Towards Automatic Full Heart Segmentation in Computed-Tomography Images

O Ecabert[1], J Peters[1], C Lorenz[2], J von Berg[2], M Vembar[3],
K Subramanyan[3], G Lavi[4], J Weese[1]

[1]Philips Research Laboratories, Aachen, Germany
[2]Philips Research Laboratories, Hamburg, Germany
[3]Philips Medical Systems, Cleveland, OH, USA
[4]Philips Medical Systems, Haifa, Israel

Abstract

We present a robust, fast and fully automatic approach enabling the segmentation of the main anatomical structures of the heart in CT images. The proposed method is based on the adaptation of a 3D triangulated mesh to new unknown images exploiting simultaneously knowledge of organ shape and typical gray level appearance in images, both learned from a training database made of 28 data sets. The described approach was tested on more than 50 volume images at different cardiac phases. Visual inspection by experts reveals that the proposed method is overall robust and succeeds in segmenting the heart up to minor interactive local corrections.

1.   Introduction

Automatic segmentation plays a central role when inspecting reconstructed 3D cardiac images (e.g. from CT or MR scanners) [1]. An accurate classification of the different cardiac regions is usually the first step of subsequent tasks like: Visualization, coronary artery inspection, measurement of the ejection fraction for the left and right ventricles, wall motion analysis, intervention planning (e.g. for electro-physiology treatment).

Because a significant amount of information about global functional analysis is delivered by the left ventricle, much work has been dedicated in extracting this structure from medical images. However, volumetric images acquired by emerging imaging techniques like e.g. multislice CT offer much more information, motivating the need for methods that are able to segment the other anatomical structures of the cardiac region.

In this paper, we present a model-based approach capable of extracting the main anatomical structures of the heart. Unlike other works on model-based segmentation, we also concentrated our efforts in procedures enabling the initialization of the model without any click required from the user.

During the training phase, typical boundary appearances are learned from representative images. This allows a very robust boundary detection when segmenting new unknown images. Moreover, although we adapt a single mesh, special attention has been paid in allowing each anatomical part of the model to be free to globally deform in an independent fashion.

2.   Shape-constrained deformable models framework

Deformable models have been widely used for the segmentation of medical images [2]. However, they might have too much flexibility when adapting towards the organ of interest and be sensitive to image artifacts. Using *a priori* information about shape variability to constrain the deformation flexibility has been recognized to improve the robustness of the segmentation process [3]. In this paper, we will use the approach introduced in [4], which constrains the model to remain close to the trained shape while allowing local non–learned deformations to account for the individuality of each patient.

The input of the *shape–constrained* deformable model approach is a mesh with vertices v_j $(j = 0...V)$ connected in T triangles. The mesh is adapted to a new image minimizing an energy function, which is usually made of two contributions. The first one (called external energy) attracts the model to the object boundaries in the image, whereas the internal energy penalizes deviations of the deformed model from the learned shape:

$$E = E_{\text{ext}} + \alpha\, E_{\text{int}}\,, \tag{1}$$

where α is a parameter which balances the contribution of each energy.

Briefly, the external energy can be expressed as

$$E_{\text{ext}} = \sum_{i=1}^{T} w_i \left(\frac{\boldsymbol{\nabla} I(\boldsymbol{x}_i^{\text{target}})}{\parallel \boldsymbol{\nabla} I(\boldsymbol{x}_i^{\text{target}}) \parallel} \cdot (\boldsymbol{x}_i^{\text{target}} - \boldsymbol{c}_i) \right)^2 \tag{2}$$

where the sum is performed over the mesh triangles. The triangle centers c_i are attracted towards target points x_i^{target} detected at the object boundary. The projection of $(x_i^{\text{target}} - c_i)$ onto the normal vector $\nabla I / \parallel \nabla I \parallel$ at the target point makes the energy invariant to movements of the triangle within the object tangent plane, avoiding that the point keeps stuck at the target position. The target points are detected in the image maximizing a feature function as explained in subsection 3.1 in more details. Finally, the weights w_i are large for reliably detected target points and small otherwise.

