



Segmentation of the heart and great vessels in CT images using a model-based adaptation framework

Olivier Ecabert^a, Jochen Peters^a, Matthew J. Walker^b, Thomas Ivanc^b, Cristian Lorenz^c, Jens von Berg^c, Jonathan Lessick^d, Mani Vembar^b, Jürgen Weese^{a,*}

^a Philips Research Europe – Aachen, X-ray Imaging, 52062 Aachen, Germany

^b Philips Healthcare, CT Clinical Science, Cleveland, OH 44143, USA

^c Philips Research Europe – Hamburg, Digital Imaging, 22335 Hamburg, Germany

^d Philips Healthcare, CT Engineering, 31004 Haifa, Israel

ARTICLE INFO

Article history:

Received 20 July 2010

Received in revised form 4 June 2011

Accepted 7 June 2011

Available online 16 June 2011

Keywords:

Heart segmentation

Model-based segmentation

Shape variability

Computed tomography

ABSTRACT

Recently, model-based methods for the automatic segmentation of the heart chambers have been proposed. An important application of these methods is the characterization of the heart function. Heart models are, however, increasingly used for interventional guidance making it necessary to also extract the attached great vessels. It is, for instance, important to extract the left atrium and the proximal part of the pulmonary veins to support guidance of ablation procedures for atrial fibrillation treatment. For cardiac resynchronization therapy, a heart model including the coronary sinus is needed.

We present a heart model comprising the four heart chambers and the attached great vessels. By assigning individual linear transformations to the heart chambers and to short tubular segments building the great vessels, variable sizes of the heart chambers and bending of the vessels can be described in a consistent way. A configurable algorithmic framework that we call adaptation engine matches the heart model automatically to cardiac CT angiography images in a multi-stage process. First, the heart is detected using a Generalized Hough Transformation. Subsequently, the heart chambers are adapted. This stage uses parametric as well as deformable mesh adaptation techniques. In the final stage, segments of the large vascular structures are successively activated and adapted. To optimize the computational performance, the adaptation engine can vary the mesh resolution and freeze already adapted mesh parts.

The data used for validation were independent from the data used for model-building. Ground truth segmentations were generated for 37 CT data sets reconstructed at several cardiac phases from 17 patients. Segmentation errors were assessed for anatomical sub-structures resulting in a mean surface-to-surface error ranging 0.50–0.82 mm for the heart chambers and 0.60–1.32 mm for the parts of the great vessels visible in the images.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Numerous algorithms for heart segmentation have been published in the past (see Frangi et al. (2001) for a review). Most of the algorithms have been designed for the segmentation of the left and/or right ventricle. Respective methods comprise image driven approaches using thresholds and clustering (Redwood et al., 2005), fitting of simple geometric models to images (Cauvin et al., 1993), active shape models (Fritz et al., 2005; van Assen et al., 2006, 2008), active appearance models (Mitchell et al., 2001, 2002; Stegmann and Pedersen, 2005; Zhang et al., 2010), and flexible deformable models that allow local deformations subject to a

smoothness constraint (McInerney and Terzopoulos, 1995; Kaus et al., 2004; Montagnat and Delingette, 2005). Recently, approaches have been proposed that allow the segmentation of the whole heart with its four chambers using a geometric model (Ecabert et al., 2008; Zheng et al., 2008) or registration-based techniques (Lötjönen et al., 2004; Zhuang et al., 2010; Kirisli et al., 2010).

An important application of these methods is the characterization of heart function. The geometric heart models resulting from segmentation have, in addition, been increasingly used to guide invasive therapeutic procedures. It has been proposed, for instance, to overlay heart models onto live fluoroscopy data to support ablation procedures for treating atrial fibrillation (Rhode et al., 2005; Knecht et al., 2008) or stem cell injection for myocardial repair (Gutiérrez et al., 2007; Lehmann et al., 2009). For these applications, a heart model comprising only the four chambers is often insufficient. It is, for instance, important to extract the left

* Corresponding author. Tel.: +49 40 5078 1466; fax: +49 40 5078 2510.

E-mail address: juergen.weese@philips.com (J. Weese).

atrium and the proximal part of the pulmonary veins to support guidance of ablation procedures for atrial fibrillation treatment. For cardiac resynchronization therapy, a heart model including the coronary sinus is needed to facilitate implantation of the pacemaker lead.

We present a method for segmentation of the heart and the attached great vessels. This method builds upon our previously published method for the automatic segmentation of the heart in computed tomography angiography (CTA) images (Ecabert et al., 2006; Ecabert et al., 2008) which was successfully applied for the automatic characterization of global heart function (Ghersin et al., 2009; Coche et al., 2010; Plumhans et al., 2009; Abadi et al., 2010). Briefly, the method detects the heart with an adapted Generalized Hough Transformation (GHT) and performs segmentation by adapting a generic heart model represented by a surface mesh to images. This is done by successively adjusting the global orientation and scaling, optimizing the parameters characterizing shape variability of the heart, and performing a deformable adaptation that roughly preserves shape. Robustness is achieved by using optimal boundary detection functions extracted from reference annotations using Simulated Search (Peters et al., 2010).

For the combined segmentation of the heart and attached great vessels, a properly extended geometric heart model and an empirical parametrization for approximating the heart shape of different individuals is presented. In particular, shape variations are described by assigning individual linear transformations to the heart chambers and to short tubular segments building the great vessels. In that way, variable sizes of the heart chambers and bending of the vessels can be described in a single consistent framework. Reliable adaptation of this complex heart model is achieved by first adapting the heart model without the great vessels (Peters et al., 2008). Afterwards the tubular segments of the large vascular structures are successively activated, initialized, and adapted. To optimize the computational performance, the mesh resolution is varied and already adapted mesh parts are frozen. Implementation has been done in a way that the adaptation mechanisms build a configurable framework that we call adaptation engine. The sequence of different steps and control parameters are encoded in an extended model.

In the following section, the heart model with the attached great vessels is presented. The section describes mesh generation and geometry, the parametrization of shape variability, and the assignment of boundary detection functions. Section 3 describes the adaptation engine, the specific techniques used in the engine, and the succession of the model adaptation steps as encoded in the extended model. Section 4 describes ground truth generation and results for the segmentation accuracy. Respective values are given for individual structures. In addition, specific metrics have been used to characterize segmentation performance for the great vessels. Summary and conclusions are presented in the final section.

2. Heart and vessel model

We start by providing information about the image data that have been used to build the heart and vessel model. Then we present the heart and vessel model itself. We describe how the model geometry has been generated. For adapting the model to images it is important to have a parametric description of its shape variability. For that purpose we extend the description of shape variability by multiple linear transformations (Ecabert et al., 2008) and show how it can be used to model bending and diameter variations of large vessels. Finally, information for boundary detection during model adaptation is attached to each triangle. In this context, we rely on the concept of Simulated Search (Peters et al., 2010).

2.1. Image data for model building

Allover, 35 data sets from 20 patients were used for model building. The data sets were retrospectively reconstructed at various phases of the cardiac cycle. They represent typical images that can be expected in clinical routine for patients indicated for cardiac CTA. The images were acquired with 16-, 40-, and 64-slice CT scanners (Brilliance CT, Philips Healthcare, Cleveland, OH) and standard CTA protocols with in-plane voxel resolution between 0.30×0.30 and $0.78 \times 0.78 \text{ mm}^2$ (typically $0.49 \times 0.49 \text{ mm}^2$), spacing between the slices between 0.4 to 2.0 mm (typically 0.4–0.5 mm) and slice thickness varying from 0.6 to 3.0 mm (typically twice the spacing between slices). The data sets comprise the 28 data sets of Ecabert et al. (2008) and 7 additional data sets that were particularly suited to construct the mesh geometry of the great vessels.

2.2. Mesh model

A geometric mesh model of the heart and the major vessels can be generated in various ways. Starting with an annotation of the different structures in an image, the mesh could be constructed from scratch. For us it was, however, important to re-use previously generated ground truth annotations and corresponding reference meshes. In addition, we wanted to have the meshes of the vessels made of regular rings as this facilitates definition of the shape variability. Hence, we used the 4-chamber heart model (Fig. 1) of Lorenz and von Berg (2006) and Ecabert et al. (2008) as starting point. This model consists of $V = 7286$ vertices combined in $T = 14,771$ triangles and represents the average heart shape of the 28 data sets used to build the model. We extended this model by generating and attaching structured meshes of the great vessels.

For that purpose, 7 CTA data sets were used in which the desired vessels were properly contrasted. These data sets also include the aortic arch and substantial portions of the superior and inferior vena cava. The centerlines of the aorta, the superior vena cava, the inferior vena cava, the coronary sinus, and the four pulmonary veins (up to the first bifurcation) were interactively delineated in these data sets and represented by equidistantly distributed points. These centerlines were complemented with radius information using a search-ray based approach. This approach casts 20 radial search rays, detects the position along a ray where the intensity drops below a threshold, discards the 10% smallest and 10% largest values, and estimates a mean radius from the

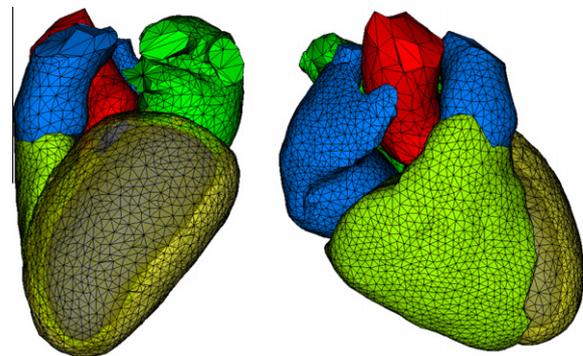


Fig. 1. 4-chamber heart model. The surface describes the endocardium except for the left ventricle where the epicardial wall is also modeled. The colors show the subdivision into different anatomical regions (epicardium of the left ventricle (yellow), endocardium of the left ventricle (gray), right ventricle (light green), left atrium (green), right atrium (light blue), aorta (red), trunk of the pulmonary artery (light blue)).

remaining values. We manually adjusted the threshold in this model building step in a way that the resulting radii describe the vessel lumen.

In the second step, we generated the mean geometry of the great vessels. For that purpose, the 4-chamber heart model was adapted to the 7 CTA data sets and the pose of the undeformed 4-chamber model was determined by point-based registration to the adapted heart model. After transforming the centerlines into the coordinate system of the undeformed 4-chamber model, average radii and centerlines were generated. For the centerline, averaging was done by averaging the difference vectors between two successive centerline points. This approach allows to take all centerline information into account, though the centerlines for a given vessel extracted from the 7 data sets varied considerably in length.

