A Skeleton-Based Inflation Model for Myocardium Segmentation

André Neubauer and Rainer Wegenkittl
VRVis Research Center, Vienna, Austria

Abstract

Analysis of four-dimensional cardiac MRI data sets has become a popular tool in modern cardiology. Semi-automatic segmentation of objects of interest, like the myocardium or the left ventricle, helps to increase effectiveness in the cardiologist's work. This paper presents a new algorithm for semi-automatic segmentation of the myocardium in a 3D or 4D MRI data set. The user manually approximates the center-line (skeleton) of the myocardial cross-section in one slice. The skeleton is used to initialize a simple deformable model which is then iteratively aligned with myocardial contours. It is primarily attracted to sharp boundaries fragments of the myocardium and fine-tuned by aligning with remaining image gradients. The segmentation result is finally automatically propagated to all other slices.
Keywords: myocardium, segmentation, inflation, deformable model

1 Introduction and Related Work

Four-dimensional Magnetic Resonance Imaging (MRI) has become an important tool in modern cardiology. A patient's cardiac activity can be recorded by generating MRI data volumes of the heart at a number of time points evenly distributed over one cardiac cycle. In order to be able to automatically extract important properties of the patient's heart function, objects of interest, like the heart muscle (the myocardium) and the left ventricle of the heart, must be segmented in each data volume. Segmentation of cardiac MR images is often problematic due to poor image quality which has various reasons: Patient movement, turbulent blood flow and physical and chemical disturbances of the imaging process lead to a number of artifacts which have to be compensated by the segmentation technique.

Segmentation of MRI data is still an active field of research, although a lot of different methods have already been proposed: Most techniques are based on the concept of deformable models [6, 11]. A boundary is initialized and iteratively aligned to object contours by applying internal and external forces. Internal forces keep the boundary in a smooth volumetric shape and enforce shape-based constraints often stemming from a-priori knowledge about object contours. External forces fit the boundary to image properties by, for instance, driving it towards significant image gradients. Many different concepts for initialization and refinement of deformable models in the field of medical image segmentation have been published: One of the most popular is the method known as 'snakes' [3]. A snake is an energy-minimizing spline which is driven by internal forces, constrained (user-defined) forces and image forces towards the correct object boundaries.

Most deformable models, including snakes, are very sensitive with respect to initialization. Proper initialization, preferrably as close to the correct boundary as possible, is often crucial for the avoidance of getting stuck in local minima of the underlying energy functional. One popular concept is to initialize the model somewhere inside the segmented object and apply inflation forces to expand the boundary until it fits the object boundaries. Deformable models using this technique are often called 'balloons' [1]. Jones and Metaxas use pixel-affinity [10] (a measure of probability that two neighboring pixels belong to the same object) and balloon forces to initialize a two-dimensional deformable model [2]. McInerney and Terzopoulos use a three-dimensional finite element surface model which is also pushed towards the final object boundaries by balloon forces and then refined by data-driven forces [5]. These approaches cannot segment the myocardium right-away, due to the concave shape of the heart muscle. The left ventricle must be segmented first, then the shape of the epicardium (the outer boundary of the myocardium) can be initialized outside the ventricle and further refined [8]. However, image noise, MRI artifacts and papillary muscles

(small muscles inside the LV) can lead to initializations far from the correct boundaries preventing the desired quality of result.

Techniques which combine the deformable model approach with anatomical atlases [4, 7] have been proposed. Data sets acquired from a number of patients or volunteers form a training set. They are abstracted to common features which are collected to an atlas and aid the automatic segmentation of new data sets. Research in this area, however, is still young. Current systems are either restricted to 2D images or have trouble segmenting data sets where structures are significantly different to those encountered in the training set.

This paper introduces a new approach based on a deformable model which is initialized around an approximated center-line of the cross-section of the myocardium in each slice of the MRI volume and then iteratively refined. This algorithm reduces the probability of the model locking in a local minimum by detecting highly significant boundary voxels prior to the process of model refinement. For model refinement, an inflation force and image- and shape-based forces are applied to drive the deformable boundary towards these detected boundary fragments and attach it to the real boundaries of the myocardium.

