Automatic 3D ASM construction via atlas-based landmarking and volumetric elastic registration

A.F. Frangi[1], D. Rueckert[2], J.A. Schnabel[3], W.J. Niessen[1]

[1] Image Sciences Institute, University Medical Center Utrecht (UMC)
Room E.01.334, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands
{alex,wire}@isi.uu.nl
[2] VIPG, Department of Computing, Imperial College, London, UK
dr@doc.ic.ac.uk
[3] CISG, Radiological Sciences, Guy's Hospital, King's College London, UK
julia.schnabel@kcl.ac.uk

Abstract. A novel method is introduced that allows for the generation of landmarks for three-dimensional shapes and the construction of the corresponding 3D Active Shape Models (ASM). Landmarking of a set of examples from a class of shapes is achieved by (i) construction of an atlas of the class, (ii) automatic extraction of the landmarks from the atlas, and (iii) subsequent propagation of these landmarks to each example shape via a volumetric elastic deformation procedure. This paper describes in detail the method to generate the atlas, and the landmark extraction and propagation procedures. This technique presents some advantages over previously published methods: it can treat multiple-part structures, and it requires less restrictive assumptions on the structure's topology. The applicability of the developed technique is demonstrated with two examples: CT bone data and MR brain data.

1 Introduction

Statistical models of shape variability [5] or Active Shape Models (ASM) have been successfully applied to perform segmentation and recognition tasks in two-dimensional images. In building those statistical models, a set of segmentations of the shape of interest is required as well as a set of landmarks that can be defined in each sample shape.

Manual segmentation and determining point correspondences are time consuming and tedious tasks. This is particularly true for three-dimensional applications where the number of slices to analyze and the amount of landmarks required to describe the shape increases dramatically with respect to two-dimensional applications. This work aims at automating the landmarking procedure while we still rely on the existence of a manual segmentation of the shapes.

Several authors have proposed techniques to find point (landmark) correspondences but only a few of them have indicate or investigated their applicability in the field of statistical shape models. Wang *et al.* [16] use a surface registration technique to find 3D point correspondences based on a metric matching surface-to-surface distance, and surface normals and curvature. The authors suggest that

this technique could be used to build 3D ASMs but they do not report any results on statistical model building. Kelemen *et al.* [9] report on the construction of 3D ASMs of neuroradiological anatomical structures. In this method the authors used 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 3D ASMs, no explicit volumetric or surface registration between shapes takes place.

To our knowledge, little work has been done on the automatic construction of 3D ASM using elastic registration [7, 8, 4, 3]. The frameworks proposed by Brett and Taylor [4, 3] are most closely related to this paper. In these approaches, each shape is first converted into a polyhedral representation. In the first approach [4], shape pairs are matched using a symmetric version of the Iterative Closest Point (ICP) algorithm of Besl and McKay [2]. Using this method, the authors were able to build 3D ASMs by automatically finding corresponding landmarks between surfaces. Surfaces are represented by means of dense triangulations that are matched via sparse triangulations (obtained by triangle decimation from the dense triangulations). The nodes of this 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 does not incorporate connectivity constraints (purely local registration). In Brett and Taylor [3] this problem is overcome by transforming the surface to a planar domain by means of 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 multi-resolution elastic 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 this, in this work we use a free-form elastic registration technique based on maximization of normalized mutual information (volume-based technique). In addition, we provide experiments giving empirical evidence of the convergence of the atlas generation procedure that is not analyzed in [7, 8].

In this work a technique is introduced that addresses the shortcomings of point-based registration where no overall connectivity constraints are imposed. It uses a free-form elastic registration technique based on maximization of normalized mutual information (volume-based technique). Our method introduces global constraints by modifying the pairwise shape corresponder from a point-based registration technique into a volume-based elastic registration technique. By construction, the deformation field is enforced to be smooth and the regularization term of the deformation will further penalize folding. In addition, our method can be applied to multiple-part shapes.

The paper is organized as follows. In Section 2, our approach is described. In Section 3, results are presented that show the applicability of the method

to modeling the radius in volumetric Computed Tomography (CT) data and the caudate nucleus in Magnetic Resonance Imaging (MRI); empirical evidence is given on convergence properties and reconstruction errors. Finally, Section 4 closes the paper with some conclusions and directions for future research.