In the final step, meshes of the great vessels were constructed and attached to the 4-chamber heart model. After selecting a point on the mean centerline within the respective vessel trunk of the 4-chamber heart model, the normal plane at the selected point was used to cut off the distal mesh part of the vessel trunk. We resampled the vessel centerlines with a distance of 0.8 times the local vessel radius and constructed meshes composed of regularly triangulated rings. At the cut plane, the contour of the vessel trunk and the respective centerpoint do not coincide with the corresponding vessel contour and centerpoint, and a smooth transition between both mesh geometries must be generated. Starting from the cut plane we updated the centerline points of the great vessels with the difference vector between respective centerpoints multiplied by a weight factor exponentially decreasing towards the distal vessel part. Analogously, the radial vectors pointing from the centerline to the contour have been modified to generate a smooth transition. Further details can be found in Peters et al. (2008).

Fig. 2a and b show the resulting 4-chamber mesh model with the attached great vessels. The mesh model has $V = 8506$ vertices and $T = 17,324$ triangles. The 8 attached great vessels are composed of 99 individual rings and describe the vessel lumen. In addition, Fig. 2c and d show a model variant with reduced mesh resolution. This model has been generated by subsequently removing 4819 (56.7%) vertices building the heart chambers in a way that the resulting triangles have similar size. The vertices of the low-resolution heart model are by construction a subset of the vertices of the high resolution model.

2.3. Shape variability

To derive a parametric description of the shape variability, we subdivide the reference model represented by the mesh vertices $\mathbf{m}_1, \dots, \mathbf{m}_V$ into K individual parts. The subdivision is described by the weights $w_{i,k}$, which could be defined by

$$w_{i,k} = \begin{cases} 1 & \text{if vertex } i \text{ belongs to part } k \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

if all parts would be independent. To each part a linear transformation $\mathcal{T}_k(\mathbf{q}_k)[\cdot]$ (e.g. similarity transformation, affine transformation) is assigned that depends on the parameters \mathbf{q}_k . Accordingly, the entire model can be transformed by

$$\mathcal{T}_{\text{multi-linear}}(\mathbf{q})[\mathbf{m}_i] = \sum_{k=1}^K w_{i,k} \cdot \mathcal{T}_k(\mathbf{q}_k)[\mathbf{m}_i] \quad (2)$$

and modified shapes can be generated by varying the parameters $\mathbf{q} = (\mathbf{q}_1, \dots, \mathbf{q}_K)$ of the transformations. As this approach would generate discontinuities between the individual parts, transition regions between two or more parts are introduced where the transformations are interpolated. Practically, this is done by assigning weights $0 < w_{i,k} < 1$ with

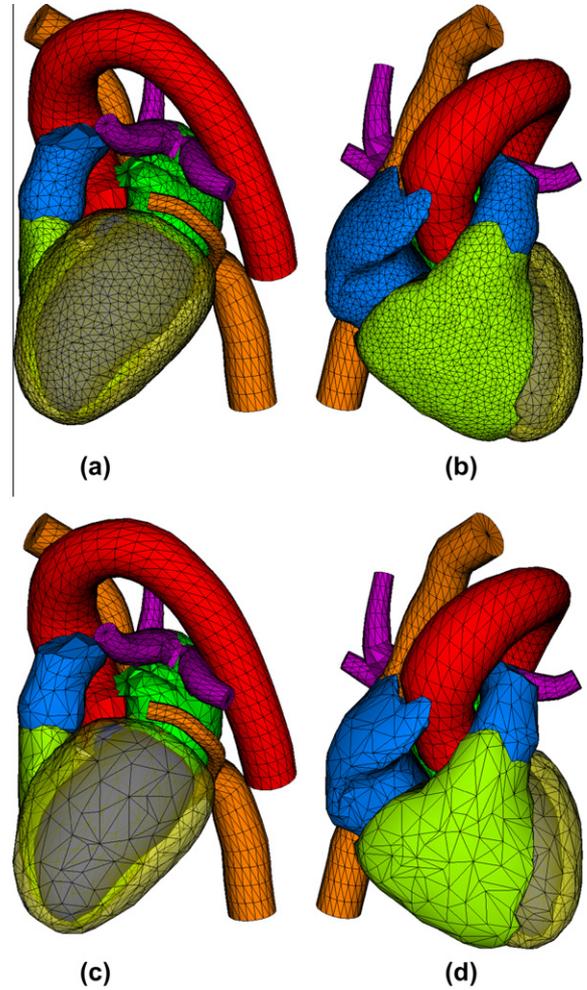


Fig. 2. Heart model with the great vessels ((a and b): fine resolution; (c and d): low resolution). As for the 4-chamber heart model, the surface describes the endocardium except for the left ventricle where the epicardial wall is also modeled. The tubular structures describe the vessel lumen (aorta (red), pulmonary veins (purple), coronary sinus (orange), inferior & superior vena cava (orange)) and are composed of regularly triangulated rings. For the low resolution mesh model only the chambers are sub-sampled.

$$\sum_{k=1}^K w_{i,k} = 1 \quad \forall i \quad (3)$$

in transition regions that depend on the geodesic distance to the border between the parts.

Our approach to model shape variability by a multi-linear transformation provides an empirical parametrization of an anatomical shape. This parametrization is not derived by statistical analysis from an ensemble of shapes as it is, for example, done in the context of principal component analysis (Cootes et al., 1994). There is also no mechanism that constrains the parameters to enforce a realistic shape or that prevents overlaps of different parts. Rather, it is assumed that accurate adaptation to images leads to a realistic shape and limits or resolves issues with overlapping structures. A prerequisite is that this approach provides enough flexibility to approximate anatomical shapes of different individuals. For the 4-chamber heart model this has been shown and a more accurate approximation of the heart shape for different individuals was obtained than with eigenmodes resulting from principal component analysis (Ecabert et al., 2006; Ecabert et al., 2008).

The subdivision of the heart chambers is done as described in Ecabert et al. (2008). An affine transformation is assigned to the left

atrium, the right atrium, the epicardium around the left ventricle, the left ventricle together with the aortic trunk, and the right ventricle together with the trunk of the pulmonary artery. For the attached vessels, we combine between 3 and 6 successive structured rings into short tubular segments. To each segment, we assign an individual similarity transformation. Between two tubular segments, there is a transition region (7 mm for the aorta and 5 mm for the other vessels) to preserve a smooth mesh geometry. Fig. 3 shows the heart model with the different tubular segments represented in different colors.

The description of shape variability by linear transformations assigned to model subparts is a simple and intuitive approach, though subdivision of the model is done manually. In the context of our application it has several advantages. First, shape variations of the heart chambers as well as bending and diameter variations of the vessels can be described in a single, consistent framework. Second, since shape variability is described by an empirical parametrization and not derived by statistical analysis, 7 data sets are sufficient to define geometry and shape variability of the great vessels. Third, parameter variations have by construction only a local influence. That is, when changing the diameter of the descending aorta, the heart chambers are not influenced. This is a difference to approaches based on principal component analysis where modification of a single parameter usually influences all parts of the model.

2.4. Assignment of boundary detection functions

For boundary detection, characteristic gray-value transitions are searched along profiles parallel to the triangle's normal \mathbf{n} and passing through its center \mathbf{c} :

$$\hat{\mathbf{x}} = \mathbf{c} + \arg \max_{i=-l, \dots, +l} [F(\mathbf{c} + i\delta\mathbf{n}) - D^2 \delta^2] \cdot \delta \cdot \mathbf{n} \quad (4)$$

with $\hat{\mathbf{x}}$ denoting the detected boundary point, $(2l + 1)\delta$ being the profile length, δ defining the sampling step size, and D biasing boundary detection towards nearby points. The boundary detection functions that we use are defined by

$$F(\mathbf{x}) = \begin{cases} \pm G_{\text{proj}}^{\text{limit}}(\mathbf{x}) & \text{for } Q_k(\mathbf{x}) \in [\min_k, \max_k] \quad \forall k \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

with

$$G_{\text{proj}}^{\text{limit}}(\mathbf{x}) = (\mathbf{n} \cdot \nabla I(\mathbf{x})) \cdot \frac{g_{\text{max}}(g_{\text{max}} + \|\nabla I(\mathbf{x})\|)}{g_{\text{max}}^2 + \|\nabla I(\mathbf{x})\|^2}. \quad (6)$$

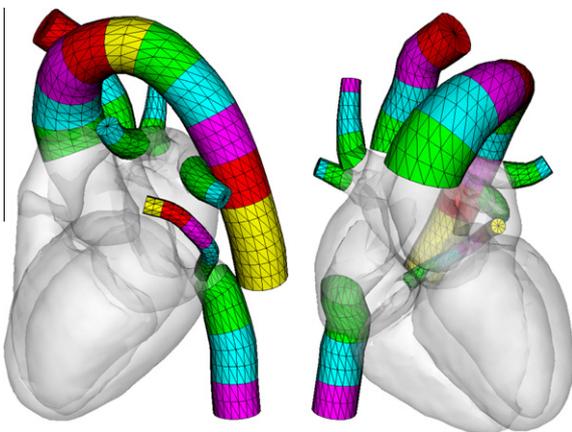


Fig. 3. Color-coded visualization of the combination of successive structured rings into short tubular segments. A separate similarity transformation is assigned to each group of rings. This multi-linear transformation enables the modeling of both global bending and local diameter variation.

Here, $I(\mathbf{x})$ denotes the gray-value at the point \mathbf{x} , ∇ is the gradient operator and g_{max} is a heuristic damping factor. The quantities $Q_k(\mathbf{x})$ are, for instance, the gray-value $I(\mathbf{x} + \delta\mathbf{n})$ on one side of the investigated position, and allow to discriminate different gray-value transitions. For the parameters we use the same values as in Ecabert et al. (2008) ($\delta = 1.0$ mm; $D = 0.5$ HU mm⁻² and $l = 10$ for a short search profile; $D = 0.125$ HU mm⁻² and $l = 20$ for a long search profile; $g_{\text{max}} = 50$ HU mm⁻¹).

Individual boundary detection functions are defined for each triangle by using triangle specific intervals $[\min_k, \max_k]$. In order to assign boundary detection functions and define the intervals $[\min_k, \max_k]$ for each triangle, we use a number of annotated reference images with corresponding meshes. At the mesh surface, vectors comprising values for all criteria Q_k are generated. By applying k -means clustering to these vectors, we derive reasonable intervals for the values of the criteria Q_k and generate a large number of boundary detection function candidates. Assignment of a boundary detection function candidate to a triangle is done by Simulated Search (Peters et al., 2010). Within this process, the triangle under investigation is displaced, boundary detection is performed and the boundary detection error (i.e. distance of the detected point to the tangent plane at the triangle position) is recorded. This is done for numerous displacements of the triangle, all reference images and all boundary detection function candidates. At the end of the process, the boundary detection function leading to the smallest boundary detection error is assigned to the triangle.

