ELSEVIER

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Medical Image Analysis 8 (2004) 371–386

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

Statistical shape model of atria, ventricles and epicardium from short- and long-axis MR images

J. Lötjönen [a,*], S. Kivistö [b], J. Koikkalainen [c], D. Smutek [a], K. Lauerma [b]

[a] *VTT Information Technology, P.O. Box 1206, Sinitaival 6, FIN-33101 Tampere, Finland*
[b] *Helsinki Medical Imaging Center, Helsinki University, P.O. Box 281, FIN-00029 HUS, Finland*
[c] *Laboratory of Biomedical Engineering, Helsinki University of Technology, P.O. Box 2200, FIN-02015 HUT, Finland*

Available online 21 July 2004

Abstract

We describe a new 3-D statistical shape model of the heart consisting of atria, ventricles and epicardium. The model was constructed by combining information on standard short- and long-axis cardiac MR images. In the model, the variability of the shape was modeled with PCA- and ICA-based shape models as well as with non-parametric landmark probability distributions and a probabilistic atlas. The statistical atlas was built from 25 healthy subjects. The shape model was evaluated by applying it to image segmentation. The probabilistic atlas was found to be superior to the other shape models ($p < 0.001$) in this study.

Keywords: Shape analysis; Statistical shape models; Cardiac models; Atlas; Model-based image analysis

1. Introduction

Cardiac imaging has an established role in the diagnosis of cardiac diseases. Currently modern imaging modalities produce detailed anatomic and functional information of the heart. Automated analysis of huge amounts of image data is usually a pre-requisite for comprehensive quantitative analysis of cardiac images (Frangi et al., 2001).

Many different types of atlases as well as different techniques to generate them have been published for the brain (Thompson and Toga, 1997; Chen et al., 1999; Caunce and Taylor, 2001; Hill et al., 2002; Rueckert et al., 2003) and for the heart (Frangi et al., 2002; Mitchell et al., 2002). In many atlases, the shape variability has also been modeled. The most popular approach for modeling the shape changes is the point distribution model (PDM), also referred to as active shape model (ASM) (Cootes et al., 1995) while the corresponding technique for modeling the appearance is the active appearance model (Cootes et al., 1998). PDM defines a mean model and its typical deformation modes on the basis of a training set. The deformation modes are the eigenvectors of the covariance matrix determined for point-to-point correspondence between the different models in the training set. In other words, principal component analysis (PCA) is performed. Alternatively, the shape can be modeled by independent component analysis (ICA) (Hyvärinen et al., 2001). ICA-based shape models have recently been applied into medical image segmentation (Uzümcü et al., 2003; Koikkalainen and Lötjönen, 2004). In both studies, it was found that the use of ICA deformation modes, instead of or in addition to PCA modes, improves the segmentation accuracy. Parametric Gaussian shaped distributions are another approach to model the spatial variability (Thompson and Toga, 1997; Chen et al., 1999). These distributions give a probability $p_i(x, y, z)$ that a landmark point i is found from the location (x, y, z).

The motivation for building the shape models is twofold. (1) The shape models can be used in segmentation, which is needed before any quantitative measures, such as the ejection fraction or the left ventricle mass, can be computed from image data. The use of the shape modes or probability distributions in segmentation

* Corresponding author. Tel.: +358-50-346-3250; fax: +358-3-317-4102.
E-mail address: jyrki.lotjonen@vtt.fi (J. Lötjönen).

1361-8415/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2004.06.013

regulate the deformation process. For example, only shapes similar to the training set are allowed in PCA- or ICA-based models. In addition, statistical shape models provide a compact parameterization of the deformation, because only the weights for different shape modes need to be defined. These properties make the segmentation task more robust to noise and incomplete data. (2) The shape model itself can be used as a source of quantitative information, not only as a tool for segmentation. For example, it can be used to define whether the shape of a specific structure in a patient is abnormal (assuming that the shape model is representative enough).

This paper concentrates on the shape modeling of the heart. Typically only short-axis (SA) images have been used in modeling and in automated segmentation of the left and right ventricles. Since the slice thickness is normally several times larger than the pixel size in cardiac MR images, the resolution is poor in the direction orthogonal to the slice. Therefore, the accurate localization of the ventricles in the basal and apical levels from SA images has been difficult. In practice, the models of the left ventricle are simply cut by a plane at some basal level. This certainly introduces some inaccuracies to the quantitative measures computed from the heart.

To our knowledge, only Lelieveldt et al. (2003), in addition to our own work (Lötjönen et al., 2003a), have used multiple views for creating statistical models from the heart. They used three 2-D frames taken from different orientations for building the active appearance model of the left ventricle. The shape and appearance vectors from each view were concatenated before the PCA was computed in order to preserve the coherence between the views. In our approach, we reconstructed the 3-D geometry of the heart using information from short- and long-axis (LA) images of each subject, and computed the PCA-modes for the 3-D shapes. LA images provide comprehensive and more accurate information on the modeling of the basal and apical levels. In addition to the ventricles, the atria were included in the model. The variability of the shape was modeled by PCA- and ICA-based shape models, as well as by non-parametric spatial landmark probability distributions (LPDs) and probabilistic atlas. Mean shape and gray-scale models were also constructed.

Before the 3-D geometry can be reconstructed from SA and LA views, the movement artifacts, e.g. from breathing, need to be corrected. The problem has been discussed and reported very little in the literature. Moore et al. (2003) recently built a high resolution dynamic heart model from coronal slices acquired from a healthy volunteer. They corrected the breath-hold misalignment by registering a 3-D volume with sagittal and axial scout images. A line-by-line mean squared difference was minimized. In this work, we extended this idea to the registration of two volumes. We optimized the locations of SA slices based on data from LA slices and vice versa.

The volume changes of the ventricles and atria during the cardiac cycle, both during systole and diastole, indicate the cardiac function. Volumetry of all four cardiac chambers throughout the entire cycle without tools for segmentation is so time consuming that a lot of valuable functional data is overlooked. Analysis of atrial and right ventricular function can indicate a subtle cardiac disease in patients that have normal systolic left ventricular function. Therefore it is essential to provide segmentation for cardiac chamber analysis to screen for early signs of myocardial disease. In (Lötjönen et al., 2003b), an a priori model was non-rigidly registered to target data using simultaneously both SA and LA images. In this work, we applied the technique to segmentation, instead of motion tracking, and used the statistical information about the shape variability to regulate the deformation of the model.

This paper makes the following contributions:

- A novel technique to reconstruct the 3-D geometry from multiple views is demonstrated.
- A segmentation technique, utilizing simultaneously both SA and LA images, and different statistical shape models, is proposed.
- A novel method is presented for correcting movement artifacts from MR images.
- Non-parametric LPD is proposed for modeling the shape changes.
- The shape model of the atria is built for the first time.
- The automatic segmentation of the atria was demonstrated.

In Section 2 we present all the steps needed to construct shape models from the SA and LA images, and a segmentation technique to evaluate the shape models. Section 3 visualizes the shape models generated and presents the segmentation results. Finally, we conclude with a discussion in Section 4.

2. Methods

The whole model construction protocol has been summarized in Fig. 1.

2.1. Materials

MR images of 25 healthy adult subjects formed the database. The subjects were the control persons of two separate clinical studies. The images were obtained with a 1.5 T Siemens Vision and Siemens Sonata imagers and with a phased array coil (Siemens, Erlangen, Germany) at the Helsinki Medical Imaging Center (Helsinki University Central Hospital, HUCH). Breath hold gradient echo turbo flash cine images were acquired in the LV short-axis planes with 6–7 mm slice thickness and 7–8 mm gap for the volumetric study of ventricles (Alfakih et al., 2003). Cine images in the four chamber planes

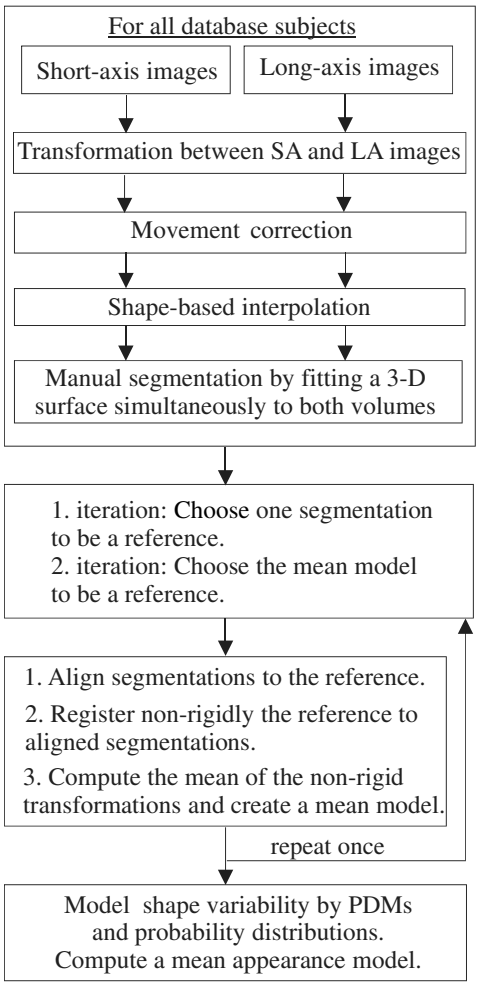

Fig. 1. The model construction protocol.

with 6–7 mm slice thickness and 3–4 gap were obtained for atria (Järvinen et al., 1994a,b). The following imaging parameters were used: repetition time 30 or 40 ms, echo time 4.8 ms, flip angle 20°, matrix size 256×256 and field of view 250–300 mm. In practice, a slight tilt compared to the standard LA view is used in HUCH in order to get the mitral and tricuspidal valves to the same image slice. The pixel size for SA and LA images was either 1.0×1.0 or 1.4×1.4 mm. The number of SA and LA slices was 4–5 and 4–7, respectively, depending on the size of the subject's heart. In this work, the 3-D shape model was built from the first image of each time series, i.e. from the end-diastolic phase. Fig. 2 shows SA and LA slices from one subject.