2 Method

2.1 Background

Suppose that we have n shapes described as vectors, $\{\mathbf{x}_i; i = 1 \cdots n\}$. Each shape consists of l 3-D landmarks, $\{\mathbf{p}_j = (p_{1j}, p_{2j}, p_{3j}); j = 1 \cdots l\}$ that represent the nodes of a surface triangulation. How to obtain those l 3-D landmarks is not a trivial issue and is precisely the topic of this paper. Each vector is of dimension $3l$ and is made up of the concatenation of the landmarks, *i.e.* $\mathbf{x}_i = (p_{11}, p_{21}, p_{31}, p_{12}, p_{22}, p_{32}, \cdots, p_{1l}, p_{2l}, p_{3l})$. Moreover, it is assumed that the positions of the landmarks of all shapes are in the same coordinate system. These vectors form a distribution in a $3l$-dimensional space. The goal is to approximate this distribution with a linear model of the form

$$\mathbf{x} = \hat{\mathbf{x}} + \mathbf{\Phi}\mathbf{b} \tag{1}$$

where $\hat{\mathbf{x}} = \frac{1}{n}\sum_{i=1}^{n} \mathbf{x}_i$ is the average landmark vector, $\mathbf{b}$ is the shape parameter vector of the model, and $\mathbf{\Phi}$ is a matrix whose columns are the principal components of the covariance matrix $\mathbf{S} = \frac{1}{n-1}\sum_{i=1}^{n}(\mathbf{x}_i - \hat{\mathbf{x}})(\mathbf{x}_i - \hat{\mathbf{x}})^T$. 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 $\mathbf{\Phi}$ contains the t eigenvectors corresponding to the largest eigenvalues, then we can approximate any shape of the training set, $\mathbf{x}$, using Eqn. (1) where $\mathbf{\Phi} = (\phi_1 | \phi_2 | \cdots | \phi_t)$ and $\mathbf{b}$ is a t dimensional vector given by $\mathbf{b} = \mathbf{\Phi}^T(\mathbf{x} - \hat{\mathbf{x}})$.

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 Eqn. (1). Under the assumption that the cloud of landmark vectors follows a multi-dimensional 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 the shapes contained in the training class.

2.2 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 3D shape. Therefore, we will consider pseudo-landmarks, *i.e.* landmarks lying on the shape's surface and determining its geometry.

In our framework, automatic landmarking is carried out by mapping the landmarks of an atlas that is representative of a set of training shapes. Let us

assume that n segmented shapes (3D binary images) are available, $\mathcal{T}_n = \{B_i\}$ where $i = 1 \cdots n$. To generate the landmarks for the n shapes, the task is to build an atlas A, landmark it, and propagate its landmarks to the n shapes (Fig. 1). In the following we will describe these three steps in detail.

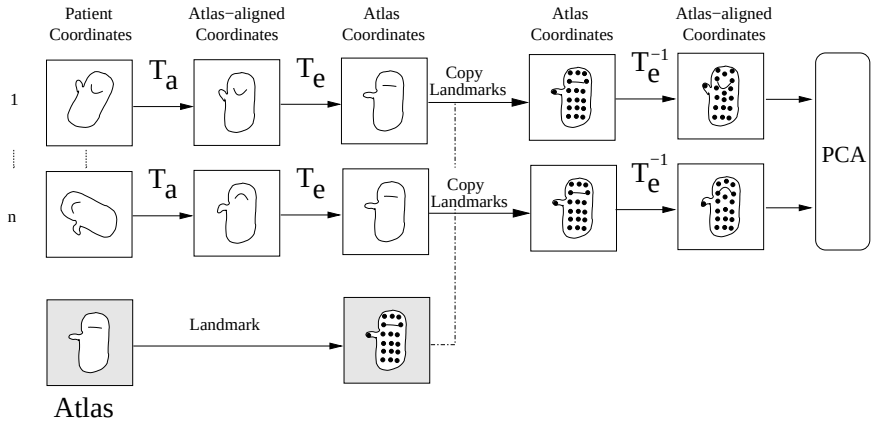

Fig. 1. Overview of the automatic landmarking framework. All individual data sets are matched to an atlas via an quasi-affine transformation (T_a) and an elastic transformation (T_e). The landmarks in the atlas can then be copied to the individual patients. The elastic deformation is subsequently reversed. Thus, Principal Component Analysis (PCA) is carried out in a space where all shapes are aligned with the atlas (atlas-aligned coordinates). The principal modes of variation will therefore account for elastic deformations and not for pose or size changes.

Atlas building. In the context of this paper, an *atlas* is an average representation of the shape of a structure inferred from a set of training shapes $\mathcal{T}_n$. In order to build the atlas, three issues have to be addressed: the selection of a *pairwise corresponder* to match two different shapes, a strategy to *blend shapes* that are represented as binary volumes in a common coordinate frame, and a scheme to obtain an average or *mean shape* with marginal bias towards a particular individual.

