

Automatic Construction of Multiple-Object Three-Dimensional Statistical Shape Models: Application to Cardiac Modeling

Alejandro F. Frangi*, Daniel Rueckert, Julia A. Schnabel, and Wiro J. Niessen

Abstract—A novel method is introduced for the generation of landmarks for three-dimensional (3-D) shapes and the construction of the corresponding 3-D statistical shape models. Automatic landmarking of a set of manual segmentations from a class of shapes is achieved by 1) construction of an atlas of the class, 2) automatic extraction of the landmarks from the atlas, and 3) subsequent propagation of these landmarks to each example shape via a volumetric nonrigid registration technique using multiresolution B-spline deformations. This approach presents some advantages over previously published methods: it can treat multiple-part structures and requires less restrictive assumptions on the structure's topology. In this paper, we address the problem of building a 3-D statistical shape model of the left and right ventricle of the heart from 3-D magnetic resonance images. The average accuracy in landmark propagation is shown to be below 2.2 mm. This application demonstrates the robustness and accuracy of the method in the presence of large shape variability and multiple objects.

Index Terms—Atlas, cardiac models, model-based image analysis, nonrigid registration, statistical shape models.

I. INTRODUCTION

STATISTICAL models of shape variability or active shape models (ASMs) [1] have been successfully applied to perform segmentation and recognition tasks in two-dimensional (2-D) images. In building statistical models, a set of segmentations of the shape of interest is required, as well as a set of corresponding landmarks defined over the set of training shapes.

Manual segmentation and determination of point correspondences are time-consuming and tedious tasks. This is particularly true for three-dimensional (3-D) applications, where the

amount of image data to analyze and the amount of landmarks required to describe the shape increase dramatically in comparison to 2-D applications. This work aims at automating the landmarking procedure. That is, we still rely on the existence of a manual segmentation of the shapes, but the process of landmark extraction will be automated.

Several authors have proposed techniques to find point (landmark) correspondences, but only a few of them have indicated or investigated their applicability in the field of statistical shape models. Wang *et al.* [2] use a surface registration technique to find 3-D point correspondences based on a metric matching surface-to-surface distance, surface normals, and curvature. The authors suggest that this technique could be used to build 3-D ASMs, but they do not report any results on statistical model building. Kelemen *et al.* [3] report on the construction of 3-D ASMs of neuroradiological anatomical structures. In this method, the authors use a correspondence-by-parameterization approach to establish surface landmarks. The landmark correspondence is defined in the parameter domain of an underlying spherical harmonic parameterization. Although this approach has been used to build ASMs, no explicit volumetric or surface registration takes place. Lorenz and Krahnstover [4] show an improved method for building dense surface models that is similar to the one proposed in this paper. All the input meshes corresponding to several examples are warped onto the hand-placed landmarks of a single example; then a coating procedure is used to resample each surface to solve the correspondence problem. A mesh regularization step is included to ensure that folds in the surface introduced by the coating procedure do not appear in the final model. Unfortunately, this technique does not handle multiple-part objects as a whole; therefore, it cannot be guaranteed that it is free from inconsistencies arising from collisions between objects when treating the parts separately.

To our knowledge, only a few authors have addressed the problem of automatic construction of 3-D ASMs using nonrigid registration [5]–[9]. The frameworks proposed by Brett and Taylor [5], [6] are closely related to this work. In these approaches, each shape is first converted into a polyhedral representation. In the first approach [5], shape pairs are matched using a symmetric version of the iterative closest point (ICP) algorithm by Besl and McKay [10]. Using this method, the authors were able to build 3-D ASMs by automatically finding corresponding landmarks between surfaces. Surfaces are represented by dense triangulations that are matched to sparse triangulations (obtained by triangle decimation). The

Manuscript received February 22, 2001; revised July 6, 2002. The work of A. F. Frangi was supported by The Netherlands Ministry of Economic Affairs under Grant IOP Beeldverwerking IBV97009 and by EasyVision Advanced Development, Philips Medical Systems BV, Best, The Netherlands. He is currently supported by Spanish Ministry of Science and Technology under a Ramón y Cajal Research Fellowship, Grant TIC2002-04495-C02 and Grant FIT-070000-2002-935. The work of D. Rueckert was supported by the EPSRC under Grant GR/N24919. The work of J. A. Schnabel was supported by EasyVision Advanced Development, Philips Medical Systems BV, Best, The Netherlands. *As-terisk indicates corresponding author.*

*A. F. Frangi is with the Division of Biomedical Engineering, Aragon Institute of Engineering Research, University of Zaragoza, María de Luna 1, Centro Politécnico Superior, E-50018 Zaragoza, Spain (e-mail: afrangi@unizar.es).

D. Rueckert is with the Visual Information Processing Group, Department of Computing, Imperial College of Science, Technology and Medicine, SW7 2BZ London, U.K.

J. A. Schnabel is with the Computational Imaging Science Group, Division of Radiological Sciences, Guy's, King's, and St. Thomas' School of Medicine, Guy's Hospital, King's College London, SE1 9RT London, U.K.

W. J. Niessen is with the Image Sciences Institute, University Medical Center Utrecht, 3584 CX Utrecht, The Netherlands.

Digital Object Identifier 10.1109/TMI.2002.804426

nodes of the sparse triangulation become the final landmarks. One problem acknowledged by the authors is the possibility of obtaining shape models with surface folding due to some landmark groups (triples) matched in different order between training examples. This is a consequence of the use of the ICP technique, which is a purely local registration technique and does not incorporate any connectivity constraints. In the work by Brett and Taylor [6], this problem is addressed by transforming the surface to a planar domain using harmonic maps where connectivity constraints can be explicitly enforced. This technique avoids invalid cross-correspondences but is only applicable to single-part shapes that are topologically isomorphic to a disk. The work by Fleute and Lavallée [7], [8] is also closely related to our work. They use a multiresolution nonrigid registration technique based on octree-splines. This approach is a surface-based technique that registers shapes by minimization of a distance measure. In contrast to their work, we use a free-form volumetric nonrigid registration technique based on maximization of two novel similarity measures, coined *label consistency* and κ *statistic*. In addition, we provide experiments that give empirical evidence of the convergence of the atlas generation procedure, which is missing in [7], [8]. Finally, Subsol *et al.* [9] reported on a method for automatically constructing 3-D morphometric anatomical atlantes that is based on the extraction of line and point features and their subsequent nonrigid registration. This method is very attractive for applications where line features are evenly spread over the whole object, as is the case in the skull for which the authors illustrated their method. However, these features can be insufficient in modeling the heart since the cardiac chambers are rather smooth. Moreover, the extraction of line features is not trivial and may hamper the use of this technique in large-scale databases due to lack of robustness.

Recently, Davies *et al.* [11] published a method to automatically extract a set of optimal landmarks using the minimum description length (MDL) principle. Although this is a promising method that has a sound mathematical formulation, it is still required to elucidate whether the optimal landmarks in the MDL sense are also optimal in the sense of anatomical correspondence. On the other hand, although the authors sketch how to extend this method to 3-D, the extension does not seem trivial. Experiments with 3-D shapes still need to be performed to show its usefulness and computational feasibility, in particular with multiple-part objects.

In this paper, a technique is introduced that tries to address the shortcomings of point- or surface-based registration where no overall connectivity constraints are imposed. Our method introduces global constraints by matching shapes via a volumetric nonrigid registration technique using multiresolution B-spline deformation fields [12], [13]. Owing to the multiresolution nature of the deformation field, the mappings between matched shapes have been shown to be smooth [13]. Although in our experience this is not essential, an additional smoothness penalty can be incorporated, which further prevents folding when recovering the deformation field. An important feature of our approach is that the same method for establishing correspondences can be simultaneously applied to all the parts of a composite object. This has not been previously investigated in [5]–[8].

This paper makes two main contributions. The first contribution is methodological and lies in providing a generic framework for automatic selection of corresponding landmarks in 3-D shapes. The second contribution is of a more practical value and lies in applying this technique to the construction of 3-D shape models from 3-D cardiac magnetic resonance imaging (MRI) as well as in addressing a number of application-specific issues.

Model-based analysis of cardiac images is an active area of research—see e.g., a recent literature survey on 3-D model-based techniques for functional analysis of cardiac images [14]. Model-based approaches have been used for several applications ranging from computation of global cardiac functional parameters, mainly ejection fraction and left ventricular mass [15]–[21], to local estimation of cardiac motion [22]–[26] and stress [27]–[31]. Some authors have applied cardiac models to cardiac scan planning [32] or to derive new descriptors of cardiac function [26], [33]–[39]. A common denominator of these methods is that shape description is performed using standard modeling primitives such as, for instance, superquadrics [15], [26], [39], [40], series expansions [17], [40], [41], constructive solid geometry [32], or polyhedral representations [21], [23], [35], [42], [43]. Shape constraints are either explicitly enforced by the selected primitive (e.g., superquadrics) or via *ad hoc* regularization terms in the shape recovery process.

In this paper, the use of 3-D statistical shape models to describe the shape of the left and right ventricle of the heart is proposed. Although shape representation is based on standard triangulated meshes, the model contains information of the average location of each node in the mesh, its variability, and the correlation between node locations inferred from a set of training shapes. This information can be used to regularize the problem of model recovery or image segmentation without the need of any *ad hoc* smoothing constraints. However, this paper focuses on the automatic construction of statistical shape models and does not address the problem of shape recovery or segmentation, which is one of our future research issues.

This paper is organized as follows. Section II provides a summarized background on statistical shape models. Section III focuses on a multiresolution free-form registration algorithm that will be used in the landmarking procedure. In Section IV, our proposed method for automatic landmarking is described. In Section V, the method is applied to construct statistical models of the left and right ventricle of the heart from 3-D MRI; empirical evidence is given on convergence properties and reconstruction errors. Finally, Section VI closes the paper with some conclusions and directions for future research.

II. STATISTICAL SHAPE MODELS

Let $\{\mathbf{x}_i; i = 1 \dots n\}$ denote n shapes. Each shape consists of m 3-D landmarks, $\{\mathbf{p}_j = (p_{1j}, p_{2j}, p_{3j}); j = 1 \dots m\}$, that represent the nodes of a surface triangulation. Obtaining those m 3-D landmarks is a nontrivial task and presents the main topic of this paper. Each vector \mathbf{x}_i is of dimension $3m$ and consists of the landmarks $(p_{11}, p_{21}, p_{31}, p_{12}, p_{22}, p_{32}, \dots, p_{1m}, p_{2m}, p_{3m})$. Moreover, assume that the positions of the landmarks of all shapes are expressed in the same coordinate system. These vectors form

a distribution in a $3m$ -dimensional space. The goal is to approximate this distribution with a linear model of the form

$$\mathbf{x} = \hat{\mathbf{x}} + \Phi \mathbf{b} \quad (1)$$

where $\hat{\mathbf{x}} = (1/n) \sum_{i=1}^n \mathbf{x}_i$ is the average landmark vector, \mathbf{b} is the shape parameter vector of the model, and Φ is a matrix whose columns are obtained by performing a principal component analysis (PCA) of the covariance matrix $\mathbf{S} = (1/(n-1)) \sum_{i=1}^n (\mathbf{x}_i - \hat{\mathbf{x}})(\mathbf{x}_i - \hat{\mathbf{x}})^T$. The principal components of \mathbf{S} are calculated as its eigenvectors ϕ_i , with corresponding eigenvalues λ_i (sorted so that $\lambda_i \geq \lambda_{i+1}$). If Φ contains the $t < \min\{m, n\}$ eigenvectors corresponding to the largest nonzero eigenvalues, we can approximate any shape of the training set \mathbf{x} using (1), where $\Phi = (\phi_1 | \phi_2 | \dots | \phi_t)$ and \mathbf{b} is a t -dimensional vector given by $\mathbf{b} = \Phi^T(\mathbf{x} - \hat{\mathbf{x}})$.