2 Myocardium Segmentation

This section gives a detailed description of the new algorithm for segmenting the myocardium from cardiac data sets obtained from MRI.

The result of a cardiac MRI scanning process is a set of images of cross-sections (slices) of a patient's heart. The algorithm presented here works on a slice-by-slice basis. The core structure of the technique is outlined in figure 1. Segmentation of the myocardium from a 3D data set starts with the user manually approximating the center line (the skeleton) of the cross-section of the myocardium in an arbitrary slice by placing a poly-line onto the slice (step 1). Utilizing this user-defined skeleton, a presegmentation step indentifies data gradients which are likely to represent parts of the boundary (steps 2 and 3). Next, a simple deformable boundary is initialized tightly enclosing the skeleton (step 4) and iteratively refined ("inflated", step 5) to finally match the object contours. This progressive adaptation is achieved by applying a set of forces to the deformable model, which attract it towards the previously identified (step 5) gradients, and at the same time enforce constraints which keep the boundary in the correct shape. Once the final contours have been found, the center line of the segmented myocardial cross-section is calculated using skeletonization (step 6). This center line is then copied to the next slice (step 7), slightly adapted to better fit the data (step 8) and then used to initialize a new two-dimensional segmentation step.

The following paragraphs describe those steps in detail:

2.1 Manual Skeleton Placement

In the first step of the segmentation algorithm, the user selects a slice and places a poly-line on it, which approximates the center line of the myocardial cross-section on the slice (see the leftmost image in figure 2).

2.2 Detection of Boundary Fragments

One of the key parts of this algorithm is the identification of significant gradients which represent parts of the object boundary with high probability and can be used to attract the deformable boundary. This edge detection technique is based on voxel connectivity: For each voxel in the slice, a value is calculated which denotes its connectivity to the skeleton. This connectivity value is high, if there exists a path from the voxel to any part of the skeleton, which traverses similar data values and no significant gradients.

In MRI data sets, a set of artifacts can be observed, which can cause the data level to vary quite significantly within the myocardium. Care has to be taken to compensate these possible data inhomogeneities. This is done by using all intensity values that the skeleton passes through locally as reference values for the calculation of connectivity:

The connectivity values of all voxels of the slice, which are not part of the skeleton, are initialized to 0, those of skeleton voxels are 1. Each voxel is also assigned a reference value. This reference value is initially 0 for non skeleton voxels and equal to the intensity value for skeleton voxels. A sorted list of voxels is generated. Entries are sorted by descending connectivity values. The list initially holds all skeleton voxels. In order to calculate the remaining connectivity values, the following steps are iteratively performed:

1. Retrieve the voxel with highest connectivity from the list. If the connectivity is below 0, stop execution of the algorithm. Let c be the connectivity value and r the reference value of the obtained voxel.