Pairwise shape corresponder. Given a shape B_i, it is matched to the atlas, A, using an *quasi-affine* registration algorithm with nine degrees of freedom (rigid transformation plus anisotropic scaling) adapted from [14]. This algorithm matches shapes using a criterion based on normalized mutual information [15]. Since the shapes are binary images, we have experimented with several other registration measures (sum of squared differences and cross-correlation) but normalized mutual information was found to be superior.

After registration, the shape B_i is expressed in the coordinate system of A. The coordinate system of A will be referred to as the atlas-aligned coordinate system.

Shape blending. Once we have found the quasi-affine transformations that map each of the B_i shapes into atlas-aligned coordinates, these shapes have to be combined to form an average shape (binary image).

Let B_i' and $\mathcal{DT}(B_i')$ denote the shape in atlas coordinates and its Euclidean distance transform [6] respectively, with the convention that inner points have a negative distance while outer points have a positive distance. Then, an average shape can be obtained in the distance transformed domain by computing $\mathcal{DT}(B_{av}) = \frac{1}{n}\sum_{i=1}^{n}\mathcal{DT}(B_i')$. A binary representation of the shape B_{av} can be obtained by thresholding the distance transform map to its zero-level set (Figure 2(a)).

Mean shape. To generate the mean shape it is necessary to register all $\mathcal{T}_n$ shapes into a common reference frame (atlas-aligned coordinates). However, the atlas is not initially known. To solve this problem an iterative algorithm was developed. One training shape is randomly selected as the initial atlas, A_0, and all remaining shapes are registered to it using the pairwise shape corresponder. After this step, all shapes $\mathcal{T}_n$ are expressed in the canonical system of A_0 and can be blended to generate a new atlas A_1. This procedure is iterated I times to reduce the effect of the initial shape. Any metric of similarity between the atlases of two consecutive iterations can be used to monitor the convergence of the procedure. The final atlas is A_I. This iterative algorithm is summarized in the flow diagram of Figure 2(b). To check for the influence of the randomly selected training shape, atlases with different start shapes have been quantitatively compared.

Atlas landmarking. By means of the iterative procedure of the previous subsection a binary atlas, A, has been obtained. In order to landmark this atlas the *marching cubes* [11] algorithm is used which generates a dense triangulation of the boundary surface. This triangulation can be decimated to obtain a sparse set of nodes that keeps the geometry of the original triangulation to a desired degree of accuracy. The number of nodes in this decimated triangulation corresponds to the number of landmarks. The use of different triangle densities (decimation ratios) has been investigated to observe their influence in the statistical models generated with our technique (see results section). The decimation strategy applied in this paper is the one proposed by Schroeder *et al.* [13]. Note that, as an alternative to marching cubes, an expert could manually pinpoint anatomical landmarks in the atlas. Anatomical landmarks, however, may be too sparse to accurately represent the shape of the structure. By using marching cubes, a dense and approximately even distribution of landmarks is obtained.

Landmark propagation. Once the atlas is constructed and landmarked, its landmarks can be propagated to the individual shapes. This is carried out

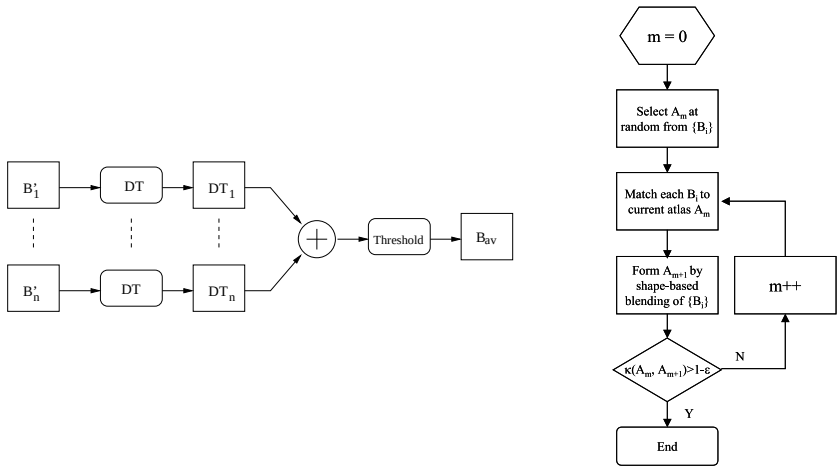

Fig. 2. (a) Shape-based blending of n registered binary shapes based on distance transforms (DT). By convention, the inside of the shape has negative distance and the outside positive distance. (b) Flow diagram of the iterative atlas construction algorithm.

by warping each sample binary volume into the atlas with a transformation, $T = T_a + T_e$, that is composed of an quasi-affine (T_a) and an elastic (T_e) transformation. The transformation T_a accounts for pose and size differences between the atlas and each sample volume while the transformation T_e accounts for shape differences.

