Generation of point based 3D statistical shape models for anatomical objects

Cristian Lorenz

Philips Research Hamburg, Röntgenstraße 24-26, D-22335 Hamburg, Germany

C.Lorenz@pfh.research.philips.com

Nils Krahnstöver

Department of Computer Science and Engineering, Pennsylvania State University, 220 Pond Lab, University Park, PA 16802, USA

krahnsto@cse.psu.edu

3D models for anatomical objects

Cristian Lorenz,

Philips Research Hamburg, Röntgenstraße 24-26, D-22335 Hamburg,

T. ++49-40-5078-2063, Fax -2510, C.Lorenz@pfh.research.philips.com

ABSTRACT

A novel method that allows the development of surface point based three-dimensional statistical shape models is presented. Given a set of medical objects, a statistical shape model can be obtained by Principal Component Analysis. This technique requires that a set of complex shaped objects is represented as a set of vectors that uniquely determine the shapes of the objects, and at the same time are suitable for a statistical analysis. The correspondence between the vector components and the respective shape features has to be identical in order for all shape parameter vectors to be considered. We present a novel approach to the correspondence problem for arbitrary three-dimensional objects, which involves developing a template shape and fitting this template to all objects to be analyzed. The method is successfully applied to obtain a statistical shape model for the lumbar vertebrae.

Keywords: 3D, statistical, shape model, point distribution model, medical imaging, anatomical objects, triangulation

$\Re$: German type "R"

$\subset$: is subset of

$\vec{v}$: small "v", vector

α : Greek alpha

λ : Greek lambda

σ : Greek sigma

1. INTRODUCTION

In recent years, a rapidly increasing portion of research in the field of medical image analysis has begun to focus on shape as an anatomical object property. Shape representations and shape models are used in connection with: three-dimensional visualization of anatomical objects; segmentation of 3D medical images [1]; diagnosis (e.g., based on deformable atlases) [2]; surgical simulation [3]; motion analysis (e.g., of the heart) [4] and radiotherapy treatment planning [5]. The desired properties of shape representations and shape models depend largely on the medical application to be supported. With reference to shape, one has to distinguish between shape classes (e.g., heart-shaped), shape instances (the shape of certain heart), shape properties (such as local curvature, extension), shape representations (e.g., bi-cubic-interpolation) and shape parameterization (e.g., the control points and the (u,v) surface coordinates of Bézier patches).

Shape representation and parameterization are crucial for computerized processing and the manipulation of shapes. The choice of representation determines to a great extend the flexibility, processing speed, and amount of user interaction provided by a given application. One of the most important issues is the number of parameters needed to describe an object using a given representation. The shape of a sphere-like object, for example, can be described at one extreme using the set of surface points or at the other extreme using the four parameters of center and radius. The identification of a 'good' shape representation usually requires minimizing the number of necessary parameters to represent a shape of given complexity. Depending on the application of the shape model, it may also be required that the shape representation and parameterization are meaningful to the user, to ease user interaction.

In the domain of medical applications, the capture of object variability is especially important. A shape model for a given anatomical object should be able to resemble a large portion of the shape instances that can be found either in different individuals (inter-patient variability) or within the same individual (intra-patient variability). This effect can be achieved by a statistical analysis of a set of shape instances of the object of interest [6,7]. The parameter vector of the shape parameterization is

interpreted as a random vector. The parameter vectors of a set of shape instances define a multivariate distribution that can be described by the infinite series of moments of the random vector. If the distribution is Gaussian, it is already determined by the first (mean) and second (covariance) moments. In the case of a non-Gaussian distribution, the first and second moments can still be used as an approximation. Thus, in the Gaussian approximation, the shape variability is captured by the covariance matrix of shape vectors representing a set of objects. The eigenvectors of the covariance matrix can be interpreted as fundamental modes of variation, or short "eigenmodes". A subset of these modes can be used as a set of basis vectors defining an approximated "shape-space" of the set of objects. A weighted sum of these basis vectors and the mean shape is well suited for the generation of new shape samples in order to support tasks such as shape recognition or image segmentation[7]. The appealing property of this approach is the ability to create models that depend on a small number of parameters, (i.e., on a small number of eigenmodes), even for intricately shaped objects.

Different types of shape representations and parameterizations have been used for statistical shape models. Point distribution models with parameter vectors being composed of the coordinates of contour points, have been frequently applied [8,9]. If an object contour is represented using a Fourier decomposition, the statistical analysis can be applied to the vector of Fourier coefficients [10] instead. Wavelet decompositions can be treated in a similar manner [6]. In general, the eigenmode analysis can be based on any kind of object representation.

A crucial requirement for a statistical analysis is the one-to-one correspondence between the elements of parameter vectors for different shape instances. All vector elements with the same index must represent corresponding shape information. Even in two dimensions, it can be difficult to find this correspondence for complicated objects [11]. In case of a single closed contour per object correspondence can be achieved by simple ordering schemes. In case of a point distribution model, it is sufficient to define the starting point and a direction of the contour and add equidistant key-points along the contour. Another possibility is to define a set of meaningful landmarks along the contour and to fill the remaining segments with a fixed number of equidistant key-points. In case of Fourier or wavelet decompositions, only a starting point and a contour direction has to be defined as the decomposition coefficients are already intrinsically ordered. For 3D shapes however, the

correspondence requirement is in general difficult to fulfill. In case of one-connected objects, the intrinsic ordering of coefficients of a decomposition in spherical harmonics can be used [10]. This assumes a definition of a spherical coordinate-system on the surface of the object. In order to achieve a homogenous parameterization, in the sense that the size of surface elements is roughly the same all over the object surface, optimization schemes for the definition of a spherical coordinate system have been proposed [10]. Another possibility is the 3D extension of point distribution models using a stack of 2D contours [8,9]. This approach is well-suited for roughly cylindrical objects. Difficulties can arise in case of object branching. An undesirable effect however, is the dependency on the slice orientation.

The work presented herein, adopts a new approach to achieve parametric correspondence. We use a point distribution model representing the object shape by a set of surface points. The point-set results from a point density reduction procedure of the object surface in the voxel image. A triangulation of the point-set defines a surface mesh and can be used to cover the surface area in-between the points. The point density reduction performed on two objects will in general produce non-corresponding point sets. Therefore, the point density reduction is performed only for one object instance and the resulting point set is used subsequently as a template to be "coated" onto the surfaces of the remaining object instances. The coating procedure preserves the required one-to-one correspondence. Applied to all object instances, this procedure results in a set of parameter vectors, the dimensions of which are three times the number of surface points. The final shape model consists of the mean shape parameter vector and variation modes in the form of the eigenvectors of the covariance matrix. The model can serve as a low-parametric organ specific shape model and can be used, for example, for segmentation or registration of medical images. We used abdominal CT scans of lumbar vertebrae for a first application of the model generation procedure. The obtained model is based on 31 object instances and contains roughly 600 surface points and 1200 triangles.

The following section provides an overview of the model creation scheme. Section 3 describes the creation of the shape template and Section 4 describes the coating procedure. The statistical analysis and an overall application of our method are presented in section 5 and 6, followed by a discussion of the procedure in section 7.

