Automatic Segmentation of the Left Ventricle in 3D SPECT Data by Registration with a Dynamic Anatomic Model

Lars Dornheim[1], Klaus D. Tönnies[1], and Kat Dixon[2]

[1] Institut für Simulation und Graphik, Fakultät für Informatik,
Otto-von-Guericke-Universität Magdeburg, Germany
[2] Medical Imaging Research Group, University of British Columbia, Canada

Abstract. We present a fully automatic 3D segmentation method for the left ventricle (LV) in human myocardial perfusion SPECT data. This model-based approach consists of 3 phases: 1. finding the LV in the dataset, 2. extracting its approximate shape and 3. segmenting its exact contour.

Finding of the LV is done by flexible pattern matching, whereas segmentation is achieved by registering an anatomical model to the functional data. This model is a new kind of stable 3D mass spring model using direction-weighted 3D contour sensors.

Our approach is much faster than manual segmentation, which is standard in this application up to now. By testing it on 41 LV SPECT datasets of mostly pathological data, we could show, that it is very robust and its results are comparable with those made by human experts.

1 Introduction

Heart diseases, especially heart infarcts, are one of the most common causes of death in the western world. Modern functional medical imaging techniques can contribute significantly to diagnosis and particularly to the quantitative assessment of these diseases.

Myocardial perfusion SPECT provides functional, three-dimensional images of the human left ventricle (LV) (figure 1a), where infarct-affected regions can be seen as gaps within the well-perfused LV (figure 1b). For the qualitative and quantitative assessment of the damage of the LV, anatomical knowledge is needed (figure 1c).

LV segmentation is often performed manually, which is very time-consuming. Automatic registration of an anatomical model with functional data either requires extensive training based on segmented data, or a flexible model supplementing information, if perfusion and anatomy do not coincide.

2 Related Work

A lot of literature exists on the segmentation of the LV, but only little of it deals with SPECT data. SPECT data is noisy and the level of detail is low. Anatomy may be extracted from it only if an external model is supplied.
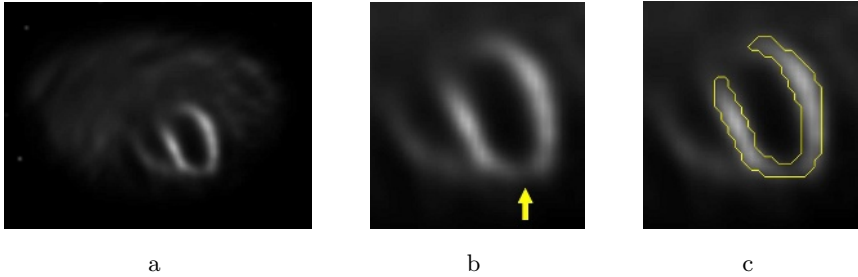
Fig. 1. (a) thorax 3D SPECT data slice. **(b)** enlarged LV section with gap due to infarct (arrow). **(c)** manual segmentation.

[1] gives an overview on previous approaches to segment the LV in SPECT data. They conclude, that many approaches (threshold-based, region-based, etc.) lead to bad results due the lack of model knowledge. Existing model-based techniques either need user interaction, or they use unflexible and simple models, so that no adequate segmentation of the LV is possible. No method exists, to cope with perfusion gaps, which are typical for infarct patients.

The approach in [2] tries to solve this problem by using superquadrics, followed by free form deformation. These implicit dynamic 3D shape models have already been used in [3] for tracking the segmented LV in 4D-SPECT data. In [2], a very simple model is used for an initial segmentation. Subsequent fine-tuning via free-form deformation does not use model knowledge, but searches the data slicewise for points supporting the interpolation. A bounding box of the LV must be provided by the user for initial placement of the implicit mode.

The most promising of these methods all use some kind of model knowledge. However, up to now, no appropriate automatic approach exists for the segmentation of the LV in 3D-SPECT data.

3   Requirements

Our approach to automatic segmentation of the LV in 3D-SPECT data registers a dynamic model with the data. We use *mass spring models*, known from the linear elasticity theory ([4]), which are a special kind of deformable models.

Unlike active contours and snakes ([5]) these models are shape models in that they model a specific shape intrinsically. They are governed by *internal forces* (from springs) and *external forces* (from sensors at the mass points). They also differ from mass spring models used in surgery simulation in that they do not try to simulate a physically correct world. Our models use forces only in order to find good segmentations.

Furthermore, the mass points do not represent exact locations with special semantic meaning like in ASMs [6]. They are (beside their function in the model dynamics) placeholders for sensors, searching for certain image features.

These mass spring models have several advantages. They can easily be applied to 3D data. No training is necessary (in contrast to stochastic models). Local

adjustment is easy, making interaction possible in a intuitive way. But also two major disadvantages exist:

− Mass spring models are already very instable in a 2D environment, which is even more so in 3D.
− A new mass spring model needs to be generated for every specific segmentation task. Its topology and shape must be adequate.