The global transformation is obtained using an quasi-affine registration algorithm adapted from [14]. Registration of binary volumes was carried out using normalized mutual information [15]. The elastic transformation is expressed as a volumetric free-form deformation field computed with the method of Rueckert *et al.* [12] that also uses normalized mutual information as a registration measure.

Once the global transformation T has been found, the landmarks of the atlas could be propagated to the atlas-aligned coordinate system by applying the inverse of the elastic transformation (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 atlas-aligned coordinates, pose and size variations are explicitly eliminated from further analysis. These transformed landmarks are subsequently used as input for Principal Component Analysis (PCA) as indicated in Figure 1.

Figure 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.

3 Results

3.1 Data sets

In order to exemplify the methodology, two case studies were analyzed. The first case study consists of a set of 14 manual segmentations of the head of the radius, a bone of the wrist, extracted from CT scans (voxel dimensions $1 \times 1 \times 2$ mm^3). The second is a set of 20 manual segmentations of the caudate nucleus, a deep structure of the brain, from MR scans (voxel dimensions $1 \times 1 \times 1.2$ mm^3).

In building the model of the caudate nucleus each hemisphere of the structure was treated independently. This was done because this particular two-part structure has an almost specular symmetry with respect to the sagittal plane separating the left and right brain hemispheres. Such symmetry would be difficult to capture with a single quasi-affine transformation. After the landmarks of each side (sub-atlas) are extracted and propagated, Principal Component Analysis (PCA) is applied to the concatenation of the landmarks of both sides. In this way, inter-hemisphere relationships are included in the statistical analysis.

3.2 Atlas construction

Convergence properties. As a metric to measure convergence we have used the κ statistic [1]. This statistic measures the similarity between two binary images, $\kappa(A_m, A_{m-1})$, in a way that is independent of the structure's volume. Figure 3 shows the evolution of the κ statistic, $\kappa(m)$, as a function of the iteration number, m. This statistic ranges between 0.0 and 1.0 and a value above 0.9 is usually regarded as an excellent agreement [1]. The $\kappa(m)$ statistic compares the similarity between the atlases A_m and A_{m-1}. Figure 3(a) corresponds to the atlas of the radius. Two curves are shown for two different initial shapes used in the initialization procedure. Similar curves are drawn in Figure 3(b) for the left and right caudate nucleus atlases. The atlas of each subpart (left/right caudate nucleus) was obtained independently. The trend of these plots is similar to that observed in the atlas of the radius. Figure 3 indicates that after five iterations the shape of the atlas stabilizes ($\kappa > 0.97$).

Effect of initial shape. We investigated whether the atlases generated with the two different initializations are comparable in shape, *i.e.* similar up to an quasi-affine registration. This was done in the following way. For each individual shape, two quasi-affine transformations can be found that map it to each of the two atlases, A and B. Let us call these transformations T_{A_i} and T_{B_i}, respectively. Let T_{AB} be the quasi-affine transformation that maps the atlas A into the atlas B. In this situation, the transformation $T_i = T_{Bi}^{-1} T_{AB} T_{Ai}$ should be equal to the identity transformation, T_I. It is possible now to measure the average and the standard deviation of the difference $T_i - T_I$. These two measures will provide the bias and dispersion introduced by using two different initial shapes to build the atlas. The results of this analysis are shown in Table 1 for each atlas and

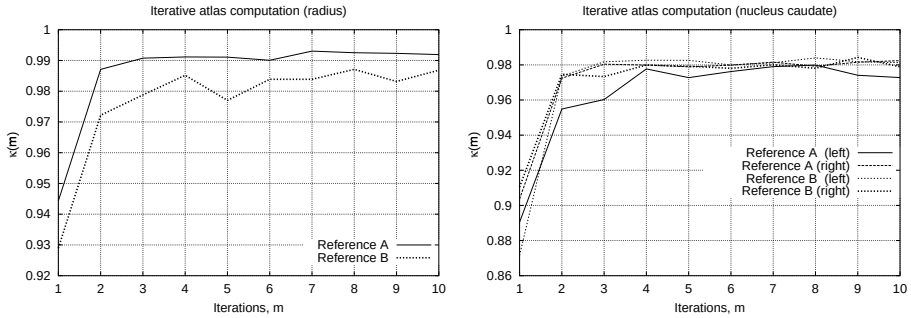

(a) (b)