2. SHAPE-MODEL GENERATION SCHEME

An overview of the shape-model generation scheme is given in figure 1. The input is a set of 3D images. The assumption is made that the images are segmented with respect to the object of interest. The first step is the definition of an object-related coordinate system. In case of the vertebrae, we chose the center of mass as the origin and the main axes of the second momentum matrix of the binarized objects as the coordinate axes. It turns out, that these axes are roughly aligned along the patient left/right, anterior/posterior and head/foot axes. The second step is a re-sampling of the voxel images along the object coordinate axis. This step results in isotropic voxel images of the objects with predefined resolution, which are the basis for all further processing. The third step involves the generation of a point distribution shape template, which incorporates a surface-curvature guided point density reduction and a triangulation procedure. This procedure is performed only once. In step four, the shape template is "coated" to all remaining object instances. First, a small set of corresponding landmarks is defined interactively by the user on the template object and the object to be coated. Using an image warping approach, the template is then deformed such that corresponding landmarks are in alignment. As a result, the template roughly matches the object to be coated. Nevertheless, in general the template points will not be located on the surface of the object. To achieve this, the landmark based deformation is followed by a mesh relaxation that on the one hand, draws the template points to the object surface and, on the other hand, tries to retain mesh properties such as distance relations between neighboring surface points. For the relaxation, we use a spring-mass model based on the surface point set and the connecting edges, as defined by the triangulation. The final step of the model generation scheme contains the statistical analysis of the shape parameter vectors that are produced in the previous step. The mean shape is calculated as well as the eigenvectors and eigenvalues of the co-variance matrix of the set of shape vectors.

3. GENERATION OF A 3D POINT DISTRIBUTION TEMPLATE

One of the objectives of a shape model is the reduction of parameters needed to describe a shape instance. In a first step towards this goal, the set of surface voxels of an object in a volume data set is thinned out to obtain a lower number of surface points. The method and parameters of this point density reduction procedure determine the resolution of the shape representation. The reduction is guided by the local surface curvature because it is desirable that the surface point density at a certain surface patch correlates with the local "degree of detail" of the object. In addition, the *selection* of the actual points can be guided by the local surface curvature.

After the point density reduction procedure, a triangulation of the resulting point set is performed. A triangulation is necessary because it contains the information about the surface that the point distribution actually represents. Furthermore, the triangulation allows fast visualization of the shape model.

The reduction procedure is described in section 3.1 followed by the description of the surface curvature estimation in section 3.2. The method that was developed to obtain a surface triangulation is outlined in section 3.3.

3.1 Point density reduction

For the generation of a point-distribution model of a surface, it is necessary to select a sub-set of surface voxels as *surface points*. In a typical CT data-set, the number of surface voxels of a lumbar vertebra is roughly 10^4. A surface voxel is defined as an object voxel that is 26-connected to a background voxel. This selection procedure, referred to as the point density reduction procedure, works as follows [13]: Initially all surface voxels are candidates and can be selected by the reduction algorithm. The algorithm iteratively selects surface points from the candidate list to be included in the

point distribution model and eliminates all candidates that are within a certain neighborhood of the selected point. This selection and elimination of surface voxel is repeated until no candidates are left. The point density reduction is schematically shown in figure 2. Details about the procedure by which the surface voxels are selected and eliminated from the candidate set are two important aspects of the reduction procedure. The surface voxels can be selected either be at random or selection can curvature guided. It is of advantage to select voxels from those parts of the surface where the local curvature is high because those points lie on edges and tips of the surface. The set of surface voxels to be eliminated is determined as follows: Given a selected surface voxel and a reduction radius r, all surface voxels that have a *surface* distance less than r are eliminated around the voxel. Surface distance between two surface voxels is defined as the shortest path along neighboring surface voxels, where the neighborhood relation is defined as 26-connectedness [13]. The length of a path is given by the sum of the Manhattan distances between voxels along the path. The reduction radius r defines the resolution of the point distribution and can either be a constant or dependent on the surface curvature, as previously used. In the latter case, it is chosen to be small in regions where the surface curvature is high. Making the reduction radius variable leads to an anisotropic point distribution. The reduction radius can be related linearly to the surface curvature c or non-linearly such that small reduction radii are given more weight using the following function: Denoting with $c_{\min}$ and $c_{\max}$, the minimal and maximal surface curvature found for the object and with $r_{\min}$ and $r_{\max}$, the minimal and maximal reduction radius that shall be used for the reduction process, the reduction radius is given by:

$$r = \left(\frac{c - c_{\min}}{c_{\max} - c_{\min}} \right)^{1/p} \cdot \left(r_{\min} - r_{\max} \right) + r_{\max} \qquad (1)$$

The parameter p steers the nonlinear behavior. For $p = 1$, the relation between c and r is linear, for $p > 1$, the fraction of small radii used is being enlarged. For $p \to \infty$ only $r_{\min}$ is used. The parameters p, $r_{\min}$ and $r_{\max}$ have to be chosen by the user. We hope that, in the future, the above relation between c and r can be replaced by a relation needing only one user defined parameter such as the maximal allowed distance between a triangle in the surface triangulation and the actual object surface.

3.2 Surface curvature estimation

As shown above, it is recommended to perform a surface curvature dependent point density reduction of the surface voxels. The curvature estimation used here is described in [12] and is referred to as "Coordinate Transformation Method". The procedure consists of the following steps for every surface voxel:

1. Estimate the surface normal.
2. Select a surface patch of a predefined size around the voxel.
3. Calculate the coordinates of the patch voxels in the coordinate system with x- and y axis in the tangential plane and z-axis along surface normal.
4. Perform a parabolic approximation of the patch with $z = f(x, y) = a \cdot x^2 + 2 \cdot b \cdot x \cdot y + c \cdot y^2$ using the pseudo inverse method and estimate the main curvatures k_1 and k_2 of the surface patch as the eigenvalues of the matrix $2 \cdot \begin{vmatrix} a & b \\ b & c \end{vmatrix}$. As a mean curvature we use $c = \sqrt{k_1^2 + k_2^2}$.

For the surface normal estimation, we used a method described in [13]. The normal is approximated by the difference vector between the center of the voxel of interest and the center of mass of all volume voxels in a sphere around the voxel of interest. The radius of the sphere determines the degree of smoothing applied to the surface normal estimation. The diameter of the surface patch on the other hand determines the smoothing applied to the curvature estimation.

The result of the procedure for a cube and a vertebra is shown in figure 3.

3.3 Delaunay Triangulation

To obtain a complete parameterization of the object surface and to ease visualization, the surface point set resulting from the point density reduction procedure is triangulated. To do this, a Voronoi graph of the surface point set is used to obtain the Delaunay triangulation [14]. An elementary multiple source

shortest path graph algorithm is used to simultaneously expand regions of surface voxels around each surface point selected in the point density reduction procedure. As the regions increase in size, the borders of the regions start to collide. The points where the fringes collide represent the points that have equal surface distance from two or more surface points and approximately represent the edges and nodes of the Voronoi graph on the surface of the object. The points where more than two regions collide, i.e., the nodes of the Voronoi graph, represent the faces of the triangulation.