The internal energy is given by

$$E_{\text{int}} = \sum_{j=1}^{V} \sum_{k \in N(j)} ((v_j - v_k) - (\mathcal{T}[m_j] - \mathcal{T}[m_k]))^2 \quad (3)$$

with $N(j)$ the set of indices of the neighbor vertices of vertex v_j, and m_j the vertex coordinates of the reference mesh undergoing a geometric transformation $\mathcal{T}[.]$. This transformation describes allowed global shape variations like e.g. rigid body motions or deformations given by a principal component analysis (PCA) learned from a training set of shapes.

3. Heart model

The mesh representing the heart model was built using simple geometric entities and a set of basic operations to combine them as explained in [5]. The result of these operations is a mesh made of 14771 triangles including seven distinct anatomical parts, namely the atria and ventricles, the trunk of the aorta and pulmonary artery as well as the myocardium around the left ventricle. Segments of the superior and inferior vena cava as well as of four pulmonary veins were also modeled at the respective atria.

The model was thereafter semi-automatically adapted by experts to 8 different patients at different cardiac phases yielding a training data set with a total of 28 mesh configurations. The final model geometry could be computed as the mean mesh of these different representations (fig. 1).

3.1. Learning boundary appearances

Images acquired by CT scanners are calibrated, and the different types of tissues can be mapped into characteristic gray–value intervals that can be learned from the training data sets along with their corresponding meshes.

The target points in eq. (2) are usually detected by maximizing a feature function along a profile perpendicular to the mesh triangles. Since organ boundaries are related to gray-level transitions, a common feature function used is the gradient operator. However, the gradient operator alone will return a high response at every gray–level transition in the image irrespective of the adjacent organs.
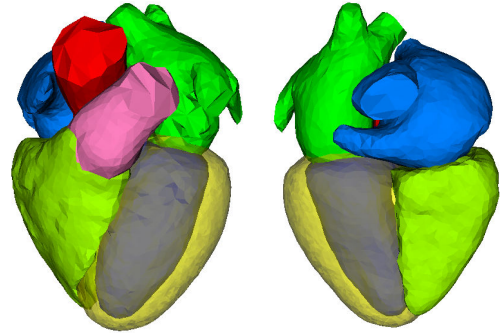

Figure 1. Surface rendering of the heart model used for segmentation. Each anatomical part is coded with another color.

In order to make the feature search much more discriminative, we propose to use a feature function of the type [6]:

$$F(c_j, x) = \begin{cases} \pm n(c_j) \cdot \nabla I(x) & \text{if } Q_i \in [\min_i, \max_i] \\ & \text{for all } Q_i \\ 0 & \text{if } Q_i \notin [\min_i, \max_i] \\ & \text{for any } Q_i \end{cases}$$

The core of this feature function is the image gradient $\nabla I(x)$ projected onto the triangle normal n, which intrinsically distinguishes between dark-to-bright and bright-to-dark transitions (together with sign $\pm$). To distinguish boundaries between different types of tissues we may add one or more clipping criteria by requiring that some quantities Q_i lie within certain trained clipping intervals $[\min_i, \max_i]$ (e.g. gray value at x or mean gray values on both sides of the mesh).

The basic idea to assign the optimal function to each triangle is to *simulate* the boundary detection for the given 3D images and to select those features that minimize the distance between the detected position and the desired object boundary [6].

3.2. Modeling shape variability

Instead of modeling the shape variability for the whole model with a PCA, we propose to assign an affine transformation to each of the anatomical parts of the mesh. The trunk of the arteries and veins are sufficiently short to be still efficiently modeled by a linear transformation. In order for the mesh to stay smooth at the transitions between the anatomical regions, we linearly interpolate the respective contributions of the affine transformations. The geometric transformation of eq. (3) is therefore given by

$$\mathcal{T}_{\text{Affine}}[m_j] = \sum_k \lambda_k \cdot (A_k \cdot m_j + b_k)$$

with A_k and b_k the linear matrix and translation vector of the affine transformation for the anatomical part k, and

λ_k the transformation weights such that $0 \le \lambda_k \le 1$ and $\sum_k \lambda_k = 1$.

This kind of shape variability modeling is motivated by the fact that variations in size and/or shape of the chambers can be described independently either for a single case over the cardiac cycle or among different patients.

4. Segmentation chain

The fully automatic segmentation procedure can be split into three steps, building up a coarse-to-fine segmentation chain. First, the location of the heart is roughly estimated in the image and the model is translated to that position. Second, the initialization can be improved allowing the whole model to deform only according to global transformations. Finally, fine segmentation can be reached finding the vertex configuration that minimizes eq. (1).