Fig. 3. Convergence of the atlas construction algorithm. The κ statistic between two consecutive atlases as a function of the iteration number. Iteration zero corresponds to the reference (initial) shape used in the iterative algorithm. The $\kappa(m)$ statistic compares the agreement between the atlases A_m and A_{m-1}. Curves for different initial shapes (A and B) are shown.

each transformation parameter. This table indicates that the deviation from an identity transformation depends on the type of shape. For the very elongated and thin structure of the caudate nucleus the error standard deviations (SDs) are larger compared to the radius. As a consequence, the influence of the initial shape on the final atlas will depend on the shape itself. Translation and rotation error SDs are below 3.3 mm and 0.1°, respectively. Scaling error SDs are below 14.5%. From a practical point of view Table 1 indicates that the atlas does indeed depend on the initial shape and that the effect is has to do with the class of shapes being modeled. In the applications presented in this chapter, this effect is not critical. After performing an quasi-affine registration of the atlases generated with two different initializations, the average boundary-to-boundary distance between the two atlases was 1.3 mm and 0.6 mm for the radius and the

Table 1. Mean (standard deviation) of the error in each transformation parameter (translation, rotation and scaling) of the transformation T_i with respect to the identity transformation for three different atlases.

Parameter	Units	Radius	Caudate (L)	Caudate (R)
t_x	[mm]	-0.72 (1.68)	+1.25 (3.28)	+0.62 (1.42)
t_y	[mm]	-1.20 (1.32)	-0.20 (0.71)	-0.14 (0.57)
t_z	[mm]	+0.64 (1.99)	-0.25 (0.54)	+0.06 (0.17)
r_x	[°]	+0.01 (0.02)	-0.01 (0.03)	+0.02 (0.03)
r_y	[°]	-0.01 (0.02)	-0.04 (0.09)	+0.10 (0.05)
r_z	[°]	-0.01 (0.02)	+0.01 (0.08)	-0.02 (0.06)
s_x	[%]	-0.57 (1.99)	+3.45 (14.51)	-5.60 (8.20)
s_y	[%]	-1.48 (1.78)	-2.12 (6.28)	-1.47 (3.92)
s_z	[%]	+1.57 (6.08)	-3.22 (7.23)	-1.98 (4.12)

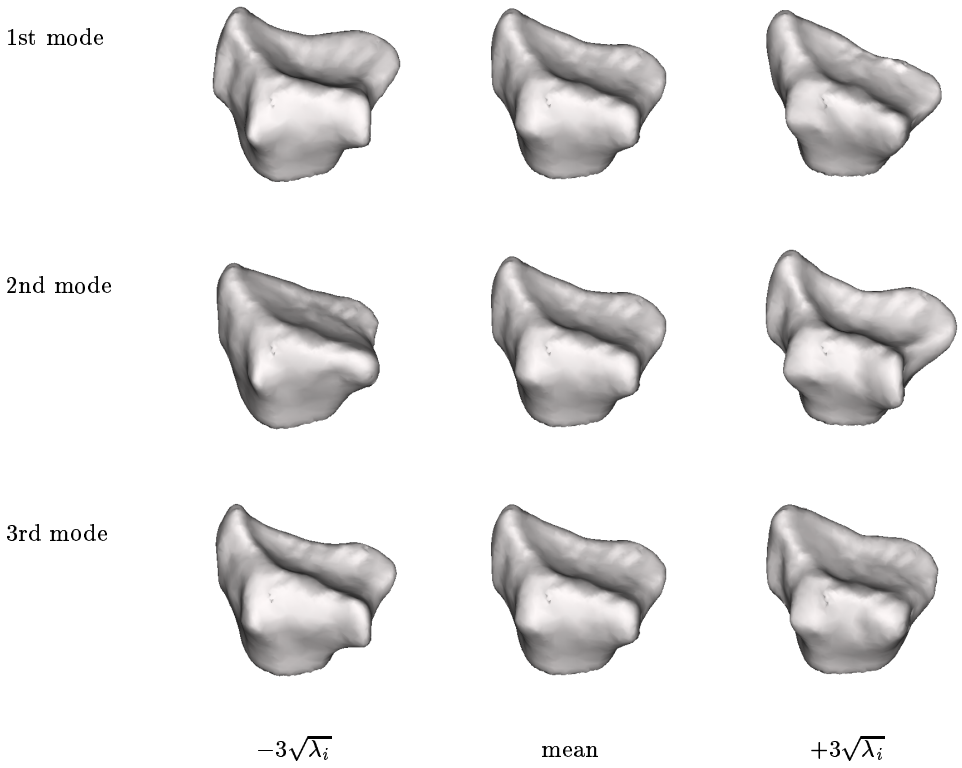

Fig. 4. Shape instances generated using the 3D model from 14 data sets of the radius. The instances are generated by varying a single shape parameter, fixing all others constant at zero standard deviations from the mean shape. Each instance of the model consists of 2500 nodes.

two caudate nucleus atlases, respectively. These errors are on the order of, and slightly smaller than the voxel dimensions, respectively.