Figure 4 shows three triangulations of a cube. Fig. 4 (left) shows clearly that a random selection of surface voxels is ignorant of edges and corners while the curvature dependent triangulation manages to capture the shape of the cube while using less vertices (Fig 4 (middle)).

It must be kept in mind that a shape template does not have to be obtained by the above procedure. Other triangulation algorithms as well as alternative surface representations could be used. For this work the template shape was selected randomly from the set of shape instances. This approach showed to be sufficient for this study. To avoid that the statistical analysis potentially depends on the choice of the template shape, a template that globally depends on all shape instances could be developed. One global template could be the logical union of all voxel images. One could also perform the statistical analysis twice: Once with a randomly selected template to obtain a mean shape and a second time with the obtained mean shape as template.

4. ADAPTATION OF THE TEMPLATE TO A SET OF OBJECTS

For a statistical shape analysis, it is necessary to obtain parameter vectors for a set of objects. The procedures described in the previous section resulted in a triangulation $T = (V, F)$ with vertices $V = \{\vec{v}_1, \vec{v}_2, \vec{v}_3, ..., \vec{v}_n\} \subset \Re^3$ and faces $F \subset V^3$. The set of vertices constitutes the point distribution model. The set of vertices can be used to obtain a 3n-dimensional parameterization $p = (\vec{v}_{1x}, \vec{v}_{1y}, \vec{v}_{1z}, \vec{v}_{2x}, \vec{v}_{2y}, \vec{v}_{2z}, ..., \vec{v}_{nx}, \vec{v}_{ny}, \vec{v}_{nz}) \in \Re^{3n}$ of the object. Given a set of object samples, performing a surface point density reduction and a triangulation of the objects independently would result in a set of parameter vectors with non-corresponding components. This is due to the fact that the i-th surface point of one object will in general not correspond to the i-th surface point of another object. In addition, the number of surface points and triangles (and thus the dimension of the parameter vector) will in general be different for each individual object. Those parameter vectors would be unsuitable for a statistical shape analysis. This is the well known correspondence problem. For a statistical analysis it is necessary that a given component of a set of parameter vectors correspond to a well-defined surface property, for example, a well-defined part of the surface in the case of a point distribution (see Fig. 5).

To overcome this problem, we propose to generate, for one object sample, a shape template and to "coat" this template onto the remaining set of object samples guided by user defined corresponding landmarks. The coating should preserve the one-to-one correspondence of the components of the parameter vectors. In this report, we propose the following "coating-scheme" to adapt a point distribution shape template (source) to another object of the same type (destination):

1. Define a set of corresponding landmarks on the source and destination object.
2. Deform the surface point distribution of the source object in such a way, that corresponding landmarks of the source and destination object match. The surface points will be mapped close to but not exactly onto the target surface.

3. Perform a mesh relaxation that moves the surface points on the target object in such a way that the point distribution resembles, as much as possible, the source mesh while keeping the points on the target surface.

The method described here is not completely automatic since the user has to manually define a set of landmarks on each object. This is, on the one hand, a disadvantage because the land mark definition is a time consuming task, especially in the case of a large training set. It is on the other hand an advantage, because it enables the user to define a point to point correspondence possibly based on anatomical knowledge in addition to the image information. A completely automatic locally affine surface registration method is described in [15]. The algorithm is based on the assumption that corresponding points have similar surface curvatures. However, surface curvature can vary strongly between anatomic object instances and is only one of several possible criteria for point correspondence.

4.1 Landmark-based mesh deformation

A user has to define pairs of landmarks $(\vec{r}_i, \vec{s}_i)$ in the source and the destination object space. This set of landmarks is used to obtain a mapping from the source object space to the destination object space. We used the scattered point thin-plate spline interpolation technique, which was introduced into the medical image processing domain by Bookstein [16]. The obtained spline function smoothly maps the source landmarks to the destination (see Fig. 6). The purpose of this procedure is to move the points in the neighborhood of a landmark similar to the way the landmark moves towards its corresponding target landmark.

Denoting $\vec{r}$ as a coordinate in the source object space, $\vec{F}(\vec{r})$ as the transformed coordinate and $\vec{r}_i$ as the i-th landmark position, the resulting coordinate transformation can be written as [17]:

$$\vec{F}(\vec{r}) = \vec{a}_0 + A \cdot \vec{r} + \sum_i \vec{w}_i \cdot U\left(\left|\vec{r} - \vec{r}_i\right|\right) \qquad (2)$$

The first part of this equation is a translation and the second part an affine transformation (defined by a matrix A). The third part sums all transformations of all landmarks and consists of the radial basis function $U(r)$ and the weight vectors $\vec{w}_i$.

The above transformation must map all landmarks from their source to their destination position $\vec{F}(\vec{r}_i) = \vec{s}_i$. These equations can be used to obtain the parameters of the above mapping.

The name of the method originates from the physical problem of a thin plate having minimal bending energy while being constrained at a number of points. The appropriate radial basis functions solving the thin-plate problem are $U(r) = r^2 \ln r$ and $U(r) = r \ln r$ [17] in the 2D and 3D case, respectively (Duchon functions). Since in our case there is no real physical problem to be solved, other functions could be chosen, as for example the often used functions of type $U^H{}_i(r) = \left(|r - r_i|^2 + s_i{}^2\right)^a$ with $a > 0$ and s_i being a "stiffness-radius", introduced by Hardy [18]. We have implemented both Duchon- and Hardy-type radial basis functions. The Duchon-type function has a larger long-distance influence and seems to work better for our purpose, as the aim is to influence the position of many vertices with only a few landmarks.

For an efficient, user-guided development of a shape model, it is desirable that the number of landmarks to be defined is as low as possible. It is not yet clear what the optimal number of landmarks is. In the case of the vertebra, we used 15 landmarks. In addition, the identification of suitable landmarks is dependent on the shape to be investigated and currently involves user-knowledge. It is imaginable that in the future some landmarks could be detected automatically, for example, based on surface curvature information in the binary image or based on gray-value curvature in the gray-value image [19].

4.2 Mesh relaxation

The landmark-guided warping leads to a rough adaptation of the template to the target object. However, the vertices of the triangulation are not necessarily mapped onto the destination object surface. To achieve a precise mapping, we use a distance map [20] to project every vertex to the nearest surface point. The projection might be omitted in the future, because the mesh relaxation step that follows the projection includes a surface term, pulling the vertices to the surface of the target object.