4.1. Heart detection

The first step consists of finding the center of mass of the heart as well as a coarse estimate of its size casting rays through the volumetric image and analyzing the gray value properties along these rays as introduced in [7].

First, rays are cast along x and only those rays that encounter a lung-muscle transition and have no Hounsfield units (HU) below -500 for at least 60 mm are selected for subsequent processings. This procedure is then repeated for y and z. Finally, for each of the intersection point, a sphere is fit to image such that none of the included intensities is lower than -500 HU. The center of mass of the sphere with maximum radius is interpreted as the heart location. Moreover, the maximum radius delivers information about the size of the heart present in the image.

4.2. Model initialization

To improve the model initialization and account for pose, as well as relative sizes of the model parts, we propose to constrain the vertex to adapt according to global deformations only. The internal energy term in eq. (1) becomes therefore superfluous and can be omitted. The problem can be re-written as

$$E = E_{\text{ext}}(\boldsymbol{v}(\boldsymbol{q})) \qquad (4)$$

with $\boldsymbol{q}$ a vector of parameters describing the global transformation. The goal is now to find the parameter vector $\boldsymbol{q}$ such that eq. (4) is minimized. Efficient implementation can be done with a Gauss-Newton optimization method [8].

To perform a robust, but nevertheless fast initialization, we propose to first minimize (4) assuming a single rigid transformation and then multiple affine transformations as

introduced in sub-section 3.2. Again, the boundary detection is performed as explained in sub-section 3.1. The rigid transformation will compensate for slight misalignments, rotations and scaling errors. The multiple affine transformations will resize and deform globally each piece of the model to the actual anatomy.

4.3. Segmentation

Finally, at that point, the model is usually well enough initialized to proceed with the minimization of eq. (1) until the mesh converges to a steady state. For this last adaptation, all vertex coordinates are considered as degrees of freedom. Feature detection and shape variability modeling are still implemented as described in sub-sections 3.1 and 3.2.

5. Results

This section shows the results of the proposed segmentation algorithm for some data sets. The parameter α was set to 30 HU/mm^3 and 10 iterations were performed for both the rigid and multiple affine transformations during initialization (sub-section 4.2). The feature functions were learned as described in [6].

Fig. 2 and 3 illustrate the segmentation results of the proposed fully automatic method for a wide range of images. Please note that none of the images shown in this section was part of the training database.
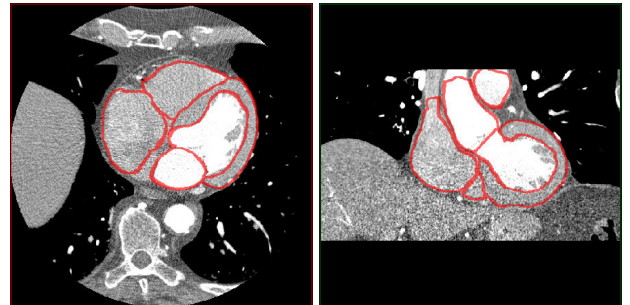

Figure 2. Fully automatic segmentation of a noisy cardiac data set in axial (left) and coronal (right) views.

Fig. 2 shows an axial and coronal views of a rather noisy data set. The proposed approach could successfully segment this image, even in presence of a large amount of noise. The influence of the model was especially noticeable at the bottom of the heart (fig. 2 right) in regions of low contrast.