2.2. Transformation between short- and long-axis images

Since both SA and LA images were utilized to build the shape model, voxel-by-voxel based correspondence need to be defined between the image volumes. Standard registration techniques were not required, because the SA and LA image series were acquired during one imaging session and the image headers contained all necessary information to compute the location of each voxel relative to the imaging device. In general, the transformation between the volumes in the voxel-space is not rigid because of anisotropic voxels, different size of voxels and movement artifacts. The effect of possible movement artifacts is discussed in Section 2.3.

In the following, the co-ordinates of a voxel in the source volume, denoted by (X, Y, Z), are defined in the co-ordinate system of the destination volume, denoted by (X', Y', Z'). The co-ordinate system of the imaging device is denoted by (X^*, Y^*, Z^*). Next, the parameters for the source volume are defined. Corresponding symbols for the destination volume have a dash. The pixel size of the source volume is (s_x, s_y, s_z). The pixel size in the z-direction is defined to be the distance between neighboring slices, i.e. slice separation. The slice separation is equal to the sum of the slice thickness and the possible gap between the slices. The position of the first voxel of the source volume in the scanner's co-ordinate system is denoted by (o_x, o_y, o_z). In addition, the orientation of x- and y-directions (row and column) are denoted by (r_x, r_y, r_z) and (c_x, c_y, c_z). The normal vector (n_x, n_y, n_z) of the slices can be computed, for example, by subtracting the image positions of the second and the first slices of the volume, and normalizing its length.

The location of a source voxel in the co-ordinate system of a scanner is computed as follows:

$$X^* = s_x X r_x + s_y Y c_x + s_z Z n_x + o_x, \tag{1}$$

$$Y^* = s_x X r_y + s_y Y c_y + s_z Z n_y + o_y, \tag{2}$$

$$Z^* = s_x X r_z + s_y Y c_z + s_z Z n_z + o_z. \tag{3}$$

The location of the voxel in the co-ordinate system of the destination volume is computed as follows:

$$X' = \frac{(X^* - o'_x)r'_x + (Y^* - o'_y)r'_y + (Z^* - o'_z)r'_z}{s'_x}, \tag{4}$$

$$Y' = \frac{(X^* - o'_x)c'_x + (Y^* - o'_y)c'_y + (Z^* - o'_z)c'_z}{s'_y}, \tag{5}$$

$$Z' = \frac{(X^* - o'_x)n'_x + (Y^* - o'_y)n'_y + (Z^* - o'_z)n'_z}{s'_z}. \tag{6}$$

2.3. Movement correction

Typically, several MR image series are acquired during one imaging session. If a subject moves during the imaging session, the relation between image series, derived from the image headers, is lost and image registration is needed to realign the images. A subject may move because of several reasons, e.g. coughing, breathing or change of inconvenient pose. Breathing is a major source of movement artifacts in cardiac imaging,

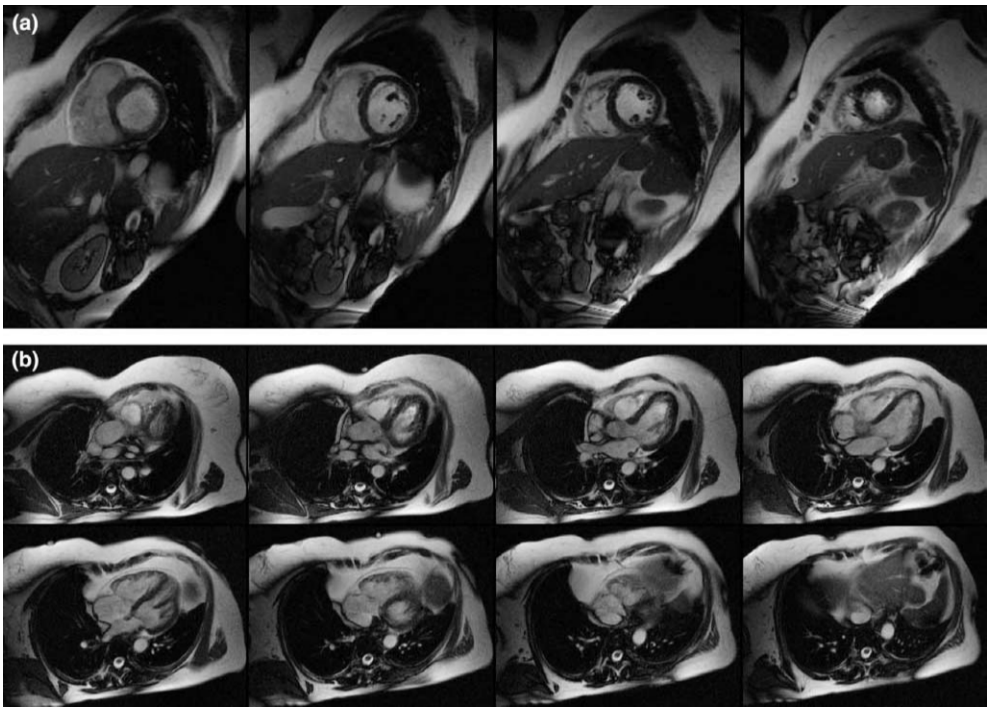

Fig. 2. Original MR images of one subject: (a) SA images and (b) LA images.

as the heart's own movement is handled by the ECG gating. McLeish et al. (2002) studied the movements of the heart due to respiration. They observed translations up to 23.5 mm in the heart because of breathing. If modern imaging techniques, such as SENSE, were used to acquire the whole volume during a breath hold, all slices would be in correct locations relative to each other. However, when the cine sequences are used to track the cardiac motion, an image series produced during a breath hold contains typically slices from several time points but only from one spatial location. If the phase of the breathing cycle is not similar during all acquisitions, slices from different image series will be misaligned relative to each other, and a volume built from the image series does not represent the real anatomy of the subject.

The movement artifacts can be visually observed by forming an image volume from the cine series and computing a cross-section of the volume. Figs. 3(a) and (b) show one SA and LA slice, respectively, from a subject having severe movement artifacts. The horizontal lines superimposed on the images indicate the cross-section planes. The cross-sections computed from the SA (six slices) and LA volumes (eight slices) are shown in Figs. 3(c) and (d), respectively. The cross-sections have been interpolated to isotropic voxel size using nearest neighbor interpolation for better visualizing the shifts. The dark, almost vertical stripe in the middle of the images represents the septum, which is naturally a smooth and continuous object. In the cross-sections, the septum is not smooth nor even continuous. For example, the shift between slices three and four in the SA cross-section, i.e. the gap in the septum, corresponds to 13 mm. Figs. 3(e) and (f) show the cross-sections after applying the automatic correction algorithm, which is described next.

The criterion of the smooth and continuous septum could be used to shift the slices in the SA and LA volumes separately. The problem in making the corrections separately is that the comprehensive shape information

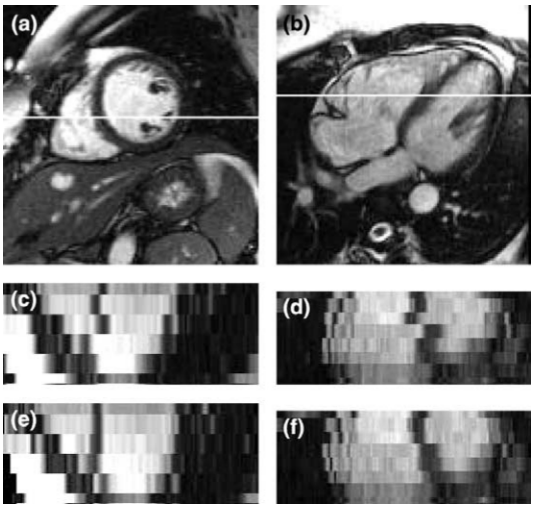

Fig. 3. Movement correction. Original (a) SA and (b) LA slices. The horizontal line represents the location from which cross-sections are computed. The cross-sections of the (c) SA and (d) LA volumes before movement correction. The cross-sections of the (c) SA and (d) LA volumes after movement correction.

available in the other volume is not used. In other words, mapping the structures from the SA volume to the LA volume, using Eqs. (1)–(6), could be inaccurate although the shape of the heart may appear visually realistic in the both volumes separately.

The basic idea of our automatic registration algorithm is presented in Fig. 4. The normalized mutual information (NMI) (Studholme et al., 1999) is maximized between the SA and LA data,

$$E_{\mathrm{NMI}} = \frac{H(\mathrm{SA}) + H(\mathrm{LA})}{H(\mathrm{SA}, \mathrm{LA})}, \tag{7}$$

where $H(\mathrm{SA})$ and $H(\mathrm{LA})$ are marginal entropies and $H(\mathrm{SA}, \mathrm{LA})$ a joint entropy of the data. The voxel-by-voxel correspondence is calculated as described in Section 2.2. The displacement of each slice is assumed to be independent on the displacements of the other slices. One slice is randomly chosen from the SA and LA stacks and it is moved to the direction that maximizes the NMI. The slice locations are iterated until the NMI does not increase.

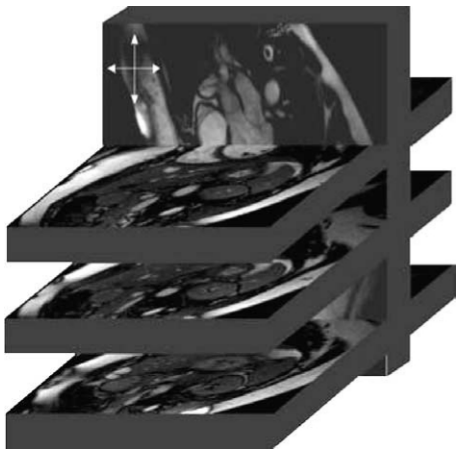

Fig. 4. The optimization of a LA slice relative to a SA volume.