Training of the boundary detection functions was done using all 35 data sets, the 28 data sets of Ecabert et al. (2008) and the 7 additional data sets that were used to construct the mesh geometry of the great vessels. In a first step, the boundary detection functions of the 4-chamber model were mapped onto the high resolution heart model with the great vessels. The 7 data sets were used to train boundary detection functions of the great vessels. In a second step, the resulting model was adapted to 5 out of the 28 data sets where all great vessels were inside the field-of-view and well contrasted. The segmentations were thoroughly corrected and the boundary detection functions of the great vessels were trained using the 7 + 5 data sets. In a third step, the resulting model was used to segment all 35 data sets. Again, the segmentations were thoroughly corrected. For the aorta, the inferior vena cava, and the coronary sinus, reference segmentations could only be obtained for a subset of the 35 data sets (26, 12, and 16, respectively), because these structures were either not completely in the field-of-view or weak contrast made a reliable segmentation impossible.

The resulting reference annotations were used to train the boundary detection functions for the high and the low resolution heart model (4 chambers, trunk of the pulmonary artery, pulmonary veins and superior vena cava). Separate trainings were performed for the aorta, inferior vena cava, and the coronary sinus to handle the varying number of reference annotations. In all cases, the gray values on both sides of the detected edge and the first and second order Taylor coefficient of the local gray value profile were used as criteria Q_k . Respective Q_k values in the reference images have been clustered in $C = 5$ and 10 clusters. Acceptance intervals $[\min_k, \max_k]$ were derived by rejecting the low and high 5% or 10% percentiles. In that way, 400 boundary detection function candidates were generated per training that use one or two of the criteria Q_k and associated acceptance intervals $[\min_k, \max_k]$. Assignment of a specific boundary detection function to an individual triangle was done using Simulated Search with 147 positions (21 positions on each profile of 20 mm length; 1 profile passing through the triangle center and parallel to the triangle normal leading to displacements of $-10, -9, \dots, -1, 0, 1, \dots, 9, 10$ mm; three profiles laterally shifted by 10 mm with respect to the first profile; three profiles passing through the triangle center and tilted by 30 deg with respect to the first profile).

3. Model adaptation

Automatic, accurate and fast adaptation of the heart model with the great vessels requires various techniques. Table 1 summarizes the sequence of processing steps: The heart chambers are first localized and the low-resolution heart model is positioned. Pose misalignment is corrected by matching the heart chambers and vessel trunks to the image using a global similarity transformation. Afterwards the descending aorta is detected. Adaptation of the heart chambers and the most distal segment of the descending aorta is further refined by parametric adaptation using a multi-linear transformation. In the following steps, the model is accurately adapted using deformable adaptation while the tubular segments are successively activated. After activating the 1st generation of tubular segments that are directly connected to the vessel trunks and the most distal segment of the descending aorta, the high-resolution heart model is initialized. Deformable adaptation is continued while the remaining tubular segments are successively activated. Already well-adapted mesh parts are frozen in this step to speed-up adaptation. Fig. 4 shows the heart mesh at various stages of the adaptation pipeline for an example data set.

Implementation has been done in a way that these techniques build a configurable algorithmic framework that we call adaptation engine. The order of adaptation steps, the technique used in each step, respective parameters, information for boundary detection as well as the heart shape itself are part of the extended heart model that controls the engine. Fig. 5 illustrates the overall architecture. First, we describe the techniques provided by the engine. Afterwards, we provide the details about the processing order and specific parameters settings.

3.1. Heart localization and adaptation

The basic methods for automatic model adaptation comprise heart detection, parametric model adaptation and deformable model adaptation. The latter two approaches comprise new extensions to support partial model adaptation, incremental model adaptation as well as acceleration techniques.

3.1.1. Localization

Localization of the heart is based on the Generalized Hough Transformation (Ballard, 1981) and the 4-chamber model of Fig. 1. The approach combines the R-tables of several reference shapes to better capture variations in orientation and local shape variability within an object class. Training and detection is performed on subsampled images. The geometric transformations, that the shape can undergo, are limited to translations and isotropic scaling with a couple of factors (0.9, 1.0, and 1.1). The resulting parameters are used to position the heart model with the great vessels at a fixed orientation.

The details about the algorithm, implementation and parameter settings are given in Ecabert et al. (2008). For training, the originally used 28 datasets have been complemented by the 7 data sets used for the construction of the great vessels.

3.1.2. Parametric adaptation

A parametric description of the heart shape with vertices $\mathbf{v}_i(\mathbf{q})$ is obtained by applying a transformation \mathcal{T} defined by the parameters \mathbf{q} to the mesh model with vertices \mathbf{m}_i :

$$\mathbf{v}_i(\mathbf{q}) = \mathcal{T}(\mathbf{q})[\mathbf{m}_i] \quad (7)$$

Accordingly, the center point of triangle i is given by

$$\mathbf{c}_i(\mathbf{q}) = (\mathbf{v}_{i_1}(\mathbf{q}) + \mathbf{v}_{i_2}(\mathbf{q}) + \mathbf{v}_{i_3}(\mathbf{q}))/3 \quad (8)$$

Table 1

Sequence of processing steps for the adaptation of the heart model with the great vessels to cardiac CT images.

Step	Method	Comments	Mesh parts
1	Localization	Initialization of low-resolution model	
2	Parametric adaptation	Similarity transformation	Heart chambers and vessel trunks
3	Localization	Initialization of descending aorta	
4	Parametric adaptation	Multi-linear transformation	Heart chambers and vessel trunks + distal segment of descending aorta
5	Deformable adaptation	Activation of tubular segments	Heart chambers and vessel trunks + distal segment of descending aorta + 1st generation of tubular segments
6	Upsampling	Initialization of high-resolution model	
7	Deformable adaptation	Activation of tubular segments and freezing of mesh parts	Heart chambers and vessel trunks + distal segment of descending aorta + 1st generation of tubular segments + 2nd generation of tubular segments ...

with i_1 , i_2 and i_3 denoting the vertices building triangle i . The transformation \mathcal{T} may be a rigid, similarity, affine, or multi-linear transformation.

For adaptation, boundary points $\hat{\mathbf{x}}_i$ are detected for a subset S of triangles. Afterwards, the mesh coordinates are adjusted by minimizing the external energy

$$E_{\text{ext}}(\mathbf{q}) = \sum_{i \in S} \tilde{w}_i \left(\frac{\mathbf{V}(\hat{\mathbf{x}}_i)}{\|\mathbf{V}(\hat{\mathbf{x}}_i)\|} \cdot (\hat{\mathbf{x}}_i - \mathbf{c}_i(\mathbf{q})) \right)^2 \quad (9)$$

using a Gauss-Newton method (Gill et al., 1981). The weights

$$\tilde{w}_i = \max\{0, F_i(\hat{\mathbf{x}}_i) - D \cdot (\hat{\mathbf{x}}_i - \mathbf{c}_i)^2\}. \quad (10)$$

reflect the reliability of a detected boundary and are considered as constants during optimization. It is important to note that the transformation \mathcal{T} with associated parameters \mathbf{q} and the set of triangles S must be defined in a way that Eq. (9) has a well-defined minimum. In particular, no linear transformation \mathcal{T}_k should be associated with a model part for which no boundary detection is performed, when using a multi-linear transformation according to Eq. (2). Boundary detection and mesh adaptation are repeated for a pre-defined number of iterations.

The set S of triangles can be used in various ways. Performing boundary detection for a subset of triangles that are homogeneously distributed over the model surface can be used to accelerate model adaptation. The set may also be used to adapt model parts. In particular, we use the set S to adapt the heart chambers initially while neglecting the attached vascular structures.

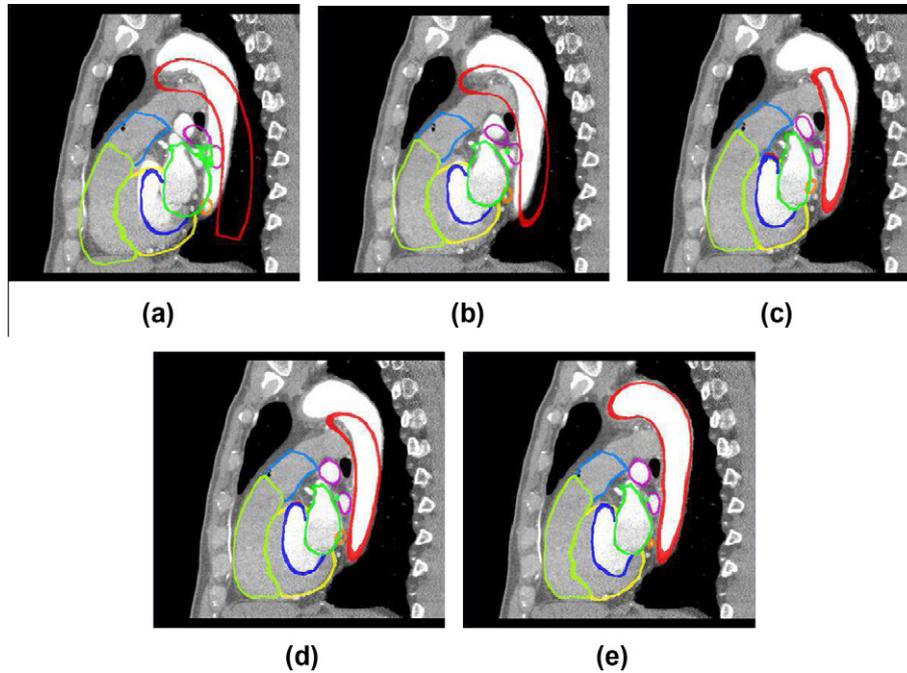


Fig. 4. Heart mesh after step 1 (a), 2 (b), 4 (c), 6 (d) and 7 (e) of the adaptation pipeline (see Table 1) overlaid onto a sagittal CTA image. The endocardium of the left ventricle is represented in blue. The colors of the other anatomical parts are chosen as in Figs. 1 and 2. As it has been shown in Ecabert et al. (2008), segmentation accuracy of the heart chambers is successively improved by parametric adaptation using a similarity transformation, parametric adaptation using a multi-linear transformation and deformable adaptation. Segmentation of the great vessels is essentially done during deformable segmentation by successively initializing tubular segments using an already well adapted structure.

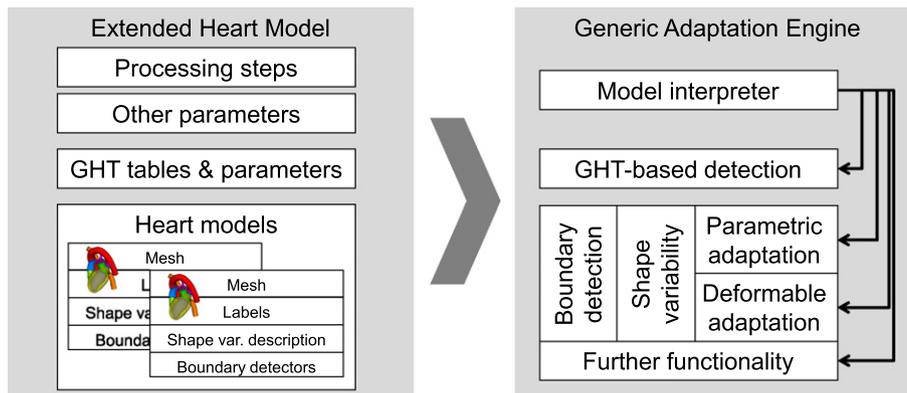


Fig. 5. Overall architecture of the segmentation framework. The engine provides generic techniques for the automatic adaptation of a mesh model to images. The extended heart model controls the engine. Apart from the mesh topology, boundary detection functions and shape variability information, it also defines the order of processing steps and related parameters required to automatically segment the heart with the attached great vessels by model adaptation.