The vector \mathbf{b} defines the shape parameters of the ASM. By varying these parameters, we can generate different instances of the shape class under analysis using (1). Assuming that the cloud of landmark vectors follows a multidimensional Gaussian distribution, the variance of the i th parameter b_i across the training set is given by λ_i . By applying limits to the variation of b_i , for instance, $|b_i| \leq \pm 3\sqrt{\lambda_i}$, it can be ensured that a generated shape is similar to those contained in the training class.

Although in this paper the statistical analysis is carried out using PCA, there are other alternatives that may be considered, such as independent component analysis [44].

III. MULTIREOLUTION NONRIGID REGISTRATION

In this section, a multiresolution free-form nonrigid registration algorithm is presented. This technique is employed in this paper to find correspondences between shapes. In contrast to most registration literature, our aim here is to match labeled rather than gray-level images, i.e., image segmentations where each voxel value indicates the type of structure this voxel belongs to.

The matching algorithm summarized here is a multiresolution free-form nonrigid registration algorithm, which was formulated by Rueckert *et al.* [12] and further developed by Schnabel *et al.* [13]. This algorithm manipulates a shape by embedding it into a subsequently refined volumetric mesh, which defines a continuous deformation field through a set of B-spline basis functions. For each location in the reference shape, the corresponding locations in the individual shapes are found to obtain an optimal match. The corresponding optimal deformation field is obtained by maximizing a voxel similarity measure on the basis of the corresponding labels. This maximization is carried out using a standard gradient ascent algorithm. The registration method and two novel similarity measures for labeled images are briefly summarized in the following paragraphs.

A. Transformation Model

Let $\mathbf{T} : (x, y, z) \mapsto (x', y', z')$ be a transformation that maps any point (x, y, z) in the source image into the corresponding target image coordinates (x', y', z') . The target image could be an *atlas* or average shape defining the coordinate system in which all other images will be expressed.

To accommodate for nonrigid deformations, \mathbf{T} will consist of a global transformation \mathbf{T}_g and a nonrigid transformation ob-

tained in a coarse-to-fine manner \mathbf{T}_e^H , where H is the number of mesh subdivisions that take place in the multiresolution strategy

$$\mathbf{T}(x, y, z) = \mathbf{T}_g(x, y, z) + \mathbf{T}_e^H(x, y, z). \quad (2)$$

B. Global Transformation

The global transformation describes the pose and size of the transformed shape with respect to the atlas. This can be accomplished with a global transformation in the form of a quasi-affine model

$$\mathbf{T}_a(x, y, z) = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} \\ \theta_{21} & \theta_{22} & \theta_{23} \\ \theta_{31} & \theta_{32} & \theta_{33} \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} \theta_{14} \\ \theta_{24} \\ \theta_{34} \end{pmatrix} \quad (3)$$

where the coefficients Θ parameterize the 12 degrees of freedom of a generic affine transformation. In a quasi-affine transformation only nine parameters are independent: translation $\{t_x, t_y, t_z\}$, rotation $\{r_x, r_y, r_z\}$, and anisotropic scaling $\{s_x, s_y, s_z\}$. The computation of these parameters is accomplished with a modified version of the method by Studholme *et al.* [45]. The main modification is that we do not use a registration measure based on mutual information but either *label consistency* or the κ *statistic* similarity measure, as described in Section III-D. However, we employ the same multiresolution strategy for computing the registration parameters.

C. Local Transformation

To accommodate for detailed shape differences, the global deformation field has to be supplemented with a local deformation model. The local deformation field is represented by a free-form deformation (FFD) based on B-splines. The basic concept of FFDs is to deform an object by embedding it into a volumetric mesh of control points and subsequently manipulating the nodes of the mesh. The mesh is deformed in a multiresolution fashion by subsequently subdividing the mesh (i.e., decreasing the mesh spacing). At any mesh resolution level, the control points are smoothly interpolated by a set of B-spline basis functions that define a continuous deformation field.

At each level of mesh resolution h , the FFD is represented by a tensor-product B-spline. Assume that we denote the domain of the image volume as $\Omega = \{(x, y, z) | 0 \leq x < X, 0 \leq y < Y, 0 \leq z < Z\}$. Let Φ^h denote an $n_x \times n_y \times n_z$ mesh of control points $\phi_{i,j,k}^h$ at level h and with uniform spacing $\delta_h = \delta_o/2^h$, where δ_o is the initial mesh spacing. Then, the FFD at level h can be written as the 3-D tensor product of the familiar one-dimensional cubic B-splines

$$\mathbf{T}_e^h(x, y, z) = \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 B_l(u) B_m(v) B_n(w) \phi_{i+l, j+m, k+n}^h \quad (4)$$

where $i = \lfloor x/n_x \rfloor - 1$, $j = \lfloor y/n_y \rfloor - 1$, $k = \lfloor z/n_z \rfloor - 1$, $u = x/n_x - \lfloor x/n_x \rfloor$, $v = y/n_y - \lfloor y/n_y \rfloor$, $w = z/n_z - \lfloor z/n_z \rfloor$, and where B_l represents the l th basis function of the B-spline [46].

The control points Φ^h act as parameters of the B-spline, and the degree of nonrigid deformation that can be modeled depends on the resolution level h . The smaller the associated spacing

δ_h , the more local the deformation is allowed to be. On the other hand, the computational complexity of the algorithm is approximately inversely proportional to the spacing parameter. The tradeoff between deformation flexibility and computational complexity is mainly an empirical choice, which is determined by the accuracy required to deform the atlas into the individual shapes. In our experiment with cardiac data sets, we have used a deformation field with three mesh subdivisions ($H = 3$) and an initial spacing $\delta_o = 20$ mm.

D. Similarity Measure for Labeled Image Registration

Since the correspondences of structures across both images are encoded explicitly in the labeling, we are only interested in maximizing the overlap of structures denoted by the same label in both images. Therefore, we have used two registration measures that favor the mapping of identical labels. We have coined these measures *label consistency* and κ *statistic*.

1) *Label Consistency*: Assume that $P_{AB}(i, j)$ is the joint probability of labels i and j in the source (A) and target (B) shapes, respectively. This can be estimated as the number of voxels with label i in image A and label j in image B divided by the total number of voxels in the overlap region of both images. The label consistency measure is then defined as

$$C_{LC}(A, B) = \sum_{i=1}^l P_{AB}(i, i). \quad (5)$$

Label consistency is a measure of how many labels of all the labels in the source shape are correctly mapped into the target shape. When the label consistency measure is zero, none of the source labels has been correctly mapped into the target shape. If all reference labels are correctly matched, the label consistency measure yields one.

2) *The κ Statistic*: The measure is inspired in a statistic used frequently in biomedical research to assess the agreement between two raters measuring the same quantity or performing a classification task [47]. For a given voxel, each image can be considered as an observer who assigns a class label to it. Therefore, comparing labeled images is equivalent to comparing the agreement between two observers. This statistic is defined as follows:

$$\kappa = \frac{p_a - p_c}{1 - p_c} \quad (6)$$

$$p_c = \sum_{i=1}^l P_A(i)P_B(i) \quad (7)$$

$$p_a = \sum_{i=1}^l P_{AB}(i, i) \quad (8)$$

where $P_{AB}(i, j)$, $P_A(i)$, and $P_B(i)$ are the joint probability density and marginal probability densities for the labels in the images A and B. The κ statistic is a measure of agreement between two classifications p_a that is corrected for chance agreement p_c . To interpret the values of the κ statistic, one can refer to standard tables provided in the literature [47]. This statistic has a maximum value of one, and values above 0.9 are usually regarded as very good agreements.

IV. AUTOMATIC LANDMARKING

A. Overview

Ideally, a landmark is an anatomically characteristic point that can be uniquely identified on a set of shapes. However, anatomical landmarks are usually too sparse to accurately describe a 3-D shape. Therefore, we will consider pseudo-landmarks, *i.e.*, landmarks lying on the shape's surface and determining its geometry.

Assume that n segmented shapes of the structure of interest are available. These segmentations constitute the training set $\mathcal{S}_n = \{L_i\}$, where $i = 1 \dots n$. Each shape in the training set is represented by a *labeled image* whose voxel values belong to a label set \mathcal{L} . To generate the landmarks for the n shapes, a shape atlas A is constructed and landmarked; these landmarks are finally propagated back to the n shapes (Fig. 1). The following sections describe these three steps in detail.

B. Preprocessing

The automatic landmarking algorithm can be applied to a set of previously segmented structures. In our application, three main structures of the heart have been segmented from a number of volunteer scans. Each segmentation consists of manual identification of the background (BG), the left ventricular blood pool (LV_{bp}) and myocardium (LV_{myo}), and right ventricular blood pool (RV_{bp}) (see Fig. 2). In the segmentation of the left ventricle, the papillary muscles are considered part of the blood pool, as is customary in functional cardiac analysis. Each segmentation is represented as a labeled image, where each voxel is labeled according to the tissue type to which it belongs ($\mathcal{L} = \{\text{BG, LV}_{\text{bp}}, \text{LV}_{\text{myo}}, \text{RV}_{\text{bp}}\}$).

Owing to the large voxel anisotropy in MR short-axis acquisitions of functional cardiac data sets, manual segmentations have significant staircase artefacts in the direction of the long axis of the heart. To facilitate image registration and to smooth out those artefacts, shape-based interpolation [48], [49] has been applied to obtain labeled images of isotropic voxel size (Fig. 3). To preserve the diameter of the ventricles in the basal segmentation slices, this slice was copied (repeated) again toward the base prior to computing the distance transform required to perform shape-based interpolation. Once the distance transform is linearly interpolated, it is thresholded to obtain a binary volume and subsequently masked to remove the extra slice. In this way, the segmentations before and after interpolation have the same basal diameter. In our opinion, this is a conservative way of respecting the ventricular dimensions in the absence of label information beyond the most basal slices.

C. Atlas Construction

Given a set of labeled images that are instances of an anatomical structure, a shape atlas can be viewed as an average labeled image representation of the shape under consideration. In the next section, a method to obtain an atlas for bivalued labeled images is introduced, which will be subsequently extended to labeled images containing multiple, and possibly nested, structures.