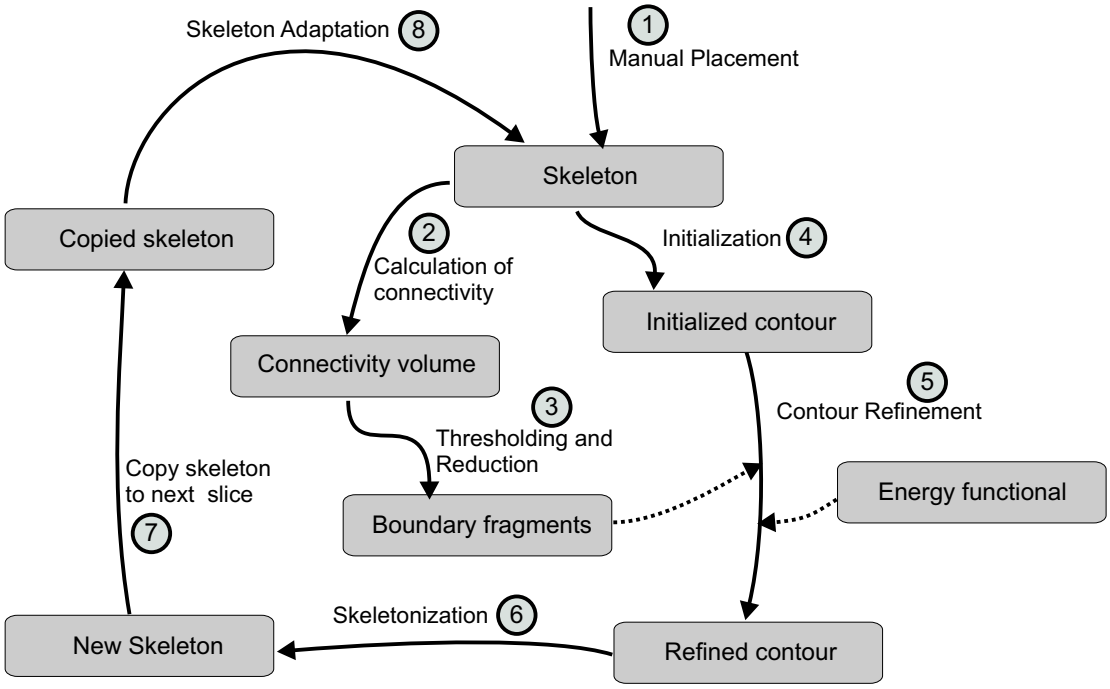

Figure 1: algorithm overview

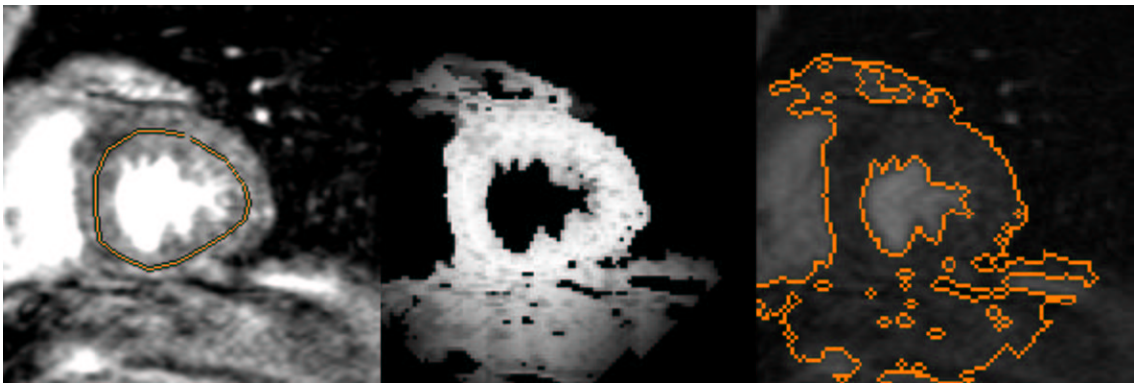

Figure 2: left: manually placed skeleton, center: voxel connectivity, right: preliminary boundaries

2. For each immediate neighbor voxel (4-neighborhood), calculate the value t with:

$$t = c - C_1 - C_2(1 - S)^{C_3}$$

S is a measure of similarity of the intensity value v of the neighbor voxel with the reference value r:

$$S = e^{\frac{((v-r)/C_4)^2}{-2}}$$

The constant values C_1, C_2, C_3 and C_4 control the dynamics of connectivity propagation: C_1 determines the degree of penalization of longer paths. C_2 determines the degree of penalization of higher gradients. C_3 controls the steepness of the function assigning decrease of connectivity to gradient size and C_4 controls the similarity relation between data values. For the example illustrated in the figures of this paper, the values $C_1 = 0.02$, $C_2 = 28.5$, $C_3 = 2$ and $C_4 = 14$ were used.

3. If t is higher than the current connectivity value of the neighbor voxel, the connectivity value is changed to t and the reference value is changed to r.

4. return to 1.