3.1.3. Deformable adaptation

In the same way as parametric adaptation, deformable adaptation is done by iterating boundary detection and mesh adjustment for a pre-defined number of iterations. To support adaptation of mesh parts, boundary detection can be done for a set \mathcal{S} of triangles. In addition, only a subset \mathcal{V} of vertices which we call active vertices are modified during mesh adjustment. The other vertices are either frozen and belong to the set \mathcal{F} or inactive and belong to the set \mathcal{I} . Frozen vertices belong to already well adapted mesh parts and serve as boundary condition during mesh adjustment. Inactive vertices are not taken into account during mesh adjustment and are adapted in later stages of the segmentation process.

In particular, mesh adjustment is done by minimizing

$$E(\mathbf{q}, \{\mathbf{v}_i; i \in \mathcal{V} \cup \mathcal{F}\}) = E_{\text{ext}}(\{\mathbf{v}_i; i \in \mathcal{S}_v\}) + \alpha E_{\text{int}}(\mathbf{q}, \{\mathbf{v}_i; i \in \mathcal{V} \cup \mathcal{F}\}) \quad (11)$$

with respect to the parameters \mathbf{q} and the active vertices $\{\mathbf{v}_i; i \in \mathcal{V}\}$ using the two-step approach described in Ecabert et al. (2008). The set $\mathcal{S}_v \subseteq \mathcal{V} \cup \mathcal{F}$ in the above equation denotes the vertices associated with the triangles in set \mathcal{S} . The parameter α balances the external versus the internal energy and we use $\alpha = 60 \text{ HU mm}^{-3}$. The internal energy is defined by

$$E_{\text{int}}(\mathbf{q}, \{\mathbf{v}_i; i \in \mathcal{V} \cup \mathcal{F}\}) = \sum_{i \in \mathcal{V} \cup \mathcal{F}} \sum_{j \in \mathcal{N}(i), j \neq i} \sum_{k=1}^K w_{i,k} ((\mathbf{v}_i - \mathbf{v}_j) - (\mathcal{T}_k(\mathbf{q}_k)[\mathbf{m}_i] - \mathcal{T}_k(\mathbf{q}_k)[\mathbf{m}_j]))^2, \quad (12)$$

where $\mathcal{N}(i)$ is the set of indices of the first-order neighbor vertices of vertex i . Again, the transformations \mathcal{T}_k with associated parameters \mathbf{q}_k and the sets \mathcal{S} , \mathcal{V} , \mathcal{I} and \mathcal{F} must be defined in a way that Eq. (11) has a well-defined minimum.

3.2. Adaptation of the great vessels

Adaptation of the great vessels assumes that an initial segmentation of the heart chambers is available. Afterwards the tubular segments are successively initialized, activated, and adapted. This approach allows to properly initialize the tubular segments for deformable adaptation except for the case when the aortic arch is not in the field-of-view. For that reason, a dedicated localization algorithm for the descending aorta has been introduced.

3.2.1. Descending aorta localization

The descending aorta is localized within an axial slice that is determined from the most distal tubular segment of the descending aorta after positioning the heart model and optimizing its global orientation and scaling. Localization is done with a modified 2D Hough transformation (Hough, 1962) for circles. Candidate edges are detected using 3D methods and additional criteria are applied to reduce the number of candidates. For instance, a candidate edge is only considered, if its 3D gradient is sufficiently parallel to the axial plane (i.e. the angle between the 3D gradient vector and the axial plane should be smaller than 26.6 deg). In addition, edges with gradient magnitude or gray-value difference outside pre-defined ranges are skipped. During accumulation, the homogeneity within the associated circle is taken into account. To ensure that circular arcs are suppressed and only almost complete circular structures are detected, the circle is subdivided into 4 wedges and the associated counts are accumulated separately. The final vote is obtained by multiplying the results of four wedges building a circle and the circle with the maximum final vote represents the localization result for the descending aorta.

3.2.2. Adaptation of tubular segments

During deformable adaptation, the set \mathcal{V} of active vertices consists of the vertices of the heart structures and tubular segments that are currently being adapted (i.e., these vertices belong to a triangle out of the set \mathcal{S} for which boundary detection is performed) and of the vertices building the directly connected tubular segments. The weights $w_{i,k}$ of the multi-linear transformation $\mathcal{T}_{\text{multi-linear}}$ are defined in a way that respective tubular segments use the same transformation as the mesh part they are connected to. With this mechanism, the shape of these tubular segments is defined by the internal energy and their orientation is defined via the transformation $\mathcal{T}_{\text{multi-linear}}$. This mechanism ensures a smooth transition between mesh parts that are currently being adapted and the directly connected tubular segments, for which deformable adaptation can be activated in a subsequent iteration.

In order to activate the adaptation of a tubular segment, boundary detection is enabled by adding its triangles to the set \mathcal{S} and the multi-linear transformation $\mathcal{T}_{\text{multi-linear}}$ is complemented by the transformation describing its shape variability. If a further inactive tubular segment is connected, its vertices are added to the set \mathcal{V} of active vertices and the weights $w_{i,k}$ are defined in a way that both tubular segments use the same transformation.

3.3. Acceleration techniques

Computation time can be crucial, especially in an interventional context. For that reason, acceleration techniques are important for rapidly adapting the heart model with the attached great vessels to images. Multi-resolution meshes as well as mesh freezing have been introduced before and we outline the realization of both techniques in our adaptation engine.

3.3.1. Multi-resolution meshes

To support multi-resolution meshes, the actual mesh geometry corresponding to a low resolution mesh model can be used to ini-

tialize the mesh geometry corresponding to a high resolution mesh model. The implemented mechanism assumes that for each vertex of the first model there is a corresponding vertex in the second model. The new vertices are interpolated by minimizing the internal energy.

Let the quantities $\mathbf{v}_{1,i}$, $\mathbf{v}_{2,i}$, $\mathbf{m}_{1,i}$, $\mathbf{m}_{2,i}$, N_1 and N_2 ($N_1 < N_2$) denote the vertices \mathbf{v} of the actual mesh geometry, the vertices \mathbf{m} of the reference model, and the number N of vertices referring to the low resolution and high resolution model, respectively. For simplicity of the notation we assume that corresponding vertices have corresponding indices, i.e. vertices with indices $\leq N_1$ are corresponding vertices in the low resolution and high resolution model, whereas vertices with index $> N_1$ are the additional vertices of the high resolution model. In addition, the sets \mathcal{V}_1 and \mathcal{V}_2 ($\mathcal{V}_1 \subseteq \mathcal{V}_2$) as well as \mathcal{F}_1 and \mathcal{F}_2 ($\mathcal{F}_1 \subseteq \mathcal{F}_2$) denote corresponding mesh parts. In a first step, the parameters \mathbf{q} of the transformation $\mathcal{T}_{\text{multi-linear}}$ associated with the high resolution mesh are computed by minimizing

$$E_{\text{reg}}(\mathbf{q}) = \sum_{i \in \mathcal{V}_1 \cup \mathcal{F}_1} (\mathbf{v}_{1,i} - \mathcal{T}_{\text{multi-linear}}(\mathbf{q})[\mathbf{m}_{2,i}])^2 \quad (13)$$

resulting in the parameters \mathbf{q}^* . In a second step, the vertices $\mathbf{v}_{2,i}$ are initialized. The corresponding vertices are copied ($\mathbf{v}_{2,i} = \mathbf{v}_{1,i}$ for $i \in \mathcal{V}_1 \cup \mathcal{F}_1$). After minimization of the internal energy $E_{\text{int}}(\mathbf{q}^*, \{\mathbf{v}_{2,i}; i \in \mathcal{V}_2 \cup \mathcal{F}_2\})$ with respect to the additional vertices $\{\mathbf{v}_{2,i}; i \in \mathcal{V}_2 \cup \mathcal{F}_2 \wedge i > N_1\}$, initialization of the high resolution mesh is completed.

3.3.2. Mesh freezing

Mesh parts can be frozen after a pre-defined number of iterations. For that purpose, vertices are removed from the set \mathcal{V} of active vertices and added to the set \mathcal{F} of frozen vertices. Savings in computation time result from two effects. First, it is not necessary to perform boundary detection for triangles where all vertices are frozen and the set \mathcal{S} can be adapted accordingly. Second, less variables need to be considered, when adjusting the mesh coordinates by minimization of Eq. (11).

3.4. Adaptation of the heart and vessel model

Having introduced the different techniques for model adaptation, we now describe the sequence of processing steps (see Table 1) in more detail. In addition, we provide information about the parameter settings used in the different steps.

3.4.1. Localization and parametric adaptation

After GHT-based heart localization, the low resolution heart model (Fig. 2c and d) is placed in the image volume (Fig. 4a). In the next step, location, orientation and scaling are adjusted using parametric model adaptation (Eq. (9)) and a global similarity transformation as transformation \mathcal{T} . Boundary detection is performed for the heart chambers and vessel trunks, and the associated triangles build the set \mathcal{S} . Altogether, 20 iterations of boundary detection and mesh adaptation are performed.

The descending aorta is poorly initialized in this stage (Fig. 4b). Using the resulting similarity transformation, the gravity center of the most distal tubular segment of the descending aorta is computed and aorta detection is carried out in the corresponding axial image slice. Model adaptation is further refined by parametric adaptation using a multi-linear transformation (Fig. 4c). For the heart chambers and vessel trunks, the multi-linear transformation is defined as in Ecabert et al. (2008). An additional similarity transformation is used for the most distal tubular segment of the aorta. Furthermore, the triangles of the most distal tubular segment of the descending aorta are added to the set \mathcal{S} . Parametric adaptation with the multi-linear transformation is finalized after 20 iterations.

3.4.2. Deformable adaptation and vessel segmentation

For deformable adaptation, the number of iterations is defined per structure. Additionally, the number of iterations, after which a tubular segment is activated, is defined. In this way, simple as well as complex adaptation schemes can be realized.

Adaptation of the low resolution mesh is done in 10 iterations. Within the first 9 iterations, the heart chambers with the attached vessel trunks and the most distal segment of the descending aorta are adapted. Then, the first generation of tubular segments is activated and one additional iteration is performed. Afterwards, the resulting mesh is up-sampled (Fig. 4d).