It can happen, that the projection of the vertices to the destination surface causes folds of the triangular mesh. In addition, irregular large or small triangles or misplaced vertices may appear. We introduced as a last step a mesh-relaxation based on a mass-spring model, to "unfold" the surface and to regularize the mesh, (see fig. 7). Each vertex is associated with a certain mass and the triangle-edges with springs. Denoting the current distance between vertex i and j with $d_{i,j}$, the distance between vertex i and j in the template mesh with $d_{i,j}{}^{0}$, the unit vector pointing from vertex i to j with $\vec{e}_{ij}$, and the current distance vector between vertex i and the target surface with $\vec{d}_{DM}$ (resulting from the distance map), the force acting on the i-th vertex is defined as:

$$\vec{F}_i = \sum_j k_1 (d_{i,j} - d_{i,j}{}^{0}) \vec{e}_{ji} + k_2 \vec{d}_{DM}(\vec{r}_i) \qquad (3)$$

The first term of this equation models the edges between vertices as springs of length $d_{i,j}$ and spring constant k_1 according to Hooks law [21]. Correspondingly, the second term resembles a spring force towards the surface of the object resulting from a spring of zero length with spring constant k_2. The equation of motion is given by $\vec{F}_i = m \cdot \vec{a}_i + g \cdot \vec{v}_i$ with $\vec{a}_i$ being the acceleration, $\vec{v}_i$ the velocity, and g the damping coefficient. The damping term is necessary for the mesh to assume an equilibrium position after a finite time. Solving the resulting differential equation for the coordinates of the vertices smoothes out the mesh and pulls the vertices towards the surface. Folds in the mesh can be detected using the difference between the triangle normals and the surface normals of the binary target

object. If a fold is detected, the surface force is switched off for the vertices of that particular triangle to allow the necessary distance from the surface during unfolding. The parameters k_1 and k_2 steer the interplay between surface and mesh force. A value of $k_2 > k_1$ assures that the vertices assume their equilibrium positions on the surface of the destination object. Figure 8 shows the template mesh, the landmark-based deformed mesh, and the relaxed mesh together with an axial cross-section of a vertebra. It is clearly visible, that the relaxation step is necessary for a good resulting fit between mesh and voxel surface.

5. STATISTICAL ANALYSIS

The output of the procedures described above, is a set of n shape parameter vectors $p_i = (p_{i,1}, p_{i,2}, ..., p_{i,3N}) \in \Re^{3N}$ with $i = 1,2,...n$. Each shape vector contains the coordinates of N surface points. All vectors are derived from one shape template and have thus the same dimension $3N$. The coordinates of the surface points are given in the object-related coordinate system defined by the second moment matrix and the center of mass of the binary voxel object. No other normalization, such as the 'Procrustes Analysis' [7], has been applied. Therefore, scaling is expected to be part of the natural variability and will be part of the model variability. After the coating procedure there is a one-to-one correspondence between the vector elements of a given index, in the sense that for all instances this vector element represents corresponding shape information. If the object instances used so far are representative for the class of objects under consideration, it should be possible to extract statistical information about the object variability. More specifically, we are interested in the mean shape vector

$$\overline{p} = \frac{1}{n}\sum_{i=1}^{n} p_i \tag{4}$$

and in the correlation between vector elements, as given by the empirical covariance matrix

$$S = \frac{1}{n}\sum_{i=1}^{n} (p_i - \overline{p})(p_i - \overline{p})^T \ . \tag{5}$$

If we write P for the $3N \times n$ matrix of centered parameter vectors, S can be written as $S = PP^T$. The principal component analysis of S yields the eigenvalues $l_1 \geq l_2 \geq ... \geq l_{3N} \geq 0$ and the corresponding eigenvectors q_i, $i = 1,2,...,3N$. The eigenvectors corresponding to the largest eigenvalues describe the most significant modes of variation in the parameter space. Here, it is assumed, that a small number of modes account for a large portion of the total variability. If the number of shape samples n is smaller than the shape vector dimension $3N$, only the first n eigenvalues of S are non-zero. In this case, instead of the $3N \times 3N$ eigenvalue problem, the associated $n \times n$ eigenvalue problem of the implicit covariance matrix $\tilde{S} = P^T P$ can be solved [22]. The first n eigenvalues and eigenvectors of S can be determined by the eigenvectors of $\tilde{S}$:

$$l_i = \tilde{l}_i\,, \quad q_i = \frac{P\,\tilde{q}_i}{|P\,\tilde{q}_i|}\,, \quad i = 1,...,n\,. \tag{6}$$

The model allows the generation of new shape samples by adding a weighted sum of modes to the mean parameter vector:

$$p = \bar{p} + Qb \tag{7}$$

with $Q = (q_1,...,q_m)$, $m \le n$ being the matrix of the first m modes and $b = (b_1,...,b_m)$ the weight vector. This can be used to explore the object shape space with a low parametric representation of the object, for example during a segmentation procedure. The function

$$n_{s^2}(k) = \sum_{i=1}^{k} \frac{s_i^2}{\sum_{j=1}^{3N} s_j^2} \tag{8}$$

evaluated for the variances s_i^2 of the coordinates of the surface points and for the eigenvalues l_i of the covariance matrix can be used to determine the concentration of variability that has been achieved by the eigenmode analysis (see also fig. 9).

6. RESULTS

Segmented Abdominal CT scans of lumbar vertebrae have been used for a first application of the model generation procedure. The model is based on 31 object instances (9xL1, 11xL2, 9xL3 and 2xL4) and contains roughly 600 surface points and 1200 triangles. For the landmark-based mesh deformation, 15 landmarks distributed over the vertebra have been used. The definition of the 15 landmarks took approximately 2 to 3 minutes per vertebra.

In Fig. 9, the magnitude of the eigenvalues (left) and the concentration of variability in the first n eigenvalues (right) according to (8) is shown. A variability of 90 % is already captured by the first 10 eigenvectors. From this it can be concluded the that the variability of the objects within the training set can considerably be captured by the linear approximation based on the eigenvectors of the co-variance matrix.

Fig. 10 shows axial, sagittal and coronal views of the mean shape and the first two eigenmodes of the generated model. The shapes are generated according to (7). For illustration the first and second entry of the weight vector are set to twice the natural standard deviation: $b_{1,2} = 2\sqrt{l_{1,2}}$, respectively. The first eigenmode is to some degree a size variation mode. This is due to the fact that size variation was not removed by a normalization procedure, but was instead included in the natural variability of the anatomical object. Both modes are more or less symmetric about the mid-sagittal plane. Other modes show asymmetric properties.

7. DISCUSSION

The results presented in the previous section are of preliminary nature. For a thorough statistical analysis, it would be desirable to base the evaluation on a larger set of object instances. Especially interesting is the question whether the eigenmodes of variation converge for an increasing number of shape samples. It is also not clear yet which portion of the inter- and intra-patient variability can be captured in one model. Is it reasonable to generate a shape model that includes all human vertebrae, or all lumbar vertebrae or just for example L2? Other criteria could be age, sex or specific diseases.

Despite of the preliminary results, it can be stated that the model generation scheme has been proven to work very satisfactory. The coating procedure succeeded in establishing a correspondence between the shape parameterization of the sample set. No significant variability (i.e. a variability introduced by the model generation scheme and not by the shape variability within the set of object samples) has been introduced into the model. To some extent this can be seen in figure 10. An artificial variability introduced for example by a mismatch of corresponding surface areas would result in deformation modes with local vertex movements being in contradiction to the large-scale shape deformation. In figure 6 and 7 corresponding triangles are colored equally. This allows to check if corresponding triangles match to corresponding surface areas. The deformation captured by the eigenvectors of the covariance matrix are smooth and describe the variability of the sample set accurately.