Although this work concentrated in building a 3-D shape model, the MR-images acquired were 4-D cine images (Section 2.1) containing information about the motion of the heart. Because all time points related to one spatial location were acquired during the same breath-hold, data from all time instants can be used to compute the NMI. In other words, as a slice is displaced, all time points of the slice are displaced. Although the use of several time points is not essential, we have found that it improves the registration accuracy.

Fig. 5 shows the SA and LA volumes visualized using a chessboard visualization technique before the motion correction (the top row) and after the motion correction (the bottom row). The first two columns show the volumes as the LA volume is transformed to the SA coordinate system. The third and fourth columns visualize the result as the SA volume is transformed to the coordinate system of the LA volume. As can be noticed from the images, the edges are more continuous across the chess-boxes with the movement correction than without it. The black boxes indicate that no data is available for that location in the volume, e.g. the SA volume does not contain any data from atria. The algorithm was evaluated more carefully in (Lötjönen et al., 2004).

In practice, a subject can normally keep almost a similar breathing phase between the acquisitions and no clear misalignment can be visually detected from the volumes. However, when movement artifacts existed, the automatic correction performed well. A good indication of successful movement correction was that the final manual segmentation results fitted well to both SA and LA volumes.

2.4. Manual segmentation

Atria, ventricles and epicardium were manually segmented from the images by a medical specialist. A

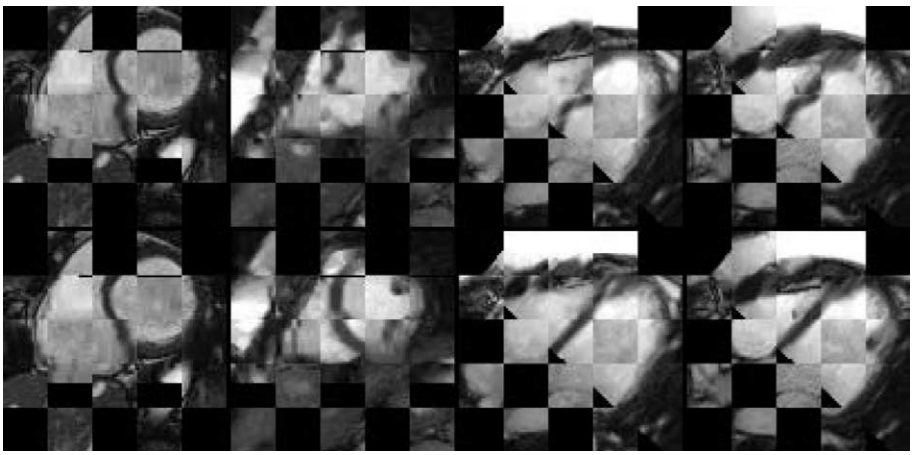

Fig. 5. Movement correction result. The top and bottom rows show the SA and LA volumes before and after, respectively, the movement correction. The chessboard visualization technique was used.

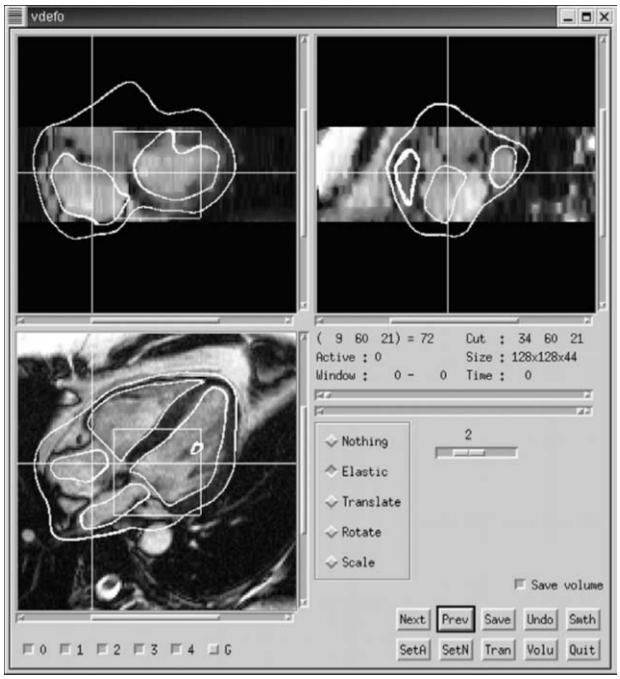

Fig. 6. A software tool for manual segmentation.

software tool, designed for interactively making 3-D deformations, was used (Fig. 6) (Lötjönen, 2003). With it, a triangulated surface model (white contours superimposed on the MR images in Fig. 6) was registered manually and non-rigidly to edges in the SA and LA image volumes. The user can drag-and-drop the surface locally in 3-D. The surface is elastically deformed inside a user defined box of interest (the white 3-D box on the MR images in Fig. 6): the deformation is zero on the borders and outside the box, and as dragged by the user in the center of the box. The tool shows three cross-sections of the volume and the deformations can be made in any of these cross-sections. We found, that the orthogonal cross-sections to the original LA slices were very useful in segmenting atria (see the right atrium under the white cross in all images in Fig. 6; the left atrium is also visible in the image on the right, i.e. the right-most round shape).

The SA and LA volumes were shape-based interpolated before the segmentation in order to obtain an isotropic voxel size (Grevera and Udupa, 1996). Because no severe gray-scale variation was detected between slices in our dataset, the intensities were not normalized. Figs. 7(a) and (b) show original slices on the top and bottom, and an interpolated slice in the middle. Although the slice separation is large compared to the pixel size, i.e. about 10 slices need to be interpolated between the original slices in the SA volumes, the results are visually satisfactory, because the heart is geometrically a relatively smooth object. In addition, the shape model built in this work is even smoother, because only the main structures were modeled. For these reasons, we

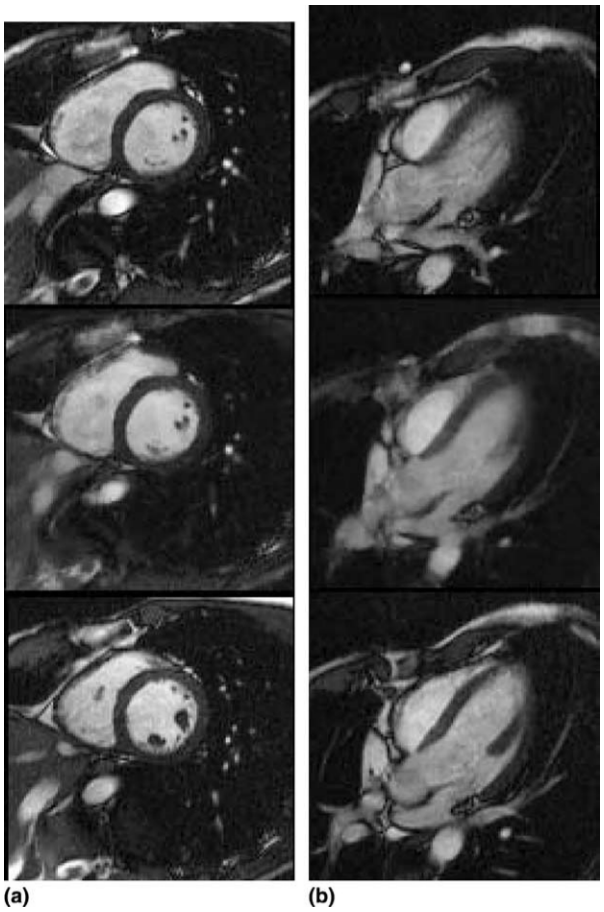

Fig. 7. Shape-based interpolation of (a) SA and (b) LA slices. The original slices are shown on the top and bottom, and an interpolated slice in the middle.

believe that geometrical errors from interpolation do not play an important role in the modeling. The total computation time for interpolating the SA and LA volumes was 49 s (1.7 Gz Pentium processor).

Thereafter, the prior model was manually deformed to the SA volume. After finishing the deformation in the SA volume, the deformed model was transformed to the co-ordinate system of the corresponding LA volume and the matching was continued in the LA volume. The deformed model was transferred iteratively between the volumes until no deformations were needed. The transformation used was described in Section 2.2. The whole segmentation procedure is shown in Fig. 8.

A segmentation result for one subject is shown in Fig. 9. The following guidelines were used in segmentation:

- The basal regions of the right and left ventricles were defined using the LA volume. The tricuspid and mitral valves, visible in the images, form an interface between the ventricles and atria. Because our final target in the future is to automatically segment all time points, the valves are not good landmarks, as they open after the systole and make the definition

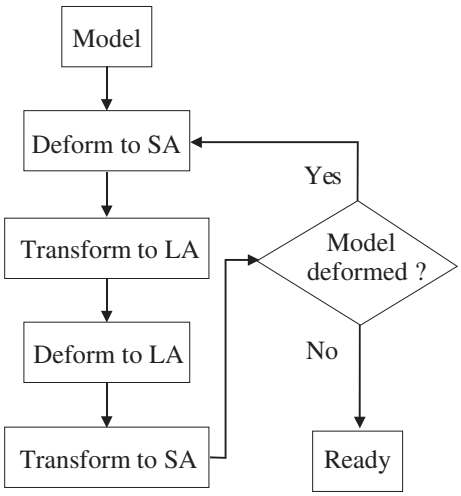

Fig. 8. Segmentation flow-chart.

of the boundary between ventricles and atria impossible. Therefore, we defined the most basal points of the myocardium (the points where the valves attach to the myocardium). The ventricle-atrium boundary was approximated by a straight line between the points in each slice.