Segmentation of transplanted heart images could also be performed successfully as illustrated on fig. 3 (top left). The remaining pictures demonstrate that the variability of the gray level distributions at the boundaries could be effectively learned during the training phase. Indeed, there
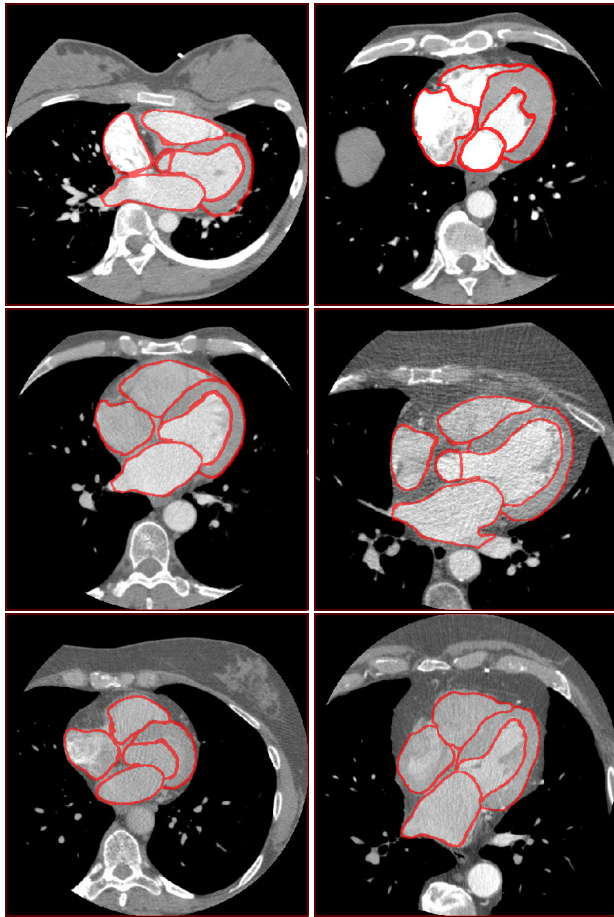
Figure 3. Fully automatic segmentation of six cardiac data sets. Top left: Transplanted heart, top right: Thick myocardium, middle: Medium contrast images, bottom left and bottom right: Low contrast images.

are strong variations in contrast for the right chambers between the top and bottom rows. Contrast variability can also be present within the right atrium for a single patient (fig. 3 bottom left).

The described approach was tested on more than 50 data sets at different cardiac phases. Visual inspection by experts shows that the proposed method is overall robust and succeeds in segmenting the heart up to minor interactive local corrections. The results are generally accurate enough for numerous applications, but the segmentation might still require some user interactions at complex and variable structures like the pulmonary veins.

Most of the failures are attributed to inaccuracies during steps 1 and 2, which could be solved by a manual initialization procedure. The whole segmentation chain runs in about 20 seconds on a workstation with Dual-Xeon Hyper-Threading Intel processors (2 × 2.4 GHz) and 1 GByte of RAM.

6.   Discussion and conclusions

In this article, we introduced a model-based approach to automatically and robustly segment the heart in CT images. The model is deformed under the influence of forces which pull it towards boundaries in the image, while deviations from a pre-defined shape are penalized. The typical boundary appearances are automatically learned from a training data base as well as the mean mesh representation. The shape variability is simply encoded in the model, letting each anatomical part globally deform according to a linear transformation.

We especially present a combination of three methods which allow to automatically segment the main cardiac structures in an image. In the first step, the model is translated to the potential heart position in the image. In the following steps, it is first adapted globally, then locally by progressively relaxing constraints on the allowed deformations. Although we used a model-based approach, it showed to be able to cope well with large deformation and complex structures in images.

Although the proposed method yields excellent visual results, it should be quantitatively evaluated. A leave-one-out evaluation based on our training data sets is planned.

References

[1] Frangi AF, Niessen WJ, Viergever MA. Three-dimensional modeling of functional analysis of cardiac images: A review. IEEE Trans Medical Imaging Jan. 2001;20(1):2–25.
[2] McInerney T, Terzopoulos. D. Deformable models in medical image analysis. Medical Image Analysis 1996;1(2):91–108.
[3] Cootes TF, Taylor CJ, Cooper DH, Graham J. Active shape models – their training and application. Computer Vision and Image Understanding Jan. 1995;61(1):38–59.
[4] Weese J, Kaus M, Lorenz C, Lobregt S, Truyen R, Pekar V. Shape constrained deformable models for 3D medical image segmentation. In Proc. IPMI. 2001; 380–387.
[5] von Berg J, Lorenz C. Multi-surface cardiac modelling, segmentation, and tracking. In Proc. FIMH. 2005; 1–11.
[6] Peters J, Ecabert O, Weese J. Feature optimization via simulated search for model-based heart segmentation. In Proc. CARS. 2005; 33–38.
[7] Lorenz C, von Berg J. Fast automated object detection by recursive casting of search rays. In Proc. CARS. 2005; 230–235.
[8] Gill PE, Murray W, Wright MH. Practical Optimization. Academic Press, 1982.

Address for correspondence:

Olivier Ecabert
Philips Research Laboratories
Weisshausstr. 2, 52066 Aachen, Germany
olivier.ecabert@philips.com