Available online at www.sciencedirect.com
ScienceDirect
Medical Image Analysis 11 (2007) 99–109
MEDICAL IMAGE ANALYSIS
www.elsevier.com/locate/media

Statistical deformable bone models for robust 3D surface extrapolation from sparse data

Kumar T. Rajamani [a], Martin A. Styner [b], Haydar Talib [a], Guoyan Zheng [a], Lutz P. Nolte [a], Miguel A. González Ballester [a,*]

[a] *MEM Research Center, Institute for Surgical Technology and Biomechanics, University of Bern, Stauffacherstrasse 78, CH-3014 Bern, Switzerland*
[b] *Departments of Computer Science and Psychiatry, Neurodevelopmental Disorders Research Center, University of North Carolina at Chapel Hill, USA*

Received 19 August 2005; received in revised form 25 April 2006; accepted 5 May 2006

Abstract

A majority of pre-operative planning and navigational guidance during computer assisted orthopaedic surgery routinely uses three-dimensional models of patient anatomy. These models enhance the surgeon's capability to decrease the invasiveness of surgical procedures and increase their accuracy and safety. A common approach for this is to use computed tomography (CT) or magnetic resonance imaging (MRI). These have the disadvantages that they are expensive and/or induce radiation to the patient. In this paper we propose a novel method to construct a patient-specific three-dimensional model that provides an appropriate intra-operative visualization without the need for a pre or intra-operative imaging. The 3D model is reconstructed by fitting a statistical deformable model to minimal sparse 3D data consisting of digitized landmarks and surface points that are obtained intra-operatively. The statistical model is constructed using Principal Component Analysis from training objects. Our deformation scheme efficiently and accurately computes a Mahalanobis distance weighted least square fit of the deformable model to the 3D data. Relaxing the Mahalanobis distance term as additional points are incorporated enables our method to handle small and large sets of digitized points efficiently. Formalizing the problem as a linear equation system helps us to provide real-time updates to the surgeons. Incorporation of M-estimator based weighting of the digitized points enables us to effectively reject outliers and compute stable models. We present here our evaluation results using leave-one-out experiments and extended validation of our method on nine dry cadaver bones.

Keywords: Medical image analysis; Shape analysis; Statistical shape models; Principal component analysis; Deformable models

1. Introduction

In surgery, the computerized visualization of three-dimensional patient data models has both pre- and intra-operative purposes. Pre-operatively, simulators may be used to train practitioners in basic surgical tasks as well as complete interventions. Patient specific models allow the practice of complex procedures prior to working with the patient directly. This could be used for effective diagnosis and procedural planning methods. Intra-operatively, it presents opportunities in navigation by augmenting the surgeon's view of the operating field with computer-generated data. The common approach to obtain 3D models is to use imaging techniques such as CT or MRI scans. These have the disadvantage that they are expensive and/or induce high radiation doses to the patient. Additionally a number of orthopaedic surgeries such as total hip arthoplasty (THA) and total knee arthoplasty (TKA) do not warrant a pre or intra-operative scan. The alternative is to build a statistical deformable model and adapt it to the patient anatomy.

Statistical shape analysis (Dryden and Mardia, 1998; Kendall, 1989; Small, 1996) is emerging as an important tool for understanding anatomical structures from medical

* Corresponding author. Tel.: +41 31 631 5950; fax: +41 31 631 5960.
E-mail address: Miguel.Gonzalez@MEMcenter.unibe.ch (Miguel A. González Ballester).

1361-8415/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2006.05.001

images. Statistical models give an efficient parameterization of the geometric variability of the anatomy. Model based approaches are popular (Kelemen et al., 1999; Turk and Pentland, 1991; Cootes et al., 1994) due to their ability to robustly represent objects. Intra-operative 3D anatomical visualization can be potentially achieved through the use of statistical shape models. Statistical model building consists of establishing legal variations of shape from a training population. The statistical model is then adapted, or *fitted* to the patient anatomy using intra-operatively digitized bone surface points. Thus the aim of statistical shape model fitting is to extrapolate from an extremely sparse set of 3D points a complete and accurate anatomical surface representation. This is particularly interesting for minimally invasive surgery (MIS), largely due to the operating theater setup. Statistical modeling technologies allow minimal intrusion on the surgical environment as 3D computerized models may be directly injected into the scene, enabling enhanced visualization.