1) *Atlas Construction in Single Object Shapes*: Let us assume that a set of n training shapes \mathcal{S}_n is available and that each shape is represented by a bivalued labeled image B_i . That

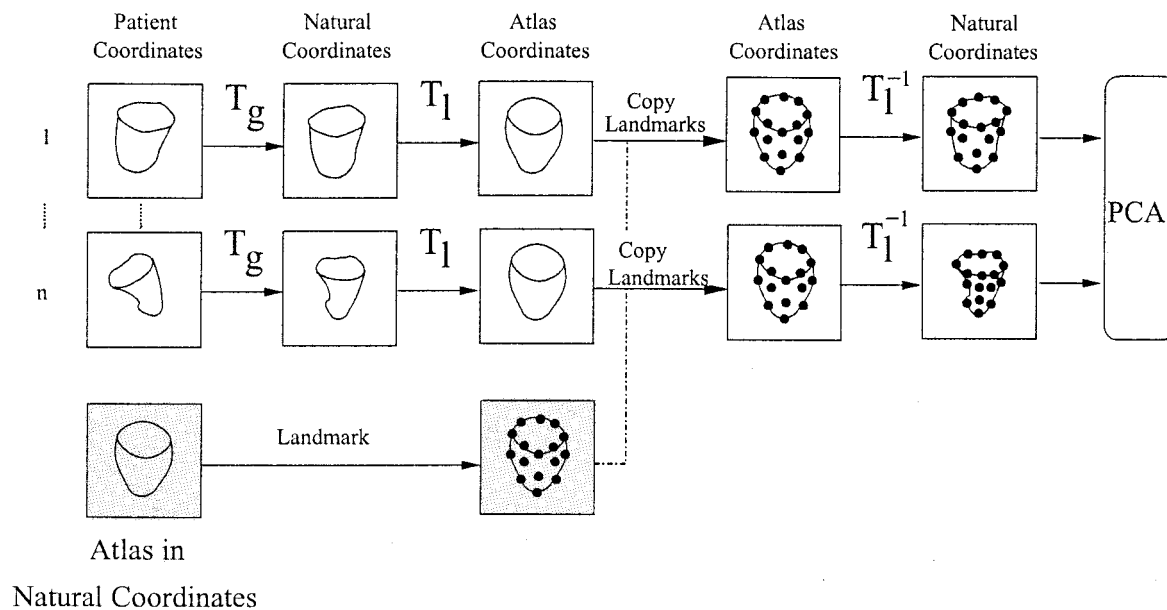


Fig. 1. Overview of the automatic landmarking framework. All individual data sets are matched to an atlas via a quasi-affine transformation (T_n) and a nonrigid transformation (T_l). The landmarks in the atlas can then be copied to the individual patients. The nonrigid deformation is subsequently reversed. Thus, PCA is carried out in a space where all shapes are aligned with the atlas (the *natural coordinate system*, explained in Section IV-C3). The principal modes of variation will therefore account for nonrigid deformations and not for pose or size differences.




Fig. 2. Example MR image and manual segmentation from the training set (short-axis and simulated long-axis view). The labeled images in (b) were manually extracted from 3-D cardiac MR scans in (a) and subsequently interpolated using shape-based methods. In each slice, the LV blood pool (dark gray), LV myocardium (light grey), and RV blood pool (white) were manually outlined. The papillary muscles were included in the segmentation of the blood pool, as is customary in clinical practice. The labeled image shown in (b) has been resampled using shape-based interpolation.

is, the label set contains only two labels corresponding to the object and background segmentations. For the sake of simplicity, we assume that S'_n is the set of shapes S_n after they have been aligned to a reference coordinate system.

Let B'_i and $DT(B'_i)$ denote the shape in atlas-aligned coordinates and its signed Euclidean distance transform [50], respectively. The signed distance transform of an average shape B_{av} can be obtained by computing $DT(B_{av}) = (1/n) \sum_{i=1}^n DT(B'_i)$. The average shape can be retrieved by thresholding the distance transform map to its zero-level set. We coin this averaging procedure *shape-based blending*.

As was mentioned earlier, prior to shape-based blending, all shapes have to be aligned into an atlas-aligned coordinate system. As the atlas is still to be constructed, the initial coordinate system can be chosen arbitrarily to coincide with that of any of the shapes in the training set S_n . To reduce the bias introduced by the selection of the initial reference shape, an iterative algorithm has been developed. In the first iteration, one shape of the training set is randomly selected to be the atlas. Subsequently, all other shapes in the set are registered to the current atlas using the global registration technique presented in Section III-B and the label-based similarity measures introduced in Section III-D. After registration, all shapes are blended and a new atlas is generated. This new atlas then becomes the current atlas, and the process is iterated until the difference between the current and new atlas falls below a certain threshold. This can be monitored, for instance, by defining a suitable measure of label agreement between two images. In the results section, the convergence properties of this algorithm and the influence of the randomly selected initial reference are investigated.

2) *Atlas Construction in Multiple-Part Shapes*: Let us assume that a set of n training shapes S_n is available and that each shape is represented by a labeled image L_i , in which l objects are represented by distinctive labels.

The atlas construction algorithm as described above applies to bivalued labeled images for which the distance transform is

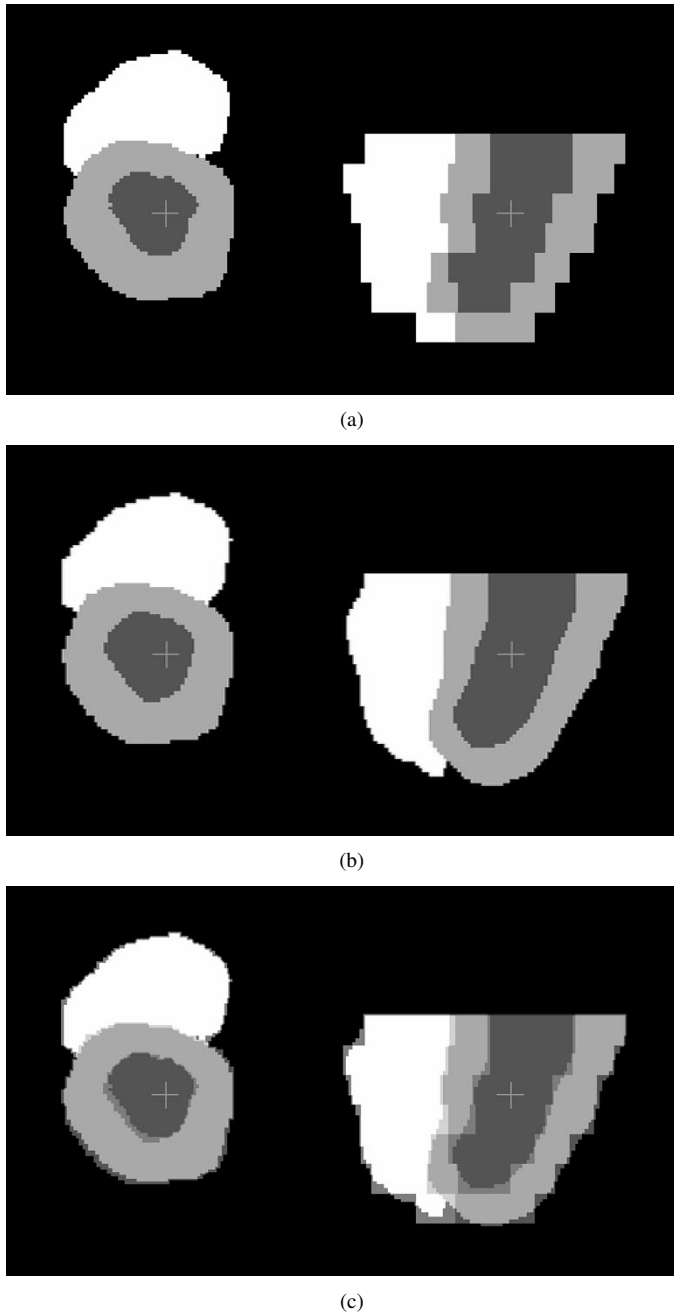


Fig. 3. Owing to voxel anisotropy, the labeled images were resampled to create isotropic geometric models. (a) shows an original labeled image with large anisotropy in the axial direction. (b) shows the corresponding shape-based interpolated image. (c) overlays the two previous images for comparison.

defined. However, it is possible to generalize the method to multiple-part shapes, which is needed for cardiac modeling. In fact, to extend the method of the previous section, we only need to specify a method for shape-based blending of multivalued images. The atlas alignment procedure used in this paper readily extends from single-object to multiple-object labeled images. In the following, we discuss our method for shape-based blending in cardiac labeled images.

Fig. 4 schematically represents a typical arrangement of LV_{bp} , LV_{myo} , and RV_{bp} as extracted manually from cardiac MR images. This figure also indicates how the cardiac shape can be decomposed into three subshapes: LV_{bp} , $LV_{bp} \cup LV_{myo}$,

and RV_{bp} . By construction, $LV_{bp} \cup LV_{myo}$ always embeds LV_{bp} , and both of them are nonoverlapping structures with respect to RV_{bp} . Each one of these subshapes can now be represented as a separate bivalued labeled image. This transforms the problem of shape-based blending of a multivalued labeled image into l problems of shape-based blending of bivalued labeled images. After the l average subshapes have been generated, all of them are combined into a new labeled image by taking the initial arrangement and labeling into account. An example of shape-based blending with two shapes consisting of three objects is shown in Fig. 5.

3) *Natural Coordinate System:* The atlas construction phase is iterative in order to reduce the bias toward the initial shape used as the initial “atlas.” However, the pose and size of the atlas are still biased toward those of this initial shape. This fact can have a negative influence on the landmark propagation, as we will explain in the following. An atlas whose pose is biased toward a specific subject may lead to a privileged situation of that subject in the nonrigid warping phase. To remove any remaining bias, we use a method similar to that presented in Rueckert *et al.* [51] to define a *natural coordinate system* for statistical deformation models.

Suppose that we have n deformation fields T_i . They map the atlas constructed with the iterative algorithm of Section IV-C2 into each individual in S_n . These deformation fields can be recovered with the nonrigid registration algorithm described in Section III.

We define now a natural coordinate system, which will be the coordinate system requiring the least residual nonrigid deformation to explain the anatomical variability across all individuals. Based on a point \mathbf{x} in the space of the reference subject, we can find the corresponding point \mathbf{x}' in its natural coordinates by applying the average deformation \hat{T} to \mathbf{x}

$$\mathbf{x}' = \mathbf{x} + \hat{T}(\mathbf{x}). \quad (9)$$

An interpretation of this natural coordinate system is shown in Fig. 6. Applying the average deformation vector \hat{T}_A , a point \mathbf{x}_A in the coordinate system of subject A should map the point \mathbf{x}_A into the point \mathbf{x}' in natural coordinates. Applying the average deformation vector \hat{T}_B to the corresponding point \mathbf{x}_B in the coordinate system of subject B should also map it to the point \mathbf{x}' in natural coordinates. Thus, assuming a perfect nonrigid registration that establishes a one-to-one correspondence between the anatomies of different subjects, the choice of the reference subject becomes irrelevant when constructing a natural coordinate system, since points \mathbf{x}' are independent of the choice of the reference subject.

In conclusion, before proceeding with the actual autolandmarking, the atlas is warped into its natural coordinate system, where there is no influence of the initial shape. This unbiased atlas is the one used as a template to extract landmarks and subsequently propagate them according to Fig. 1.

D. Landmark Extraction

After an atlas has been constructed and mapped into natural coordinates, it needs to be landmarked. In this paper, we shall




Fig. 4. Two-chamber model decomposition. To apply shape-based blending, the heart is decomposed into three bivalent labeled images representing, from left to right, the right ventricular blood pool (RV_{bp}), the left ventricular interior ($LV_{bp} \cup LV_{myo}$), and the left ventricular blood pool (LV_{bp}), respectively.




Fig. 5. Shape-based blending of label images. (d) illustrates the blending result of the two shapes of (a) and (b). (c) shows the two blended images mixed for comparison. All figures contain short-axis and simulated long-axis views.




Fig. 6. Natural coordinate system. In this space, the coordinates of the atlas are not biased toward the initial shape under the assumption of a perfect nonrigid registration.