Deformable adaptation is continued with the high resolution heart model. In particular, segmentation of the heart chambers is slightly refined and the tubular structures are successively activated and adapted to the image (Fig. 4e). The adaptation scheme is complex. Adaptation of the heart chambers and vessel trunks is performed for 2 iterations. Afterwards respective vertices are frozen by removing them from the set of active vertices \mathcal{V} and adding them to the set of frozen vertices \mathcal{F} . Deformable adaptation of the tubular segments is performed for 8–14 iterations. Again, respective vertices are afterwards frozen. Activation of the next tubular segment takes place after 4–10 iterations. The exact numbers differ for different vessels and tubular segments. To guarantee a smooth transition between two connected segments, at least four iterations are performed where both segments are adapted to the image. Overall, deformable adaptation of the high resolution heart model with the great vessels is completed after 20 iterations.

3.4.3. Parameter selection

A considerable number of parameters must be defined for training and model adaptation, and we provided the parameter values in the preceding sections. In general, we use the same values as Ecabert et al. (2008), where also the stability with respect to specific parameters of the basic concepts of our model adaptation approach has been discussed. New values were introduced for the number of iterations during deformable adaptation as different mesh resolutions are used and the number of iterations are defined per part. These iteration numbers are empirically selected in a way that they limit computation time while leading visually to a good segmentation. Larger iteration numbers will slightly change the segmentation, but are not expected to lead to clear improvements without modifying other parameter such as the sampling distance for boundary detection ($\delta = 1$ mm) at the same time. In general, further improvements may be achieved by systematically testing parameter value combinations and selecting those values that minimize a metric of the segmentation accuracy as used during validation. For instance, Hautvast et al. (2006) performed full factorial experiments and analyzed the results using the technique of analysis of variances (ANOVA).

4. Validation

We first outline the material used for validation. Specifically, we describe how ground truth segmentations have been derived in a practicable and consistent manner. We then introduce several error metrics to validate the distinct types of anatomical structures composing the heart model. Finally, we present and discuss numerical results along with segmentation examples.

4.1. Data

The validation of the presented method was carried out on CT images reflecting the variability in terms of quality, noise, field of view, etc., which can be expected in clinical routine. We selected

$N_I = 37$ data sets retrospectively reconstructed at various phases of the cardiac cycle from 17 adult patients clinically indicated for cardiac CTA. To avoid a bias, these images were from different patients than the images used for model building and training. The images were acquired with 16-, 40-, and 64-slice CT scanners (Brilliance CT, Philips Healthcare, Cleveland, OH, USA) and standard coronary CTA protocols with in-plane field of view varying from 144×144 to 255×255 mm² and craniocaudal coverage ranging from 110 to 268 mm. The in-plane voxel resolution was between 0.28×0.28 and 0.49×0.49 mm² (typically 0.49×0.49 mm²). The spacing between the slices was between 0.33 and 0.50 mm (typically 0.4–0.5 mm) and the slice thickness was in all cases twice as large as the spacing. The gray level distribution of the blood pool in the chambers and vessels may vary due to differences in arrival times of the contrast agent bolus. This can strongly reduce the contrast between the most distal parts of the vessels and the surrounding structures. As for the training data, the aortic arch was not covered in all images.

4.2. Ground truth segmentation

A multi-step approach was conducted to obtain accurate ground truth segmentations of the test images. The extended heart model and adaptation method were applied to the new images. The resulting meshes were then converted into annotation masks of same resolution as the underlying images, and each voxel enclosed by the mesh was assigned a distinct organ code. This annotation was represented as a multi-color mask overlaid on the original image to enable detailed local corrections, and the masks were thoroughly manually edited by a clinical expert. These expert-edited masks are considered as ground truth annotations in our experiments.

In order to generate ground truth meshes, for mesh-based comparisons, the ground truth masks were converted into gray level images by assigning a distinct gray value to each color code. The boundary detection functions of the heart model (see Section 2.4) were then modified to detect the expected transitions in these *gray-plateau* images and the model was adapted following the scheme described in Section 3. The resulting meshes were inspected by the clinical expert and corrected where needed. These expert-edited meshes are considered as ground truth meshes in our experiments.

With this approach to generate ground truth, it is measured during validation how much the automatic segmentation deviates from a segmentation that is acceptable for a clinical expert. A disadvantage of this approach is that the segmentation error is underestimated if the manual corrections are not thoroughly performed. A completely independent segmentation method for ground truth segmentation would prevent this bias, but manual segmentation leading to accurate and consistent 3D segmentations of complex anatomical structures such as the heart with the great vessels is a complex research topic in itself.

4.3. Error metrics

The different parts composing the extended vessel model can be divided into two structurally different groups. The chambers have a rounded shape, whereas the vessels are elongated tubular structures. To account for these structural differences, several error metrics have been used to measure the quality of the segmentation results.

4.3.1. Constrained point-to-surface distance

To characterize the accuracy of surface detection, we use the concept of constrained point-to-surface (CPS) distance as introduced by Ecabert et al. (2006). Let

$$\varepsilon_{\Pi(r)}(\mathbf{c}_i^{\text{adapt},j}, \mathbf{c}_i^{\text{ref},j}) = \min_{\mathbf{x} \in \Pi(r, \mathbf{c}_i^{\text{ref},j})} \|\mathbf{c}_i^{\text{adapt},j} - \mathbf{x}\| \quad (14)$$

be the smallest Euclidean distance between the triangle center \mathbf{c}_i of the adapted mesh in image j and a surface patch Π of geodesic radius r surrounding the corresponding triangle center in the reference mesh. For $r=0$, $\varepsilon_{\Pi(0)}$ corresponds to the point-to-point distance between corresponding triangle centers and for $r \rightarrow \infty$, $\varepsilon_{\Pi(\infty)}$ corresponds to the distance between the triangle centers and the surface of the other mesh. We set $r = 10$ mm in the experiments. This radius gives a good approximation of the point-to-surface distance while preventing meaningless comparisons between close, but non-corresponding surfaces.

As this distance is not symmetric (from one mesh to the other or vice versa), we define the mean CPS distance as

$$\varepsilon_{\text{mean}}^{\text{CPS}} = \frac{1}{N_j \cdot M} \sum_{j=1}^{N_j} \sum_{i=1}^M \frac{1}{2} \left(\varepsilon_{\Pi(r)}(\mathbf{c}_i^{\text{adapt},j}, \mathbf{c}_i^{\text{ref},j}) + \varepsilon_{\Pi(r)}(\mathbf{c}_i^{\text{ref},j}, \mathbf{c}_i^{\text{adapt},j}) \right). \quad (15)$$

The inner summation adds up the unsigned distances or unsigned errors for the whole mesh (i.e., $M = T$) or for a specific substructure. The outer summation is performed over the available data sets.

4.3.2. Constrained point-to-centerline distance

Centerline-to-centerline distances are good indicators of the accuracy of the vessel course. Similarly to the CPS distance, we introduce the constrained point-to-centerline (CPC) distance. The centerlines are defined as sequences of straight line segments connecting the centers of successive vertex rings of the tubular meshes. Let

$$\varepsilon_{A(l)}(\mathbf{p}_i^{\text{adapt},j}, \mathbf{p}_i^{\text{ref},j}) = \min_{\mathbf{x} \in A(l, \mathbf{p}_i^{\text{ref},j})} \|\mathbf{p}_i^{\text{adapt},j} - \mathbf{x}\| \quad (16)$$

be the smallest Euclidean distance between a centerline point \mathbf{p}_i from an adapted centerline in image j and an interval A of length $2 \cdot l$ surrounding the corresponding point in the reference centerline. We set $l = 10$ mm for the experiments. We define the mean CPC distance similarly to the previous metric as

$$\varepsilon_{\text{mean}}^{\text{CPC}} = \frac{1}{N_j \cdot N_s} \sum_{j=1}^{N_j} \sum_{i=1}^{N_s} \frac{1}{2} \left(\varepsilon_{A(l)}(\mathbf{p}_i^{\text{adapt},j}, \mathbf{p}_i^{\text{ref},j}) + \varepsilon_{A(l)}(\mathbf{p}_i^{\text{ref},j}, \mathbf{p}_i^{\text{adapt},j}) \right). \quad (17)$$

where N_s is the number of segments composing a centerline.

4.3.3. Volume overlap

The third metric is based on volume overlap. Volume measurements are commonly used in clinical routine to determine global functional parameters such as stroke and ejection fraction (Frangi et al., 2001). The adapted and reference meshes are converted to annotation masks with voxel resolution $0.2 \times 0.2 \times 0.2$ mm³. This resolution is finer than the resolution of the actual data to reduce inaccuracies caused by the discretization process.

Let $\varepsilon^{\text{overlap}}(R^{\text{adapt},j}, R^{\text{ref},j})$ be the relative volume overlap between a region R enclosed by the adapted mesh in image j and the corresponding region in the ground truth annotation defined as

$$\varepsilon^{\text{overlap}}(R^{\text{adapt},j}, R^{\text{ref},j}) = \frac{|R^{\text{adapt},j} \cap R^{\text{ref},j}|}{|R^{\text{adapt},j}|} \quad (18)$$

where $|\cdot|$ is the cardinality operator. The mean volume overlap is defined by

$$\varepsilon_{\text{mean}}^{\text{overlap}} = \frac{1}{N_j} \sum_{j=1}^{N_j} \frac{1}{2} \left(\varepsilon^{\text{overlap}}(R^{\text{adapt},j}, R^{\text{ref},j}) + \varepsilon^{\text{overlap}}(R^{\text{ref},j}, R^{\text{adapt},j}) \right). \quad (19)$$

According to Eq. (9), the triangles are allowed to slide along image boundaries at no cost. For that reason, we use a slightly different definition of the volume overlap for the vessels. Let $\varepsilon_e^{\text{overlap}}(R_i^{\text{adapt},j}, R_{[-e, i+e]}^{\text{ref},j})$ be the overlap between the region defined by ring R_i of the adapted mesh in image j and the region $R_{[-e, i+e]}$ defined by the concatenation of neighboring rings of the reference mesh, i.e.

$$\varepsilon_e^{\text{overlap}}(R_i^{\text{adapt},j}, R_{[-e, i+e]}^{\text{ref},j}) = \frac{|R_i^{\text{adapt},j} \cap R_{[-e, i+e]}^{\text{ref},j}|}{|R_i^{\text{adapt},j}|}. \quad (20)$$

For the vessels, the mean volume overlap is defined as

$$\varepsilon_{\text{mean}}^{\text{vessel overlap}} = \frac{1}{N_j \cdot N_s} \sum_{j=1}^{N_j} \sum_{i=1}^{N_s} \frac{1}{2} \left(\varepsilon_e^{\text{overlap}}(R_i^{\text{adapt},j}, R_{[-e, i+e]}^{\text{ref},j}) + \varepsilon_e^{\text{overlap}}(R_i^{\text{ref},j}, R_{[-e, i+e]}^{\text{adapt},j}) \right) \quad (21)$$

where N_s is the number of rings composing a vessel. In the experiments we set $e = 2$ corresponding approximately to the interval $A = 20$ mm used for the constrained point-to-centerline distance. Note that setting $e = 0$ in Eq. (21) leads to Eq. (19) used for the chambers.