To overcome these difficulties, we have developed a new kind of stable 3D mass spring model that stabilizes itself by the use of an additional internal force, the *torsion force* ([7]), even if it is only sparsely interconnected. Stability in traditional models stems from a dense spring meshing that makes them inflexible, expensive to compute and difficult to control.

Furthermore, we developed a technique to generate complex 3D mass spring models automatically from an arbitrary sample segmentation ([8]). Models are composed of substructures representing the features of the sample segmentation; the level of abstraction can be controlled.

4   Segmentation Process

We now describe the automatic segmentation process. Its strategy is first searching the object in the dataset, then segmenting its region roughly and finally tracing its exact contour. So, the segmentation process consists of 3 phases:

1. *Initial placement of the model:* This step is crucial for the success of the following phases, since dynamic models can only adapt locally.
2. *Shape adaptation of the model:* Here, shape, size and direction of the model are fitted to the data.
3. *Contour adaptation of the model:* Now the model will be adapted to the contour of the object to segment. The shape fitting of phase 2 as a starting point will give this process a stable context.

4.1   Initial Placement

In the literature, the first and important phase of initially placing the model is usually performed interactively or discussion is omitted completely. However, shape models provide excellent possibilities to automate this step.

We generate a voxel template P from the intrinsic shape of our LV model, which we employ for a flexible pattern matching on the 3D-SPECT dataset D to find the approximate position of the LV in the data as a good initial position for our model. This approach is possible, as the size and direction of the LV never change dramatically in our datasets. Besides, there is always only one LV-shaped object in the dataset, so that no ambiguities need to be resolved. Since our model is very robust with respect to initial placement errors, a rough pattern matching should provide sufficient information.

a   b   c
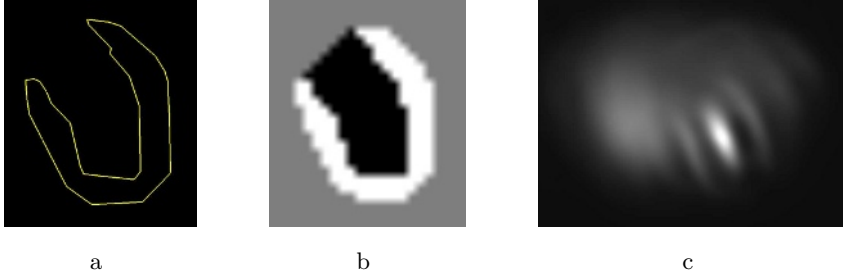
Fig. 2. (a) contour cross section of the relaxed model. **(b)** according slice of the generated voxel template. **(c)** correlation image of the slice: bright means high correlation.

The template P is constructed from the relaxed model with no external forces acting. The contour substructure of the model is treated as a polygonal mesh, which is converted to a 3D voxel template (figure 2b) by 3D polygon filling.

For the pattern matching, we use two alternative techniques. One is based on *correlation*, which can computed very fast by FFT. The other technique is based on the *pattern vector angle*. It is reasonably fast, but much more robust to brightness artefacts in the data.

Correlation. The SPECT dataset has only non-negative values, which represent the counts in a voxel. We apply the generated LV template P to the data. P is a 3D voxel matrix with value 1 at the bright LV wall, value -1 in the dark LV interior and value 0 at unspecified outer regions (figure 2b). So, both the strong signal LV wall values and the weak LV interior values contribute to high correlation values (see figure 2c).

The similarity measure e_{Corr} we use is proportional to a 3D discrete correlation. Correlation is susceptible to intensity artefacts (figure 3a). In such cases, the second technique described in the following section is a better option.

Pattern vector angle (PVA). A fairly fast pattern matching technique that does not rely on intensity levels can be realized by the use of the *pattern vectors*. These are comprised of all the voxel values from a region in a fixed order. If one measures the similarity of two equally-sized patterns via the angle between them, it does not depend on the intensity level of the data but on the shape ([9]). The shape is represented by the direction of the pattern vectors, whereas their lengths depend on this intensity level. Equation 1 shows the calculation of this angle at position p, i.e. the similarity measure e_{PVA}.

$$e_{\mathrm{PVA}}(p) = \cos \angle(\boldsymbol{D}(p), \boldsymbol{P}(p)) = \frac{\langle \boldsymbol{D}(p), \boldsymbol{P}(p) \rangle}{\|\boldsymbol{D}(p)\| \cdot \|\boldsymbol{P}(p)\|} \tag{1}$$

The method is less susceptible to artefactual intensity peaks (compare figures 3a and 3b) at the expense of higher computation times.