Fig. 7. Atlas construction and landmarking. (a) shows the result of the iterative atlas construction algorithm as three orthogonal views. (b) visualizes a surface rendering of the atlas with the wire frame joining the extracted landmarks. The landmarks have been extracted using marching cubes with subsequent mesh decimation (95% decimation factor). The examples shown were generated using the κ statistic measure, although similar results were obtained for the label-consistency registration measure.

consider surface landmarks only. For cardiac modeling, landmarks for the endocardial and epicardial wall of the left ventricle and for the endocardial wall of the right ventricle can be extracted. To landmark the atlas automatically, the *marching cubes* [52] algorithm is applied. This algorithm generates a dense triangulation of the boundary isosurfaces that can be further decimated to reduce the amount of triangle nodes. The decimation process can be implemented in such a way that it preserves the accuracy of the surface representation in the original triangulation to a desired tolerance. We use the method by Schroeder *et al.* [53], which takes into account the curvature of the surface to keep the density of the triangles higher around relevant edges but decreases the number of triangles in flat areas. The nodes in the decimated triangulation form the landmarks of the shape. In our implementation of Schroeder's algorithm, we specify a target decimation rate (95%) subject to the constraint that the error of the decimated mesh is kept to a given tolerance with respect to the original mesh. The tolerance we used was the in-plane resolution of our images (1.2 mm).

By using marching cubes, a dense and approximately even distribution of landmarks is obtained. Alternatively, any other automatic algorithm for mesh extraction from a binary (labeled) image can be applied, such as *surface nets* [54] or the *wrapper algorithm* [55]. Note that, as an alternative to these algorithms, an expert could manually localize anatomical landmarks in the atlas. Anatomical landmarks, however, may be too sparse to represent the shape of 3-D structures and may also be difficult to accurately localize.

E. Landmark Propagation

Once the atlas is constructed, mapped to the natural coordinate system, and landmarked, its landmarks can be propagated to the individual shapes. This is carried out by warping each sample labeled volume into the atlas with a transformation $\mathbf{T} = \mathbf{T}_g + \mathbf{T}_e$ that is composed of a quasi-affine (\mathbf{T}_g) and a nonrigid (\mathbf{T}_e) transformation. The transformation \mathbf{T}_g accounts for global (pose and size) differences between the atlas and each sample volume, while the transformation \mathbf{T}_e accounts for local shape differences. The recovering of \mathbf{T} is carried out using the algorithms described in Section III.

Once the full transformation \mathbf{T} has been found, the landmarks of the atlas can be propagated to the natural coordinate system by applying the inverse of the nonrigid transformation (\mathbf{T}_e^{-1}). This process is repeated for each sample shape. As a result, a set of landmarks is obtained that describes shape variations with respect to the atlas. Since these landmarks are now in natural coordinates, pose and size variations are explicitly eliminated from further analysis. These transformed landmarks are subsequently used as the input for PCA, as indicated in Fig. 1.

Fig. 1 suggests that each sample shape is warped to the atlas. In this case, the inverse of the deformation field has to be computed to propagate the landmarks. However, this mapping does not necessarily exist. This was illustrated for the sake of conceptual simplicity only. From a computational point of view, it is more convenient to warp the atlas to each sample shape and use the direct deformation field for landmark propagation.

V. RESULTS

A. Data Sets and Preprocessing

Fourteen adult subjects, free of clinical cardiovascular disease, were scanned on a 1.5-T MR scanner (Philips ACS-NT, PowerTrak 6000 Gradient System, Philips Medical Systems, Best, The Netherlands) using an ECG-triggered Echo Planar Imaging (FFE-EPI) sequence. Cine acquisitions consisting of eight to ten short-axis slices of the heart in 18–20 phases of the cardiac cycle were performed. Scan parameters were repetition time: $TR = 632\text{--}857$ ms; echo time: $TE = 8.9\text{--}9.1$ ms; flip angle: $\alpha = 20^\circ$; slice thickness: 10 mm; a 256×256 image matrix; and a $300 \times 300\text{-mm}^2$ field-of-view.

From the acquired temporal sequence of each volunteer, the end-diastolic frame was manually segmented, as indicated in Section IV-B. Subsequently, all segmentations were resampled to isotropic voxels with size equal to the in-plane resolution (1.2 mm) using shape-based interpolation [48], [49], as described in Section IV-B.

B. Atlas Construction

Fig. 7 shows the result of the atlas building process described in Section IV-C. The shape-based blending procedure captures

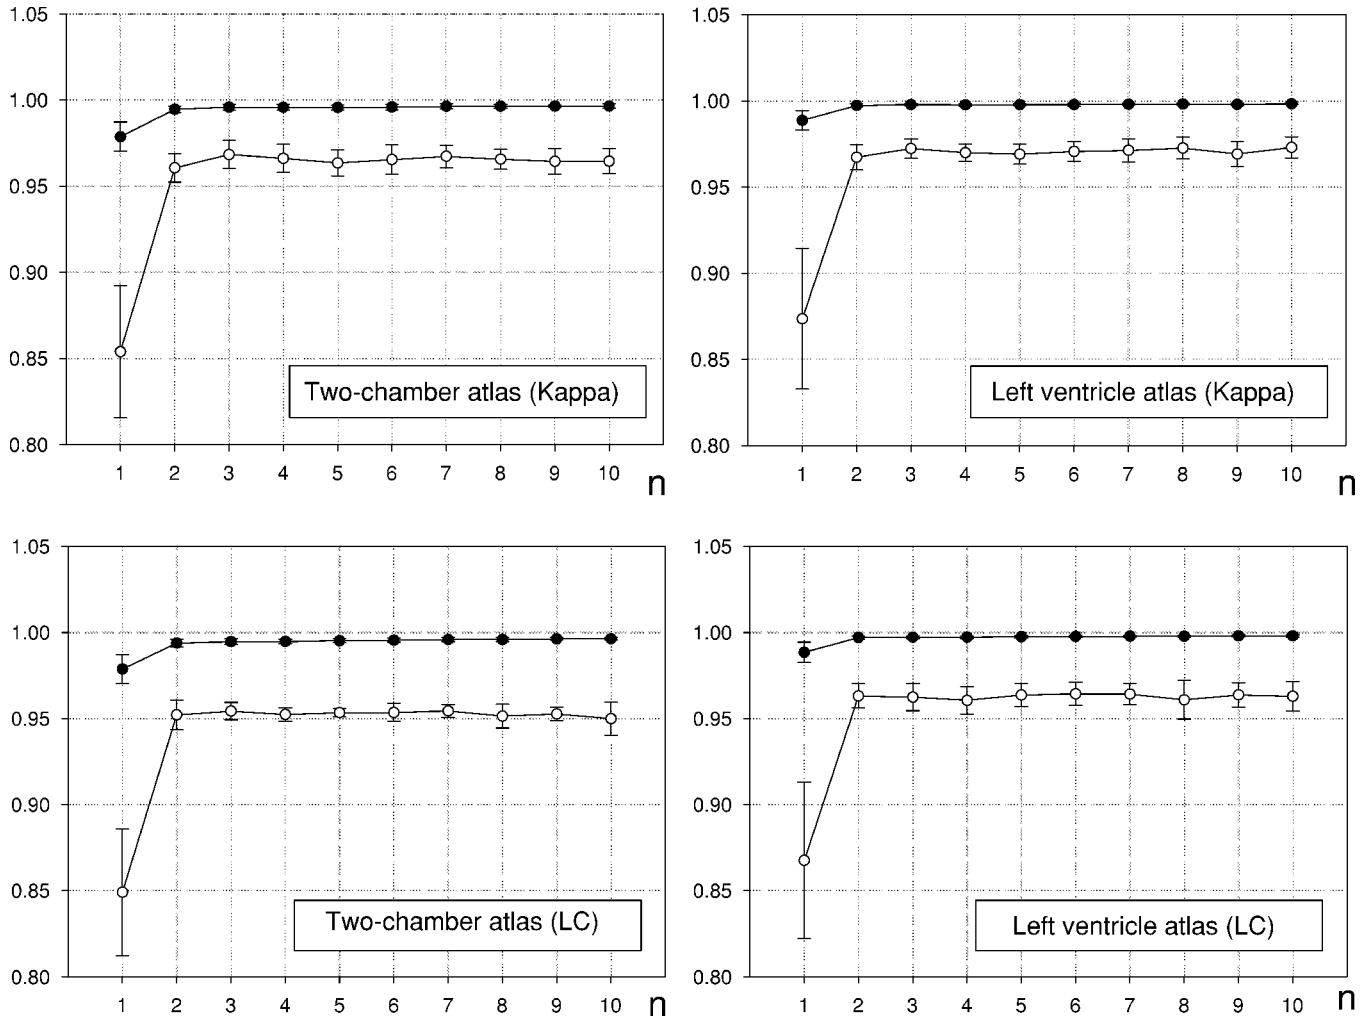


Fig. 8. Convergence of the atlas construction algorithm. κ statistic and the *label consistency* similarity measures between two consecutive atlases as a function of the iteration number n . In the inset of each plot, the similarity measure is indicated. The graph denoted with \bullet corresponds to the evolution of the *label consistency* measure, while the graph with \circ corresponds to the evolution of the κ statistic. Each curve is the average of the convergence curve resulting from initializing the atlas building procedure with one of the 14 shapes.

the global shape of the ventricles without an apparent bias in the wall thickness at any particular sector.

The convergence of the iterative atlas construction and the effect of different initial shapes in the generation of the atlas have also been studied. Fig. 8 plots the κ statistic between two successive iterations of the algorithm of Section IV-C. The same curves were generated for the 14 atlases obtained by using each of the 14 subjects as initial shapes. For each atlas and both registration measures used in the atlas building procedure, we show the evolution of both metrics as a function of the iteration number. After two iterations, the agreement is excellent ($\kappa > 0.94$) and the label consistency very high ($LC > 0.99$). It is notable that both metrics seem to have a different sensitivity to matching, as their dynamic range is different. The plots also indicate that the convergence is not monotonous but tends to level out after a few iterations. This could be caused by marginal changes in the registration measure due to small mismatches at boundaries when working with labeled images. In fact, there is hardly any visual difference between atlases in successive iterations after the third iteration. Given that this process is reasonably fast and to avoid any residual bias, we used the atlas of the fifth iteration in this study.

Another way of assessing the influence of the initial shape is to create atlases with all possible initial shapes, map all of them into natural coordinates, and compute the corresponding shape-based average atlas. Subsequently, the mean, standard deviation, and maximum distance between the boundaries of each atlas and those of the average atlas can be computed and averaged (Fig. 9). Table I shows the average mean, standard deviation, and maximum error for each part of the model. The average maximum error is always below 4 mm and, therefore, is on the order of the voxel dimensions ($1.2 \times 1.2 \times 10 \text{ mm}^3$).

C. Statistical Shape Models

To construct a statistical model from the cardiac atlas, landmarks were extracted automatically from the atlas using marching cubes and subsequent mesh decimation (95% decimation ratio). This procedure yielded 1352 (Kappa) and 1304 (LC) landmarks for the left ventricular epicardial surface, 679 (Kappa) and 631 (LC) landmarks for the endocardial surface, and 1841 (Kappa) and 1693 (LC) landmarks for the right ventricular endocardial surface.

This table reports on the errors of a left ventricle model and a combined left and right ventricle model, respectively.




Fig. 9. Contours in three planes of the atlantes created starting with all possible initial shapes after mapping into natural coordinates overlaid on the atlas image. There is very little difference between shapes, as indicated by the close agreement between contours.

In the former, only the left ventricle segments (LV_{bp} and LV_{myo}) were taken into account in the nonrigid registration and landmark propagation procedures. In the latter, both the left ventricle (LV_{bp} and LV_{myo}) and the right ventricle (RV_{bp}) were nonrigidly registered to the atlas, and the landmarks were subsequently propagated.