Extrapolation via principal component analysis (PCA) based statistical shape models has been explored by several scientists. Fleute and Lavallée (1998) fit the deformable model surface to intra-operatively digitized point data via jointly optimizing deformation and pose. This technology developed by Fleute et al. has been clinically evaluated and these results have been published (Stindel et al., 2002). Chan et al. (2003) optimize deformation and pose separately using an iterative closest point (ICP) method. In our prior work (Rajamani et al., 2004b) we proposed to iteratively remove shape information coded by digitized points from the PCA model. The extrapolated surface is then computed as the most probable surface in the shape space given the data. Unlike earlier approaches, this approach was also able to include non-spatial data, such as patient height and weight. It is applicable for very small set of known points.

We present here a novel bone deformation method that can seamlessly handle both small and large sets of digitized points and provide real time interactivity. We have formulated the problem as a least squares error minimization with additional regularization terms that computes the Mahalanobis distance of the predicted model (Rajamani et al., 2004a). We solve for the shape parameters that minimize the residual errors between the reconstructed model and the cloud of random points. The novelty is that the Mahalanobis distance term enables stable prediction with minimal number of known surface points. In addition, the computation is performed in real time as shape parameters are determined by solving a single linear system. The formalization also enables the incorporation of the complete set of eigenvectors for the shape estimation. This scheme was then improved to have better convergence behaviour by having an additional parameter in the objective function that relaxes the Mahalanobis distance term as additional points are digitized (Rajamani et al., 2004c). As more information in terms of additional digitized points is received we relax the constraint on the surface to remain close to the mean and allow it to deform so that the error between the predicted surface and the set of digitized points is minimized as far as possible. Finally, the usage of M-estimators based weighting enables for a smart estimation mechanism that is robust to outliers.

Surface points are typically acquired by use of a tracked digitizing pointer. It can, due to limited surgical access, be difficult to acquire a set of points that sufficiently spans the patient's anatomy to ensure accurate shape prediction of a given statistical model. Hence we explored using ultrasound imaging for non-invasive intra-operative surface points digitization. We briefly illustrate the application of our deformable bone models concurrently with automatic segmentation of 2D B-mode ultrasound contours (Kowal et al., 2003), to provide for a rapid, automatic intra-operative visualization for navigation and planning especially in minimally invasive orthopaedic surgery.

This paper is structured as follows. Section 2 briefly describes model construction using principal components and outlines the method we chose for building our model. In Section 3 we describe in detail the evolution of our bone deformation algorithm. In Section 4 we present our evaluation results using leave-one-out experiments, extended validation of our method on plastic and dry cadaver bones and finally application of our deformation algorithm in conjunction with ultrasound contours resulting in rapid, automatic intra-operative visualization. We conclude by discussing the results and limitations and possible extensions of our work.

2. Statistical model construction

The first step is to build a statistical shape model from a training database. Several different geometric representations have been used to model anatomy. Bookstein (1986) uses landmarks to capture the important geometric features. The active shape model (ASM) of Cootes and Taylor (Cootes et al., 1995) represents an object's geometry as a dense collection of boundary points. Cootes et al. (1998) have augmented their statistical models to include the variability of the image information as well as shape. Kelemen et al. (1999) use a spherical harmonic (SPHARM) decomposition of the object geometry. Recent researchers are also exploring methods towards constructing a statistical shape model using nonrigid deformation of a template mesh (Heitz et al., 2005).

For our model building we have employed the representation of shapes using point distribution models (PDM). The basic idea is to compute the mean shape and to establish from the training set the pattern of legal variations in the shapes for a given class of images. This is achieved using principal component analysis (PCA) (Jolliffe, 1986). PCA defines a linear transformation that decorrelates the parameter signals of the original shape population by projecting the objects into a linear shape space spanned by a complete set of orthogonal basis vectors. The axes of the

shape space are oriented along directions in which the data has its highest variance. If the parameter signals are highly correlated, then the major variations of shape are described by the first few basis vectors. Furthermore, if the joint distribution of the parameters describing the surface is Gaussian, then a reasonably weighted linear combination of the basis vectors results in a shape that is similar to the existing ones.

