

Multi-surface Cardiac Modelling, Segmentation, and Tracking

Jens von Berg and Cristian Lorenz

Philips Research Laboratories, Sector Technical Systems,
Röntgenstr. 24-26, 22335 Hamburg, Germany
Jens.von.Berg@philips.com

Abstract. Multi-slice computed tomography image series are a valuable source of information to extract shape and motion parameters of the heart. We present a method how to segment and label all main chambers (both ventricles and atria) and connected vessels (arteries and main vein trunks) from such images and to track their movement over the cardiac cycle. A framework is presented to construct a multi-surface triangular model enclosing all blood-filled cavities and the main myocardium as well as to adapt this model to unseen images, and to propagate it from phase to phase. While model construction still requires a reasonable amount of user interaction, adaptation is mostly automated, and propagation works fully automatically. The adaptation method by deformable surface models requires a set of landmarks to be manually located for one of the cardiac phases for model initialisation.

1 Introduction

The aim of our work is a comprehensive model of the geometry of the human heart contraction as well as its inter-individual variations. Such a model introducing a priori knowledge about typical properties of a beating heart will be highly beneficial in the whole chain of image-based cardiac diagnostics, as well as in many cardiac treatment procedures. The model covers landmarks, the coronary tree, and the surfaces of the large vessels [1]. The latter is the subject of the work reported here. The most valuable and practically unique source of information for the modelling process are cardiac images from clinical practice. In this paper the use of multi-slice computed tomography (MSCT) images is reported that have a voxel size of about 0.5 mm in each direction and a temporal resolution of 10 volumes per cardiac cycle. Mostly, cardiac MRI were used previously for this purpose [2, 3, 4, 5, 6]. MSCT may provide even better insight into the morphology of the human heart [7]. Extracting the relevant information from these images is hardly feasible without a priori knowledge [8]. Many approaches to cardiac segmentation were based on manually segmented images, which is a good means to both tune parameters by automated supervised learning, and to finally prove their performance in comparison to human expertise. However, manually segmenting an MSCT series means delineating each object

of interest in about two thousand images. This dilemma led us to a bootstrap approach with a consecutive refinement of the model during successive analysis of new images.

2 Model Construction

The model covers the blood pool of both the left and the right heart. The blood pools of the ventricles should be distinguishable from those of the atria. All attached vessels should also be modelled, i.e. the aorta, the pulmonary artery, the vena cava, and the pulmonary vein trunks. As it is clearly visible, and diagnostically relevant, also the left myocardium should be represented in the model. Including adjacent volumetric entities required a surface modelling scheme beyond two-dimensional manifolds. In the discrete case with triangular faces this means that there are faces with more than three neighbours wherever multiple surfaces share an edge. In order to enable multi-scale / multi-resolution approaches or to just find an ideal trade-off between accuracy and complexity, a multi-resolution representation of the surface discretisation was desired.

The initial step was the construction of single basic shapes like spheres (atria), tubes (attached vessels), and opened ellipsoids (ventricles). Each one modelled an anatomical entity. These shapes were then positioned in the training image and adapted to the corresponding entities. A re-sampling closed this step to get a defined level of granularity. The third step was the most important one that combines the single basic two-dimensional manifolds to form the multi-surface model. The method used in this third step is explained in some detail below. In the resulting model, each face should be assigned a label that indicates the anatomical structure it belongs to. This information was derived by storing which of the initial basic shapes a face originates from. The basic shapes were *left atrium* a_l , *left ventricle endocardium* (inner part v_i), *left ventricle epicardium* (outer part v_o), *aorta* a , *vena cava superior* v_s , *vena cava inferior* v_l , *right atrium* a_r , *right ventricle* v_r , *pulmonary artery* a_p (right branch only), and the pulmonary vein trunks (v_1, v_2, v_3, v_4) that drain into the left atrium.