a   b
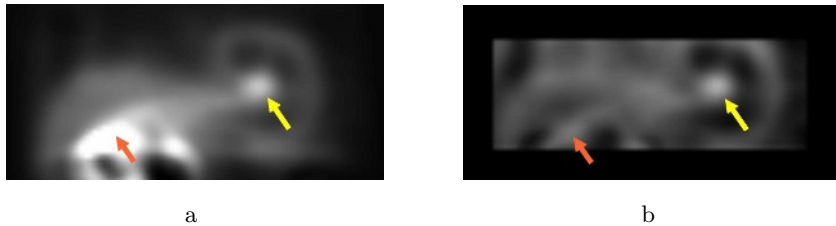
Fig. 3. (a) correlation image and (b) PVA image in Y direction. Bright regions depict high similarity values. An artefact in the dataset (red, left arrow) is only in (a) calculated with a higher similarity value than the correct position (yellow, right arrow), which has in (b) the highest similarity value.

4.2   Shape Adaptation

After initial placement, the model starts adapting to the dataset. The primary goal in this step is a rough adaptation of the model shape.

This is achieved by running the model with the external forces from its *intensity sensors* only. These sensors drive the mass points towards certain intensity levels in the data. They are evenly spread over the model and intensity information is present all over the dataset. The model therefore has good data support to adapt its rough shape quickly to it. Gaps in the LV are not critical in this phase, because they are surrounded by enough intensity information on which the model can rely. *Contour sensors*, which act on mass points towards contours, are not used in this phase. This way, a fast, flexible and robust model adaptation is achieved without disturbing the model by too many details too early.

The model movement is simulated until the model dynamics reach an equilibrium. The criterion for this is that every mass point stays within a small neighborhood around its position for a fixed number of simulation steps (equation 2, where $\boldsymbol{m}_{i_t}$ is the position vector of the ith mass point at time step t).

$$\forall i \ \forall 0 \leq j < n-1 : \ \left\| \boldsymbol{m}_{i_{t-j}} - \boldsymbol{m}_{i_{t-(n-1)}} \right\| \leq \varepsilon \quad \text{for} \quad t \geq n \qquad (2)$$

4.3   Contour Adaptation

The goal of this last step is the exact model adaptation to the data. This phase is similar to the previous phase, but now contour sensors are also active. Due to the previous step they should find the correct contour very close to their starting point. This leads to a fast and robust segmentation in this final phase.

In the case of gaps in the bright LV wall, where no such contour can be found, the still active intensity sensors in this region will stabilize the model. This way, these gaps are interpolated by the model.

Because of the high complexity in 3D, there are normally many potential contours close to a sensor. We therefore developed a directionally sensitive sensor, which takes the gradient direction of the 3D contour into account ([10]).

5   Evaluation

For our tests, we developed a prototype, which was running on a standard PC (1,3 GHz Pentium M, 512 RAM) with a modern 3D graphics card (NVidia GeForce-Go5200 64MB). The concept of dynamic cached sensor cards (see [10] for that) for each sensor in the model allows for a very fast model dynamics simulation, so that the model motion can be visualized in real time. This way it is much easier to assess the model behavior on the data.

We tested our method on 41 3D-SPECT datasets of human LV's from 25 different patients (128 x 128 voxel slices, 25 to 38 slices per dataset, cubic voxels with 4.795 mm edge length). For evaluation, we used a subset of 7 datasets, which were mostly (5 of 7) affected by infarcts (up to 25 %). We used manual segmentations from a medical professional as gold standard.

For a quantitative assessment, we used several measures (relative segmentation error, Hausdorff distance and average surface distance) to calculate the similarity between the gold standard and the binary voxel segmentation.

The model was generated automatically from one of the 7 datasets. It consisted of 153 mass points, 630 springs and 124 contour faces. The average spring degree at a mass point was 8.24. This is not a very complex 3D model, but it is already too complex to be generated by hand.

We weighted external model energy in the simulation 10 times stronger than the internal energy and used a damping of 2 %. For a detailed review of all these mostly non critical parameters see [10].

5.1   Segmentation Process Evaluation

Table 1 shows the results of the two techniques for initial placement discussed in section 4.1. Correlation was fast (below 1 sec) and mostly correct. It failed only in one case, when the dataset had a bright artefact. PVA had no problems with this artefact, as expected, but was noticeably slower (35 sec). Thus, it was

Table 1. correctness of the initial model placement by means of correlation respectively pattern vector angle (PVA)

Similarity / Dataset	1	2	3	4	5	6	7
Correlation	correct	correct	correct	wrong	correct	correct	correct
PVA	correct	correct	correct	correct	correct	correct	correct

Table 2. effect of phase count on the average automatic segmentation quality on the datasets 1 to 3

Phase Count	False Seg.	Hausdorff.	Aver. Dist.
2-phases approach	27.3 %	1.79 vox	0.410 vox
3-phases approach	26.4 %	1.71 vox	0.398 vox

only selected in datasets with bright artefacts being recognizeable by abnormal high count values.