- Because of a large slice thickness and slice separation, the segmentation of the apex using the SA volume is inaccurate. The apex was located mainly based on information from the LA volume.
- The pulmonary artery was left outside the right ventricle in the model. The detection of the interface between the artery and the ventricle would be easy if the pulmonary valve was visible in the images. This was not the case in our data. The pulmonary artery was usually clearly visible on the most basal SA slices but the valve level was difficult to locate as moved toward the apex. A compromise was made and a smooth boundary was created from the slice where the artery was visible to a more apical slice which definitely contained only the right ventricle. The pulmonary artery was eliminated from the right ventricle because of volumetric measures. As volumetric measures, such as the ejection fraction, are computed, the results will be skewed if the right ventricle contains the pulmonary artery.
- The inferior and superior vena cavas were left outside the right atrium in the model. The vena cavas can be seen as round structures below the right atrium in Fig. 9.
- The pulmonary vein was left outside the left atrium in the model.
- The papillary muscles can be seen clearly inside the left ventricle. In cardiac volumetry, the papillary muscles are usually classified to the left ventricle making the left ventricle almost circular in the SA slices. For modeling purposes, especially for the mean gray-scale image, the delineation of the papillary muscles is beneficial. Therefore, we made two segmentations for the left ventricle, with and without papillary muscles.
- The epicardium was not segmented to any structures above the tricuspid and mitral valves but it was only set to contain the atria.

2.5. Modeling the variability of shape

The variability in the shape was modeled in three ways: with PDMs, with LPDs and with probabilistic atlas. Since the heart contains only a few anatomic landmark points, a set of semi-landmarks from the surfaces of the objects was used. A set of corresponding semi-landmarks was determined for each database subject using non-rigid registration, as was described in (Frangi et al., 2002).

Instead of registering non-rigidly gray-scale volumes, segmented data were utilized. First, intensity volumes were created from surface data. In these intensity volumes, each object was represented using one gray-scale value. One volume, volume V_0, out of N volumes was chosen to be a reference model. Then, the other volumes, $V_1, V_2, \ldots, V_{N-1}$, were aligned with the reference using translation, rotation, and isotropic scaling. The

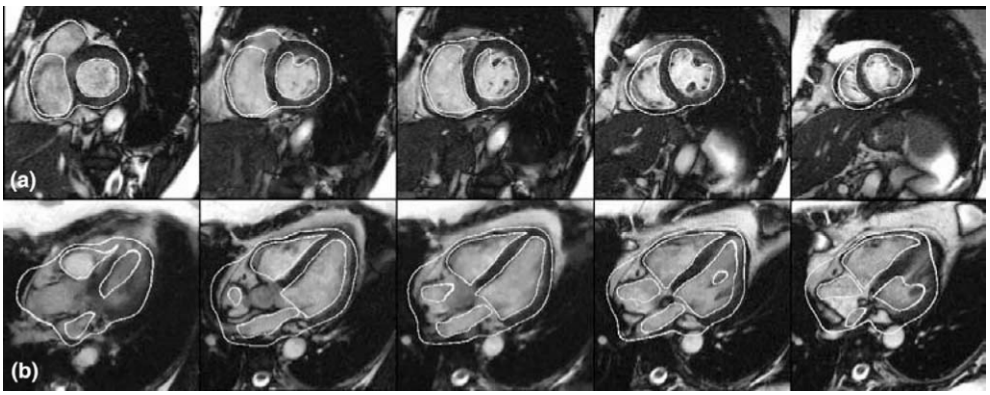

Fig. 9. A segmentation result for one subject shown on a few slices. The contours on the SA volume (a) and on the LA volume (b) are from the same triangulated surface.

gradient minimization technique was used to maximize the NMI between the volumes.

Next, the reference volume, V_0, was non-rigidly registered to the aligned volumes, $V_1, V_2, \ldots, V_{N-1}$, using a deformation sphere technique (Lötjönen and Mäkelä, 2001) which is closely related to the free-form deformation (FFD)-algorithm with radial basis functions proposed by Rohde et al. (2003). In the deformation sphere technique, smooth deformations are applied to voxels inside a sphere in such a way that the NMI is maximized. The location of the sphere is randomly chosen from the surfaces of ventricles, atria and epicardium, and it is varied during the iteration. In addition, the radius of the sphere is gradually decreased during the iteration. The program outputs a volumetric transformation $\mathbf{T} : x \mapsto x'$ from the reference model point x to the corresponding target point x'.

The nodes of the triangulated surface model of the reference subject, obtained from the manual segmentation, were considered as semi-landmarks. These semi-landmarks, $\mathbf{L}_0$, were propagated to the remaining database subjects based on the volumetric transformations $\mathbf{T}_i$, $i = \{0, 1, \ldots, N-1\}$. As a result, a set of corresponding semi-landmarks was achieved for each database subject, $\mathbf{L}_i$, $i = \{0, 1, \ldots, N-1\}$. The mean shape, $\bar{\mathbf{L}}$, was then computed from the semi-landmarks:

$$\bar{\mathbf{L}} = \frac{1}{N} \sum_{i=0}^{N-1} \mathbf{L}_i. \tag{8}$$

To reduce the bias of the mean shape towards the selected reference subject, and to give a better a priori estimate in the non-rigid registration, the preceding procedure was repeated by using the mean shape as a reference model. An intensity volume, $\bar{V}$, was constructed from the mean shape, $\bar{\mathbf{L}}$, and the volumes $V_0, V_1, \ldots, V_{N-1}$ were registered with the intensity volume of the mean shape, $\bar{V}$. Finally, the semi-landmarks of the mean shape, $\bar{\mathbf{L}}$, were propagated to the database subjects to establish the final semi-landmarks, $\mathbf{L}_i$, $i = 0, 1, \ldots, N-1$. From the semi-landmarks, a new mean model can be computed by using Eq. (8).

2.5.1. Point distribution model

In PDMs, typical modes of shape variation were defined for set of points, typically for surface points. The modes of shape variation were determined from the deviations of database subjects' landmarks from the mean shape, $\mathbf{d}_i = \mathbf{L}_i - \bar{\mathbf{L}}$. The shape variability was modeled using a linear model:

$$\mathbf{L} = \bar{\mathbf{L}} + \Phi\mathbf{b}, \tag{9}$$

where the matrix $\Phi = [\phi_1|\phi_2|\ldots|\phi_t]$ consists of the modes of shape variation, ϕ_k, and $\mathbf{b} = [b_1\ b_2\ \ldots\ b_t]^{\mathrm{T}}$ is a vector containing the weights for the modes.

In this study, the modes of shape variation were calculated in two different ways: using PCA and ICA. When PCA is used (Cootes et al., 1995), the matrix Φ consists of the eigenvectors ϕ_k of the covariance matrix defined for the displacement vectors $\mathbf{d}_i$. In PCA, the objective is to find the modes of shape variation that explain the maximal amount of the variance in the database. In ICA, on the other hand, the objective is to find such modes of shape variation that are statistically independent. The Central Limit Theorem states that the combination of independent distributions is more Gaussian than the original distributions. Hence, the problem of finding independent components is equivalent to maximizing the non-Gaussianity of the data. There are several measurements which can be used to estimate the amount of the Gaussianity, such as kurtosis and negentropy (Hyvärinen et al., 2001). In this work, ICA was performed with the FastICA MATLAB package (FastICA, 1998).

2.5.2. Landmark probability distribution

In most cases, the assumption of the ellipsoid-shaped variability of landmarks is applicable, as assumed in parametric Gaussian shaped distributions (Thompson and Toga, 1997; Chen et al., 1999). However, if, for example, the twisting motion of the heart was modeled, a Gaussian distribution would not realistically present the variability of landmarks. In this paper, the distribution of landmark points was modeled by non-parametric probability functions $p_i(\mathbf{x}_i)$, where $p_i(x, y, z)$ is a probability that the landmark i appears at the position (x, y, z), and $\mathbf{x}_i = (x_i, y_i, z_i)$ denotes the 3-D position of the landmark point.

LPDs were computed for the set of semi-landmarks, $\mathbf{L}_i$. The probability distribution of the landmark i was a volume, W_i, centered around the landmark in the mean shape. The dimensions of W_i were large enough to contain the landmark i in all volumes V_i, $i \in \Omega$. If the number of the volumes V_i were high, the probability values could be accurately estimated by computing the number of landmarks in each voxel and normalizing the value by the total number of landmarks. Because the number of volumes in our database was relatively low, we used Parzen windowing with a Gaussian kernel ($\sigma = 4$ mm) before the normalization (Duda and Hart, 1973).

2.5.3. Probabilistic atlas

A probabilistic atlas gives a spatial probability $p_s(\mathbf{x})$ that a structure s appears at location $\mathbf{x}$. A probabilistic atlas is built by (1) affinely registering all database subjects with a reference subject, (2) blurring the registered segmented edge images with a Gaussian kernel ($\sigma = 4$ mm), and (3) averaging all blurred images. Probabilistic atlases have been used in the expectation maximization (EM) based segmentation of the brain (Leemput et al., 1999). Lorenzo-Valdés et al. (2003) recently built a probabilistic 4-D cardiac atlas of ventricles

and epicardium, and used it in the segmentation of cardiac cine images. In this work, we built a 3-D probabilistic atlas separately for atria, ventricles and myocardium.

2.6. Construction of mean model

In Section 2.5, the mean model was constructed from the corresponding semi-landmarks using Eq. (8). However, this model is only mean in the sense of shape, not appearance, i.e., gray-scale. Although the emphasis in this paper was to model the variability of the shape, not appearance, a mean model for appearance was also constructed. This was performed by utilizing the volumetric transformations of the non-rigid registrations, $\mathbf{T}_i$, that were described in Section 2.5 to establish the correspondence between database subjects. The applied procedure was similar to the procedure proposed in (Guimond and Thirion, 2000), and is briefly summarized.