Figs. 10 and 11 show the first five modes of variation of the two-chamber model constructed using the kappa statistic and label consistency as registration measures, respectively. It is re-

TABLE I
AVERAGE MEAN, STANDARD DEVIATION, AND MAXIMUM DISTANCES BETWEEN THE ATLANTES CREATED STARTING WITH ALL POSSIBLE INITIAL SHAPES AND THE AVERAGE ATLAS OF ALL ATLANTES IN NATURAL COORDINATES

Model	Measure	LV endocardium			LV epicardium			RV endocardium		
		Mean [mm]	SD [mm]	Max [mm]	Mean [mm]	SD [mm]	Max [mm]	Mean [mm]	SD [mm]	Max [mm]
Left Ventricle	Kappa	0.21	0.84	2.89	0.21	0.82	2.50	-	-	-
Left Ventricle	LC	0.17	0.87	2.75	0.20	0.83	2.41	-	-	-
Two-chamber	Kappa	0.15	0.89	2.69	0.16	0.84	2.41	0.29	0.94	3.60
Two-chamber	LC	0.24	0.89	2.93	0.21	0.84	2.52	0.34	0.93	3.08

markable that there is hardly any difference in the modes. This seems to support the idea of a high independency of the model construction procedure with respect to the matching metric.

Fig. 12 indicates the percentage of the total explained variance as a function of the number of modes. Once corresponding landmarks are available, one is free to include or exclude the landmarks corresponding to a certain substructure. The more substructures that are incorporated, the larger the required number of modes to explain a given variance since the overall shape variability has been increased.

To quantitatively assess the performance of the constructed models, we have analyzed the reconstruction error by performing several leave-one-out experiments. The landmarks of all but one data set were used to build a statistical model. This model was subsequently used to reconstruct the set of landmarks not included in the PCA. The same experiment was repeated by taking out from the PCA, one in turn, each of the sets of landmarks. Finally, the average reconstruction error over the leave-one-out experiments was computed. Fig. 13 shows the mean square reconstruction error as a function of the number of modes used in shape reconstruction. Since our training set is relatively small and the shape variability is quite large, these experiments do not reveal much information on the generalization ability of the models. However, they provide a first estimate that could be refined by enlarging the database of shapes.

D. Automatic Landmark Propagation Performance

To quantify the ability of the technique to map corresponding landmarks, we have selected a set of seven landmarks in each subject and in the atlas built thereof. They are indicated in Fig. 14(a). Three observers were asked to identify these points twice for each subject in two independent sessions. The same point set was pinpointed twice in the atlas. Subsequently, the set of landmarks in the atlas was mapped into each subject using the deformation fields that were computed in Section IV-E, and a number of performance measures were computed [Fig. 14(b)]. We have calculated the intraobserver variability of the 3-D position of the landmarks placed manually in subject space. Also, we computed the intraobserver variability, measured in subject space, corresponding to the landmarks of the atlas after propagation through the deformation field. For each landmark, we calculated the average intraobserver variability over all subjects. Table II reports the mean intraobserver variability of the three observers. The variability in the automatic landmarks indicates how the uncertainty in manual placement in the atlas—of a similar magnitude to that of manual placement in subject space—propagates through the deformation fields as an uncertainty in the location of automatic landmarks.

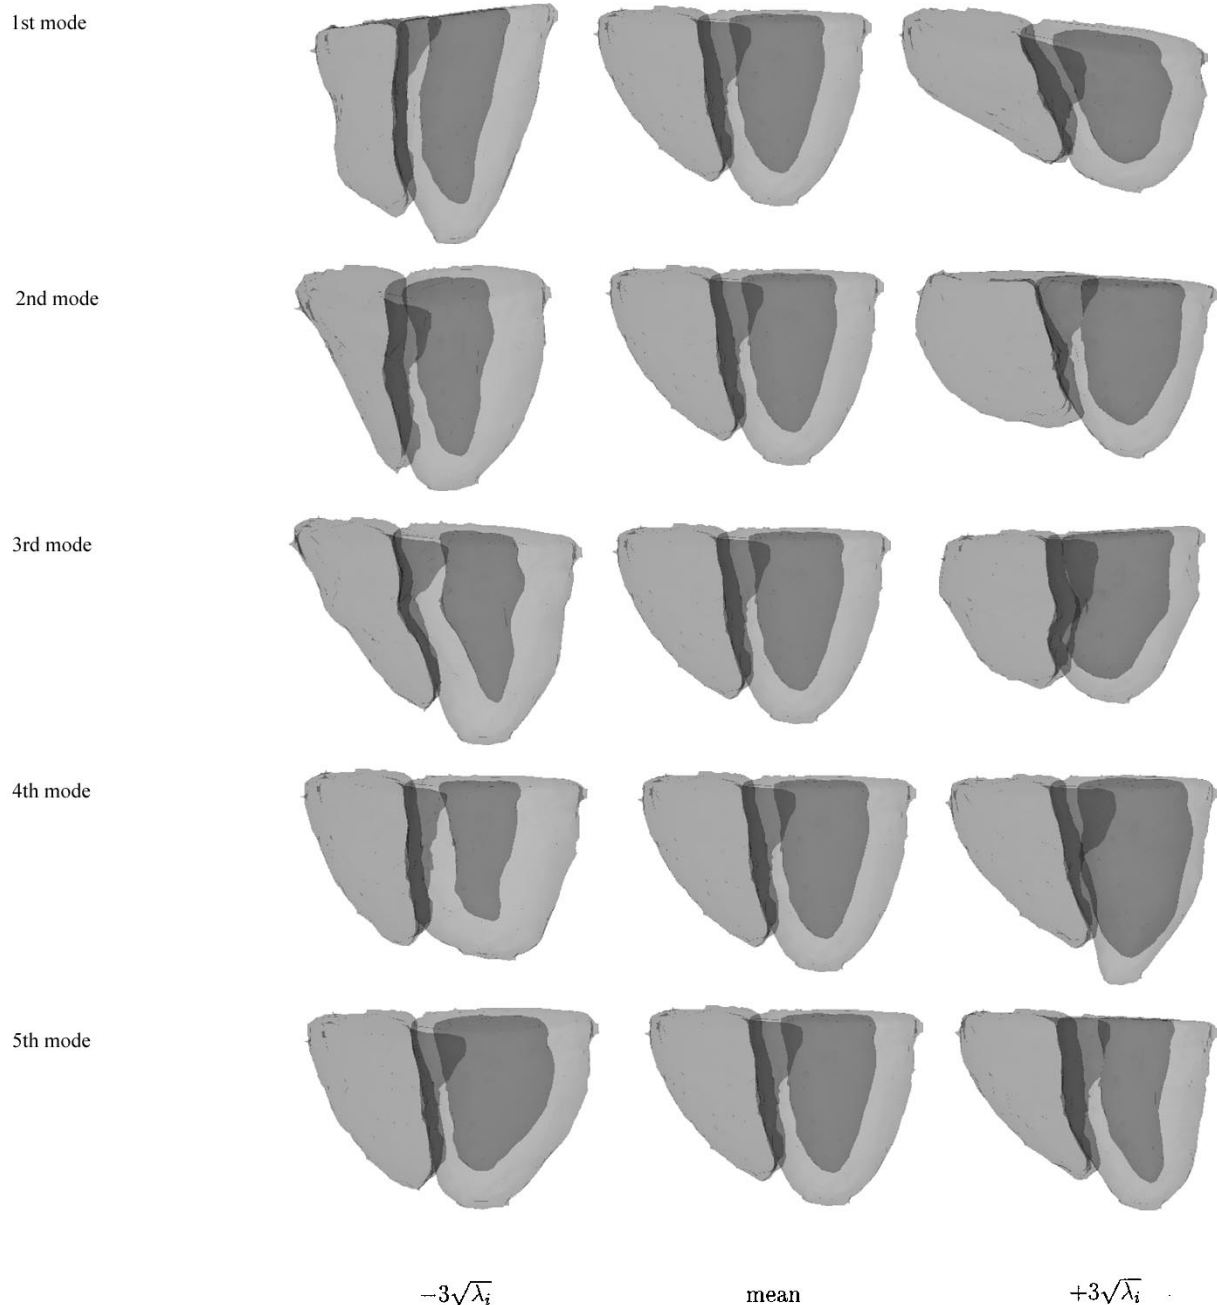


Fig. 10. Shape instances generated using the 3-D two-chamber model and the κ statistic registration measure from 14 cardiac data sets. The instances are generated by varying a single shape parameter, fixing all others at zero standard deviations from the mean shape. This two-chamber model consists of 679 LV endocardial nodes, 1352 LV epicardial nodes, and 1841 RV endocardial nodes.

Table III reports on the interobserver variability for manual landmarking in subject space and that of the landmarks manually placed in the atlas and then propagated into subject space. The latter variability measures the effective uncertainty in the position of the automatic landmarks when considering that they have been obtained from the propagation of already uncertain (manual) landmarks.

Finally, it is interesting to compute an estimate of the average distance (error) between the landmarks placed manually in subject space and those propagated to subject space from the atlas. To this end, the average position of the landmarks of all observers and all sessions has been used as the “gold standard” [textured bullet in Fig. 14(b)]. Table IV gives the average dis-

tance in subject space between the gold-standard landmarks in subject space and in atlas space after being mapped into subject coordinates. From these experiments, it is possible to state that the error in the automatic landmark placement is, on average, ± 2.2 mm. This is about two pixels with the current in-place image resolution and only moderately worse than the interobserver variability in manual landmarking.

VI. DISCUSSION AND CONCLUSION

This paper has presented a method for the automatic construction of 3-D statistical shape models. The technique is based on the automatic extraction of a dense mesh of landmarks in an