Several of the parameters for the automatic segmentation process were determined on a subset of the 7 datasets (dataset 1 to 3), so that their values could be verified on the remaining 4 datasets. These tests prove the 3-phases approach to be reasonable, since without the second phase for the approximate model adaptation (section 4.2), the segmentation quality decreases (table 2).
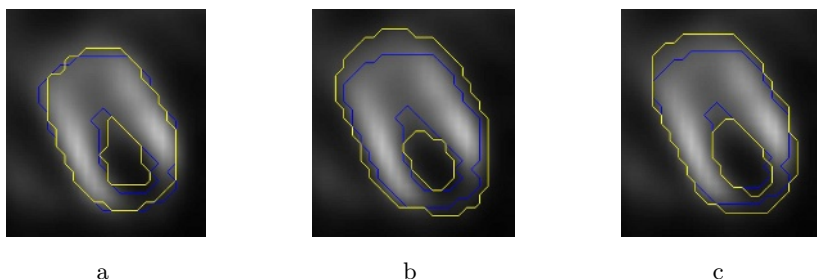
5.2   Segmentation Results

Table 3 shows the results of the automatic segmentation of the 7 datasets. They are compared with the results of the manual segmentations by two experienced users. It can easily be seen, that the automatic technique is much faster than the manual process. Its quality is slightly better. In all cases, it lies within in the inter-observer variance. This fact is illustrated by figure 4, which gives a visual impression of the segmentation quality.

Table 3. average segmentation results of the automatic process and the manual segmentations by two experienced users (A and B)

Segmentation	Duration	False. Seg.	Hausdorff.	Aver. Dist.
Auto. Seg.	25.3 sec	29.4 %	2.21 vox	0.436 vox
Seg. by A	15.6 min	36.5 %	2.81 vox	0.667 vox
Seg. by B	21.4 min	30.6 %	2.39 vox	0.534 vox

Qualitative inspection of the results on the remaining datasets, where no gold standard existed, confirmed these results: Phase 1 always detected a correct initial model position and the next phases adapted the model visually correct to the LV (similar to figure 4).

a   b   c

Fig. 4. results (yellow, bright) of automatic and manual segmentation in comparison to the gold standard (blue, dark). **(a)** automatic segmentation process. **(b)** manual segmentation of user A. **(c)** manual segmentation of user B.

6   Discussion

Our automatic segmentation method found and segmented the LV in all datasets correctly. It needs only a small fraction of the time required to achieve an equivalent result by hand. The 3-phases approach and the direction-weighted contour sensors have proven effective for the successful segmentations. Furthermore, the good results of phase 1 make the interactive initial model placement obsolete.

Our stable 3D mass spring model and the technique to generate it automatically from a sample segmentation were crucial for the good results. It enabled bridging gaps caused by reduced perfusion while being able to adapt to the individual shape of each LV.

The segmentation method was fast and robust on noisy, low-resolution data. In the next step, we will investigate its performance on more complex data.

References

1. Fernandez-Maloigne, C., Rakotobe, R.H., Langevin, F., Fauchet, M.: 3D segmentation and visualization of cardiac SPECT studies. In: 28th AIPR Workshop. Volume 3905 of Proceedings of SPIE. (2000) 222–231
2. Pohle, R., Wegner, M., Rink, K., Tönnies, K., Celler, A., Blinder, S.: Segmentation of the left ventricle in 4D-dSPECT data using free form deformation of super quadrics. In: Medical Imaging: Image Processing. Volume 5370 of Proceedings of SPIE. (2004) 1388–1394
3. Bardinet, E., Cohen, L.D., Ayache, N.: Tracking and motion analysis of the left ventricle with deformable superquadrics. Medical Image Analysis **1** (1996) 129–149
4. Gould, P.L.: Introduction to Linear Elasticity. Springer (1994)
5. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. IJCV **1** (1988) 321–331
6. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models - their training and application. CVIU **61** (1995) 38–59
7. Dornheim, L., Tönnies, K.D., Dornheim, J.: Stable dynamic 3D shape models. In: ICIP. (2005)
8. Dornheim, L., Tönnies, K.D.: Automatische Generierung dynamischer 3D-Modelle zur Segmentierung des linken Ventrikels in 3D-SPECT-Daten. In: Bildverarbeitung für die Medizin. (2005)
9. Kohonen, T.: Self-Organization and Associative Memory. Springer (1987)
10. Dornheim, L.: Generierung und Dynamik physikalisch basierter 3D-Modelle zur Segmentierung des linken Ventrikels in SPECT-Daten. Diplomarbeit, Fakultät für Informatik, Otto-von-Guericke-Universität Magdeburg (2005)