The center image of figure 2 shows the resulting field of connectivity values (black means no connectivity, white means high connectivity). A preliminary boundary is now extracted as an iso-contour in the connectivity field using a threshold T (step 3, see the rightmost image of figure 2). Since T is very small (close to zero), this preliminary boundary contains sharp edges of the myocardium as well as less significant edges well outside the heart muscle. The significant myocardium edges can now be identified by reducing the preliminary boundaries to the parts with largest connectivity gradient: Let the neighborhood of voxel v be all voxels within a Manhattan-Distance of 3 from v. The value s_{high} is the sum of the connectivity values of all neighbor voxels which have a higher connectivity than v and n_{high} the number of those voxels. Values s_{low} and n_{low} are analogous. The connectivity gradient of v is:

$$c_{high} \cdot (1 - c_{low})$$

with

$$c_{high} = \frac{s_{high}}{n_{high}} \cdot \frac{10}{max(n_{high}, 10)}$$

and c_{low} being calculated analogously. Values c_{high} and c_{low} are the average connectivity values of neighbor voxels with higher, respectively lower, connectivity, if $n_{high} = n_{low} = 10$ and lower otherwise. Higher gradient values are therefore assigned to voxels whose connectivity values are the medians in their neighborhoods. This scheme effectively reduces the impact of noise. The sizes of connectivity gradients of voxels belonging to the preliminary boundary are depicted in the leftmost image of figure 3. Now, important boundary fragments can be extracted from the preliminary boundary by removing all voxels whose connectivity gradients are below a threshold T_2. A suitable value for T_2 can be found by analyzing the function f which maps to each connectivity value c the number of preliminary boundary voxels with a connectivity $c' > c$. Threshold T_2 should be near the lowest point of the first derivative of f. The center image of figure 3 shows the resulting boundary fragments.

2.3 Boundary Model Initialization

Two deformable models are initialized as polygons whose vertices are adjacent to the skeleton, one model outside and one model inside the skeleton. The initialized boundary polygons should not intersect the skeleton.

2.4 Boundary Model Refinement

The key part of the presented algorithm is to refine the initialized boundaries and align them with the real contours of the myocardium. The boundary, initially located tightly around the skeleton, is exposed to forces which expand it until it touches the precalculated boundary fragments. The iterative enlargement of the enclosed region gives the impression of the boundary being inflated until fixed barriers (the boundary fragments) prevent further growing. Refinement consists of minimizing an energy functional e, with:

$$e = e_{Bound} \cdot e_{Shape} \cdot e_{Grad}$$

The three involved terms relate to the following aspects:

- Boundary fragments: The deformable boundaries should be attracted to the boundary fragments identified in step 3. The term e_{Bound} is proportional to the average squared distance of boundary voxels to voxels belonging to precomputed boundary fragments. For calculation of e_{Bound} a chamfer distance map denoting for each voxel in the slice the distance to the nearest voxel which is part of a precomputed boundary fragment, is computed. Chamfer distance growing is always interrupted by the skeleton. This ensures that the distance descends into all directions from the skeleton and therefore the two boundary models are attracted to opposite object contours. An example distance map is depicted in figure 3. In this image, the skeleton is indicated by discontinuities in the distance field. Minimization of e_{Bound} is the primary force applied to the models. Since the probability of boundary fragments located well inside the myocardium is extremely low, the risk of getting stuck in local minima is significantly reduced.

- Boundary model shapes: Both boundaries should maintain simple shapes. The term e_{Shape} is made up of two factors: The first factor ensures the maintenance of local curvature. Cross-sections of the epicardium in MR scans are always close to circular or elliptical. This fact can be utilized by penalizing shapes deriving too much from elliptical. The other factor preserves the topologies of the models by keeping each vertex close to the centroid of its two neighbors.

- Image gradients: For fine tuning, especially in regions, where no boundary fragments have been found, the boundaries should also align

with image gradients. The term e_{Grad} penalizes parts of the boundaries moving through regions of low gradients. Instead of using real image gradients, here, gradients in the connectivity field are used. The advantage of this is that image gradients within the myocardium, which are possible due to MRI artifacts, do, with a high probability, not coincide with gradients in the connectivity field, due to the properties of the connectivity calculation algorithm discussed in section 2.2.

Progressive refinement of the deformable model is accomplished by iteratively moving each boundary vertex. The direction of the movement of a vertex is determined by two forces: an inflation force which repells the vertex from the skeleton and an energy-minimizing force which moves the vertex such that the energy functional of the complete boundary model is minimized. Using these two forces, the boundary is expanded from its initial state and snaps to the predefined boundary fragments. Figure 4 illustrates the inflation process: The boundary is depicted after one iteration of the refinement process (each vertex has been moved once), after two iterations and after twelve iterations.