4.3.4. Remarks on statistics calculation

Two effects lead to meaningless comparisons and for that reason a subset of triangles has been removed from the error measurements. The first effect is due to the fact that the tubes in the model have been truncated at arbitrary lengths. As the triangles are allowed to slide along image boundaries at no cost (see Eq. (9)), the length of the adapted vessels may vary from image to image. The metrics defined above would thus indicate an error due to a longitudinal displacement while the segmentation is actually correct. For that reason, 30 distal rings of the 8 vessels (i.e., about one tubular segment per vessel) were removed from the error calculation in all metrics.

In addition, the distal parts of the inferior vena cava and the coronary sinus are often not contrasted in the CT images used for validation. These areas were identified by the clinical expert during ground truth generation and the corresponding vessel rings in those images are removed from the statistics. The number of excluded rings varies in each image and is dependent on the contrast filling these vessels.

The statistics calculation also excludes the triangles and rings which are outside the image volume. This can occur in certain images, e.g., at the top of the aortic arch or at the distal parts of the inferior or superior vena cava.

4.4. Results

Figs. 6–8 show segmentation results for three different patients. The visual impression of the cross-sections with the overlaid contours can be misleading with regard to a 3D assessment of the segmentation result. For example, if the local mesh surface is almost parallel to the displayed cross-section, the distance between the desired visible border and the segmentation contour may appear to be much larger than the 3D distance to the desired border. In addition, borders related to the truncation of a vessel can be misleading.

The CPS distance for the whole mesh is between 0.31 and 1.16 mm with a mean of 0.70 mm for the 37 data sets used during validation. The detailed validation results for the CPS distance, CPC

distance, and volume overlap are presented in Tables 2–4, respectively. The standard deviation $s_{\text{triangle}}^{\text{CPS}}$ and $s_{\text{line}}^{\text{CPC}}$ for the CPS distance and the CPC distance refer to the individual errors of all triangles or line segments composing an anatomical region in all images and not to the standard deviation of the mean error for different images. In Table 2 the error distribution per interval is also provided. The percentages of triangles within various error ranges are derived from the errors per triangle averaged over all images. These values give an indication whether triangles are systematically off from the ground truth.

4.5. Discussion

The visual inspection of the results indicates that the proposed approach was capable of automatically segmenting the cardiac chambers and great vessels in all test images. As can be seen in Figs. 6–8, overlaps of different model parts are largely avoided by the proper adaptation of the heart model to images though there is no explicit mechanism preventing overlaps. The aorta could be successfully extracted even if part of the arch was missing in the image. This demonstrates that a single framework can be used to simultaneously adapt a mesh to both rounded and tubular structures by appropriately assembling and activating connected deformable shape models.

4.5.1. Chambers

The mean CPS distance and the standard deviation are reported in Table 2 for each anatomical region separately. This table indicates the accuracy which can be achieved by the adapted model when retrieving the surface of the anatomical regions in the image. This table shows similar numerical results as those published by Ecabert et al. (2008), which were achieved for a heart model without appended tubular structures and with no acceleration techniques. The increased model complexity and the acceleration techniques do not seem to affect the accuracy of the chamber segmentation.

It can be observed that all chambers present a similar mean distance and error distribution. The variable regions of the atrial appendages on both sides of the heart and the difficulty in defining the boundary of the left ventricular endocardium in the vicinity of the papillary muscles may be the reason for a slightly increased mean error of these structures. The volume overlap of the different

anatomical regions is reported in Table 3. These values indicate also that the cardiac chambers are well reconstructed.

These results compare favorably to recently published studies on automatic whole heart segmentation. Zheng et al. (2008) reported a mean point-to-mesh error of 1.13–1.57 mm for the different chambers by adapting a point distribution model (PDM) of the four chambers to CT images with in-plane resolution between 0.28×0.28 – $0.74 \times 0.74 \text{ mm}^2$ and a slice-to-slice distance of 0.4–2 mm. Kirisli et al. (2010) published a mean error of 1.25 mm (varying between 0.63–1.53 mm for the different chambers) using a multi-atlas segmentation approach for CT images with a resolution of $0.32 \times 0.32 \times 0.4 \text{ mm}^3$. They also reported a mean Dice similarity coefficient of 88.3–94.6%.

We can also compare our results with studies on whole heart segmentation carried out with MR images. Lötjönen et al. (2004) reported a mean point-to-surface error of 2.53 mm by simultaneously registering a 3D atlas of the four chambers to short and long axis MR images with in-plane resolution of 1.0×1.0 or $1.4 \times 1.4 \text{ mm}^2$ and 6–7 mm slice thickness. Zhuang et al. (2010) introduced a whole heart segmentation framework combining locally affine registrations followed by free-form deformations with adaptive control point status. This approach achieved a mean point-to-surface error of 2.14 mm (ranging 1.47–2.38 mm) and a mean overlap ranging 63–85% for the different chambers on MR images with a reconstructed resolution of $1.0 \times 1.0 \times 1.0 \text{ mm}^3$. For a comparison with other work on chamber segmentation, we refer to Peters et al. (2010).

4.5.2. Great vessels

The mean CPS distance for the aorta, the pulmonary veins and the superior vena cava is in the same range as for the chambers, however somehow larger for the coronary sinus and the inferior vena cava (Table 2). A similar effect can be observed for the mean CPC distance (Table 4). This large error is mainly due to missing or very faint contrast in these vessels. For instance, this effect can be observed in the images of the second column of Figs. 7 and 8 for the inferior vena cava. In particular, the proximal part of the inferior vena cava is usually contrasted and correspondingly well segmented by the model, whereas the distal part presents almost no contrast with the surrounding structures. In this latter case, prior-shape knowledge controls the adaptation. Fig. 9 illustrates these effects for the coronary sinus. In Fig. 9a and b, the model

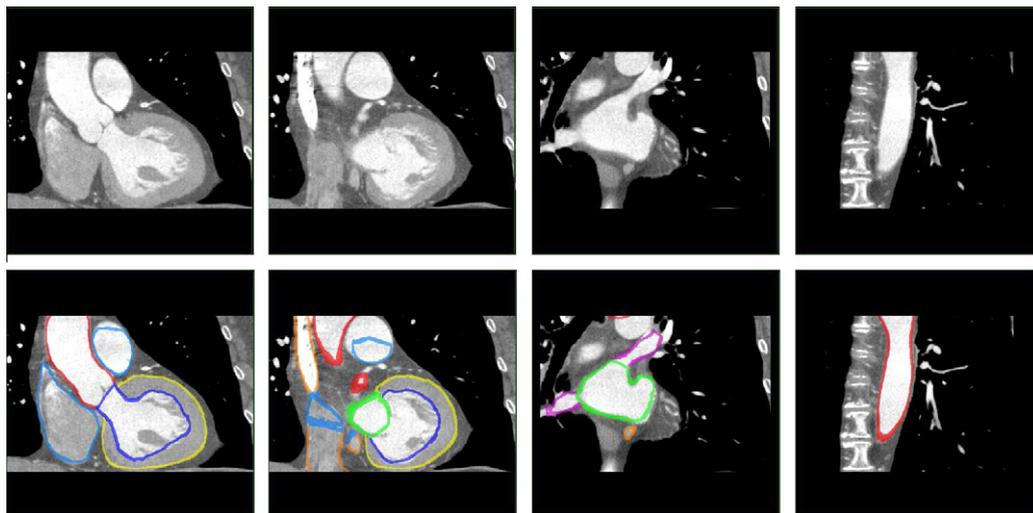


Fig. 6. Segmentation results for the heart model with the great vessels (patient 1). Four coronal views across the volume are shown without and with the segmentation results overlaid. The endocardium of the left ventricle is represented in blue. The colors of the other anatomical parts are chosen as in Figs. 1 and 2.

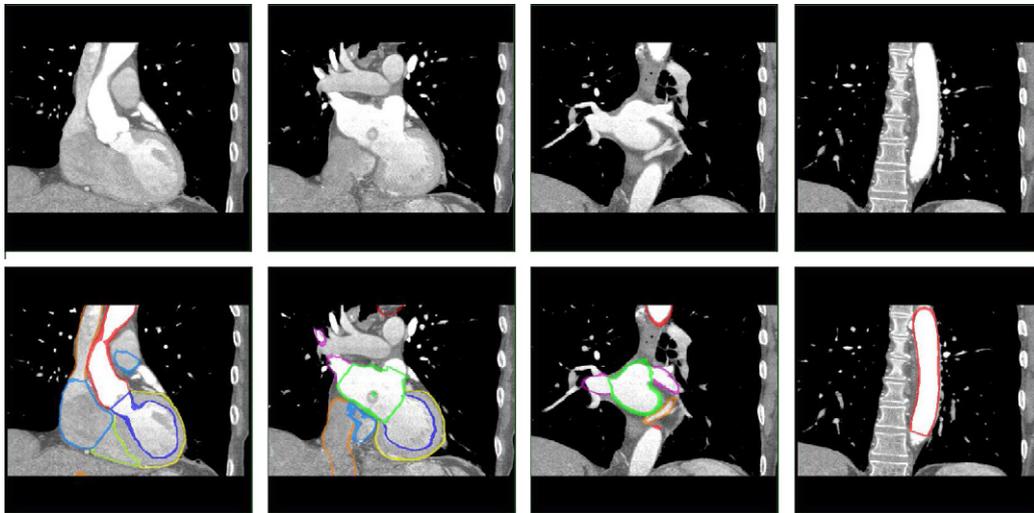


Fig. 7. Segmentation results for the heart model with the great vessels (patient 2).

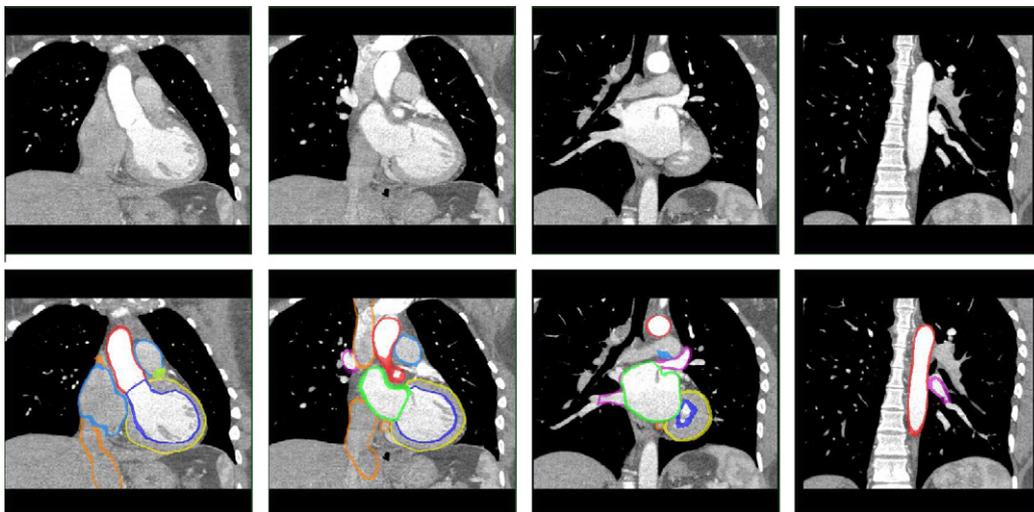


Fig. 8. Segmentation results for the heart model with the great vessels (patient 3).