2.1 Building a Multi-surface Model

The combination step was made by successive application of a handful of basic operations on surface meshes, starting with the basic meshes. There are volumetric set operations that consider the enclosed volume of two meshes, apply the union (\cup) or the difference (\setminus) operation on them, and yield the resulting surface mesh. A similar approach but with implicit surface models was proposed in [9]. Each of the present operations was defined as $\mathbf{B} \times \mathbf{B} \rightarrow \mathbf{B}$, where $m \in \mathbf{B}$ is a two-dimensional manifold mesh. As a further constraint on these operations, the intersection line between both meshes had to be closed polygons. This required open basic meshes to fully overlap with their neighbours (e.g. ventricle with atrium). The join operator (\diamond) was defined as $\mathbf{B} \times \mathbf{B} \rightarrow \mathbf{M}$, where $\bar{m} \in \mathbf{M}$ may be a non-two-dimensional manifold mesh. The join operator just unites




Fig. 1. A basic (a), some intermediate (b to d), the final (e), and a derived (f) sub mesh, all seen from left anterior. a: *left atrium* a_l , b: *left blood pool* p_l , c: *left myocardium* a_1 , e: whole heart \bar{h} , f: blood pool of the *left ventricle* only p_{v_i} . The colour-coding denotes face labels. All shown meshes are just subsets of the complete multi-surface model \bar{h}

both sets of faces and unifies corresponding faces that occur in both meshes. The unary operator c_l ($\mathbf{M} \rightarrow \mathbf{M}$) removes all edges smaller than given by parameter l and preserves the triangles' labels. It was needed to replace auxiliary triangles created by volumetric set operations. The left blood pool $p_l \in \mathbf{B}$ was build by

$$p_l = c_l(v_1 \cup v_2 \cup v_3 \cup v_4 \cup a_l \cup v_i \cup a). \quad (1)$$

In order to construct the complete multi-surface model \bar{h} , first the intermediate meshes a_1 and a_2 were constructed by

$$a_1 = c_l(v_o \setminus p_l) \quad (2)$$

$$a_2 = c_l(v_o \cup p_l), \quad (3)$$

where a_1 now exactly enclosed the left myocardium. The complete left heart model \bar{h}_l was built by

$$\bar{h}_l = a_1 \diamond a_2, \quad (4)$$

and the right blood pool was built by

$$p_r = c_l(a_r \cup v_r \cup a_p \cup v_l \cup v_s) \setminus a_1. \quad (5)$$

Left and right part were fused to

$$\bar{h} = \bar{h}_l \diamond p_r. \quad (6)$$

In Figure 1 some basic, intermediate, and the final mesh \bar{h} are rendered from the same viewing position. The edge size was set to range between 2.5 mm and

5 mm. Also other sub meshes than those required to build the final mesh may be constructed, for instance the blood pool of the left ventricle excluding the left atrium by

$$p_{v_i} = v_i \setminus a_l. \quad (7)$$

3 Model Adaptation

For adaptation of the multi-surface model to a cardiac CT image, a shape-constrained deformable surface model approach was followed as previously described in [10, 4]. The model with given vertex positions $\hat{\mathbf{v}}$ taken from a training image both served for the initialisation of the initial mesh \mathbf{v}^0 and as constraint during its adaptation to $\mathbf{v}^{1 \dots n}$. The number of triangles remained unchanged in this process.

3.1 Affine Pre-registration

In order to pose the initial mesh into the image as accurately as possible, 25 anatomical landmarks were manually located both for the image the model was built from and for each target image [1]. These landmarks are mainly centre locations of chambers, valves, and ostia. A point-based affine registration [11] was applied on the two sets of landmarks. The resulting affine matrix \mathbf{A} and translation vector \mathbf{t} were then used to pre-register the initial mesh by $\mathbf{v}^0 = \mathbf{A}\hat{\mathbf{v}} + \mathbf{t}$.

3.2 Model Deformation

In the optimisation scheme the vertex positions of the triangular surface mesh were the parameters to be varied. Mesh deformation was done by minimizing the energy term

$$E = E_{ext} + \alpha E_{int}. \quad (8)$$

The external energy E_{ext} drives the mesh towards the surface points obtained in a surface detection step. The internal energy E_{int} restricts the flexibility by maintaining the vertex configuration of a shape model. The parameter α weights the influence of both terms. A fixed number n of such minimisation steps is performed on the mesh. The different components of the deformation algorithm are described below.