2.5 Skeletonization

As already stated, only segmentation of one slice of the data set is initialized by manual skeleton placement. The segmentation result is then automatically propagated to all other slices. Propagation from one slice (the source slice) to another (the destination slice) is performed by skeletonizing the detected cross-section of the myocardium in the source slice (step 6), copying the skeleton to the destination slice (step 7), slightly adapting it (step 8), and then using it for initialization of the boundary as outlined in section 2.3 and for identification of boundary fragments (see section 2.2). There are various ways of skeletonizing a two-dimensional object. The one used in this project was proposed by Telea and van Wijk [9]. Here, each voxel which is part of the boundary of the segmented object is assigned a unique integer value, with the values monothonically increasing along the boundary. Then, using the Fast Marching Method (FMM), these values are propagated towards the center line of the object, thereby creating a field of values. The resulting skeleton consists of those voxels showing a sufficiently high gradient in this field.

After a skeleton has been found, it is copied to the next slice and slightly adapted to the data there. Then, using the adapted skeleton, a new segmentation step is initialized.

2.6 Skeleton Adaptation

Due to limited scanning speed, four-dimensional MRI data sets are usually resolved well only in two spacial dimensions. The distances between slices are usually much larger than distances between neighboring voxels of the same slice. Since there are only few slices covering the entire myocardium, there is not a high level of inter-slice coherency. The consequence of this is that the skeleton of the myocardial cross-section in one slice is generally not a good representation of the myocardial cross-sections in neighboring slices. The algorithm used for detection of boundary fragments (as described in section 2.2) is robust with respect to the exact location of the skeleton. That means that if the skeleton is completely located on the myocardium, the algorithm will work well, even if the skeleton is not a good approximation of the center line of the myocardial cross-section. However, there is the possibility that after the skeleton has been copied from the source slice, parts of it are not located on the cross-section of the heart muscle in the destination slice. So, after copying, the skeleton must be adapted such that the probability of it being completely located inside the myocardium also in the destination slice is maximized. This is done by moving the skeleton to the region where the myocardial cross-sections of source and destination slices overlap.

First each voxel of the source slice is assigned a value denoting the probability that it is contained in the overlap region. This value p is 0 for all voxels whose counterparts in the destination slice are not part of the segmented myocardial cross-section. For each other voxel V value p is calculated using the following formula:

$$p = e^{\frac{((v_d - v_s)/S_s)^2}{-2}}$$

with v_d being the intensity value of the voxel, v_s the intensity value of the voxel's counterpart in the source slice and S_s the statistical standard deviation of intensity values of the segmented cross-section in the source slice. Next, the resulting field of probability values is low-pass filtered (blurred) a number of times. This yields three results: The number of local maxima is reduced, local maxima are located on the overlap region with high probability, and values ascend smoothly towards the local maxima. The skeleton is then replaced by a polygon with some fixed number of vertices n. The polygon

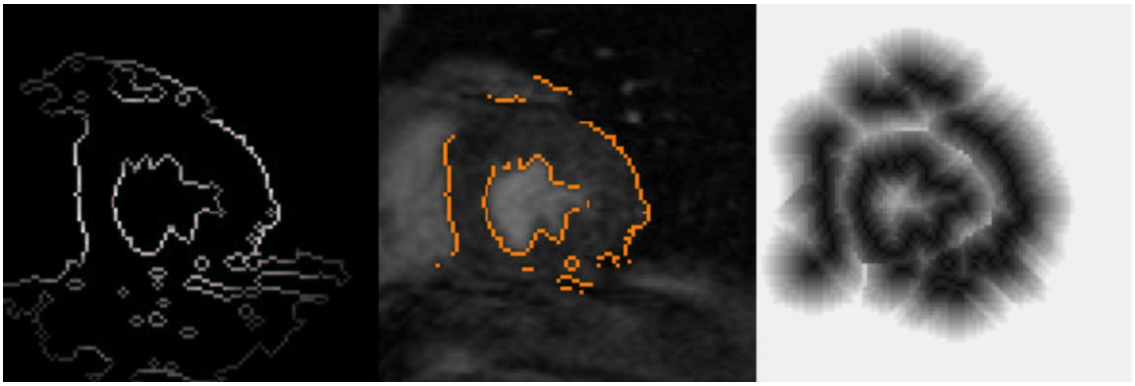

Figure 3: left: connectivity gradient sizes, center: boundary fragments, right: chamfer distance map