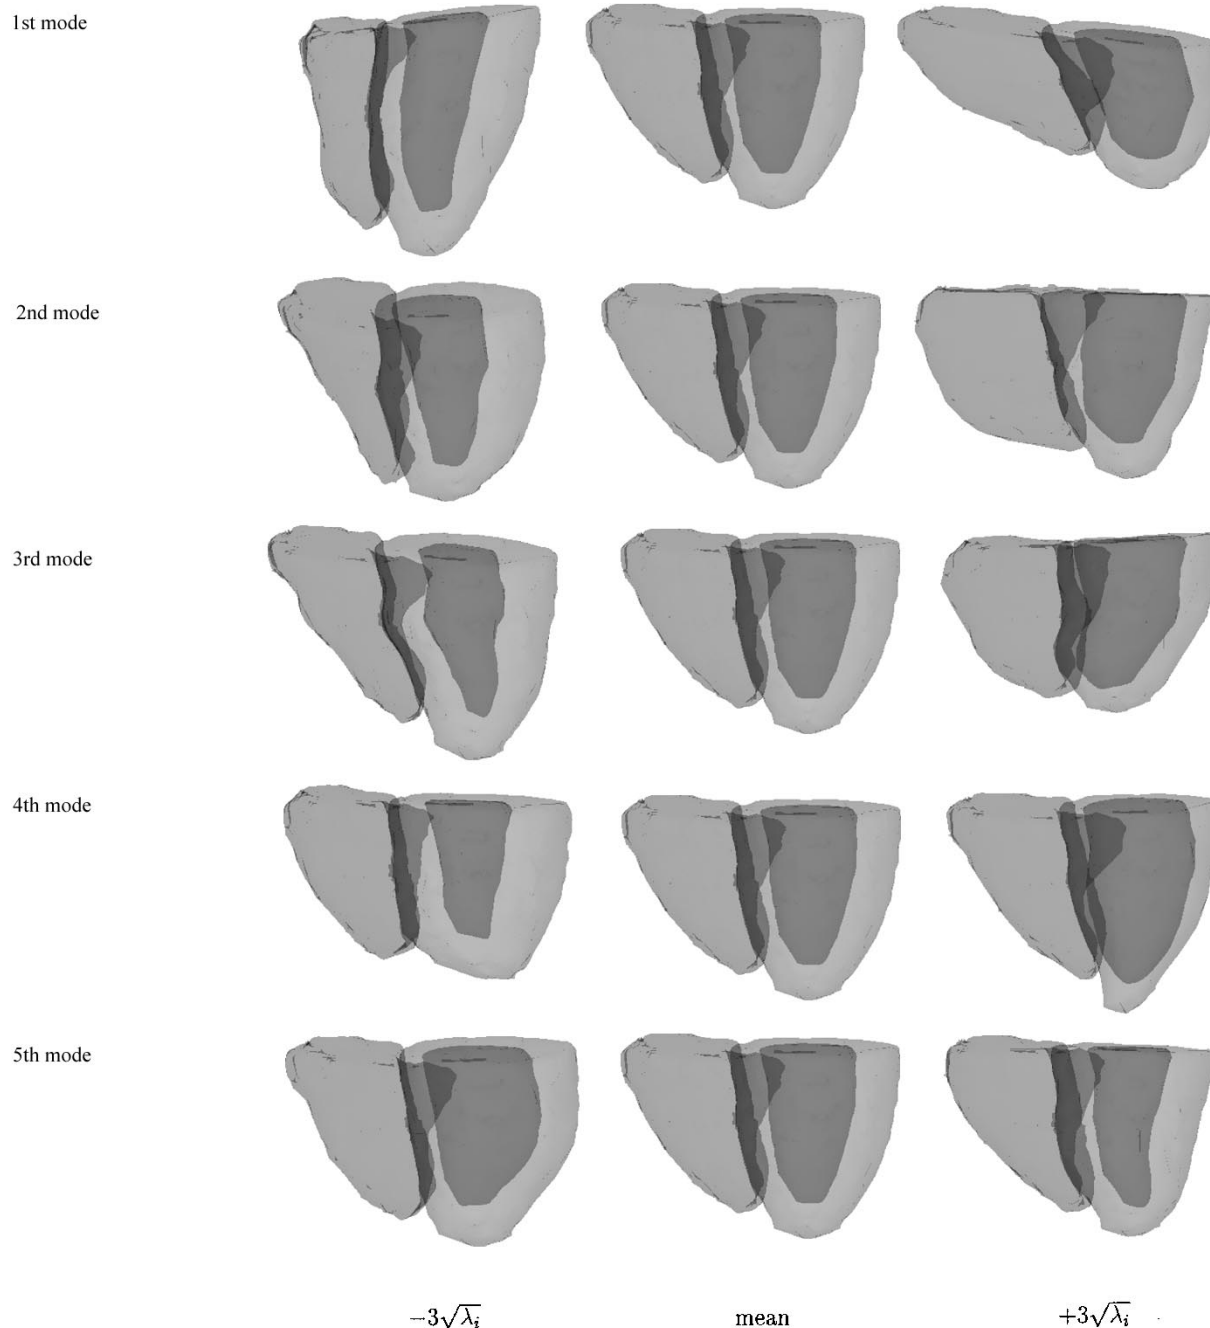


Fig. 11. Shape instances generated using the 3-D two-chamber model and the label-consistency registration measure from 14 cardiac data sets of the heart. The instances are generated by varying a single shape parameter, fixing all others at zero standard deviations from the mean shape. This two-chamber model consists of 631 LV endocardial nodes, 1304 LV epicardial nodes, and 1693 RV endocardial nodes.

atlas constructed from a set of training shapes. These landmarks are subsequently warped by a nonrigid deformation field to each shape in the training set. The method is able to treat single- and multiple-part shapes.

The first part of the proposed technique involves the building of an atlas from a set of example shapes. In Section V, we have presented experimental results supporting the hypothesis that this procedure is convergent. Moreover, different initial shapes seem to bias only marginally the final atlas when it is mapped into natural coordinates. Therefore, for practical purposes, the procedure of atlas construction can be considered to yield a unique solution. In the work by Fleute and Lavallée [7], [8],

a similar algorithm was used for building the average model (atlas). However, the atlas construction was performed on a surface representation of the shapes and required a high-resolution reference shape to initialize the iterative procedure. Finally, no experimental evidence was reported with respect to the convergence of the atlas construction algorithm nor to the selection of the initial shape (uniqueness).

An alternative to our iterative method of atlas construction is the tree-based approach presented by Brett and Taylor [5]. This hierarchical strategy is attractive since it gives a unique (noniterative) way to build an atlas from a given set of examples. However, one problem of Brett's method is that the training shapes

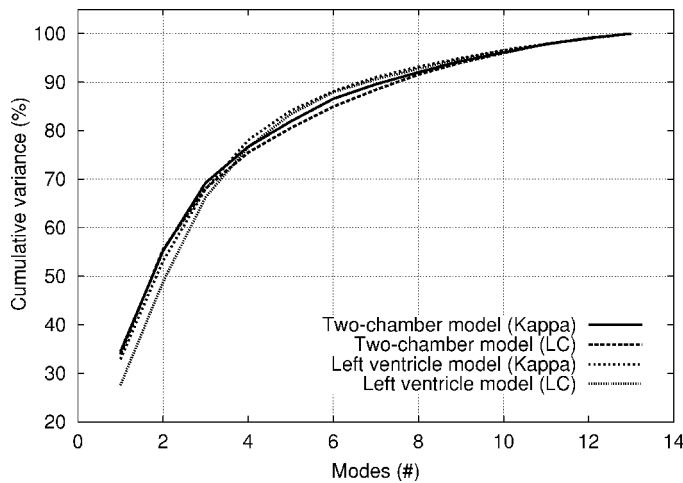
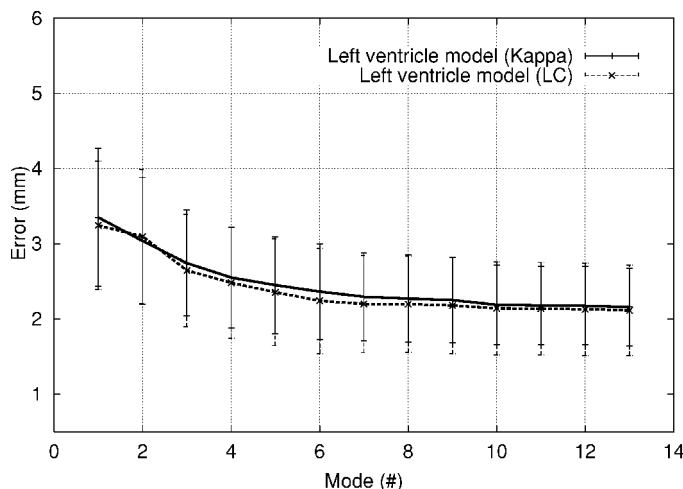
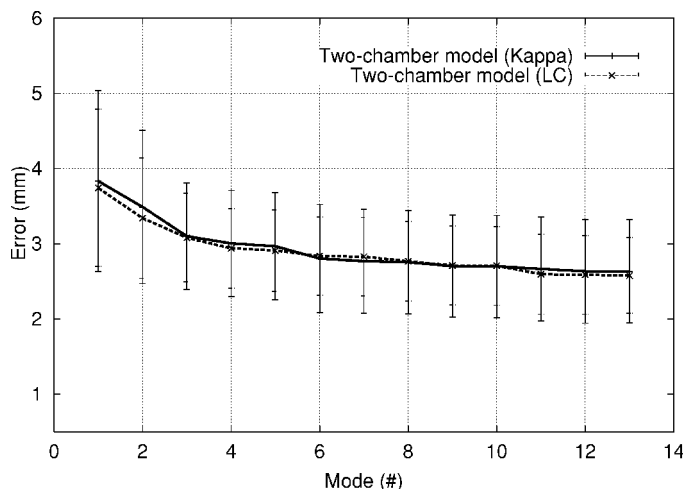


Fig. 12. Percentage of total shape variance versus the number of modes used in the 3-D ASM. The total number of landmarks was 3872 (kappa) and 3628 (LC) landmarks for the two-chamber heart model and 2550 (kappa) and 2324 (LC) landmarks for the left ventricular model, respectively. Note that as the two-chamber and LV models were constructed (landmarked) independently, marching cubes does not yield exactly the same number of nodes for a given substructure (e.g., LV epicardial surface) in both models.



(a)



(b)

Fig. 13. Reconstruction error in the leave-one-out experiments for the left ventricle and two-chamber models and different registration measures used in model building.

have to be ranked according to a pairwise match quality. This requires that all possible pairs have to be matched and scored before the tree is built. Brett presented results with only eight shapes [5], but ordering the examples according to the matching quality would be cumbersome for a more realistic amount of training shapes. For a total number of n shapes, it is necessary to compute $N = (n - 1)^2 \approx \mathcal{O}(n^2)$ pairwise matches to build the average shape. Our approach obtains the average shape in $N = nI \approx \mathcal{O}(n)$ matches, where I is the total number of iterations required for convergence. Section V provides experimental evidence that after about five iterations, the atlas shape stabilizes.

Our method for building the mean shape model is based on averaging shapes in the domain of their distance transforms. A similar strategy was proposed by Leventon *et al.* [56] to incorporate statistical constraints into the level-set approach to image segmentation. However, in that work, PCA is applied on the distance transform domain and not on a surface representation. As a consequence, the number of degrees of freedom is considerably larger than in our method. There is an intrinsic limitation in both our method and that of Leventon *et al.* Averaging distance transforms of several shapes does not necessarily yield a valid mean shape representation. It is easy to show, for instance, that in the case of a large misalignment between the averaged shapes, this procedure can introduce topological changes. Although we did not observe this problem in our experiments, this can be a potential source of failure of the technique when building models of very complex structures.

The proposed technique could be used with any nonrigid registration algorithm. In this sense, the method is a generic framework open to future investigation. Currently, a multiresolution version of the FFD nonrigid registration of Rueckert *et al.* [12] is used to match labeled images. Two novel registration measures suitable for labeled images were used. The atlantes and models built with both measures were hardly distinguishable from each other, yielding, on average, a similar performance in terms of landmark propagation accuracy. We believe that both metrics are equally useful for automatic landmarking, although label consistency is a computationally slightly more efficient measure.

The use of nonrigid registration as a method to establish shape correspondences imposes a constraint on the type of shapes that can be handled. It is assumed that the class of shapes has a well-defined topology. If there are substructures in one image not represented in the other image to be matched, the transformation would have to “destroy” those parts. This situation could arise when building a model of normal and abnormal medical structures where some parts in the latter are missing because of a diseased state or surgical procedure. However, establishing correspondences in these mixed models also remains an ill-defined problem with any of the previously published approaches [5]–[9].

Results of the construction of a one- and two-chamber cardiac model have been presented. Experiments were carried out to establish the ability of the models to generalize to shapes not present in the training set. The average reconstruction error in the two-chamber model was below 3.8 mm when the number of nodes was sufficient to explain 90% of the shape variability.