A key step in this model building involves establishing a dense correspondence between shape boundaries over a reasonably large set of training data. In 2D, correspondence is often established using manually determined landmarks, but this is a time-consuming, error-prone and subjective process. In principle, the method extends to 3D, but in practice, due to very small sets of reliably identifiable landmarks, manual landmarking becomes impractical. Most automated approaches posed the correspondence problem as that of defining a parameterization for each of the objects in the training set, assuming correspondence between equivalently parameterized points. We compared the methods introduced by Brechbühler et al. (1995) (SPHARM), Kotcheff and Taylor (1998) (DetCov), Davies et al. (2002b) (MDL) and a fourth method based on manually initialized subdivision surfaces similar to Wang et al. (2000) (MSS). We analyzed both the direct correspondence via manually selected landmarks as well as the properties of the model implied by the correspondences, in regard to compactness, generalization and specificity. Our comparison study (Styner et al., 2003) of these popular correspondence establishing methods revealed that for modeling purposes the best among the correspondence methods was minimum description length (MDL) (Davies et al., 2002b). Based on the study, for our model building, correspondence was initialized using MSS and then optimized based on the MDL criteria.

The statistical shape model is constructed based on the established point correspondences. Each member of the training population is described by individual vectors $\vec{x}_i$ containing all 3D point coordinates. The aim of building this model is to use several training datasets to compute the principal components of shape variation. PCA is used to describe the different modes of variation with a small number of parameters. For the computation of PCA, the mean vector $\bar{\vec{x}}$ and the covariance matrix D are computed from the set of object vectors (1). The sorted eigenvalues λ_i and corresponding $\vec{p}_i$ of the covariance matrix are the principal directions spanning a shape space with $\bar{\vec{x}}$ representing its origin (2). Objects $\vec{x}_i$ in that shape space can be described as linear combinations with weights $\vec{b}_i$ calculated by projecting the difference vectors $\vec{x}_i - \bar{\vec{x}}$ into the eigenspace (3).

$$D = \frac{1}{n-1}\sum_{1}^{n}(\vec{x}_i - \bar{\vec{x}}) \cdot (\vec{x}_i - \bar{\vec{x}})^{\mathrm{T}} \tag{1}$$

$$P = \{\vec{p}_i\}; \quad D \cdot \vec{p}_i = \lambda_i \cdot \vec{p}_i; \tag{2}$$

$$\vec{b}_i = D^{\mathrm{T}}(\vec{x}_i - \bar{\vec{x}}); \quad \vec{x}_i = \bar{\vec{x}} + P \cdot \vec{b}_i \tag{3}$$

Fig. 1 shows the variability captured by the first two modes of variation of our proximal femur model varied by ±2SD.

3. Model deformation algorithm

The aim of this step is to recover the patient-specific 3D shape of the anatomy from the few available digitized landmarks and surface points. Our approach uses the shape model built earlier to infer the anatomical information in a robust way and provides the best statistical shape that corresponds to the patient. The key factor is the observation that objects in our shape space, and by our hypothesis the patient's 3D shape, can be described as the mean shape plus a weighted linear combination of eigenvectors. The problem is therefore formulated as estimating the weights for this unknown shape, such that the errors between the reconstructed model and the cloud of digitized surface points is minimized.

Our model fitting algorithm is formulated as a least squares error minimization with additional regularization terms that computes the Mahalanobis distance of the predicted model (Rajamani et al., 2004a). The Mahalanobis distance term enables stable prediction with minimal number of known surface points. Where Fleute (Fleute and Lavallée, 1998) and Chan (Chan et al., 2003) consider a truncated set, we include the complete set of eigenvectors, or shape variations, without exorbitant increase in the computation time. The method consists of two steps:

- Initially a small point-set of anatomical landmarks with known correspondence to the model is digitized. This is used to register the patient anatomy to the model. This also provides an initial estimation of the 3D shape with only a few digitized points.
- To improve the prediction additional points can be interactively incorporated via closest distance correspondence.