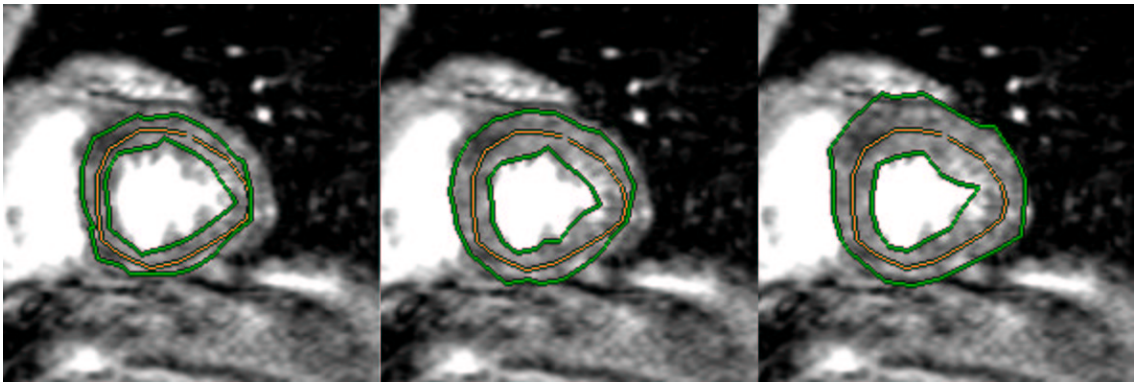

Figure 4: boundary inflation: left: 1 iteration, center: 2 iterations, right: 12 iterations

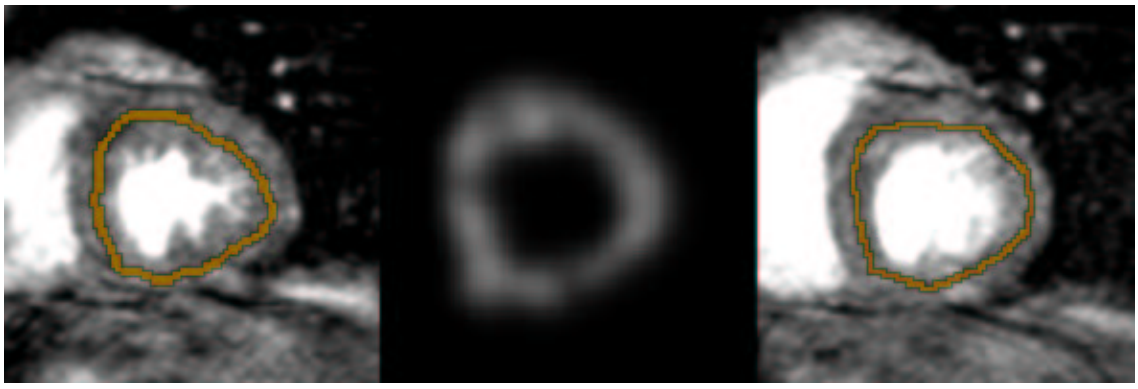

Figure 5: left: source slice with skeleton, center: probability field, right: destination slice with adapted skeleton

should be chosen such that the maximum distance between the polygon and the initial skeleton is as small as possible. Finally, each polygon vertex is iteratively moved into the direction of steepest increase in the probability field until it either reaches a local maximum or further movement would reduce its distance to a neighbor vertex below an allowed minimum. This minimal distance prevents two vertices collapsing into one. This is not allowed in order to maintain the circular shape of the skeleton as far as possible. If distances between vertices grow too large, new vertices can be added. Figure 5 shows an example of a skeleton being adapted. The image on the left shows the initial skeleton on the source slice, the center image shows the probability field and the rightmost image depicts the final adapted skeleton on the destination slice. Once all vertices have been moved to the overlap region, the resulting skeleton can be used to initiate segmentation of the destination slice.