Surface Detection. Surface detection was carried out for each triangle barycentre \mathbf{x}_i . Within a sampling grid of points \mathbf{c}_k , defined in a local co-ordinate system, that point $\tilde{\mathbf{c}}_i$ is chosen that maximizes the objective function

$$\tilde{\mathbf{c}}_i = \operatorname{argmax}_{k=-l, \dots, l} \left\{ F_i(\mathbf{x}_i + \mathbf{M}_i \mathbf{c}_k) - \delta \|\mathbf{c}_k\|^2 \right\}. \quad (9)$$

\mathbf{M}_i is a rotational matrix that rotates the z -axis of the local co-ordinate system to the triangle surface normal \mathbf{n}_i and

$$\tilde{\mathbf{x}}_i = \mathbf{x}_i + \mathbf{M}_i \tilde{\mathbf{c}}_i \quad (10)$$

is the new surface point for \mathbf{x}_i . The parameter δ controls the trade-off between feature strength and distance. The sampling grid

$$\mathbf{c}_k \equiv \mathbf{G}_L = (0, 0, k\varepsilon) : k = -l, \dots, l \quad (11)$$

was used, that results in $(2l + 1)$ equidistant sampling points along the triangle surface normal.

Feature Function. The feature function

$$F_i(\mathbf{x}) = \begin{cases} -\mathbf{n}_i^t \nabla I(\mathbf{x}) \frac{g_{max}(g_{max} + \|\nabla I(\mathbf{x})\|)}{g_{max}^2 + \|\nabla I(\mathbf{x})\|^2} & : I_{min} < I(\mathbf{x}) < I_{max} \\ 0 & : otherwise \end{cases} \quad (12)$$

was used that projects the image gradient $\nabla I(\mathbf{x})$ onto the face normal \mathbf{n}_i and damps its value so that surface points with image gradients stronger than g_{max} do not give higher response. The restriction to a dedicated intensity range may make the feature function more specific and thus makes adaptation less vulnerable to adjacent false attractors (see below).

External Energy. The external energy

$$E_{ext} = \sum_i w_i (\mathbf{e}_{\nabla I} \tilde{\mathbf{c}}_i)^2, w_i = \max \left\{ 0, F_i(\mathbf{x}_i + \mathbf{M}_i \tilde{\mathbf{c}}_i) - \delta \|\tilde{\mathbf{c}}_i\|^2 \right\} \quad (13)$$

drives each triangle barycentre \mathbf{x}_i towards the detected surface point $\tilde{\mathbf{x}}_i$. $\mathbf{e}_{\nabla I}$ is the unit vector in the direction of the image gradient at the surface point $\tilde{\mathbf{x}}_i$. Since only the projection onto $\mathbf{e}_{\nabla I}$ is penalized, this allows the triangle centre to locally slide along an iso-contour. This method proved to be superior to direct attraction by the candidate in [10] in case of intermediate false attractions.

Internal Energy. The internal energy

$$E_{int} = \sum_j \sum_{k \in N(j)} ((\hat{\mathbf{v}}_j - \hat{\mathbf{v}}_k) - s\mathbf{R}(\mathbf{v}_j - \mathbf{v}_k))^2 \quad (14)$$

preserves shape similarity of all mesh vertices \mathbf{v}_i to the model vertices $\hat{\mathbf{v}}_i$. $N(j)$ is the set of neighbours of vertex j . The neighbouring vertices are those connected by a single triangle edge. The scaling factor s and the rotational matrix \mathbf{R} are determined by a closed-form point-based registration method based on a singular value decomposition [11] prior to calculation of (14).

Optimisation. As only interdependences between neighbour vertices exist (14) and the energy terms are of a quadratic form, the conjugate gradient method [12] could be used for minimisation of (8) with a sparsely filled matrix.

Multi-surface Parameterisation. The labels assigned to each face of the multi-surface model may be used to parameterise interfaces between different anatomical entities specifically. However, dedicated parameter tuning was restricted to the epicardium border towards the lung parenchyma that differs in its appearance significantly from the other surfaces that enclose the blood pool. Thus, g_{max} was set to $60 \frac{\text{HU}}{\text{mm}}$ here instead of $120 \frac{\text{HU}}{\text{mm}}$ elsewhere. and the intensity range ($I_{min} \cdots I_{max}$) was adjusted to $350 \cdots 800$ HU instead of 1000 HU and up. The other parameters were globally set to $\alpha = \delta = 1$, $\varepsilon = 1$ mm, $l = 10$.