An important advantage of the model is the pure 3D character of the underlying shape representation. It is independent from slice orientation, voxel size and voxel anisotropy. Furthermore, it is not restricted with respect to connectedness or topology in general. Since the shape represenation is based on a set of surface points and a triangulation of this point set, it is easy to visualize and to manipulate. Due to the surface curvature driven triangulation procedure, the shape representation regards the local degree of surface detail.

It is important to note, that our method assumes a certain limit to the shape variability of the objects that are analyzed. The goal is to capture the natural variability of healthy vertebrae. Shape instances that are differ largely from the mean shape, such as pathological vertebrae, cannot and are not

intended to be captured by this scheme. A necessary condition for the method to succeed is that a sufficient number of landmarks must be well defined for all shape instances.

An important aspect of the model generation scheme is that it follows the human-in-the-loop philosophy which is especially important in the medical domain where the fault tolerance is low. The user has to define a number of landmarks in all data-sets. In addition, the input data consists of segmented data-sets, originating in our case from manual segmentations. On the other hand, an improvement of the procedure, especially in uncritical domains, would be a reduction of the necessary user interaction, especially in case of a large set of shape samples. This could be achieved by means of a bootstrap procedure. After a semi-automatic generation of a preliminary shape model, the model could be used during the coating procedure to replace the manual landmark placement by a rough automatic adaptation of the triangular mesh to the new shape instances. The mesh relaxation step would then allow to adapt the mesh to the new degree of variability.

8. CONCLUSIONS

In this article we presented a new approach for the generation of 3D statistical shape models. The method is valid for objects of arbitrary topology and is based on a triangulation of a thinned set of object surface points. From one object of a set of objects, a shape template is generated and adapted to the rest of the set. The adaptation ('coating') procedure consists of a rough landmark-based shape morphing and a mesh relaxation, and results in a set of shape parameter vectors. From this set, the mean parameter vector and the covariance matrix, as well as the eigenvector and eigenvalues of the matrix, are calculated. The eigenvectors are interpreted as modes of shape variation. The method has been successfully applied to a set of 31 lumbar vertebrae. As expected, a large portion of the total shape variability is already captured within the first few eigenvectors. However, the number of objects used so far is very small and not representative in any sense.

Future directions include the generation of a shape model based on a larger set of objects and the application of the model in a segmentation or registration procedure. An improvement of the model

generation procedure would be a reduction of the necessary number of landmarks to be identified interactively.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Dr. W. P. Th. M. Mali, Prof. Dr. B. C. Eikelboom and Dr. J. D. Blankensteijn, University Hospital Utrecht, for providing the CT images.

The algorithm was implemented on an Easy Vision workstation from Philips Medical Systems and we would like to thank Easy Vision Modules Advanced Development, Philips Medical Systems, Best, for helpful discussions and their support.

Figures:

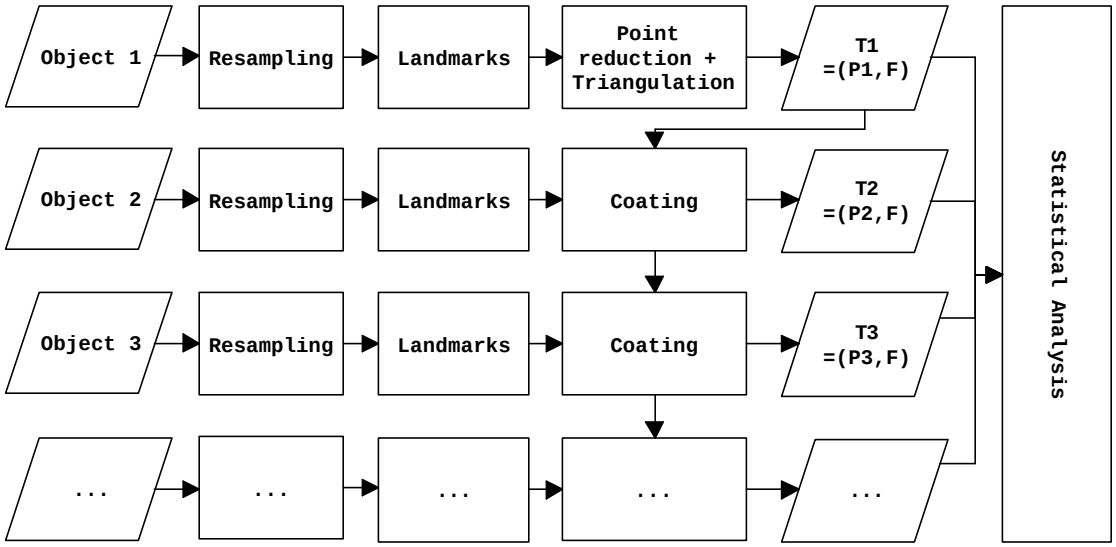

Figure 1: Shape model generation scheme. Input of the procedure is a set of binary 3D objects. The surface point density of one of the objects is being reduced and triangulated. The object shape is represented by a set of points (P1) and a face list (F) as a result of the triangulation. (P1,F) serves as a shape template for the rest of the objects. Each of the other objects are 'coated' with the template (P1,F). A statistical analysis of the point sets Pi can be used to create a statistical shape model of the set of objects.

Figure 2 : A two-dimensional schematic visualization of the point density reduction procedure.

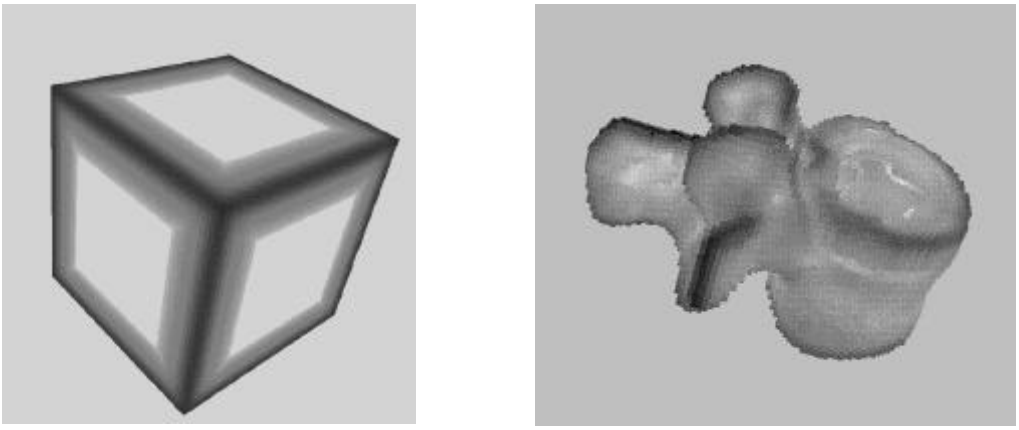

Figure 3: Gray-value encoded surface curvature of a cube and a vertebra as calculated using [12,13].