3 Results

The algorithm presented in this paper proved to be a usable technique for robust and accurate segmentation of cardiac MRI data sets. Detailed comparison of results generated by this technique with results of skilled manual segmentation revealed a deviation ratio of 0.0431 on the average (with the deviation ratio being the number of voxels which are part of only one segmentation result - either manual or automatic - divided by the number of voxels being part of one or both). As in every algorithm which propagates a segmentation result over a number of slices, an error happening in one slice can possibly accumulate during further propagation. Still, the algorithm proved to work very robustly, propagating along a distance of 10 slices or more. Significant errors usually occur only, if MRI artifacts result in too distinct data incoherencies which result in the skeleton adaptation technique (see section 2.6) placing the skeleton outside the myocardium. This, however, is observed very rarely. Furthermore, experience has shown that the algorithm usually manages to recover after segmenting one or two slices erroneously. Propagation in time dimension works in the same way as propagation in space dimension. Segmentation of a complete four-dimensional cardiac data set can therefore by accomplished after manually placing only one skeleton on one arbitrary slice in one arbitrary volume. Segmentation of a $256 \times 256 \times 12$ data set currently takes about 17 seconds on the average. The algorithm therefore works fast enough to be used in daily clinical practice. It has been successfully added to commercial PACS software (see http://www.tiani.com).

4 Conclusions

A new semi-automatic algorithm for segmentation of the myocardium from MRI volumes, using a simple deformable model has been presented. It has been shown that the technique works robustly and fast. Robustness is increased by primarilly driving the deformation of the model by attracting it towards those voxels which were identified to mark sharp boundaries of the myocardium. Robustness is further increased by the model aligning to gradients in an image made up of skeleton-voxel connectivity values instead of plain image gradients. It has further been shown that a result of segmenting one slice can be propagated to the other slices in a stable way, therefore permitting segmentation of 3D and 4D data sets.

Acknowledgements: This work has been carried out as part of the basic research project 'Virtual Reality for Scientific Applications' at the VRVis research center (http://www.vrvis.at) in Vienna, Austria, which is funded by the Austrian governmental research project Kplus. Thanks to Tiani Medgraph AG (http://www.tiani.com) for providing useful implementation facilities.

References

[1] L. D. Cohen. On active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, 1991.

[2] T. N. Jones and D. N. Metaxas. Segmentation using deformable models with affinity-based localization. Proc. CVRMed-MRCAS, pages 53–62, 1997.

[3] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Actived contour models. Proc. First International Conference on Computer Vision, 1987.

[4] M. Lorenzo-Valdés, G. I. Sanchez-Ortiz, R. Mohiaddin, and D. Rueckert. Atlas-based segmentation and tracking of 3d cardiac MR images using non-rigid registration. Proc. MICCAI, pages 642–650, 2002.

[5] T. McInerney and D. Terzopoulos. A dynamic finite element surface model for segmentation

and tracking in multidimensional medical images with application to cardiac 4D image analysis. *Computerized Medical Imaging and Gtaphics*, 19(1):69–83, 1995.

[6] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: A survey. *Medical Image Analysis*, 1(2):91–108, 1996.

[7] D. Rueckert. Geometrically deformable templates for shape-based segmentation and tracking in cardiac MR images. Proc. EMMCVPR, pages 83–98, 1997.

[8] G. I. Sánchez-Ortiz and P. Burger. Vector field analysis of the dynamics of the heart using velocity encoded cine MR imaging. 9th International Symposium on Computer Assisted Radiology (CAR), pages 228–233, 1995.

[9] A. Telea and J. J. van Wijk. An augmented fast marching method for computing skeletons and centerlines. Proc. Eurographics/IEEE VisSym, 2002.

[10] J. K. Udupa and S. Samarasekera. Fuzzy connectedness and object definition: theory, algorithms and applications in image segmentation. *CVGIP: Graphical Models and Image Processing*, 58(3):246–261, 1996.

[11] M. Worring, A. W. M. Smeulders, L. H. Staib, and J. S. Duncan. Parameterized feasible boundaries in gradient vector fields. *Computer Vision and Image Understanding*, 63(1):135–144, 1996.