4 Surface Tracking

In order to capture tissue trajectories one has to find corresponding tissue landmarks in images from different cardiac phases. This was mainly done previously either by non-linear registration [13] or by active appearance models [14]. With some modifications that rather belong to the second category and that are explained below, the adaptation method presented above was also applied for a surface tracking approach that utilizes point correspondence.

Surface detection is carried out following equation (9). In order for a surface point not only to be attracted along the surface normal, a sampling grid is used that extends into all direction. A multi-icosahedron grid

$$\mathbf{c}_{\mathbf{k}=1 \dots 37} \equiv \mathbf{G}_{\mathbf{I}} = \{(0, 0, 0), P_2, P_4, P_8\} \quad (15)$$

was used where each P_n is a set of 12 icosahedron surface points with a radius of $n\varepsilon$ mm around the origin of the local coordinate system. Individual feature functions are required for each surface point in this case to take the local image properties into account. The feature function

$$F_i(\mathbf{x}) = \frac{2l + 1}{\sum_{\mathbf{k}=-l \dots l} (I(\mathbf{x} + \mathbf{M}_i \mathbf{s}_{\mathbf{k}}) - \hat{g}_{i,k})^2} \quad (16)$$

that replaces (12) thus evaluates similarity of local appearance samples to the once learnt model $\hat{\mathbf{g}}_i$. The linear sampling grid $\mathbf{s}_{\mathbf{k}} = \mathbf{G}_{\mathbf{L}}$ from equation (11) is taken. It is applied at each sample point. The external energy is calculated by

$$E_{ext} = \sum_i \tilde{\mathbf{c}}_i^2 \quad (17)$$

instead of (13), and the internal energy is taken from (14).

In order to demonstrate the general feasibility of surface tracking with deformable models and appearance models a simple test study was carried out. A cylindrical surface mesh was posed into a cardiac CT image to roughly fit the left myocardium. The image appearance $\hat{\mathbf{g}}$ was learnt and the mesh was rotated around its main axis by $r = \pm \frac{\pi}{16}, \pm \frac{\pi}{8}, \pm \frac{\pi}{4}$, and $\pm \frac{\pi}{2}$ afterwards in a number of trials. An adaptation with $n = 80$ iterations was performed for each rotation angle. Up to $r = \pm \frac{\pi}{4}$ the mesh successfully recovered the initial position. This simple

test showed that the tracking method has a remarkable capture range, and that with the rotationally symmetric shape model, appearance alone is sufficient as driving force.

The propagation through all cardiac phases started with a phase to which the shape model was adapted successfully. From the image of this phase 1 the appearance $\hat{\mathbf{g}}_1$ was learnt at positions given by \mathbf{v}_1 . This shape \mathbf{v}_1 was used as initial mesh and further on $\mathbf{v}_i^0 = \mathbf{v}_{i-1}^n$. The same holds for the shape model $\hat{\mathbf{v}}_i = \mathbf{v}_{i-1}^n$. This was repeated until all phases were processed. For each phase the initial appearance model $\hat{\mathbf{g}}_1$ was used.

5 Results

5.1 Model Construction

A multi-surface mesh with a total number of about 7,000 vertices and 13,000 triangle faces was constructed with edge lengths ranging between 2.5 mm and 5 mm. Its shape is shown in Figure 1. The basic meshes this model was constructed from, were adapted to the anatomical entities of the end-diastolic phase of the training image. The resulting multi-surface model was then adapted to a set of five other cardiac MSCT images from different hospitals but all acquired with a Philips MX8000 IDT 16-line CT scanner. The images were contrast-enhanced as they were acquired for the purpose of coronary assessment.