The objective function that we minimize is defined as follows

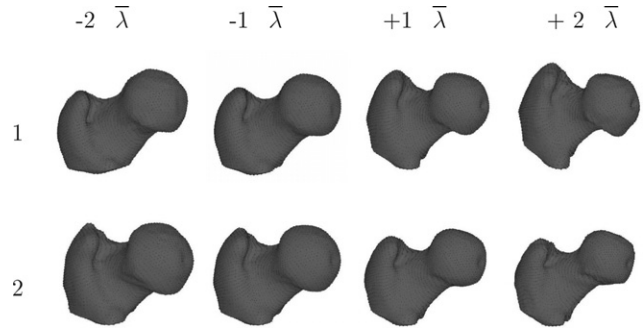

Fig. 1. The first two eigenmodes of variation of our model built from 30 segmented proximal femoral surface data. Each individual surface in the model and the shapes generated are described by a sparse triangle mesh list containing 4098 vertices. The shape instances were generated by evaluating $\bar{\vec{x}} + \omega\sqrt{\lambda_i}\vec{p}_i$ with $\omega \in \{-2, \ldots, 2\}$.

$$f = \rho\left\{\sum_{\substack{k=1\\ j=\mathrm{Index}(k)}}^{N} w_k \left\| \vec{Y}_k - \left(\vec{\overline{X}}_j + \sum_{i=1}^{m} \alpha_i \vec{p}_i(j)\right)\right\|^2\right\} + (1-\rho)\left\{\sum_{i=1}^{m} \frac{\alpha_i^2}{\lambda_i}\right\} \quad (4)$$

The first term in the function is the Euclidean distance between the N digitized points $\vec{Y}$ and the estimated shape comprising the mean $\vec{\overline{X}}$ plus a weighted sum of the eigenvectors $\vec{p}_i$. The corresponding point for each of the digitized points $\vec{Y}_k$ is computed using closest point correspondence from the current estimated shape. This is denoted as $\vec{\overline{X}}_j$, where $j = \mathrm{Index}\ (k)$ is the index of the closest point corresponding to the kth digitized point. The second term is the Mahalanobis distance of the predicted shape from the mean and controls the probability of the predicted shape. This term ensures that the predicted shapes are valid by favoring those that are closer to the mean.

The parameter ρ in the objective function enables the deformation scheme to have better convergence behaviour (Rajamani et al., 2004c). This is enabled by relaxing the effect of the Mahalanobis distance term as additional points are digitized. This makes the surface less constrained to remain close to the mean and allows it to more freely deform. Hence the error between the predicted surface and the set of digitized points is better minimized. Since the error typically decreases exponentially, we chose ρ to increase logarithmically with the number of digitized points, and was therefore defined according to the following equation

$$\rho = \begin{cases} 0.5 & N \leqslant 6 \\ \dfrac{\log\left\{\frac{N}{\mathrm{MaxN}}(ge-1)+1\right\}}{2\log(ge)} + 0.5 & N > 6 \end{cases} \quad (5)$$

where N is the number of digitized points, MaxN is the total number of points in the model, g is a factor which determines the rate of growth of ρ. To achieve faster growth rate for ρ, g was empirically set to be the number of members in the population.

The parameter w_k enables the realization of stable predictions and the robust rejection of outliers (Rajamani et al., 2005). Instead of using a box-filter based rigid threshold to reject the outliers we decided to employ a smarter and robust M-estimator based weighting mechanism that creates an inverted valley function as shown in Fig. 2. This helps us to effectively reject outliers and compute stable models. Incorporating the M-estimator based weight-function analytically in the objective function, would make the system to be solved non-linear as there is no differentiable, linear weight function that is an M-estimator. Instead we decided to include the computed weights in the objective function as single independent constants for each digitized point (computed based on its distance from the closest point in the current estimated shape). These weights could be updated iteratively

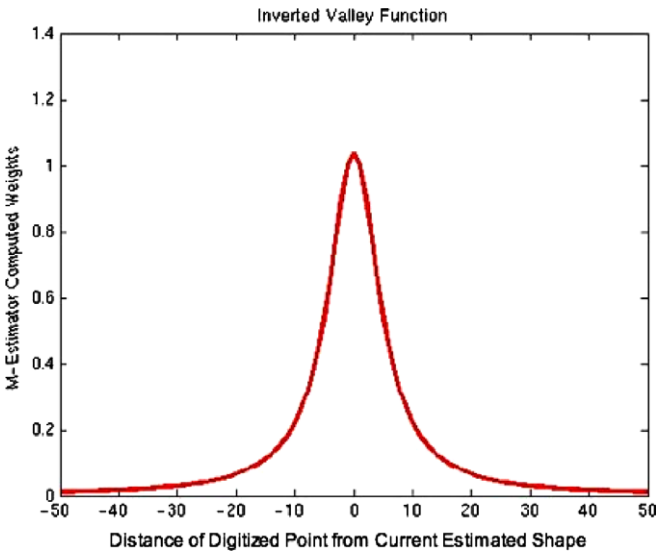

Fig. 2. The inverted valley function that was generated using the M-estimator based weights and used in the objective function to achieve outlier resistance.