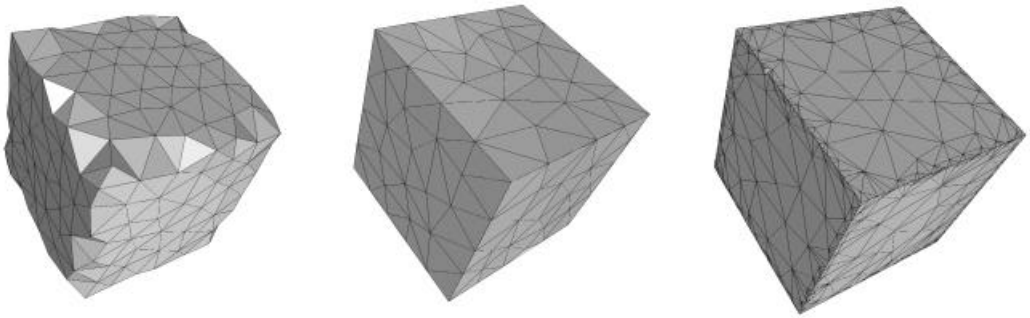

Figure 4 : Curvature guided triangulation of a cube. Left: Isotropic unguided triangulation. Middle: Isotropic triangulation, point selection curvature guided. Right: Point density and point selection curvature guided.

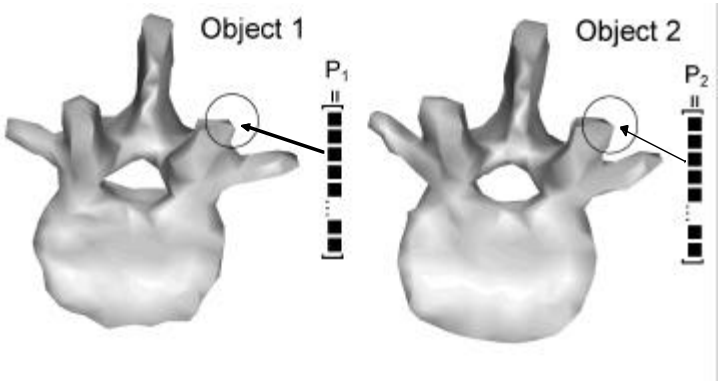

Figure 5: Illustration of the Correspondence problem of shape parameterizations. The number of parameters used to represent different shape instances as well as the type of information contained in a given parameter should be the same.

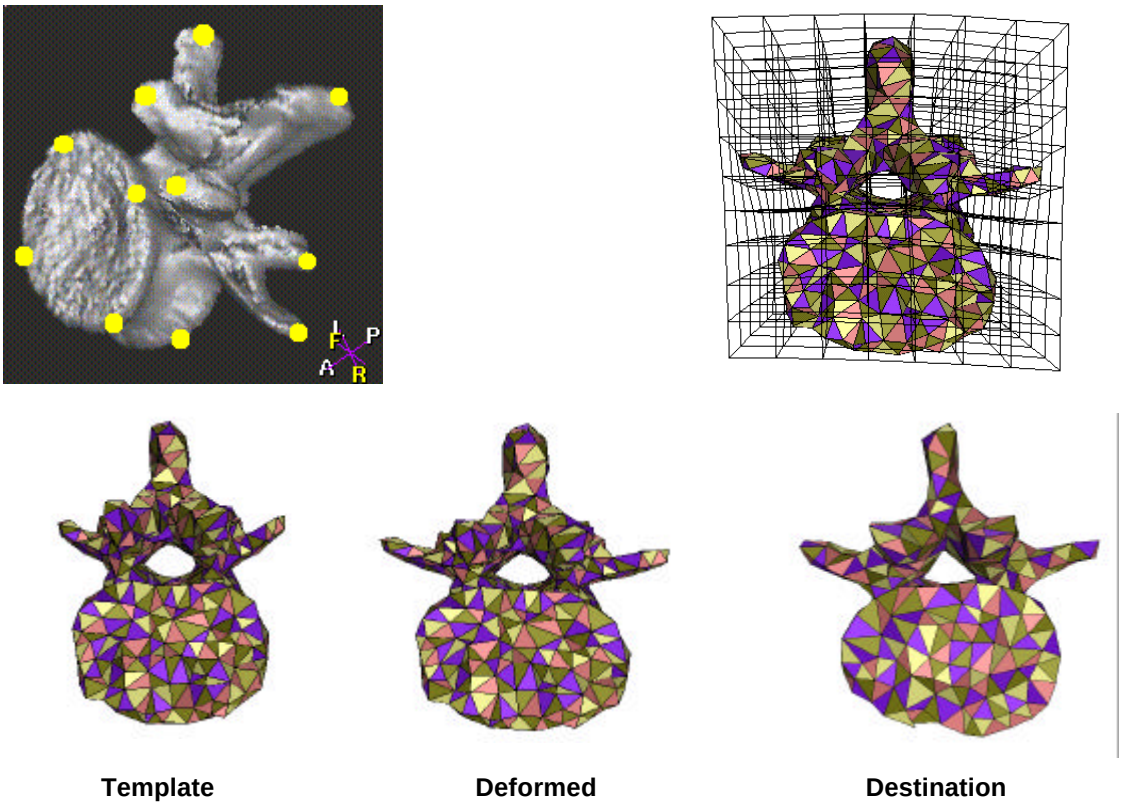

Figure 6: Landmark based mesh deformation. Top row: Based on a set of corresponding landmarks (left), a smooth deformation of the space (and the triangulation) can be achieved (right). The landmarks in the left figure show a part of the landmarks used for the shape model described in section 6.

Bottom row: The triangulated template is deformed to fit onto a destination object based on 15 landmarks. For comparison, the rightmost image shows an independent triangulation of the destination object that has no correspondence to the other two. The deformed triangulation corresponds to the template, while resembling the shape of the destination object. The coloring of the triangles is chosen in order to show the correspondence between triangulations.