5.2 Pre-registration

The affine pre-registration led to a mean (\pm standard deviation) residual landmark distance of 7.5 ± 4.3 mm, 7.0 ± 3.3 mm, 8.3 ± 3.5 mm, 5.6 ± 2.6 mm, and 13.0 ± 12.1 mm for the five images. The latter resulted in an unacceptable pre-registration, both visually and with respect to the subsequent adaptation result. An alternative rigid registration with an isotropic scale parameter




Fig. 2. Left: Pre-registration (dark mesh) and subsequent automatic adaptation (light mesh) of the multi-surface model to an unseen image. The arrow marks a local misadaptation. Right: Mean end-diastolic model of the five patients with colour-coded standard deviation (dark:0.8 mm, light:7.8 mm)

(13.0 ± 14.4 mm) resulted in an acceptable pre-registration for further processing. A typical affine pre-registered model is shown in Figure 2 in comparison to the automatic deformable adaptation based on this pre-registration.

5.3 Adaptation

The adaptation by model-based deformation significantly improved on the results of the affine pre-registration (Figure 2). Automatic adaptation with $n = 10$ iterations took about 15 seconds on a 2.6 GHz PC including real-time surface rendering. The majority of surface parts could be considered well-adapted. The reasons for remaining local mis-adaptations were mainly adaptations to false attractors e.g. of the epicardium mesh to the endocardium (see Figure 2) or to coronaries, and of the aorta mesh to the vena cava. Using the methods described in [15] manual corrections that survive subsequent automatic adaptation steps could be applied to these mis-adapted parts.

5.4 Calculating a Mean Model

The resulting individualized models were mutually registered (rigid plus isotropic scale) using a procrustes analysis of their corresponding anatomical landmarks. A mean model of the five subjects was calculated (Fig. 2).

5.5 Surface Tracking

The surface tracking method was applied to the training image sequence. The initial mesh \mathbf{v}_1 was the one that resulted from model construction and that was fit to the end-diastolic phase image of the training sequence. Each propagation step $\mathbf{v}_{i-1} \rightarrow \mathbf{v}_i$ was done with $n = 12$ iterations. Propagation was done for all nine images of subsequent phases and back to the initial image with $\mathbf{v}_1^0 = \mathbf{v}_9^{12}$. This allows for a comparison of the round-trip adaptation result \mathbf{v}_1^{12} with the initial mesh \mathbf{v}_1 . The mean (\pm standard deviation) distance of corresponding vertices between both was 1.4 ± 0.7 mm. The meshes are shown in Figure 3. The mean distance of all corresponding vertices in all phases between the forward ($\mathbf{v}_{i-1} \rightarrow \mathbf{v}_i$) and the backward ($\mathbf{v}_{i+1} \rightarrow \mathbf{v}_i$) propagation was 2.1 ± 1.3 mm. The propagation from the initial mesh ($\mathbf{v}_1 \rightarrow \mathbf{v}_i$) differs from forward propagation by




Fig. 3. Result of the consistency test: Initial mesh (white) and result of a round-trip adaptation (black) to the end-diastolic initial image




Fig. 4. Left: Trajectories of each vertex through the cardiac cycle. For visibility reasons they were scaled down by a factor 4 with respect to their initial (end-diastolic) vertex position. Centre: end-diastolic mesh. Right: end-systolic mesh

1.6 ± 1.0 mm and from backward propagation by 1.5 ± 0.9 mm. A visualisation of the moving model by a surface rendering loop gives a very natural impression of contraction (ventricles), parallel displacement (valve plane), and rather stable parts (atria). Figure 4 tries to show the results in a printed form.

6 Discussion

A method was presented that enables a widely automated construction of a multi-surface triangular mesh of cardiac chambers and vessels, mostly automatic adaptation to individual MSCT images, and automatic propagation of such an individualized model through the cardiac phases. For model construction a set of single basic shapes was adapted each to its anatomical entity. The multi-surface model resulted from their automatic combination. Some anatomical landmarks were manually located in order to pre-register this model to an unseen image by either affine or rigid registration. The subsequent deformation to fit the image boundaries was mainly gradient-based. All parameters were set explicitly during an explorative test phase resulting in a small knowledge base. Some individual surfaces of the model were parameterised specifically, which was well supported by the anatomical labels given in the multi-surface model. The surface tracking however used individual grey value profiles for each surface location learnt from the initial phase's mesh. This method was chosen in order to closely approximate the real tissue trajectories. Due to the large data volume (up to two thousand slices for a cardiac cycle) and the difficulties in manually finding reliable trajectories there was no high quality expert data available for validation. We were able to perform a capture range test and a consistency test of the method with good results. Also the animated visualisation gave a realistic impression. Only

the rotational component and the twist of the left ventricle seemed to be underestimated, which we suppose to be due to the too rigid regularisation in (14).