The volumetric transformations $\mathbf{T}_i$, $i = \{0, 1, \ldots, N-1\}$ gave the mapping from the reference model to each database subject. A mean transformation, $\bar{\mathbf{T}}$, was calculated:

$$\bar{\mathbf{T}} = \frac{1}{N} \sum_{i=0}^{N-1} \mathbf{T}_i. \tag{10}$$

The transformation from database subject i to the mean model could be computed from

$$\mathbf{T}'_i = \bar{\mathbf{T}}_0(\mathbf{T}_i^{-1}). \tag{11}$$

The mean model, also for the appearance, was computed by averaging the intensity-normalized gray-scale volumes voxel-by-voxel after applying the transformations $\mathbf{T}'_i$. The intensity normalization was done by fitting three Gaussian distributions to the gray-scale histograms using EM-algorithm. The gray-scale values in all volumes were linearly re-mapped to the gray-scale values of the reference volume by matching the histogram peaks. However, as the labeled images were used instead of the real gray-scale images in the non-rigid registrations, the correspondence between database subjects can be guaranteed only on the surfaces of the modeled objects. The transformations inside the objects are interpolated from the surfaces, which results in a blurred mean gray-scale appearance model.

2.7. Segmentation using statistical shape models

The shape model, consisting of the SA and LA volumes, was non-rigidly registered to data to be segmented. The mean gray-scale model was used as an a priori model. We also tested the shape model with gray-scale values from the reference volume, but better segmentation accuracy was achieved with the mean gray-scale model.

The NMI was used as a similarity measure:

$$E_{\mathrm{NMI}} = \frac{H(S) + H(D_i)}{H(S, D_i)}, \tag{12}$$

where $H(S)$ and $H(D_i)$ are marginal entropies and $H(S_i, D_i)$ a joint entropy of the source data S and the destination data D_i. In this study, the source data consist of both SA and LA source volumes, $S = \{S_{\mathrm{SA}}, S_{\mathrm{LA}}\}$, and the destination data consist of the SA and LA data of the destination subject, $D_i = \{D_{i,\mathrm{SA}}, D_{i,\mathrm{LA}}\}$. The absolute value of the gray-scale difference was also tested but no significant differences to the NMI was detected.

The difference in the orientation and the position of the SA and LA slices complicates the computation of the NMI. If the non-rigid transformation was defined for the SA volume, it should be aligned (Section 2.2) into the LA co-ordinates before computing the NMI for the LA volume. Making this transformation during the non-rigid registration would be very time consuming. Therefore, the LA volume was transformed to the SA co-ordinates before the non-rigid registration (Fig. 10). Because the SA and LA volumes do not contain exactly the same regions of interest, e.g. the SA volume does not contain the atria, a mask was used to define if a voxel is used to compute the NMI. Fig. 10(c) demonstrates a mask for the frames in Figs. 10(a) and (b). The SA volume information is used in the white regions, and the LA volume information in both the white and light gray regions.

Before the non-rigid registration, the location of the heart in the target data was first defined coarsely by sliding (translating) the shape model over the data and searching the maximum of Eq. (12). Then, the shape model was aligned to the target data by translation, rotation and isotropic scaling. The gradient minimization technique was used to maximize the NMI.

The non-rigid registration was regulated by utilizing the information about the variability of the shape in the database subjects. Two different approaches were used because of different characteristics of the statistical shape models. The shape variability was modeled as described in Section 2.5. However, the real gray-scale volumes were used in the initial alignment using the procedure described above.

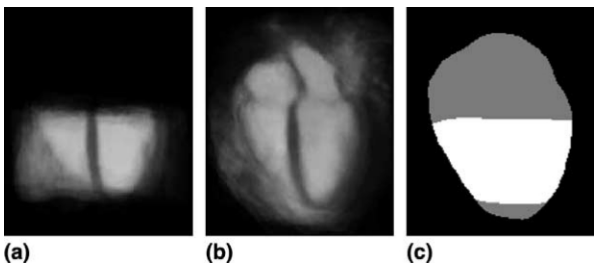

Fig. 10. Cross-sections of the (a) SA and (b) LA volumes after a transformation to a common co-ordinate system. (c) A mask used to select voxels for the computation of the NMI.

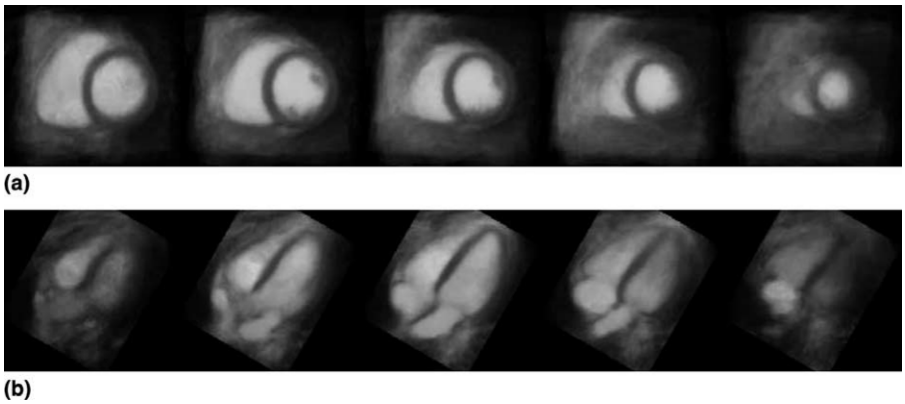

Fig. 11. Mean (a) SA images and (b) LA images.

2.7.1. Point distribution model

In PDMs, typical modes of shape variation are defined from a training set, and only shape changes consistent with the modes are allowed (Eq. (9)). In practice, the segmentation problem is reduced for defining the optimal weighting values of the modes. The conjugate gradient method was used to find the weights.

A statistical deformation model (SDM) (Rueckert et al., 2001) was built in this work. SDM contains the PDM for the non-rigid transformation instead of surface points as in ASM. The output of SDM is a 3-D transformation which can be used to transform any spatial information, such as surface points of an a priori model to produce the segmentation of destination data. Because the objective of this study is to evaluate and compare the shape models developed, the criteria to deform the models should be kept as similar as possible. In SDM, the NMI is maximized between the volumes, not only near the surfaces. Similar strategy is used also with probabilistic shape models, as described next.

2.7.2. Probabilistic shape models

In LPD and probabilistic atlas, the a priori knowledge of the shape is taken into account by maximizing the following equation:

$$E = E_{\mathrm{NMI}} + \gamma E_{\mathrm{shape}}, \tag{13}$$

where E_{NMI} is as defined in Eq. (12), E_{model} measures the consistency of the model with the probability distributions, and γ is a user-defined weight parameter.

If the probabilities between the landmarks were assumed to be independent, the likelihood of the current shape model configuration could be computed from the following equation:

$$E_{\mathrm{model}} = \frac{1}{N} \prod_{i \in 1,..,N} p(\mathbf{x}_i), \tag{14}$$

where N is the number of the landmarks, and $p(\mathbf{x})$ is the probability that the landmark i appears at location $\mathbf{x}_i$. The probability distributions are defined separately for each landmark in LPD and for each surface in probabilistic atlas. Because the direct evaluation of Eq. (14) produces enormous numerical instability, a common practice is to use the logarithmic transformation:

$$E_{\mathrm{model}} = \frac{1}{N} \sum_{i \in 1,..,N} \log[p(\mathbf{x}_i)]. \tag{15}$$

The deformation sphere approach, briefly summarized in Section 2.5, was used to produce non-rigid transformations by maximizing the measure in Eq. (13). In this work, γ was 0.05 in LPD-based and 0.1 in probabilistic atlas-based segmentations. A few different γ-values were

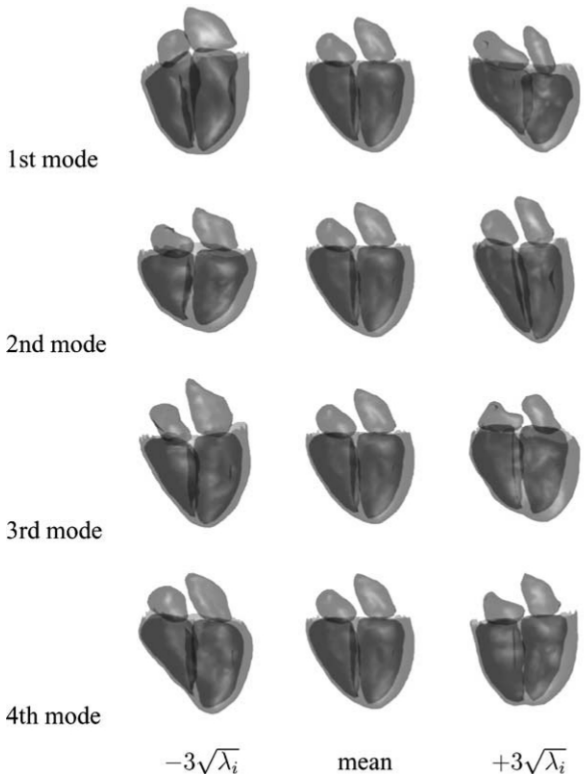

Fig. 12. Four PCA-modes.

tested and the one producing the lowest segmentation error was chosen. The radius of the deformation sphere was decreased from 40 voxels to 10 voxels during the iteration.

3. Results

3.1. Visualization of statistical shape models

In Fig. 11, slices from the mean SA and LA volumes are shown. The images represent a mean in the sense of shape and gray-scale.