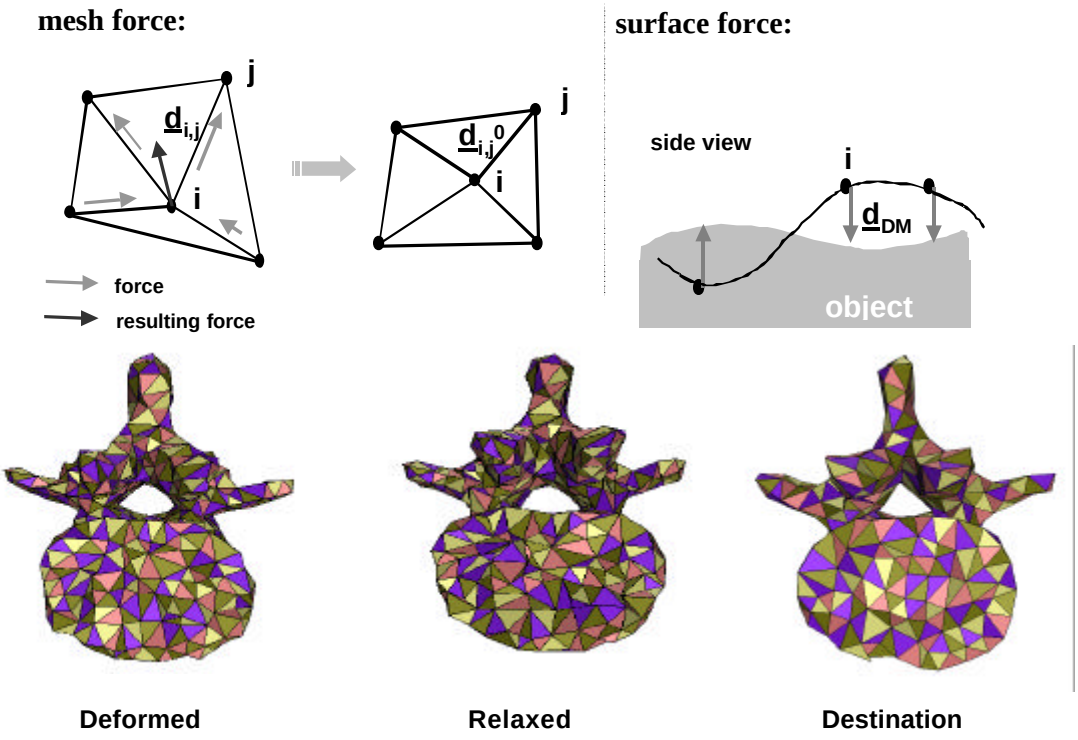

Figure 7: The mesh relaxation scheme. Input for the relaxation is the landmark based deformed mesh (bottom row, left). The forces acting on the vertices try to recover the proportional lengths of the template mesh edges and pull the mesh onto the destination surface (top row). The output of the relaxation procedure is displayed in the middle image of the bottom row. As in fig. 6, for comparison, the rightmost image of the bottom row shows an independent triangulation of the object.

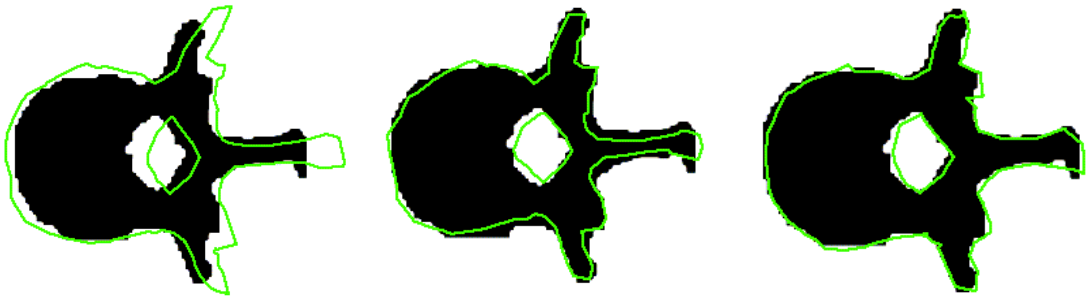

Figure 8: Axial cross-section of a segmented vertebra and the surface mesh in different stages of the coating process. The left image shows the rigid registration of the template object and the object to be coated using the second moment matrix and the center of mass of the binarized objects. The mesh after the landmark based deformation is displayed in the middle image. The surface of the object is roughly matched, but at positions far away from landmarks (see also fig. 6), deviations are visible. The right image shows the mesh after the relaxation step. The mesh closely follows the surface of the object everywhere. Slight deviations from the binary surface (ca. 10^4 voxels) are due to the triangular approximation based on approximately 600 vertices.

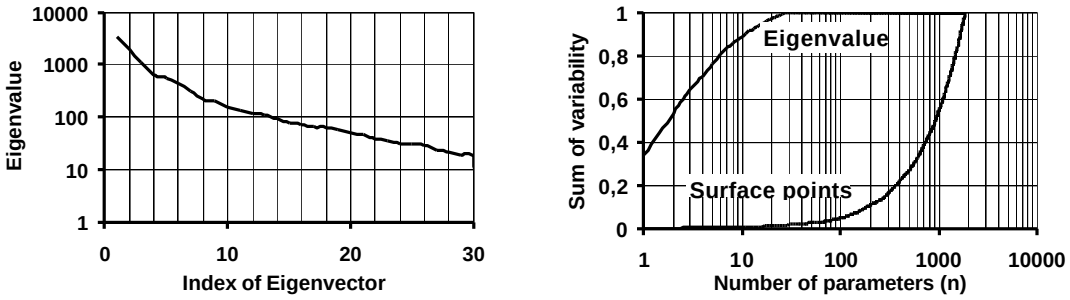

Figure 9: Variability captured by the shape model. Left: Magnitude of eigenvalues. Right: Portion of variability, captured by the first n parameters. Already the first 10 eigenvectors account for about 90 % of the total variability. Note the logarithmic scale of the diagrams.