in an ICP like mechanism until there are no more significant changes in the weights. This feature was incorporated by having the weighting parameter w_k in the objective function, adapted from Styner et al. (2000) and defined according to the following equation

$$w_k = [(1 - \mathrm{dist}^2)/(\mathrm{dist}^2 + \mathrm{sp})] + 1, \quad \text{where} \quad (6)$$

$$\mathrm{dist} = \left\| \vec{Y}_k - \left(\vec{\overline{X}}_j + \sum_{i=1}^{m} \alpha_i \vec{p}_i(j)\right)\right\|^2 \quad (7)$$

i.e. dist is the euclidean distance between the digitized surface point and its closest point in the current estimated shape and sp is a positive scalar parameter that defines the sharpness or steepness of the inverted valley function. The greater the value of sp, the sharper the valley is and the harder the outliers are disregarded. To have a gentle handling of outliers that lie within a radius of 10 mm the value of sp was empirically fixed in our application at 27.

We briefly explain here our solution strategy, where we formulate the problem as a linear equation system. To determine the shape parameters α that best describe the unknown shape, the function f is differentiated with respect to the shape parameters and equated to zero. Differentiating f with respect to α_n yields

$$\frac{\partial f}{\partial \alpha_n} = \rho \sum_{k=1}^{N} w_k \frac{\partial}{\partial \alpha_n}\left\| \vec{Y}_k - \left(\vec{\overline{X}}_j + \sum_{i=1}^{m} \alpha_i \vec{p}_i(j)\right)\right\|^2 + (1-\rho)\frac{2\alpha_i}{\lambda_n} \quad (8)$$

f is differentiated with respect to each of the α and for each of the resulting equations collating the different α terms, and dividing throughout by $2\rho w_k$ yields a linear equation system of the form $A\alpha = b$ with A being

$$\sum_{k=1}^{N}\begin{pmatrix} \vec{p_1}(j)\cdot\vec{p_1}(j)+\frac{1-\rho}{\rho w_k}\frac{1}{\lambda_1} & \cdots & \vec{p_n}(j)\cdot\vec{p_1}(j) & \cdots & \vec{p_m}(j)\cdot\vec{p_1}(j) \\ & & \vdots & & \\ \vec{p_1}(j)\cdot\vec{p_n}(j) & \cdots & \vec{p_n}(j)\cdot\vec{p_n}(j)+\frac{1-\rho}{\rho w_k}\frac{1}{\lambda_n} & \cdots & \vec{p_m}(j)\cdot\vec{p_n}(j) \\ & & \vdots & & \\ \vec{p_1}(j)\cdot\vec{p_m}(j) & \cdots & \vec{p_n}(j)\cdot\vec{p_m}(j) & \cdots & \vec{p_m}(j)\cdot\vec{p_m}(j)+\frac{1-\rho}{\rho w_k}\frac{1}{\lambda_m} \end{pmatrix} \tag{9}$$

The unknowns in our system are $(\alpha_1 \cdots \alpha_n \cdots \alpha_m)$. Collating the constant terms yields b, the right hand side of our system as follows

$$\sum_{k=1}^{N}\begin{pmatrix} (\vec{Y_k}-\vec{\overline{X_j}})\cdot\vec{p_1}(j) \\ \vdots \\ (\vec{Y_k}-\vec{\overline{X_j}})\cdot\vec{p_n}(j) \\ \vdots \\ (\vec{Y_k}-\vec{\overline{X_j}})\cdot\vec{p_m}(j) \end{pmatrix} \tag{10}$$

This results is a $m \times m$ linear equation system over α. This is solved using standard linear equations system solvers using QR decomposition. This novel solution strategy enables real time estimation of shape parameters hence facilitating incorporation of the complete set of eigenvectors for shape estimation.