Table 2

Constrained point-to-surface error $\epsilon_{\text{mean}}^{\text{CPS}}$, standard deviation $\epsilon_{\text{triangle}}^{\text{CPS}}$ and percentages of triangles within various error ranges for the different cardiovascular regions resulting from the fully automatic segmentation.

Anatomical region	$\epsilon_{\text{mean}}^{\text{CPS}}$ (mm)	$\epsilon_{\text{triangle}}^{\text{CPS}}$ (mm)	Percentage of triangles with various error ranges (mm)		
			<1.0	1.0–2.0	>2.0
Whole mesh	0.70	1.03	79.9	19.1	1.0
Left atrium	0.70	0.87	87.0	9.0	4.0
LV endocardium	0.77	1.14	76.4	23.4	0.2
LV epicardium	0.68	0.96	80.5	19.3	0.2
Right atrium	0.82	1.00	65.1	33.5	1.4
Right ventricle	0.63	0.66	87.9	12.1	4.2
Pulmonary artery trunk	0.50	0.49	100.0	0.0	0.0
Aorta	0.60	1.14	86.9	13.1	0.0
Pulmonary veins	0.60	0.72	95.3	4.7	0.0
Superior vena cava	0.64	0.74	59.7	36.7	3.6
Inferior vena cava	0.90	1.14	71.3	21.4	7.3
Coronary sinus	1.32	2.25	25.2	72.3	2.5

could track the coronary sinus which is well contrasted. Fig. 9c and d show a bad adaptation of the coronary sinus in case of poor or missing contrast. The coronary sinus also overlaps with the left atrium in the latter case, illustrating that the prevention of overlaps

relies on the proper adaptation of the heart model to the image. Nevertheless, with the exception of the coronary sinus, the reported mean CPC errors are well below the respective vascular radii which are given in Table 5.

Table 3

Volume overlap for the different cardiovascular regions resulting from the fully automatic segmentation.

Anatomical region	$\varepsilon_{\text{mean}}^{\text{overlap}}$ (%)
Left atrium	95.3 ± 1.6
LV bloodpool	95.0 ± 3.4
LV myocardium	96.2 ± 1.2
Right atrium	94.6 ± 2.0
Right ventricle	95.4 ± 1.4
Pulmonary artery trunk	94.0 ± 2.5
Aorta	95.2 ± 8.8
Pulmonary veins	83.2 ± 12.5
Superior vena cava	87.7 ± 17.4
Inferior vena cava	84.9 ± 9.0
Coronary sinus	70.4 ± 28.1

Table 4

Constrained point-to-centerline error $\varepsilon_{\text{mean}}^{\text{CPC}}$ and standard deviation $s_{\text{line}}^{\text{CPC}}$ for the great vessels resulting from fully automatic segmentation.

Anatomical region	$\varepsilon_{\text{mean}}^{\text{CPC}}$ (mm)	$s_{\text{line}}^{\text{CPC}}$ (mm)
Aorta	0.66	1.43
Pulmonary veins	0.86	1.04
Superior vena cava	1.20	2.38
Inferior vena cava	2.18	1.92
Coronary sinus	1.77	2.77

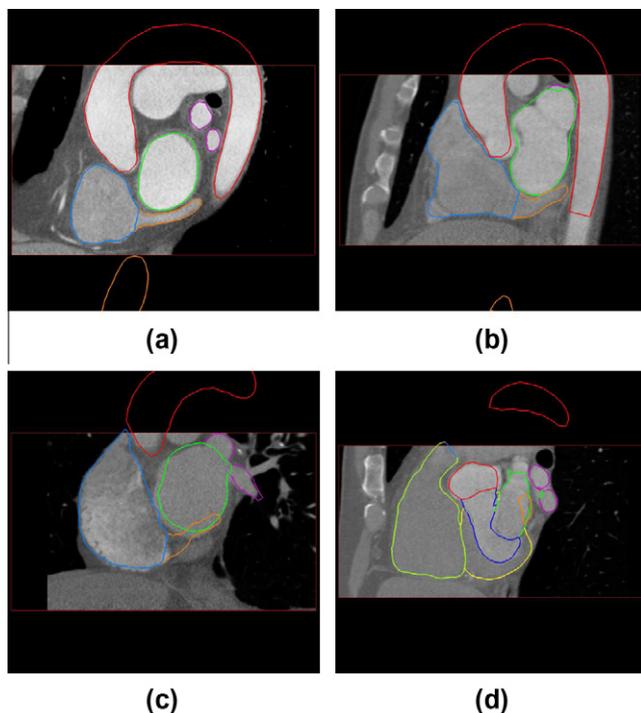


Fig. 9. Oblique views of four patients emphasizing on the coronary sinus (orange). Images (a and b) show how the proximal part of the coronary sinus model correctly tracks the corresponding structure in the image. In the images (c and d) the coronary sinus is mis-segmented due to poor or missing contrast.

Volume overlap (see Table 3) is particularly sensitive for thin structures as illustrated in Fig. 10. For the aorta, which has the largest diameter and which is usually well contrasted, we can achieve a similar performance as for the chambers. On the other hand, the coronary sinus, which is also the thinnest vessel and usually poorly contrasted, presents the lowest overlap in agreement with the errors reported in Tables 2 and 4.

Table 5

Mean vessel radius determined from the ground truth meshes.

Anatomical region	Mean radius (mm)
Aorta	13.6 ± 2.6
Pulmonary veins	7.4 ± 1.0
Superior vena cava	9.6 ± 1.1
Inferior vena cava	12.8 ± 1.4
Coronary sinus	3.7 ± 0.9

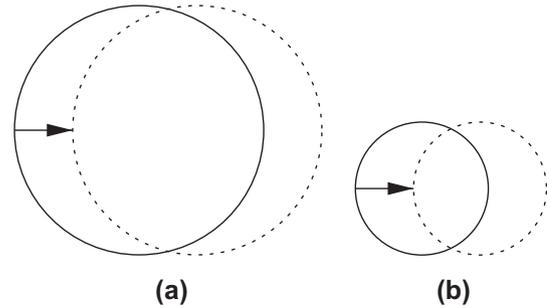


Fig. 10. Illustration of the dependence of the volume overlap on the vessel diameter. The same lateral shift reduces the overlap more strongly for a small (b) than for a large vessel cross-section (a).

To our knowledge, this is the first published approach presenting the simultaneous segmentation of the chambers and all great vessels with a single method, making a direct quantitative comparison difficult. Zhao et al. (2009) combine level-set and optimal surface segmentation algorithms in a single optimization process to extract the aorta from 4D (3D+time) MR images ($1.5 \times 1.5 \times 6.0\text{--}2.0 \times 2.0 \times 6.0 \text{ mm}^3$ voxel size) achieving a mean unsigned positioning error of 1.55 mm. We can also mention the work of de Bruijne et al. (2003) who introduced an adapting active shape model for tubular structure segmentation. They applied their technique to the segmentation of thrombus on abdominal aortic aneurysms in CT images with a resolution of $0.488 \times 0.488 \times 2.0 \text{ mm}^3$ and reported a surface root mean square error of 1.9 mm and a volume overlap of 95%.

The segmentation of the left atrium and pulmonary veins is important to support planning and guidance of ablation procedures for atrial fibrillation treatment. Several approaches have been introduced to extract the anatomy of the left atrium and pulmonary veins (John and Rahn, 2005; von Berg and Lorenz, 2005; Karim et al., 2008). However, the focus is on the segmentation of the left atrium and the pulmonary veins are extracted in a subsequent processing step. No quantitative measurements are reported for these vessels.

4.6. Computation time

The experiments were carried out on a workstation with 2 Dual-Core Intel Xeon processors (2.33 GHz, 4 GB RAM) and computation times of 12 ± 1 s were measured for the complete segmentation. The time for the different steps of the processing chain (Table 1) was 1.1 s for GHT-based heart detection (step 1), 4.3 s for parametric adaptation (steps 2–4), 2.2 s for deformable adaptation with the low resolution mesh (step 5), 0.5 s for upsampling (step 6), and 3.9 s for deformable adaptation with the high resolution mesh (step 7). Even if the computer settings are not identical, the computation time achieved for the adaption of the extended heart model is comparable to our 4-chamber model (Ecabert et al., 2008). A mean adaptation time of 4 s was reported by Zheng et al. (2008) for a model including the four chambers (without

vessels). Other approaches based on image registration for the automatic segmentation of the four chambers reported processing times of 25 min (Kirisli et al., 2010) (at reduced image resolution) and 2–4 h (Zhuang et al., 2010).

5. Summary and conclusions

An automatic model-based approach for the segmentation of the whole heart and the attached great vessels is presented. A major contribution of this article is the adaptation of a geometric model combining structurally different parts, such as rounded cardiac chambers and elongated vessels, using a single and consistent framework.

The shape variability is described by an empirical parametrization and defined using a multi-linear transformation. In particular, each chamber is assigned an affine transformation whereas the great vessels are represented by the concatenation of short tubular segments, each of them undergoing a similarity transformation. This representation enables the consistent description of inter-phase and inter-patient chamber variations, and appropriate bending of the vessels.

The adaptation of this model to an image is controlled by a flexible engine which uses different techniques such as GHT-based detection, parametric model adaptation and deformable model adaptation, and schedules when and how the different parts of the model are adapted. The adaptation engine can also control the mesh resolution and the number of iterations needed for each part of the models separately. An efficient and accurate segmentation can be achieved by dynamically activating and freezing the different parts of the model.

For validation, 37 data sets were used. These data sets originate from different patients than the data used for model building and training. The proposed approach could successfully segment the heart and the attached great vessels in all of the 37 data sets. Several metrics were used to validate the different types of anatomical structures composing the model. A mean point-to-surface error ranging 0.50–0.82 mm for the heart chambers and 0.60–1.32 mm for the visible parts of the great vessels is reported. The volume overlap ranges 94.0–96.2% for the heart chambers and 70.4–95.2% for the great vessels.

This approach is fast and promises to facilitate and accelerate quantitative image analysis for clinical diagnostics and to support the treatment of cardiac disease. Examples are guidance for ablation procedures in the treatment of atrial fibrillation or guidance for the implantation of electrodes in cardiac resynchronization therapy.

The presented architecture together with Simulated Search (Peters et al., 2010) for the training of boundary detection functions promises the adaptation to new segmentation tasks with a reasonable amount of effort. Future directions of our research are in including additional anatomical structures, such as the detailed aortic valve (Wächter et al., 2010), and in making use of the proposed architecture to address other modalities, e.g. MR (Peters et al., 2007), and other organs such as the brain (Kneser et al., 2009).