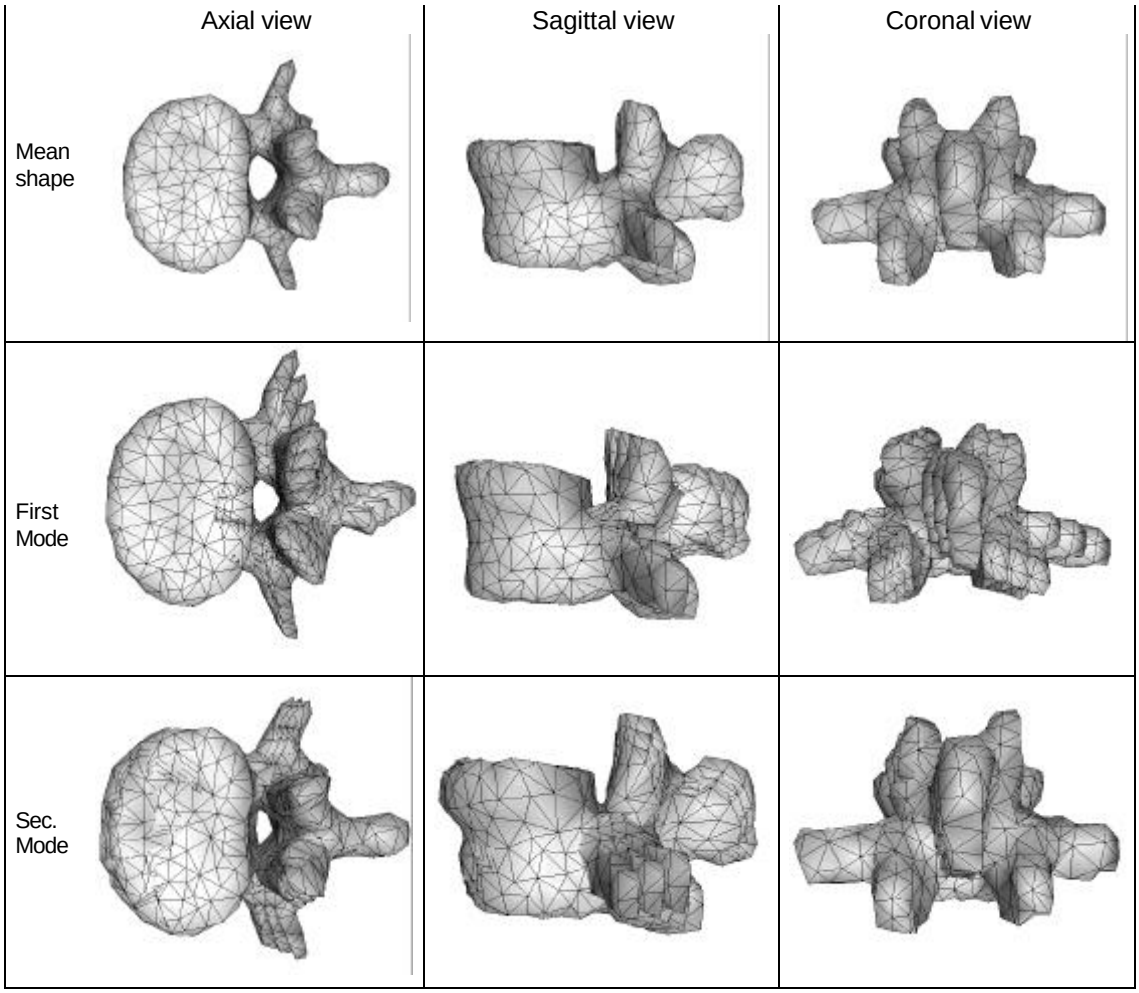

Figure 10: Visualization of the shape model. Top row: Axial, sagittal and coronal views of the mean shape . Second row: First eigenmode of the generated model visualized by an overlay of the mean shape and two deformed shapes according to the first eigenvector of the covariance matrix: $p = \overline{p} \pm 2\sqrt{l_1} \cdot q_1$. Third row: Mean shape and deformed shapes according to the second mode: $p = \overline{p} \pm 2\sqrt{l_2} \cdot q_2$. Both modes are more or less symmetric about the mid-sagittal plane. Compared to the variability of the processes the variability of the vertebral body is relatively small.

REFERENCES

1. H.Delingette. "Decimation of Isosurfaces with Deformable Models", *CVRMed-MRCAS '97, First Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, Grenoble, France, March 19-22, 1997,. Lecture Notes in Computer Science, Springer, Berlin, 1997, 83-92.
2. S.Sandor and R.Leahy. "Surface-based labeling of cortical anatomy using a deformable atlas", *IEEE Trans. Med. Imag*., 16 (1), 1993, 41-54.
3. S.Gibson, J.Samosky, A.Mor, C.Fyock, E.Grimson, T.Kanade, R.Kikins, H.Lauer, N.McKenzie, S.Nakajima, H.Ohkami, R.Osborne, and A.Sawada, "Simulating Arthroscopic Knee Surgery Using Volumetric Object Representations, Real-Time Volume Rendering and Haptic Feedback", *CVRMed-MRCAS '97, First Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, Grenoble, France, March 19-22, 1997, Lecture Notes in Computer Science. Springer, Berlin, 1997, pp. 369-378
4. D.Rueckert and P.Burger. "Shape-Based Segmentation and Tracking in 4D Cardiac MR Images", *CVRMed-MRCAS '97, First Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, Grenoble, France, March 19-22, 1997 Lecture Notes in Computer Science. Springer, Berlin, 1997, pp. 43-52
5. P.-Y. Bondiau and G.Malandain. "Eye Reconstruction and CT-Retinography Fusion for Proton Treatment Planning of Ocular Diseases", *CVRMed-MRCAS '97, First Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, Grenoble, France, March 19-22, 1997. Lecture Notes in Computer Science. Springer, Berlin, 1997, 705-714.
6. A.Neumann and C.Lorenz. "Statistical shape model based segmentation of medical images", *Computerized Medical Imaging and Graphics*, 22, 1998, 133-1413
7. L. Dryden and K. V. Mardia. *Statistical Shape Analysis*. Wiley, Chichester, 1998.

8. Hill, A. Thornham and C. J. Taylor, "Model-Based Interpretation of 3D Medical Images", *4th British Machine Vision Conference*, BMVA Press, Guildford, England, 1993, 339-348

9. D. Brett, M. F. Wilkins and C. J. Taylor, "A User Interface for 3D Statistical Model Building of Anatomical Structures", *CAR '97, Computer Assisted Radiology and Surgery, Proceedings of the 11th. International Symposium and Exhibition*, Berlin, 25-28 June 1997, Elsevier, Amsterdam, 1997, 246-251

10. Ch. Brechbühler, G. Gerig and O.Kübler, "Parameterization of Closed Surfaces for 3D Shape Description", *Computer Vision and Image Understanding*, **62** (2), 1995, 154-170

11. A. Rangarajan, H. Chui and F. Bookstein, "The Softassign Procrustes Matching Algorithm", *IPMI '97, Information processing in medical imaging, Proceedings of the 15th international conference, Poultney, Vermont, USA, June 9-13, 1997, Lecture notes in computer science*, vol. 1230, Springer, 1997

12. E. M. Stokely and S. Y. Wu, "Surface Parameterization and Curvature Measurement of Arbitrary 3-D Objects: Five Practical Methods", *IEEE Trans. On Pattern Analysis and Machine Intelligenc*, **14** (8), 1992, 833-840

13. Michael Wagner, Rekonstruktion neuronaler Ströme aus bioelektrischen und biomagnetischen Messungen auf der aus MR-Bildern segmentierten Hirnrinde, Dissertation, Technical University Hamburg-Harburg, Shaker, Aachen, 1998.

14. M. de Berg, M. van Kreveld, M.Overmars and O. Schwarzkopf, Computational Geometry, Algorithms and Applications, Springer, 1997.

15. Fred L. Bookstein, Thin-Plate Splines and the Decomposition of Deformations. *IEEE Trans. On Pattern Analysis and Machine Intelligence* **11** (6), 1989, 567-585

16. J. Feldmar and N. Ayache, "Rigid, Affine and Locally Affine Registration of Free-Form Surfaces", *International Journal of Computer Vision* 18(2), 99-119 (1996)

17. K. Rohr, H. S. Stiehl, R. Sprengel, W. Beil, T. M. Buzug, J. Weese and M. H. Kuhn, "Point-Based Elastic Registration of Medical Image Data Using Approximating Thin-Plate Splines", *VBC '96, Visualization in Biomedical Computing, Proceedings of the 4th International*

Conference, Hamburg, 22-25 Sept. 96, Lecture Notes in Computer Science 1131, K. H. Höhne and R. Kikinis (Eds.), Springer Berlin Heidelberg, pp. 297-306, 1996

18. R. L. Hardy, "Multiquadric equations of topography and other irregular surfaces", *J. Geophys. Res*, 76, 1971, 1905-1915

19. S. Frantz, K. Rohr and H. S. Stiehl, "Multi-Step Procedures for the Localization of 2D and 3D Point Landmarks and Automatic ROI Size Selection", *ECCV '98*, Lecture Notes in Computer Science 1406, H. Burkhardt and B. Neuman (Eds.), Springer Berlin Heidelberg, 1998, 1905-1915

20. G. Schmitz, J. Weese, W. Zylka and J. Sabczynski, "Registration by palpation for image guided surgery of the head", *CAR '97, Computer Assisted Radiology and Surgery, Proceedings of the 11th. International Symposium and Exhibition*, Berlin, 25-28 June 1997, Elsevier, Amsterdam, 1997, 305-310

21. H. Murakami and V. Kumar, "Efficient Calculation of Primary Images from a Set of Images", *IEEE Trans. On Pattern Analysis and Machine Intelligence*, **4** (5), 1982, 511-515

22. M. Alonso, E. J. Finn, "Physics", *Addison Wesley Publishing Company*, 1992