Fig. 12 shows the first four PCA-based deformation modes. The surfaces of the all database subjects have been visualized, for comparison, in Fig. 13. The weight b_k for the mode k was $3\sqrt{\lambda_k}$ where λ_k is the eigenvalue of the eigenvector ϕ_k. Because the database volumes were aligned before non-rigid registration using a 7-parameter transformation (rotation, translation and isotropic scaling) instead of the quasi-affine (9-parameter with anisotropic scaling) or a generic affine (12-parameter) transformation, the first modes explain mainly changes on the global size. The 9- and 12-parameter affine transformations change the shape of the objects, not only the size. Because our objective in the future is to use the shape model also in quantitative analysis, e.g. compare different populations with each other, we preferred to include all possible shape changes in the shape models.

Four ICA modes are visualized in Fig. 14. The PCA modes are sorted according to the spatial variance they explain, i.e. the larger the eigenvalue is the more variability of the input data it explains. A similar sorting is not provided directly for the ICA modes. Different criteria can, however, be developed to sort the ICA modes (Uzümcü et al., 2003). In Fig. 14, we chose a few modes where the deformations were located into atria. The modes vary the shape as well as relative size of the atria. The ICA modes are typically more local than the PCA modes (Uzümcü et al., 2003; Koikkalainen and Lötjönen, 2004).

LPDs for a set of landmark points are shown in Figs. 15(a) and (b). The figure shows the projection of

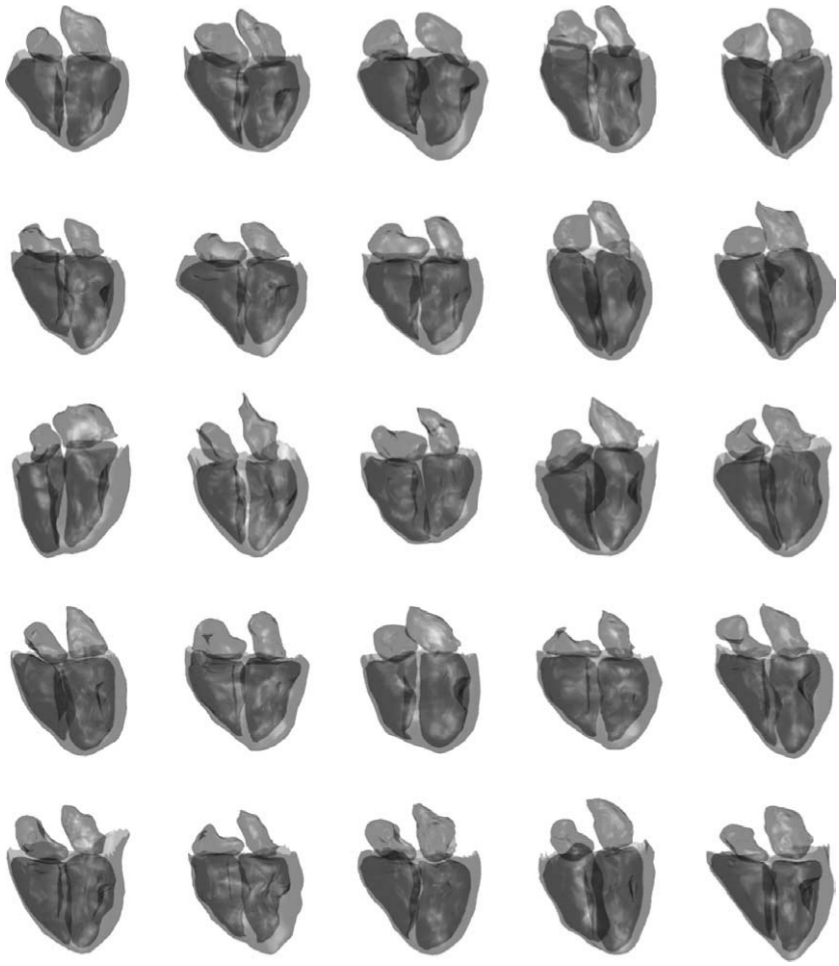

Fig. 13. Surfaces of all database subjects.

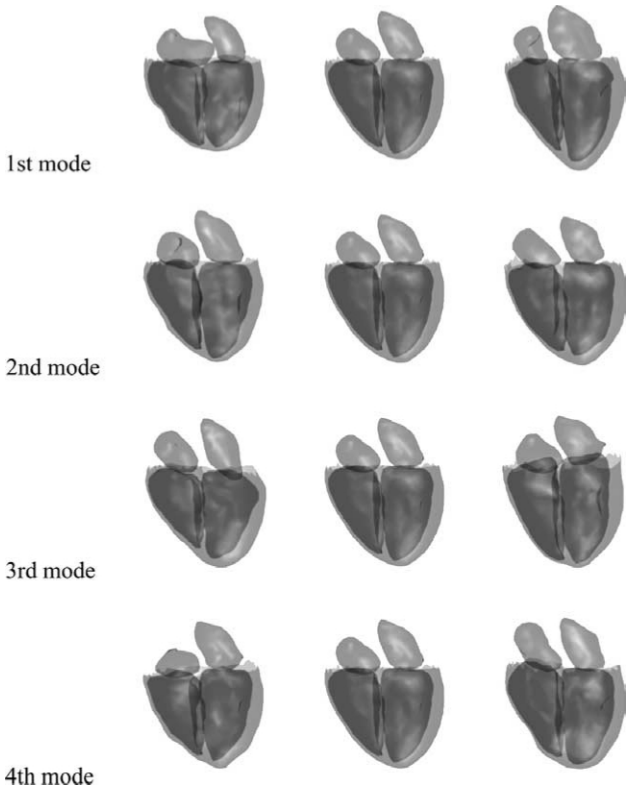

Fig. 14. Four ICA-modes.

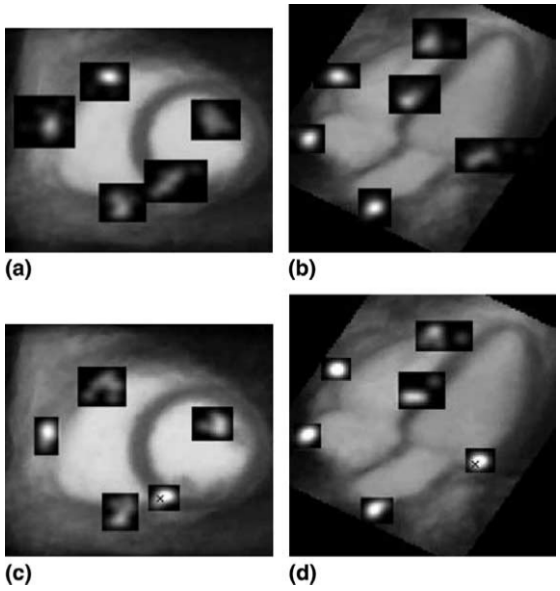

Fig. 15. Probability distributions shown on an (a) SA slice and (b) LA slice. Probability distributions considering the landmark is located on the black cross for an (c) SA slice and (d) LA slice.

the 3-D map to the 2-D slice level. We have also implemented a software tool by which a user can study the spatial relationships between landmarks. First, a user indicates a point $\mathbf{q}_i$ in the vicinity of one landmark i (black crosses in Figs. 15(c) and (d)). The maps show the spatial distribution of all landmark points assuming that the landmark i would be close to the point $\mathbf{q}_i$, formally $p(\mathbf{q}|\mathbf{q}_i)$. The landmark is considered to be close to $\mathbf{q}_i$ in the tool if its distance to $\mathbf{q}_i$ is lower than a user-defined parameter r ($r = 5$ mm voxels and contained points from seven database subjects in Figs. 15(c) and (d)). In practice, all database subjects having the landmark i close to the point $\mathbf{q}_i$ were chosen, and the probability distributions were re-computed. The relations between different landmarks can be studied by comparing the changes in the distributions as different points (black crosses) are pointed on the images. For example, if the landmark on the inferior region of the right ventricle is studied, its distribution in Fig. 15(c) remains relatively unchanged (as compared with Fig. 15(a)) indicating that its distribution is not closely correlated with the distribution of the landmark pointed by the user (black cross). Obviously, our database is still relatively small, and detailed conclusions should be avoided.

Although all database subjects were considered to be healthy, the data from one subject were clearly different. The impact of one subject is averaged out in the probability maps containing data from all subjects (Figs. 15(a) and (b)). However, the outlier can be seen on the septal landmarks in Fig. 15(d).

The probabilistic atlas was built separately for each object. Fig. 16 shows SA and LA slices from the probabilistic atlas where all five atlases have been fused.

3.2. Segmentation using statistical shape models

The segmentation results are summarized in Table 1. The segmentation error was defined as a mean distance from the surface of the shape model to the surface of the corresponding manually segmented surface. The both surfaces were presented by triangles, and the distance was computed between the nodes of the shape model and the triangles of the manually defined surface. The cross-validation was used, i.e. each database subject was once regarded as a target and the shape model was built using the other database subjects.