Acknowledgements

We would like to thank our colleagues from PMS–CT Cleveland and PMS–CT Haifa for the abundance of cardiac MSCT images. We also thank our colleagues from PMS–MIT, PMS–MR, and PMS–XRD in Best for many fruitful discussions.

References

- [1] Lorenz, C., von Berg, J.: Towards a comprehensive geometric model of the heart. In: *Functional Imaging and Modeling of the Heart*, to appear, Springer–Verlag (2005)
- [2] Frangi, A.F., Rueckert, D., Schnabel, J.A., Niessen, W.J.: Automatic construction of multiple–object three–dimensional statistical shape models: Application to cardiac modeling. *IEEE Trans. Med. Imag.* **21** (2002) 1151–1166
- [3] Lorenzo–Valdés, M., Sanchez–Ortiz, G.I., Mohiaddin, R., Rückert, D.: Atlas-based segmentation and tracking of 3D cardiac MR images using non-rigid registration. In: *Proc. of MICCAI*, Springer–Verlag (2002) 642–650
- [4] Kaus, M.R., von Berg, J., Weese, J., Niessen, W., Pekar, V.: Automated segmentation of the left ventricle in cardiac MRI. *Med. Img. Anal.* **8** (2004) 245–254
- [5] Lötjönen, J., Kivistö, S., Koikkalainen, J., Smutek, D., Lauerma, K.: Statistical shape model of atria, ventricles and epicardium from short– and long–axis MR images. *Med. Img. Anal.* **8** (2004) 371–386
- [6] Pilgram, R., Fritscher, K.D., Schubert, R.: Modeling of the geometric variation and analysis of the right atrium and right ventricle motion of the human heart using PCA. In: *Proc. of CARS*, Elsevier (2004) 1108–1113
- [7] Chen, T., Metaxas, D., Axel, L.: 3D cardiac anatomy reconstruction using high resolution CT data. In: *Proc. of MICCAI*, Springer–Verlag (2004) 411–418
- [8] Frangi, A.F., Niessen, W.J., Viergever, M.A.: Three–dimensional modeling for functional analysis of cardiac images: A review. *IEEE TMI* **20** (2001) 2–25
- [9] Leliveldt, B.P.F., van der Geest, R.J., Rezaee, M.R., Bosch, J.G., Reiber: Anatomical model matching with fuzzy implicit surfaces for segmentation of thoracic volume scans. *IEEE Trans. Med. Imag.* **18** (1999) 218–230
- [10] Weese, J., Kaus, M., Lorenz, C., Lobregt, S., Truyen, R., Pekar, V.: Shape constrained deformable models for 3D medical image segmentation. In: *Proc. of IPMI*, Springer–Verlag (2001) 380–387
- [11] Golub, G., van Loan, C.: *Matrix Computation*. 3rd edn. John’s Hopkins University Press, Baltimore, MD, USA (1996)
- [12] Gill, P., Murray, W., Wright, M.: *Practical Optimization*. Academ. Press (1981)
- [13] Wierzbicki, M., Drangova, M., Guiraudon, G., Peters, T.: Validation of dynamic heart models obtained using non–linear registration for virtual reality training, planning, and guidance of minimally invasive cardiac surgeries. *Med. Img. Anal.* **8** (2004) 387–401

- [14] Mitchell, S.C., Bosch, J.G., Lelieveldt, B.P.F., van der Geest, R., Reiber, J.H.C., Sonka, M.: 3-D active appearance models: Segmentation of cardiac MR and ultrasound images. *IEEE Trans. Med. Imag.* **21** (2002) 1167–1178
- [15] Timinger, H., Pekar, V., von Berg, J., Dietmeyer, K., Kaus, M.: Integration of interactive corrections to model-based segmentation algorithms. In: *Proc. Bildverarbeitung für die Medizin*, Springer-Verlag (2003) 11–15