4. Results

The primary application that we focus on is hip surgery such as total hip replacement (THR) and knee surgery such as total knee arthroplasty (TKA) and anterior cruciate ligament surgery (ACL). Hence we began by concentrating on the proximal femur. Our database comprised of 30 CT scans of patient hips (image resolution was $0.684 \times 0.684 \times 1.0$ mm). The datasets were anonymized when they were provided from the hospital. The details of age, height and gender were not available. There were several datasets that were acquired from the hospital and only the 30 femurs that were intact were used in the model building. The proximal femurs were segmented and surface models of the bones were extracted for the statistical model construction. Each individual surface was described by a sparse triangle mesh list containing 4098 vertices. The triangle mesh was the direct result of a octahedron based subdivision scheme after parametrically mapping the surface models onto the sphere.

Dense correspondence between points on the surface of the bones in the training database was initialized with a semiautomatic landmark-driven method and then optimized using the Minimum Description Length criterion (Davies et al., 2002a) to construct a compact optimal model. Three anatomical landmarks, the femoral notch and the upper and the lower trochanter are used as the first set of digitized points. This is used to initially register the model to the patient anatomy. This first set of points is also used for computing the bounding box of the shapes in our databse and the bounding box of the set of digitized points. This aids in scale normalization of the shapes prior to modeling and normalizing the predicted shape, making our method size invariant. The remaining points are added uniformly across the parameterization so that they occupy different locations on the bone surface. We first demonstrate proof of principle of our method using leave-one-out tests and then detail validation studies on cast and dry cadaver bones.

4.1. Leave-one-out experiments

A series of leave-one-out experiments was carried out to evaluate our method. Surface points were chosen uniformly from the surface model of the left out object so that they occupy different locations on the bone surface. Point correspondence were established by finding the closest point, in Euclidean distance terms.

Fig. 3 shows an example, with mean surface error of 1.44 mm obtained with 20 digitized points. The color-coded 3D rendering is calculated using Hausdorff's Distance to measure the distance between discrete 3D surfaces (Aspert et al., 2002). We present in Fig. 4 the reconstruction errors of this femur using different number of points picked from the surface of the bone. The maximum, 95-percentile, median and mean error with standard deviation are plotted against the number of digitized points. With 20 digitized surface points the mean surface error was 1.44 mm and the error drops down to 1.17 with 80 digitized surface points. The relatively large difference between the maximum error and 95 percentile error is explained mainly by the large distances at the bottom of the meshes, induced by different locations of the cutting plane used to define the proximal part of the femur.

The reconstruction errors of 10 different femurs using leave-one-out experiments with 10 digitized points and 90 digitized points are potrayed in Fig. 5. The maximum, 95-percentile, median and mean error with standard deviation are plotted for each femur. In our leave-one-out experiments the predicted models mean surface error ranged

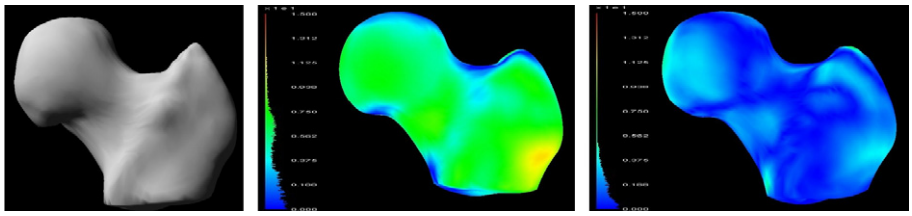

Fig. 3. Left: A typical proximal femur of the population that was used in the leave-one-out test. Middle: The average shape of the population with color coded distance map to the actual shape. The mean surface error is 3.37 mm and the median surface error is 2.65 mm. Right: The shape based on only six digitized points with color coded distance map to the actual shape. The mean surface error is 1.50 mm and the median surface error is 1.25 mm.

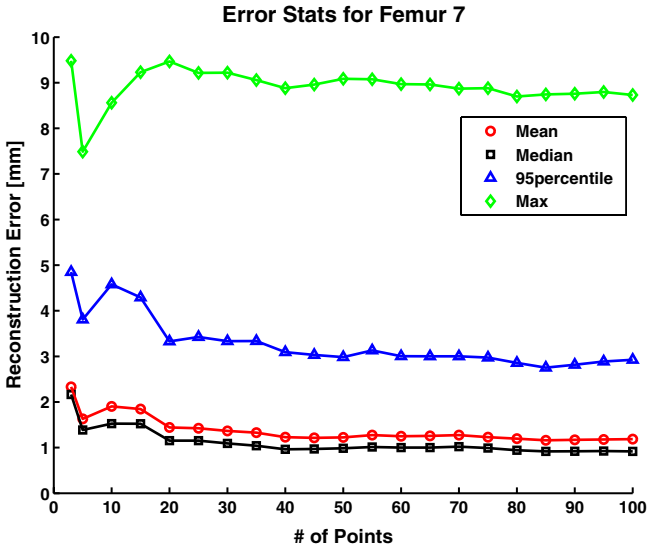

Fig. 4. Reconstruction errors for a given femur using different number of points picked from the surface of the bone. The maximum, 95-percentile, median and mean error with standard deviation are plotted against the number of digitized points.