The segmentation accuracy of the ventricles was notably better than the segmentation accuracy of the atria. Because of much larger volume, the ventricles dominated in the automatic initialization (rotation, translation and isotropic scaling). For this reason, the initialization of the atria became relatively poor in a few cases affecting also to the final segmentation accuracy. The initialization of the right atrium was very poor (9.9 mm) especially for one subject: the overlap between the aligned model and the target was less than 50%. The segmentation error of the right atrium was also high (8.2 mm using the probabilistic atlas), which was double compared with the subject having the second worst accuracy. If this case was omitted from the

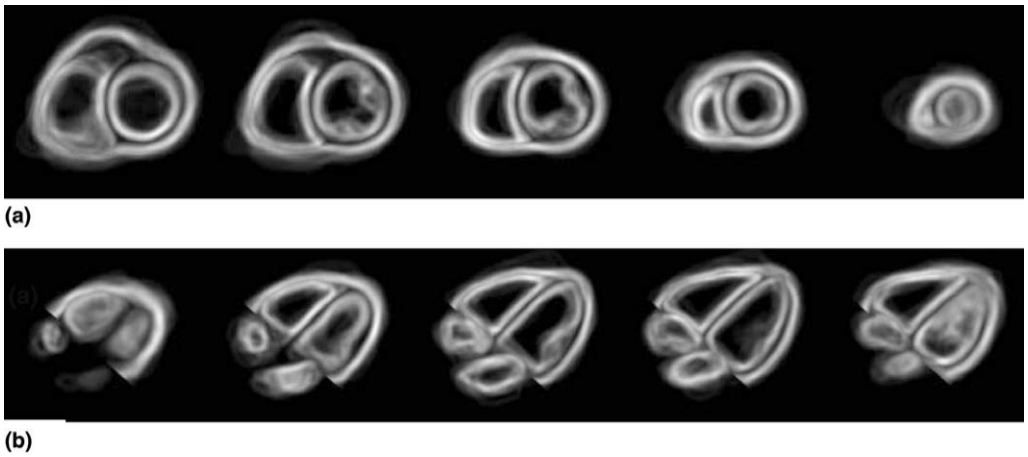

Fig. 16. (a) SA and (b) LA slices from the probabilistic atlas.

Table 1
Segmentation results

	LV	RV	LA	RA	EPI	ALL
NMI	2.32 ± 0.41	3.87 ± 1.31	3.83 ± 2.74	4.38 ± 1.99	3.61 ± 0.64	3.60 ± 1.42
LPD	2.23 ± 0.35	2.55 ± 0.55	3.00 ± 1.01	3.93 ± 1.60	2.78 ± 0.69	2.90 ± 0.84
Prob. atlas	2.01 ± 0.31	2.37 ± 0.50	2.56 ± 0.88	2.93 ± 1.30	2.77 ± 0.49	2.53 ± 0.70
PCA	2.40 ± 0.47	2.95 ± 0.84	3.78 ± 1.44	4.42 ± 3.84	3.02 ± 0.60	3.31 ± 1.17
ICA	2.47 ± 0.50	2.91 ± 0.68	3.87 ± 1.46	4.22 ± 2.41	3.11 ± 0.76	3.32 ± 0.89

The mean distance and its standard deviation ($N = 25$) from the deformed surface to the manually segmented surface in millimeters. The abbreviations used are left ventricle (LV), right ventricle (RV), left atrium (LA), right atrium (RA), epicardium (EPI) and average of all surfaces (ALL). The line NMI denotes results as no regularization was used, i.e. $\gamma = 0.0$ in Eq. (13).

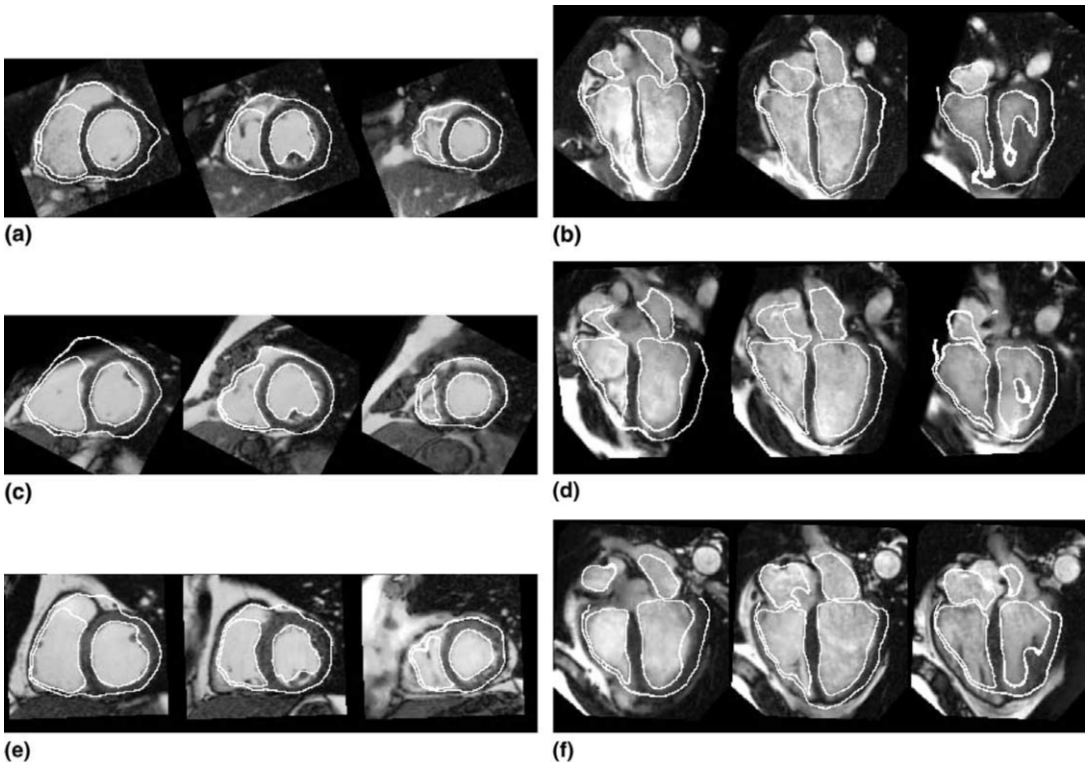

Fig. 17. Segmentation results from the cases having (a–b) the minimum (c–d) median and (e–f) maximum error value.

results, the mean error and the standard deviation for the right atrium using the probabilistic atlas would decrease to 2.71 and 0.70 mm, respectively. At least the following reasons can be found for the poor initialization in this case: (1) the older MR device with lower image quality was used, (2) only four LA images were available, and (3) the atria were not completely visible in the images.

The best results were achieved by regulating the transformation with the probabilistic atlas. The segmentation results from three cases have been superimposed on MR images in Fig. 17. The results are from the cases having the minimum (1.91 mm), median (2.41 mm) and maximum (3.26 mm, as the case with 4.05 mm error, described in the previous paragraph, was omitted) error values out of all database volumes. The segmentation accuracy of the right atrium is relatively poor (3.41 mm) in Fig. 17(d) because the superior part of the atrium was initialized about 20 mm (20 voxels) too low.

Non-parametric Wilcoxon Signed Ranks Test was used to detect statistically significant differences between the shape models. The difference to the non-regulated segmentation was significant for the probabilistic atlas ($p < 0.00001$) and for the LPD ($p < 0.001$). The probabilistic atlas was found to be superior to the LPD ($p < 0.001$), PCA ($p < 0.001$) and ICA ($p < 0.00001$). In addition, the LPD produced better accuracy than the PCA-based models ($p < 0.05$) and the ICA-based model ($p < 0.01$).

The manual segmentation is usually considered to be a golden-standard. For comparison, we manually segmented data from four subjects twice, and computed the errors between the surfaces. The average error was 1.75 mm as it was 2.57 mm for the same subjects using the probabilistic atlas-based segmentation.

4. Discussion

In this paper, we built a 3-D statistical shape model of the heart including atria, ventricles and epicardium. The new approach that combined information both from standard SA and LA cardiac images, made possible the modeling of the atria as well as the basal and apical regions of the ventricles from standard cardiac SA and LA images.

In this study, the datasets were acquired using a standard protocol for cardiac patients. The slice thickness could have been smaller to achieve a better spatial resolution and the SA images could have also contained the atria to obtain comprehensive information on atria. Although datasets were not optimal for the model construction, the advantage in using standard sets is that no specific imaging protocols are needed to increase the size of the database. All acquired datasets in the hospital could be added to the database. If a dataset contained volumes at different spatial resolutions, a high-resolution volume should be chosen to be the reference. Then, small structures, which were visible in the high-resolution images but not in the low-resolution ones, could be properly registered to other high-resolution volumes. Because the geometry of our shape model is relatively smooth, as the papillary muscles represent the finest details, increase in the resolution would not change the model considerably. However, increasing the resolution would improve the accuracy of the segmentation especially for the atria.

The shape variations were modeled in two ways: using PCA- and ICA-based shape models, and using probability distributions. A new approach of non-parametric probability distributions was proposed. The non-parametric distributions can model non-linear shape variation, which is not the case with the traditional PDM- and Gaussian models. Because of a relatively small database we could not demonstrate the advantages of the non-parametric approach to the approach using a Gaussian normal distribution. It remains still open if the generalized probability distributions would provide any advantages for the study of shape variability in SA and LA images. Nevertheless, as we extend the database and start to model the heart motion from cine data, we believe that the normal distribution is not longer general enough in modeling, for example, the twisting motion of the heart.

The shape models were evaluated in image segmentation. As the shape models were used to regulate the non-rigid registration, statistically significant improvements, compared with the non-regulated algorithm ($\gamma = 0$ in Eq. (13)), were detected for the probabilistic atlas and LPD. The best results, also in the statistical sense, were achieved with the probabilistic atlas. There are several possible reasons for this finding. The basic technical difference is that the shape variability is modeled from point distributions in PDMs and LPDs while surfaces are used in the probabilistic atlas. This means that no point-correspondences are required in the probabilistic atlas which removes one commonly known problem and error source. On the other hand, more points are used to estimate the probabilities at a specific location by using surfaces because the probability values of a point are affected also by its neighboring points (Parzen windowing was used). The results also indicated that the probability-based models performed better than the PDM-based models. Although the deformation spheres technique, used in the probabilistic atlas and LPDs, regulated the transformation according to the point distributions derived from the database, non-typical transformations were allowed if the NMI clearly improved. In PDMs, the transformations were strictly constrained to the modes of shape variation. As our database was relatively small, we believe that PDMs did not have enough degrees of freedom to model all shape

variability of the heart in 3-D. Another possible reason for the difference is that although the NMI (Eq. (12)) was used with all shape models, the different ways to deform the model and to optimize the transformation affected the result. In summary, final conclusions about the relative superiority of the shape models tested should not be made. We could conclude that the probabilistic atlas performs more robustly in segmentation than PDM- and LPD-models, if the database is relatively small. If the database is large enough, the situation might be even opposite.