3.3 Point distribution models

Figures 4 and 5 show the mean shape models and the first three modes of variation obtained from PCA for the radius and caudate nucleus test cases, respectively. The number of mesh nodes is 2500 and 1000, respectively. In both cases there are no visible surface foldings neither in the mean shape nor in the models for $\pm 3\sqrt{\lambda_i}$.

3.4 Reconstruction error

Figure 6 illustrates the relative shape variance explained with an increasing number of modes. Similar curves for different decimation ratios (number of model

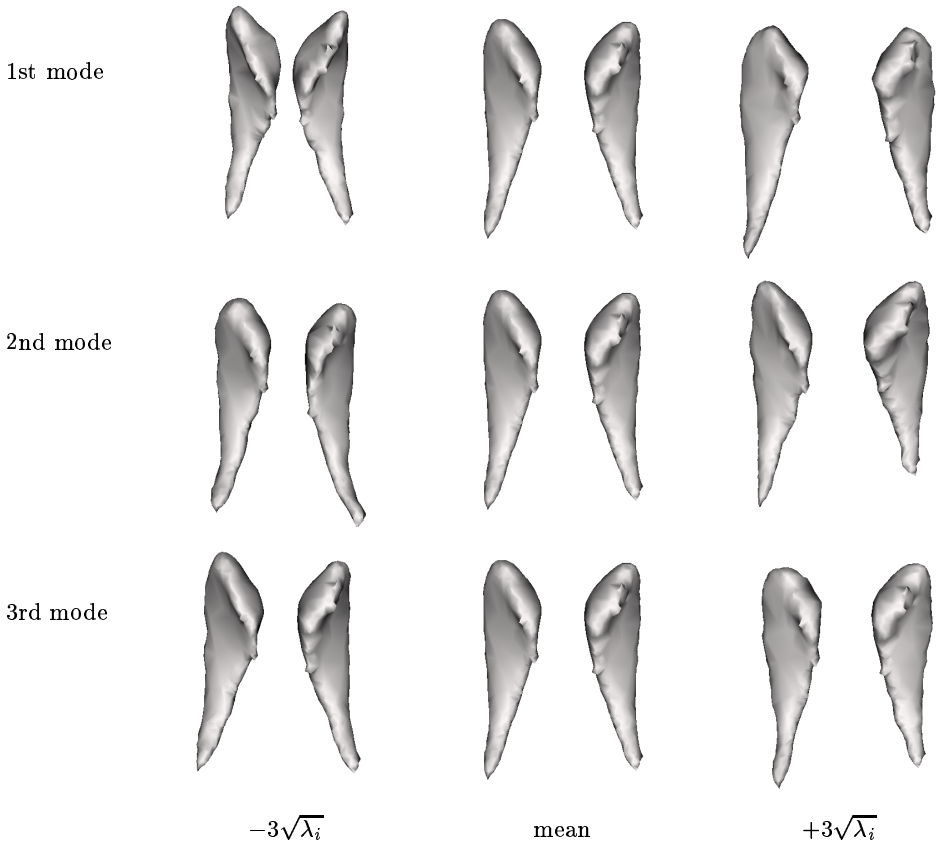

Fig. 5. Shape instances generated using the 3D model from 20 data sets of the caudate nucleus. The instances are generated by varying a single shape parameter, fixing all others constant at zero standard deviations from the mean shape. Each instance of the model consists of 1000 nodes.

triangles) are provided. These curves are only marginally dependent on this factor. From ten modes onwards, the model captures more than 90% of the shape variance. Note the steeper slope of the curves corresponding to the caudate nucleus. Over the training set there is apparently less variability in the shape of the caudate nucleus than in the shape of the radius. As a consequence, with fewer modes a larger amount of shape variation can be explained.