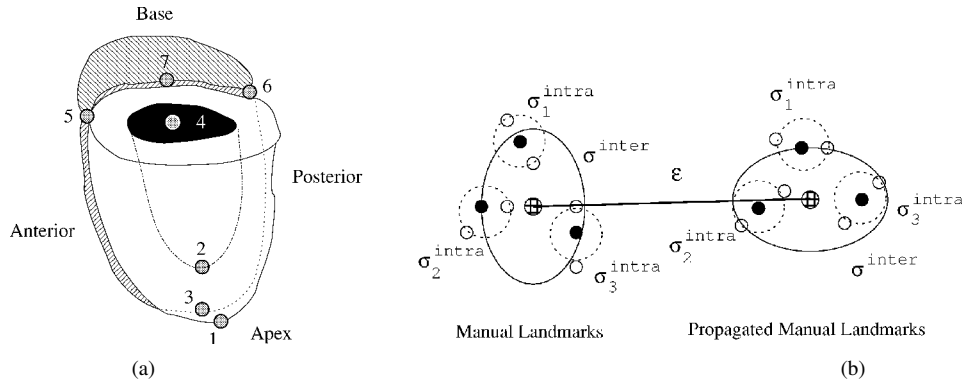


Fig. 14. Autolandmarking performance evaluation. (a) Seven landmarks were identified at the basal and apex level: LV apex epicardial (Lmk # 1), LV apex endocardial (Lmk # 2), RV apex epicardial (Lmk # 3), LV basal endocardial centroid (Lmk # 4), anterior basal septum (Lmk # 5), posterior basal RV-septum (Lmk # 6), middle basal RV-septum (Lmk # 7). (b) Definition of the computed performance measures for one of the landmarks: landmarks marked with \circ are the original manual landmarks of a subject (left) or the manual atlas landmarks propagated to subject space (right); the points marked with \bullet are the average location of two sessions for a given observer (one to three). Finally, the textured bullet is the average landmark location over all three observers. Here σ_i^{intra} and σ_i^{inter} are the intra- and interobserver variability for observer i , respectively, and ϵ is the error between the centroids of the manual and propagated landmark clouds.

TABLE II
INTRA-OBSERVER VARIABILITY DUE TO MANUAL AND
AUTOMATIC LANDMARKING

Lmk #	Manual [mm]			Automatic κ [mm]			Automatic LC [mm]		
	σ_x^{intra}	σ_y^{intra}	σ_z^{intra}	σ_x^{intra}	σ_y^{intra}	σ_z^{intra}	σ_x^{intra}	σ_y^{intra}	σ_z^{intra}
1	0.755	0.758	0.330	0.692	0.308	1.528	0.124	0.144	1.872
2	0.560	0.371	0.371	0.466	0.561	1.560	0.773	0.639	1.254
3	0.659	0.588	0.433	1.039	0.374	1.760	0.834	1.193	2.047
4	0.297	0.399	0.347	0.868	0.430	0.834	0.884	0.976	0.850
5	0.418	0.377	0.356	1.532	0.826	1.523	1.798	0.641	2.178
6	0.403	0.335	0.356	1.490	2.397	1.947	1.244	1.385	2.125
7	0.804	0.411	0.356	1.499	0.485	1.931	1.539	0.854	2.170
Average	0.556	0.463	0.364	1.084	0.769	1.583	1.028	0.833	1.785

TABLE III
INTER-OBSERVER VARIABILITY DUE TO MANUAL AND
AUTOMATIC LANDMARKING

Lmk #	Manual [mm]			Automatic κ [mm]			Automatic LC [mm]		
	σ_x^{inter}	σ_y^{inter}	σ_z^{inter}	σ_x^{inter}	σ_y^{inter}	σ_z^{inter}	σ_x^{inter}	σ_y^{inter}	σ_z^{inter}
1	1.414	0.681	0.663	0.262	0.367	1.761	0.183	0.760	2.448
2	0.755	0.548	0.616	0.478	0.235	1.019	0.372	0.408	1.872
3	1.031	1.593	1.127	3.440	0.713	1.530	1.589	0.702	3.091
4	0.484	0.623	0.319	0.668	0.853	0.713	0.858	0.373	0.594
5	0.888	0.530	0.319	2.300	1.340	0.739	2.543	0.797	2.066
6	0.395	0.518	0.319	1.413	1.655	1.117	1.086	1.191	1.967
7	1.391	0.696	0.319	1.354	0.773	1.074	2.189	0.465	1.982
Average	0.908	0.741	0.526	1.416	0.848	1.136	1.260	0.671	2.003

TABLE IV
AVERAGE DISTANCE BETWEEN MANUAL AND AUTOMATIC LANDMARKS

Lmk #	ϵ_κ [mm]	ϵ_{LC} [mm]
1	1.713	2.819
2	2.038	2.373
3	2.814	1.953
4	0.979	0.631
5	2.780	2.668
6	3.333	1.713
7	1.115	2.133
Average	2.110	2.042

In our experiments, we have not observed any problems of wrong correspondences leading to flipping of triangles or surface folding. This is an important improvement compared to the method of Brett and Taylor [5]. Also, our method is less restrictive in terms of the shapes that can be modeled. This is an important advantage over the extended method of Brett and Taylor [6], which is based on harmonic maps and therefore limited to shapes that are isomorphic to a disc. The proposed

3-D extension of the recent method by Davies *et al.* [11] has a similar problem, as it relies on a surface parameterization only applicable to shapes isomorphic to a sphere.

Our model can be classified as a statistical surface model. One of the main differences between our approach and the work by Brett and Taylor [5], [6] and Fleute and Lavallée [7], [8] is that we use a volumetric nonrigid registration algorithm as opposed to their surface-based approaches. As a consequence, after non-rigid registration, we are able to recover a dense volumetric displacement field. This could be used to propagate landmarks located inside the myocardium or blood pools producing a statistical solid model. To achieve this goal, most likely gray-level information will have to be included in the atlas construction and landmark propagation procedures. This information could provide the required textural information to match internal structures.

This paper has shown that the combination of our atlas generation method and the multiresolution FFD nonrigid registration algorithm is able to cope with the large deformations involved in intersubject matching of cardiac shapes. We had previously experimented with a single-level version of the FFD registration technique. Although with that approach we were able to propagate the landmarks of anatomical structures with moderate shape variability (deep structures of the brain and bone structures) [57], it was unsuccessful in the application presented in this paper. To cope with large shape variations, a multiresolution extension [13] of the free-form registration algorithm proposed by Rueckert *et al.* [12] was applied. In this approach, the nonrigid transformation is recovered by increasing the mesh resolution by subsequent subdivision. Gross shape warping takes place at the coarsest resolution, while shape details are captured at the finest resolution.

With this technique, we have been able to automatically place a set of specific landmarks with an average accuracy of about 2.2 mm and a precision of about 1.5 mm. In the same experiments, the precision intrinsic to manual landmarking was of about 0.8 mm. These figures are quite encouraging and indicate that automatic landmarking is feasible. To the best of our knowledge, this is the first study where the errors in automatic landmark placement are quantitatively assessed.

Our future work will focus on applying our method to model-based cardiac image segmentation and analysis. First, we are interested in enlarging the database of training shapes to capture the main modes of shape variation of cardiac chambers for a large population and not only for those of our reduced training set. Although this paper provides a proof-of-concept for our automatic landmarking algorithm, it is necessary to improve the statistics of the model to achieve an accurate segmentation. For this purpose, we have recently completed a study in which a previous version of this algorithm was applied to images of 100 subjects of the Framingham heart study,¹ of which the 14 shapes of this manuscript are a subset. This study will be published elsewhere.

All the shapes in our training set were acquired at end diastole. However, the fact that the nonrigid registration algorithm can cope with large intersubject variability suggests that the same experiments could be repeated for different phases of the cardiac cycle to build a statistical spatio temporal model of the heart. This is a line of future investigation.

Finally, a strategy has to be devised to adapt the model mesh to segment cardiac MR images. This could be achieved, for instance, by applying a method similar to the 2-D deformation procedure of active shape models [1]. For each landmark in the model, a statistical model of the intensity profile (or some other suitable image feature) along the surface normal can be computed. The model mesh could be deformed by moving the nodes along the direction of the normals to the position best matching the intensity with the statistical profile model. This method would provide an image-derived displacement for each node. The displacements applied to update the mesh can be obtained by projecting the suggested displacements onto the subspace spanned by the main modes of variation. This projection step would naturally incorporate shape constraints in the deformation of the mesh.

In conclusion, a method was presented to construct a shape atlas and to derive a statistical model of 3-D shape variability. We have demonstrated that this method is applicable to the construction of a statistical shape model of the cardiac chambers. To the best of our knowledge, this work is the first that uses 3-D statistical shape models to describe the left and right ventricle of the heart. Our future work will concentrate on the application of this model to cardiac image segmentation and analysis.

ACKNOWLEDGMENT

The authors acknowledge Dr. W. J. Manning, codirector of the Cardiac MR Center at the Beth Israel Deaconess Medical Center in Boston, MA, for kindly providing them with the functional cardiac MR data.

REFERENCES

- [1] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active Shape Models—Their training and application," *Comp. Vis. Image Understand.*, vol. 61, no. 1, pp. 38–59, 1995.
- [2] Y. Wang, B. S. Peterson, and L. W. Staib, "Shape-based 3D surface correspondence using geodesics and local geometry," in *Computer Vision Pattern Recognition*. Hilton Head, SC: IEEE Comp Society, 2000, vol. 2, pp. 644–651.