The use of patient databases in quantitative analysis is still a relatively unexamined field. Hospitals and research sites have nowadays enormous sources of image data but the data are still mostly only visually inspected and rarely compared to other cases in the database in the clinical practice. A modern university hospital may acquire tens of terabytes of image data per year. These data definitely contain much information on the variability of normal and abnormal image features, relations between different spatial regions etc., which has not been yet understood or recognized. Some information can not even be visually detected from images, because of very small differences between healthy and specific pathologic populations. However, the differences can be statistically significant. We believe that statistical modeling of different populations will be a field of increasing importance in the future as automated methods will facilitate the modeling of the variability between humans on a large scale.

In this study, we did not apply the shape model in the comparison of different patient populations. However, the local and global shape variation of the cardiac structures and the wall thickness would be natural applications for the 3-D model. In future studies, we will extend our 3-D statistical shape model to a dynamic 4-D model by using the cine and tagged MR images available in the database. The 4-D information reveals many new application areas to the statistical shape model. However, before applying the shape model to population studies, several issues related to various error sources need to be considered. The construction of the shape model consisted of the following steps: movement correction, interpolation, manual segmentation, affine registration and non-rigid registration with manual fine-tuning. All these steps cause errors to the model and skew the definitions of a 'healthy' and 'pathologic' case in the population-based studies. The relative size of these error sources and the required accuracy in the model construction (population vs. population or even individual vs. population studies) should be carefully studied in the future.

Acknowledgements

This research was supported by the National Technology Agency (Finland), Finnish Cultural Foundation, Helsinki University Central Hospital Research Funds TYH 2242 and 4233, Finnish Radiology Society, PO Klingendahl Foundation, and Biomedicum Helsinki Foundation. MR imaging was performed by two excellent radiographers, Timo Päivärinta and Aki Syrjälä.

References

Alfakih, K., Thiele, S.P., Jones, T., Ridgway, J., Sivananthan, M., 2003. Normal human left and right ventricular dimensions for MRI as assesses by turbo gradient echo and steady-state free precession imaging sequences. JMRI 107, 323–329.

Caunce, A., Taylor, C.J., 2001. Building 3D sulcal models using local geometry. Med. Image Anal. 5 (1), 69–80.

Chen, M., Pomerleau, D., Schneider, J., 1999. 3-D deformable registration of medical images using a statistical atlas. In: Taylor, C., Colchester, A. (Eds.), Medical Image Computing and Computer-Assisted Intervention, MICCAI99. Lecture Notes in Computer Science, vol. 1679. Springer, Berlin, pp. 621–630.

Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J., 1995. Active shape models – their training and application. Computer Vision Image Understand. 61 (1), 38–59.

Cootes, T.F., Edwards, G.J., Taylor, C.J., 1998. Active appearance models. In: Burkhardt, H., Neumann, B. (Eds.), Proceedings of the European Conference on Computer Vision, vol. 2, pp. 484–498.

Duda, R.O., Hart, P.E., 1973. Pattern Classification and Scene Analysis. John Wiley & Sons, New York.

FastICA, 1998. Matlab package [online]. Available: http: http://www.cis.hut.fi/projects/ica/fastica.

Frangi, A., Niessen, W., Viergever, M., 2001. Three-dimensional modeling for functional analysis of cardiac images: A review. IEEE Trans. Med. Imag. 20 (1), 2–25.

Frangi, A., Rueckert, D., Schnabel, J., Niessen, W., 2002. Automatic construction of multiple-object three-dimensional statistical shape models: Applications to cardiac modeling. IEEE Trans. Med. Imag. 21 (9), 1151–1166.

Grevera, G.J., Udupa, J.K., 1996. Shape-based interpolation of multidimensional grey-level images. IEEE Trans. Med. Imag. 15 (6), 881–892.

Guimond, A., Thirion, J.-P., 2000. Average brain models. A convergence study. Computer Vision Image Understand. 77 (2), 192–210.

Hill, D., Hajnal, J., Rueckert, D., Smith, S., Hartkens, T., McLeish, K., 2002. A dynamic brain atlas. In: Dohi, D., Kikinis, R. (Eds.), Medical Image Computing and Computer-Assisted Intervention, MICCAI 2002. Lecture Notes in Computer Science, vol. 2488. Springer, Berlin, pp. 532–539.

Hyvärinen, A., Karhunen, J., Oja, E., 2001. Independent Component Analysis. John Wiley & Sons, New York.

Järvinen, V., Hekali, P., Poutanen, V.-P., Kupari, M., 1994a. Right atrial MR imaging studies of cadaveric atrial casts and comparison with right and left atrial volumes and function in healthy subjects. Radiology 191, 137–142.

Järvinen, V., Kupari, M., Hekali, P., Poutanen, V.-P., 1994b. Assessment of left atrial volumes and phasic function using cine magnetic resonance imaging in normal subjects. Am. J. Cardiol. 73, 1135–1138.

Koikkalainen, J., Lötjönen, J., 2004. Image segmentation with the combination of the PCA- and ICA-based shape models. In: IEEE International Symposium on Biomedical Imaging: From Nano to Macro, 2004. pp. 149–152.

Leemput, K.V., Maes, F., Vandermeulen, D., Suetens, P., 1999. Automated model-based tissue classification of MR images of the brain. IEEE Trans. Med. Imag. 18 (10), 897–908.

Lelieveldt, B., Üzümcü, M., van der Geest, R., Reiber, J., Sonka, M., 2003. Multi-view active appearance models for consistent segmentation of multiple standard views: Application to long and short-axis cardiac MR images. In: Computer As-sitested Radiology and Surgery (CARS), vol. 1256. pp. 1141–1146.

Lorenzo-Valdés, M., Sancheez-Ortiz, G., Mohiaddin, R., Rueckert, D., 2003. Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm. In: Ellis, R., Peters, T. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2003. Lecture Notes in Computer Science, vol. 2878. Springer, Berlin, pp. 440–450.

Lötjönen, J., 2003. Construction of patient-specific surface models from MR images: application to bioelectromagnetism. Computer Meth. Program. Biomed. 72 (2), 167–178.

Lötjönen, J., Mäkelä, T., 2001. Elastic matching using a deformation sphere. In: Niessen, W.J., Viergever, M.A. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001. Lecture Notes in Computer Science, vol. 2208. Springer, Berlin, pp. 541–548.

Lötjönen, J., Koikkalainen, J., Smutek, D., Kivistö, S., Lauerma, K., 2003a. Four-chamber statistical 3-D heart model from short-axis and long-axis MR images. In: Ellis, R., Peters, T. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2003. Lecture Notes in Computer Science, vol. 2878. Springer, Berlin, pp. 459–466.

Lötjönen, J., Smutek, D., Kivistö, S., Lauerma, K., 2003b. Tracking atria and ventricles simultaneously from cardiac short-and long-axis MR images. In: Ellis, R., Peters, T. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2003. Lecture Notes in Computer Science, vol. 2878. Springer, Berlin, pp. 467–474.

Lötjönen, J., Pollari, M., Kivistö, S., Lauerma, K., 2004. Correction of movement artifacts from 4-D cardiac short- and long-axis MR data. In: Lecture Notes in Computer Science: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004. Springer, Berlin, in press.

McLeish, K., Hill, D., Atkinson, D., Blackall, J., Razavi, R., 2002. A study of the motion and deformation of the heart due to respiration. IEEE Trans. Med. Imag. 21 (9), 1142–1150.

Mitchell, S., Bosch, J., Lelieveldt, B., van der Geest, R., Reiber, J., Sonka, M., 2002. 3-D active appearance models: Segmentation of cardiac MR and ultrasound images. IEEE Trans. Med. Imag. 21 (9), 1167–1178.

Moore, J., Drangova, M., Wierzbicki, M., Barron, J., Peters, T., 2003. A high resolution dynamic heart model based on averaged MRI data. In: Ellis, R., Peters, T. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2003. Lecture Notes in Computer Science, vol. 2878. Springer, Berlin, pp. 549–555.

Rohde, G., Aldroubi, A., Dawant, B., 2003. The adaptive bases algorithm for intensity-based nonrigid image registration. IEEE Trans. Med. Imag. 21 (11), 1470–1479.

Rueckert, D., Frangi, A., Schnabel, J., 2001. Automatic construction of 3D statistical deformation models using non-rigid registration. In: Niessen, W., Viergever, M. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001. Lecture Notes in Computer Science, vol. 2208. Springer, Berlin, pp. 77–84.

Rueckert, D., Frangi, A., Schnabel, J., 2003. Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration. IEEE Trans. Med. Imag. 22 (8), 1014–1025.

Studholme, C., Hill, D., Hawkes, D., 1999. An overlap invariant entropy measure of 3d medical image alignment. Pattern Recogn. 32 (1), 71–86.

Thompson, P.M., Toga, A.W., 1997. Detection, visualization and animation of abnormal anatomic structure with a deformable probabilistic brain atlas based on random vector field transformation. Med. Image Anal. 1 (4), 271–294.

Uzümcü, M., Frangi, A., Sonka, M., Reiber, J., Lelieveldt, B., 2003. ICA vs. PCA active appearance models: Application to cardiac MR segmentation. In: Ellis, R., Peters, T. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2003. Lecture Notes in Computer Science, vol. 2878. Springer, Berlin, pp. 451–458.