In order to assess the ability of these models to recover shapes not used in the training set we carried out the following experiment. Reconstruction errors were computed by reconstructing the landmarks of one shape of the training set with the ASM built from the remaining shapes (leave-one-out experiment). The errors reported in Figure 7 are the average of the reconstruction errors over all shapes taking out one in turn. The same experiment was repeated for different decimation ratios and increasing number of modes of shape variation taken

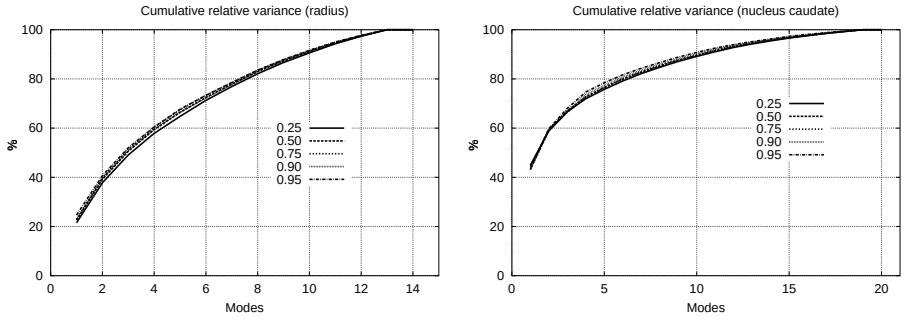

Fig. 6. Percentage of total shape variance versus the number of modes used in the 3D ASM and for various decimation ratios. The number of landmarks before decimation was 15519 for the radius, and 2320 for the caudate nucleus. The decimation ratio represents the ratio between the nodes eliminated from the triangulation of the atlas and its initial number. Note that the number of modes is at most the number of sample shapes minus one.

into the reconstruction. The reconstruction errors were computed in millimeters. For the caudate nucleus, the reconstruction error is below the voxel dimensions (10 modes). In the case of the radius, the reconstruction error is slightly larger than the slice thickness. One possible explanation to this higher error could be the fact that no image resampling was used during registration. On the other hand, in comparison to the shape of the caudate nucleus, the radius represents a more complex structure with larger shape variability in the training set. This could explain the poorer reconstruction performance in the leave-one-out experiments of the radius. The plots of Figure 7 also indicate that the reconstruction error is slightly dependent on the decimation ratio and, as expected, inversely proportional to the number of modes of variation.

4 Discussion and conclusion

This paper presents a method for the automatic construction of 3D active shape models. The technique is based on the automatic extraction of a dense mesh of landmarks in an atlas constructed from the training shapes which are propagated through an elastic deformation field to each shape of 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 3 we showed experimental results indicating that this procedure is convergent. Moreover, different initial shapes seem to contribute only marginally to the final atlas. That is, the final atlases are similar up to an quasi-affine transformation. However, we note that the influence of the initial shape depends on the class of shapes being modeled and has to be assessed on a case-by-case basis. In the work by Fleute and Lavallée [7, 8] a similar

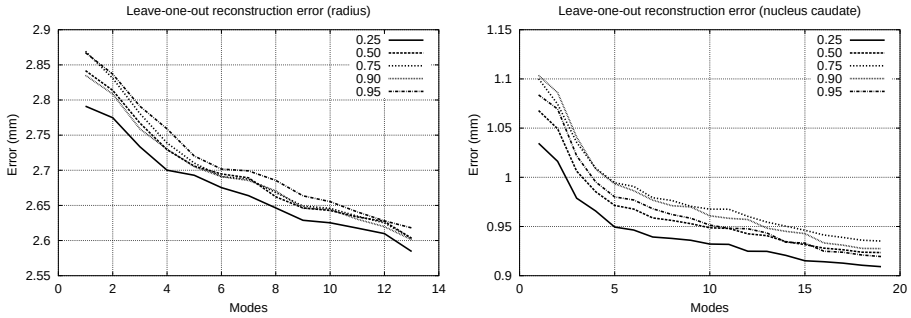

(a) (b)

Fig. 7. Reconstruction error in the leave-one-out experiments. The number of landmarks before decimation was 15519 for the radius, and 2320 for the caudate nucleus. The decimation ratio represents the ratio between the nodes eliminated from the triangulation of the atlas and its initial number.

algorithm was used to build the average model (atlas). However, no experimental evidence was reported with respect to the convergence of the atlas construction algorithm.

An alternative to our iterative method of atlas construction is the tree-based approach presented by Brett and Taylor [4]. This hierarchical strategy is attractive since it gives a unique (non-iterative) way to build an atlas from a given set of examples. However, one problem of Brett's method is that the training shapes 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 [4] but ordering the examples according to the matching quality would be cumbersome for a 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 3 shows 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.* [10] 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 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 complex structures.