- [3] A. Kelemen, G. Szekely, and G. Guerig, "Elastic model-based segmentation of 3-D neuroradiological data sets," *IEEE Trans. Med. Imag.*, vol. 18, pp. 828–839, Oct. 1999.
- [4] C. Lorenz and N. Krahnstover, "Generation of point-based 3D statistical shape models for anatomical objects," *Comp. Vis. Image Understand.*, vol. 77, pp. 175–191, 2000.
- [5] A. D. Brett and C. J. Taylor, "A method of automated landmark generation for automated 3D PDM construction," *Imag. Vis. Comp.*, vol. 18, no. 9, pp. 739–748, 2000.
- [6] —, "Automated construction of 3D shape models using harmonic maps," in *Medical Image Understanding and Analysis*, S. Arridge and A. Todd-Pokropek, Eds. London, U.K., 2000, pp. 175–178.
- [7] M. Fleute and S. Lavallée, "Building a complete surface model from sparse data using statistical shape models: Applications to computer assisted knee surgery," in *Medical Imaging Computing & Computer-Assisted Intervention*. ser. Lecture Notes in Computer Science, W. M. Wells, A. Colchester, and S. Delp, Eds. Boston, MA: Springer-Verlag, 1998, vol. 1496, pp. 879–887.
- [8] —, "Incorporating a statistically based shape model into a system for computer-assisted anterior cruciate ligament surgery," *Med. Image Anal.*, vol. 3, no. 3, pp. 209–222, 1999.
- [9] G. Subsol, J.-P. Thirion, and N. Ayache, "A scheme for automatically building three-dimensional morphometric anatomical atlases: Application to a skull atlas," *Med. Image Anal.*, vol. 2, no. 1, pp. 37–60, 1998.
- [10] P. J. Besl and N. D. McKay, "A method for registration of 3D shapes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, pp. 239–255, Feb. 1992.
- [11] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor, "A minimum description length approach to statistical shape modeling," *IEEE Trans. Med. Imag.*, vol. 21, pp. 525–537, May 2002.
- [12] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Non-rigid registration using free-form deformations: Application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, pp. 712–721, Aug. 1999.
- [13] J. A. Schnabel, D. Rueckert, M. Quist, J. M. Blackall, A. D. Castellano-Smith, T. Hartkens, G. P. Penney, W. A. Hall, H. Liu, C. L. Truitt, F. A. Gerritsen, D. L. G. Hill, and D. J. Hawkes, "A generic framework for nonrigid registration based on nonuniform multi-level free-form deformations," in *Medical Image Computing & Computer Assisted Intervention*. ser. Lecture Notes in Computer Science, W. J. Niessen and M. A. Viergever, Eds. Utrecht, The Netherlands: Springer-Verlag, 2001, vol. 2208, pp. 573–581.
- [14] A. F. Frangi, W. J. Niessen, and M. A. Viergever, "Three-dimensional modeling for functional analysis of cardiac images: A review," *IEEE Trans. Med. Imag.*, vol. 20, pp. 2–25, Jan. 2001.
- [15] E. Bardinet, N. Ayache, and L. D. Cohen, "Tracking and motion analysis of the left ventricle with deformable superquadrics," *Med. Image Anal.*, vol. 1, no. 2, pp. 129–150, June 1996.
- [16] C. W. Chen, J. Luo, K. J. Parker, and T. S. Huang, "CT volumetric data-based left ventricle motion estimation: An integrated approach," *Comp. Med. Imag. Graph.*, vol. 19, no. 1, pp. 85–100, Jan. 1995.
- [17] A. Matheny and D. M. Goldgof, "The use of three- and four-dimensional surface harmonics for rigid and nonrigid shape recovery and representation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 17, pp. 967–981, Oct. 1995.
- [18] T. McInerney and D. Terzopoulos, "A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to 4D image analysis," *Comp. Med. Imag. Graph.*, vol. 19, no. 1, pp. 69–83, Jan. 1995.
- [19] W. J. Niessen, B. M. ter Haar Romeny, and M. A. Viergever, "Geodesic deformable models for medical image analysis," *IEEE Trans. Med. Imag.*, vol. 17, pp. 634–641, Aug. 1998.
- [20] T. O'Donnell, A. Gupta, and T. Boulton, "The hybrid volumetric ventriculoid: New model for MR-SPAMM 3-D analysis," in *IEEE Computing in Cardiology*. Vienna, Austria: IEEE Computer Society Press, 1995, pp. 5–8.
- [21] L. H. Staib and J. S. Duncan, "Model-based deformable surface finding for medical images," *IEEE Trans. Med. Imag.*, vol. 15, pp. 720–731, May 1996.
- [22] S. Benayoun and N. Ayache, "Dense and nonrigid motion estimation in sequences of medical images using differential constraints," *Int. J. Comput. Vision*, vol. 26, no. 1, pp. 25–40, Jan. 1998.
- [23] W.-C. Huang and D. Goldgof, "Adaptive-size meshes for rigid and non-rigid analysis and synthesis," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, pp. 611–616, June 1993.

¹<http://www.nhlbi.nih.gov/about/framingham>.

- [24] P. Shi, A. J. Sinusas, R. T. Constable, E. Ritman, and J. S. Duncan, "Point-tracked quantitative analysis of left ventricular surface motion from 3D image sequences: Algorithms and validation," *IEEE Trans. Med. Imag.*, vol. 19, pp. 36–50, Jan. 2000.
- [25] J. Declerck, J. Feldmar, and N. Ayache, "Definition of a four-dimensional continuous planispheric transformation for the tracking and the analysis of the left-ventricle motion," *Med. Image Anal.*, vol. 2, no. 2, pp. 197–213, June 1998.
- [26] J. Park, D. N. Metaxas, and L. Axel, "Analysis of left ventricular wall motion based on volumetric deformable models and MRI-SPAMM," *Med. Image Anal.*, vol. 1, no. 1, pp. 53–72, Mar. 1996.
- [27] A. A. Young, P. J. Hunter, and B. H. Smaill, "Estimation of epicardial strain using the motions of coronary bifurcations in biplane cineangiography," *IEEE Trans. Biomed. Eng.*, vol. 39, pp. 526–531, May 1992.
- [28] C. C. Moore, W. G. O'Dell, E. R. McVeigh, and E. A. Zerhouni, "Calculation of three-dimensional left ventricular strains from biplanar tagged MR images," *J. Magn. Reson. Imag.*, vol. 2, no. 2, pp. 165–175, Mar. 1992.
- [29] M. J. Moulton, L. L. Creswell, S. W. Downing, R. L. Actis, B. A. Szabo, M. W. Vannier, and M. K. Pasque, "Spline surface interpolation for calculating 3-D ventricular strains from MRI tissue tagging," *Amer. J. Physiol.*, pt. Pt 2, vol. 270, no. 1, pp. H281–297, Jan. 1996.
- [30] W. G. O'Dell, C. C. Moore, W. C. Hunter, E. A. Zerhouni, and E. R. McVeigh, "Three-dimensional myocardial deformations: Calculation with displacement field fitting to tagged MR images," *Radiology*, vol. 195, no. 3, pp. 829–835, June 1995.
- [31] X. Papademetris, P. Shi, D. P. Dione, A. J. Sinusas, R. T. Constable, and J. S. Duncan, "Recovery of soft tissue object deformation from 3D image sequences using biomechanical models," in *Information Processing and Medical Imaging*. ser. Lecture Notes in Computer Science, A. Kuba, M. Sámal, and A. Todd-Pokropek, Eds. Berlin, Germany: Springer-Verlag, 1999, vol. 1613, pp. 352–357.
- [32] B. P. F. Lelieveldt, R. J. van der Geest, M. R. Rezaee, J. G. Bosch, and J. H. C. Reiber, "Anatomical model matching with fuzzy implicit surfaces for segmentation of thoracic volume scan," *IEEE Trans. Med. Imag.*, vol. 18, pp. 231–238, Mar. 1999.
- [33] H. Azhari, E. Grenadier, U. Dinnar, R. Beyar, D. R. Adam, M. L. Marcus, and S. Sideman, "Quantitative characterization and sorting of three-dimensional geometries: Application to left ventricles in vivo," *IEEE Trans. Med. Imag.*, vol. 36, pp. 322–332, Mar. 1989.
- [34] S. K. Mishra, D. B. Goldgof, T. S. Huang, and C. Kambhampettu, "Curvature-based nonrigid motion analysis from 3D point correspondences," *Int. J. Image Syst. Tech.*, vol. 4, pp. 214–225, 1992.
- [35] D. Friboulet, I. E. Magnin, and D. Revel, "Assessment of a model for overall left ventricular three-dimensional motion from MRI data," *Int. J. Card. Imag.*, vol. 8, no. 3, pp. 175–190, Sept. 1992.
- [36] D. Friboulet, I. E. Magnin, C. Mathieu, A. Pommert, and K. H. Hoehne, "Assessment and visualization of the curvature of the left ventricle from 3D medical images," *Comp. Med. Imag. Graph.*, vol. 17, no. 4–5, pp. 257–262, July 1993.
- [37] C. Nastar and N. Ayache, "Frequency-based nonrigid motion analysis: Application to four dimensional medical images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, pp. 1067–1079, Nov. 1996.
- [38] P. Clarysse, D. Friboulet, and I. E. Magnin, "Tracking geometrical descriptors on 3-D deformable surfaces: Application to the left-ventricular surface of the heart," *IEEE Trans. Med. Imag.*, vol. 16, pp. 392–404, Aug. 1997.
- [39] J. Park, D. N. Metaxas, A. A. Young, and L. Axel, "Deformable models with parameter functions for cardiac motion analysis from tagged MRI data," *IEEE Trans. Med. Imag.*, vol. 15, pp. 278–289, June 1996.
- [40] C. W. Chen, T. S. Huang, and M. Arrott, "Modeling, analysis and visualization of left ventricle shape and motion by hierarchical decomposition," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 16, pp. 342–356, Apr. 1994.
- [41] A. Goshtasby and D. A. Turner, "Segmentation of cardiac cine MR images for extraction of right and left ventricular chambers," *IEEE Trans. Med. Imag.*, vol. 14, pp. 56–64, Mar. 1995.
- [42] A. S. Gopal, D. L. King, J. Katz, L. M. Boxt, D. L. King, and M. Y. Shao, "Three-dimensional echocardiographic volume computation by polyhedral surface reconstruction: In vitro validation and comparison to magnetic resonance imaging," *J. Amer. Soc. Echocardiogr.*, vol. 5, no. 2, pp. 115–124, Mar. 1992.
- [43] J. Montagnat, H. Delingette, and G. Malandain, *Cylindrical Echocardiographic Image Segmentation Based on 3D Deformable Models*, ser. Lecture Notes in Computer Science. Cambridge, U.K.: Springer-Verlag, 1999, vol. 1679, pp. 168–75.
- [44] A. Hyvarinen, J. Karhunen, and E. Ojja, "Independent Component Analysis," in *Wiley Series on Adaptive and Learning Systems for Signal Processing, Communications and Control*. New York: Wiley, 2001.
- [45] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "Automated 3D registration of MR and PET brain images by multiresolution optimization of voxel similarity measures," *Med. Phys.*, vol. 24, no. 1, pp. 25–35, 1997.
- [46] S. Lee, G. Wolberg, and S. Y. Shin, "Scattered data interpolation with multilevel B-splines," *IEEE Trans. Vis. Comput. Graph.*, vol. 3, pp. 228–244, Sept. 1997.
- [47] D. G. Altman, *Practical Statistics for Medical Research*. London, U.K.: Chapman & Hall, 1991.
- [48] S. P. Raya and J. K. Udupa, "Shape-based interpolation of multidimensional objects," *IEEE Trans. Med. Imag.*, vol. 9, pp. 32–42, Jan. 1990.
- [49] G. T. Herman, J. Zheng, and C. A. Bucholtz, "Shape-based interpolation," *IEEE Comput. Graph. Applicat.*, pp. 69–79, May 1992.
- [50] P. E. Danielsson, "Euclidean distance mapping," *Comp. Graph. Imag. Process.*, vol. 14, pp. 227–248, 1980.
- [51] D. Rueckert, A. F. Frangi, J. A. Schnabel "Automatic construction of 3D statistical deformation models using nonrigid registration," W. J. Niessen, and M. A. Viergever, in *Medical Image Computing & Computer Assisted Intervention*. ser. Lecture Notes in Computer Science, W. J. Niessen and M. A. Viergever, Eds. Utrecht, The Netherlands: Springer-Verlag, 2001, vol. 2208, pp. 77–84.
- [52] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface reconstruction algorithm," in *Proc. Computer Graphics: SIGGRAPH'87 Conf.*, vol. 21, July 1987, pp. 163–169.
- [53] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen, "Decimation of triangle meshes," *Comput. Graph.*, vol. 26, no. 2, pp. 65–70, 1992.
- [54] S. Gibson, "Constrained elastic surfacenet: Generating smooth surfaces from binary sampled data," in *Medical Imaging Computing & Computer-Assisted Intervention*. ser. Lecture Notes in Computer Science, W. M. Wells, A. Colchester, and S. Delp, Eds. Boston, MA: Springer-Verlag, 1998, vol. 1496, pp. 888–898.
- [55] A. Guéziec and R. Hummel, "Exploiting triangulated surface extraction using tetrahedral decomposition," *IEEE Trans. Vis. Comput. Graph.*, vol. 1, pp. 328–342, Dec. 1995.
- [56] M. Leventon, W. E. L. Grimson, and O. Faugeras, "Statistical shape influence in geodesic active contours," in *Computer Vision Pattern Recognition*. Hilton Head, SC: IEEE Computer Society, 2000, vol. 1, pp. 316–323.
- [57] A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen, "Automatic 3D ASM construction via atlas-based landmarking and volumetric elastic registration," in *Inform. Process. Med. Imag.*, M. F. Insana and R. M. Leahy, Eds., 2001, vol. 2082, pp. 78–91.