from 0.96 to 2.59 mm with 10 digitized surface points, and the error range decreased to 0.88–2.59 mm with 90 digitized surface points. In two of the experiments the errors actually increased with more digitized surface points. The leave-one-out experiments helped us evaluate the proof-of-concept of our method and showed that we can extrapolate a three dimensional shape from sparse data.

4.2. Influence of ρ parameter

As explained earlier, the ρ parameter modulates the contribution of the Mahalanobis shape distance term. Our formulation makes ρ a factor of the number of digitized surface points and hence it is adapted automatically as more points are digitized. To evaluate the influence of the ρ factor we studied the performance of our deformation algorithm with and without the ρ parameter. Fig. 6 shows the effect of the ρ factor studied on 10 different femurs using leave-one-out experiments. The shapes are estimated with and without the ρ factor and the median errors are plotted for 30 digitized surface points. There is an improvement in accuracy using the ρ factor as can be deduced from the plots. Incorporating ρ ensured better convergence and the error factor gain in seven out of the ten femurs. For three of the femurs there was no influence, but the errors did not increase. The influence of the ρ factor is more prominent when a larger population is used to build the model. This is evident from our previous study (Rajamani et al., 2004c) of the role of the ρ parameter in Hippocampus model generated from 172 hippocampus instances, where the error factor gain was about 10%.

4.3. Effect of w_k parameter

A series of experiments was carried out to evaluate our method with regards to robustness to outliers, such as those obtained by accidental activation of the foot pedal for point digitization. A proximal cast femur with attached reference base was used for this experiment. Tracking was done by using an in-house navigation environment and maintaining an optimal distance of around 2 m to the optical tracking camera (Optotrack™, NDI, Waterloo, Canada). The accuracy for such tracking systems when used ex vitro with exposed fiducial markers is lesser than 1 mm. Fig. 7 shows screen shots of our method, when the plastic femur was estimated using surface points digitized using a calibrated navigated pointer. In the first run, 12 surface points were digitized which comprised four outliers. In spite of the large set of outliers a stable prediction was realized. The second run had fewer outliers among digitized 32 surface points. Our experiments verified that the M-estimator based weighting function was very successful in robustly eliminating outliers and enabling stable predictions.

4.4. Cadaver validation study

Nine different dry cadaver femur bones were chosen for this validation study. High-resolution CT scans of these bones were segmented (image resolution: $0.652 \times 0.652 \times 1.0$ mm) and fine 3D surface models were generated. The experiment trials were carried out in the CT coordinate system. The three anatomical landmarks and additional 51 bone surface points were digitized on the surface model of each of the cadaver bones. The deformation procedure was then employed to estimate the 3D model that best approximates the digitized set of points.