References

- Abadi, S., Roguin, A., Engel, A., Lessick, J., 2010. Feasibility of automatic assessment of four-chamber cardiac function with MDCT: initial clinical application and validation. *European Journal of Radiology* 74, 175–181.
- Ballard, D., 1981. Generalizing the hough transform to detect arbitrary shapes. *Pattern Recognition* 13, 111–122.
- Cauvin, J.C., Boire, J.Y., Zanca, M., Bonny, J.M., Maublant, J., Veyre, A., 1993. 3D modeling in myocardial 201Tl SPECT. *Computerized Medical Imaging and Graphics* 17, 345–350.
- Coche, E., Walker, M.J., Zech, F., de Crombrugge, R., Vlassenbroek, A., 2010. Quantitative right and left ventricular functional analysis during gated whole-chest MDCT: a feasibility study comparing automatic segmentation to semi-manual contouring. *European Journal of Radiology* 74, e138–e143.
- Cootes, T.F., Hill, A., Taylor, C.J., Haslam, J., 1994. The use of active shape models for locating structures in medical images. *Image and Vision Computing* 12, 355–366.
- de Bruijne, M., van Ginneken, B., Viergever, M.A., Niessen, W.J., 2003. Adapting active shape models for 3D segmentation of tubular structures in medical images. In: *Information Processing in Medical Imaging*. LNCS, vol. 2732. Springer, pp. 136–147.
- Ecabert, O., Peters, J., Weese, J., 2006. Modeling shape variability for full heart segmentation in cardiac computed-tomography images. In: *Proceedings of SPIE*, vol. 6144, pp. 61443R1–61443R12.
- Ecabert, O., Peters, J., Schramm, H., Lorenz, C., von Berg, J., Walker, M., Vembar, M., Olszewski, M., Subramanian, K., Lavi, G., Weese, J., 2008. Automatic model-based segmentation of the heart in CT images. *IEEE Transactions on Medical Imaging* 27, 1189–1201.
- Frangi, A.F., Niessen, W.J., Viergever, M.A., 2001. Three-dimensional modeling of functional analysis of cardiac images: a review. *IEEE Transactions on Medical Imaging* 20, 2–25.
- Fritz, D., Rinck, D., Unterhinninghofen, R., Dillmann, R., Scheuring, M., 2005. Automatic segmentation of the left ventricle and computation of diagnostic parameters using regiongrowing and a statistical model. In: *Proceedings of SPIE*, vol. 5747, pp. 1844–1854.
- Ghersin, E., Abadi, S., Yalonetsky, S., Engel, A., Lessick, J., 2009. Clinical evaluation of a fully automated model-based algorithm to calculate left ventricular volumes and ejection fraction using multidetector computed tomography. *Acute Cardiac Care* 11, 43–51.
- Gill, P.E., Murray, W., Wright, M.H., 1981. *Practical Optimization*. Academic Press.
- Gutiérrez, L.F., de Silva, R., Ozturk, C., Sonmez, M., Stine, A.M., Raval, A.N., Raman, V.K., Sachdev, V., Aviles, R.J., Waclawiw, M.A., McVeigh, E.R., Lederman, R.J., 2007. Technology preview: X-ray fused with magnetic resonance during invasive cardiovascular procedures. *Catheterization and Cardiovascular Interventions* 70, 773–782.
- Hautvast, G., Lobregt, S., Breuer, M., Gerritsen, F., 2006. Automatic contour propagation in cine cardiac magnetic resonance images. *IEEE Transactions on Medical Imaging* 25, 1472–1482.
- Hough, P.V.C., 1962. Method and means for recognizing complex patterns. *US Patent* 3069654.
- John, M., Rahn, N., 2005. Automatic left atrium segmentation by cutting the blood pool at narrowings. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005*. LNCS, vol. 3750. Springer, Berlin/Heidelberg, pp. 798–805.
- Karim, R., Mohiaddin, R., Rueckert, D., Miga, M.I., Cleary, K.R., 2008. Left atrium segmentation for atrial fibrillation ablation. In: *Proceedings of SPIE*, vol. 6918, pp. 69182U-1–69182U-8.
- Kaus, M.R., von Berg, J., Weese, J., Niessen, W., Pekar, V., 2004. Automated segmentation of the left ventricle in cardiac MRI. *Medical Image Analysis* 8, 245–254.
- Kirisli, H.A., Schaap, M., Klein, S., Neeffjes, L.A., Weustink, A.C., Walsum, T.V., Niessen, W.J., 2010. Fully automatic cardiac segmentation from 3D CTA data: a multi-atlas based approach. In: *Proceedings of SPIE*, vol. 7623, pp. 762305-1–762305-9.
- Knecht, S., Skali, H., O'Neill, M.D., Wright, M., Matsuo, S., Chaudhry, G.M., Haffajee, C.I., Nault, I., Gijssbers, G.H.M., Sacher, F., Laurent, F., Montaudon, M., Corneloup, O., Hocini, M., Haïssaguerre, M., Orlov, M.V., Jaïs, P., 2008. Computed tomography-fluoroscopy overlay evaluation during catheter ablation of left atrial arrhythmia. *Europace* 10, 931–938.
- Kneser, R., Lehmann, H., Geller, D., Qian, Y., Weese, J., 2009. Toward knowledge-enhanced viewing using encyclopedias and model-based segmentation. In: *Proceedings of SPIE*, vol. 7260, pp. 72601D-1–72601D-9.
- Lehmann, H., Kneser, R., Neizel, M., Peters, J., Ecabert, O., Kühl, H., Kelm, M., Weese, J., 2009. Integrating viability information into a cardiac model for interventional guidance. In: *Functional Imaging and Modeling of the Heart – FIMH 2009*. LNCS, vol. 5528. Springer, Berlin/Heidelberg, pp. 312–320.
- Lorenz, C., von Berg, J., 2006. A comprehensive shape model of the heart. *Medical Image Analysis* 10, 657–670.
- Lötjönen, J., Kivistö, S., Koikkalainen, J., Smutek, D., Lauerma, K., 2004. Statistical shape model of atria, ventricles and epicardium from short- and long-axis MR images. *Medical Image Analysis* 8, 371–386.
- McInerney, T., Terzopoulos, D., 1995. 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 Graphics* 19, 69–83.
- Mitchell, S., Lelieveldt, B., van der Geest, R., Bosch, H., Reiver, J., Sonka, M., 2001. Multistage hybrid active appearance model matching: segmentation of left and right ventricles in cardiac MR images. *IEEE Transactions on Medical Imaging* 20, 415–423.
- Mitchell, S.C., Bosch, J.G., Lelieveldt, B.P.F., van der Geest, R.J., Reiber, J.H.C., Sonka, M., 2002. 3-D active appearance models: segmentation of cardiac MR and ultrasound images. *IEEE Transactions on Medical Imaging* 21, 1167–1178.
- Montagnat, J., Delingette, H., 2005. 4D deformable models with temporal constraints: application to 4D cardiac image segmentation. *Medical Image Analysis* 9, 87–100.
- Peters, J., Ecabert, O., Meyer, C., Schramm, H., Kneser, R., Groth, A., Weese, J., 2007. Automatic whole heart segmentation in static magnetic resonance image volumes. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007*. LNCS, vol. 4792. Springer, Berlin/Heidelberg, pp. 402–410.

- Peters, J., Ecabert, O., Lorenz, C., von Berg, J., Walker, M.J., Ivanc, T.B., Vembar, M., Olszewski, M.E., Weese, J., 2008. Segmentation of the heart and major vascular structures in cardiovascular CT images. In: *Proceedings of SPIE*, vol. 6914, pp. 691417-1–691417-12.
- Peters, J., Ecabert, O., Meyer, C., Kneser, R., Weese, J., 2010. Optimizing boundary detection via simulated search with applications to multi-modal heart segmentation. *Medical Image Analysis* 14, 70–84.
- Plumhans, C., Keil, S., Ocklenburg, C., Mühlenbruch, G., Behrendt, F.F., Günther, R.W., Mahnken, A.H., 2009. Comparison of manual, semi- and fully automated heart segmentation for assessing global left ventricular function in multidetector computed tomography. *Investigative Radiology* 44, 476–482.
- Redwood, A.B., Camp, J.J., Robb, R.A., 2005. Semiautomatic segmentation of the heart from CT images based on intensity and morphological features. In: *Proceedings of SPIE*, vol. 5747, pp. 1713–1719.
- Rhode, K.S., Sermesant, M., Brogan, D., Hegde, S., Hipwell, J., Lambiase, P., Rosenthal, E., Bucknall, C., Qureshi, S.A., Gill, J.S., Razavi, R., Hill, D.L.G., 2005. A system for real-time XMR guided cardiovascular intervention. *IEEE Transactions on Medical Imaging* 24, 1428–1440.
- Stegmann, M.B., Pedersen, D., 2005. Bi-temporal 3D active appearance models with applications to unsupervised ejection fraction estimation. In: *Proceedings of SPIE*, vol. 5747, pp. 336–350.
- van Assen, H.C., Danilouchkine, M.G., Frangi, A.F., Ordás, S., Westenberg, J.J., Reiber, J.H., Lelieveldt, B.P., 2006. SPASM: a 3D-ASM for segmentation of sparse and arbitrarily oriented cardiac MRI data. *Medical Image Analysis* 10, 286–303.
- van Assen, H.C., Danilouchkine, M.G., Dirksen, M.S., Reiber, J.H.C., Lelieveldt, B.P.F., 2008. A 3-D active shape model driven by fuzzy inference: application to cardiac CT and MR. *IEEE Transactions on Information Technology in Biomedicine* 12, 595–605.
- von Berg, J., Lorenz, C., 2005. Accurate left atrium segmentation in multislice CT images using a shape model. In: *Proceedings of SPIE*, vol. 5747, pp. 351–360.
- Wächter, L., Kneser, R., Korosoglou, G., Peters, J., Bakker, N., Boomen, R., Weese, J., 2010. Patient specific models for planning and guidance of minimally invasive aortic valve implantation. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2010*. LNCS, vol. 6361. Springer, Berlin/Heidelberg, pp. 526–533.
- Zhang, H., Wahle, A., Johnson, R.K., Scholz, T.D., Sonka, M., 2010. 4-D cardiac MR image analysis: left and right ventricular morphology and function. *IEEE Transactions on Medical Imaging* 29, 350–364.
- Zhao, F., Zhang, H., Wahle, A., Thomas, M.T., Stolpen, A.H., Scholz, T.D., Sonka, M., 2009. Congenital aortic disease: 4D magnetic resonance segmentation and quantitative analysis. *Medical Image Analysis* 13, 483–493.
- Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., Comaniciu, D., 2008. Four-Chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features. *IEEE Transactions on Medical Imaging* 27, 1668–1681.
- Zhuang, X., Rhode, K., Razavi, R., Hawkes, D.J., Ourselin, S., 2010. A registration-based propagation framework for automatic whole heart segmentation of cardiac MRI. *IEEE Transactions on Medical Imaging* 29, 1612–1625.