tracking of 3D cardiac MR images using non-rigid registration. In *Proc. of MICCAI*, p. 642–650. Springer Verlag, 2002.
10. J. Montagnat and H. Delingette. Space and time shape constrained deformable surfaces for 4D medical image segmentation. In *Proc. of MICCAI*, p. 196–205. Springer Verlag, 2000.
11. J. Weese, M.R. Kaus, C. Lorenz, S. Lobregt, R. Truyen, and V. Pekar. Shape constrained deformable models for 3D medical image segmentation. In *Proc. of IPMI*, p. 380–387. Springer Verlag, 2001.
12. X. Zeng, L. H. Staib, R. T. Schultz, and J. S. Duncan. Segmentation and measurement of the cortex from 3D MR images using coupled surfaces propagation. *IEEE Trans. Med. Imag.*, 18(10): 927–937, 1999.