The proposed technique could be used with any elastic registration algorithm. In this sense, the method is a generic framework open to future research. Currently, the volumetric elastic registration of Rueckert *et al.* [12] is used to match binary images. The use of elastic 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 sub-structures 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 surgery. However, establishing correspondences in these mixed models also remains an ill-defined problem with any of the previously published approaches [7, 8, 3].

Results of the construction of models of two anatomical structures 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 was below 2.65 mm (radius) and 0.95 mm (caudate nucleus) when the number of nodes used was sufficient to explain 90% of the shape variability. These errors are on the order of, and slightly smaller than the voxel dimensions, respectively. In our experiments we have not observed problems of wrong correspondences leading to flipping of triangles and surface folding. This is an important improvement compared to the initial method of Brett and Taylor [4]. Also, our method is less restrictive in terms of the shapes that can be modeled. This is an important feature with respect to the improved method of Brett and Taylor [3] that is based on harmonic maps and therefore limited to shapes that are isomorphic to a disc.

Finally, it would be interesting to perform a comparison between the models built with different methods. In order to carry out a quantitative comparison it is necessary to define a measure of model quality. The definition of such a measure is in itself an interesting issue. Obviously, different methods will yield different sets of landmarks which precludes a landmark-based comparison. If one defines a given segmentation task, a comparison could be established on the basis of the segmentation accuracy. Although these measures can have a prominent practical value to determine the best model-building technique for a given problem, the conclusions will remain task-dependent. Possibly, other more task-independent criteria related to the compactness and generalizability of the built models could be within the interesting candidate measures to explore.

Acknowledgements

This research was sponsored by the Dutch Ministry of Economic Affairs (IOP Beeldverwerking IBV97009) and EasyVision Advanced Development, Philips Medical Systems BV, Best, The Netherlands. Dr. Maarten Hoogbergen pro-

vided us with the radius segmentations, and the Department of Psychiatry of the University Medical Center Utrecht with the caudate nucleus segmentations.

References

1. D.G. Altman. *Practical Statistics for Medical Research.* Chapman & Hall, 1991.
2. P.J. Besl and N.D. McKay. A method for registration of 3D shapes. *IEEE Trans Pattern Anal Machine Intell*, 14(2):239–55, February 1992.
3. A.D. Brett and C.J. Taylor. Automated construction of 3D shape models using harmonic maps. In S. Arridge and A. Todd-Pokropek, editors, *Medical Image Understanding and Analysis*, pages 175–78, London, July 2000.
4. A.D. Brett and C.J. Taylor. A method of automated landmark generation for automated 3D PDM construction. *Imag Vis Comp*, 18(9):739–48, 2000.
5. T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Active Shape Models - their training and application. *Comp Vis Image Underst*, 61(1):38–59, 1995.
6. P.E. Danielsson. Euclidean distance mapping. *Comp Graph Imag Proces*, 14:227–48, 1980.
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 W.M. Wells, A. Colchester, and S. Delp, editors, *Medical Imaging Computing & Computer-Assisted Intervention*, volume 1496 of *Lect Notes Comp Science*, pages 879–87, Boston, USA, September 1998. Springer Verlag.
8. M. Fleute and S. Lavallée. Incorporating a statistically based shape model into a system for computer-assisted anterior cruciate ligament surgery. *Med Image Anal*, 3(3):209–22, 1999.
9. A. Kelemen, G. Székely, and G. Guerig. Elastic model-based segmentation of 3-D neuroradiological data sets. *IEEE Trans Med Imaging*, 18(10):828–39, October 1999.
10. M. Leventon, W.E.L. Grimsom, and O. Faugeras. Shape-based 3D surface correspondence using geodesics and local geometry. In *Comp Vis Patt Recogn*, volume 1, pages 316–23, South Carolina, USA, June 2000. IEEE Computer Society.
11. W.E. Lorensen and H.E. Cline. Marching cubes: a high resolution 3D surface reconstruction algorithm. *Computer Graphics: SIGGRAPH'87 Conference Proceeding*, 21:163–69, July 1987.
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 Imaging*, 18(8):712–21, August 1999.
13. W.J. Schroeder, J.A. Zarge, and W.E. Lorensen. Decimation of triangle meshes. *Comp Graphics*, 26(2):65–70, 1992.
14. 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*, 24(1):25–35, 1997.
15. C. Studholme, D.L.G. Hill, and D.J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recogn*, 32(1):71–86, 1998.
16. Y. Wang, B.S. Peterson, and L.W. Staib. Shape-based 3D surface correspondence using geodesics and local geometry. In *Comp Vis Patt Recogn*, volume 2, pages 644–51, South Carolina, USA, June 2000. IEEE Computer Society.