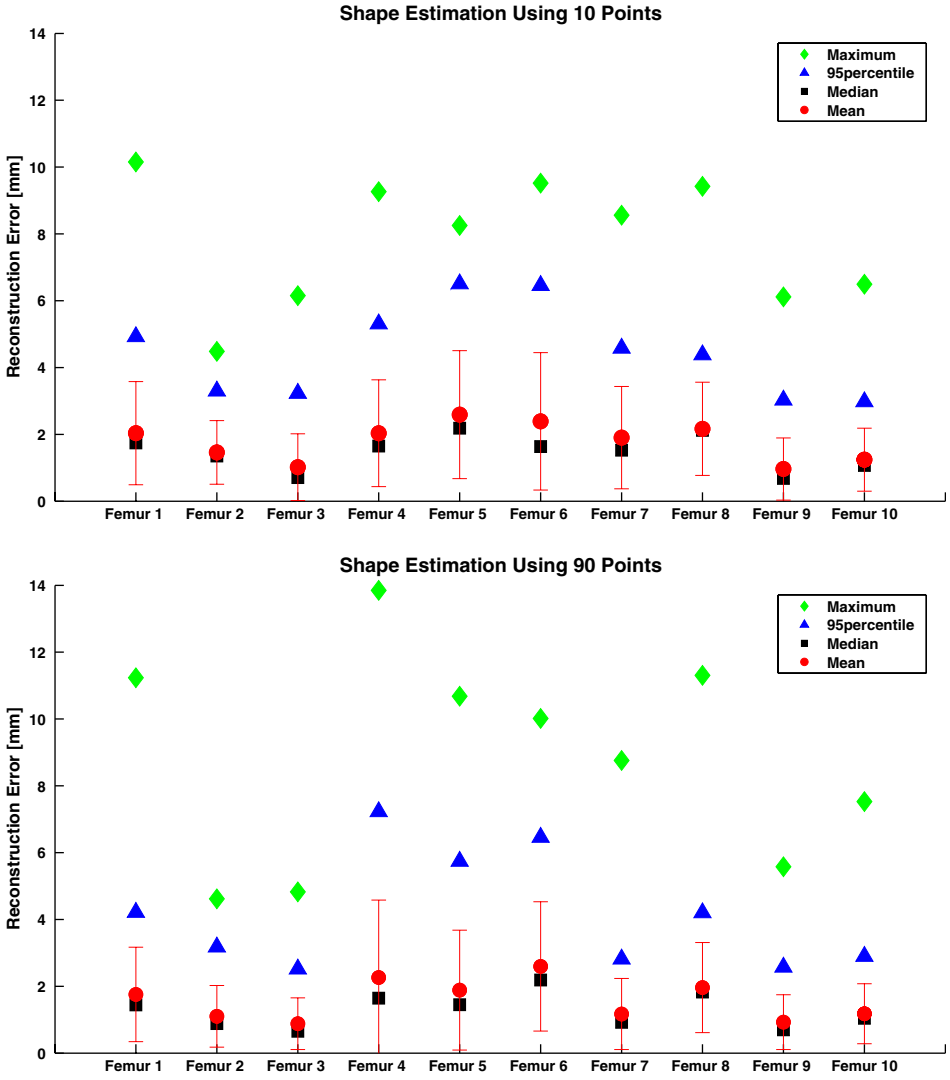

Fig. 5. Reconstruction errors of 10 different femurs using leave-one-out experiments with 10 digitized points (top) and 90 digitized points (bottom). The maximum, 95-percentile, median and mean error with standard deviation are plotted for each femur.

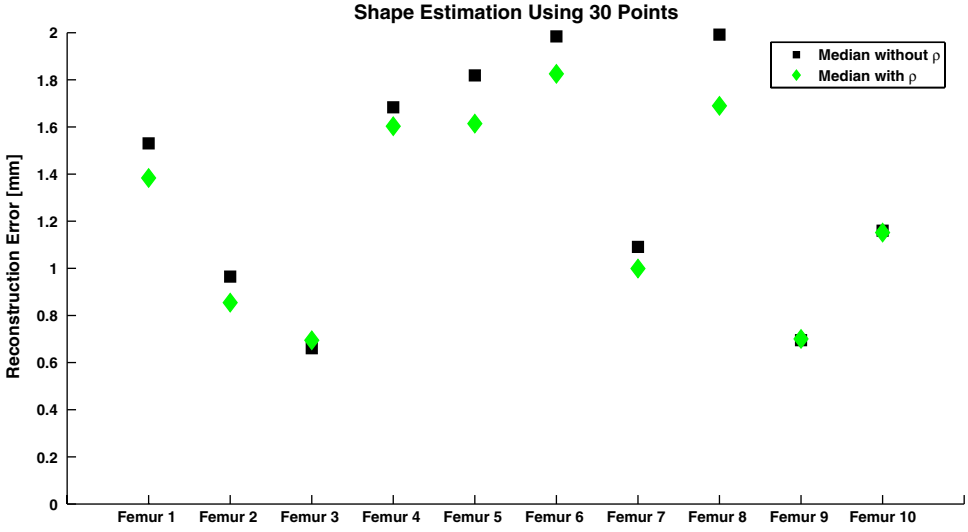

Fig. 6. Effect of the ρ factor studied on 10 different femurs using leave-one-out experiments. The shapes are estimated with and without the ρ factor and the median errors are plotted for 30 digitized surface points.

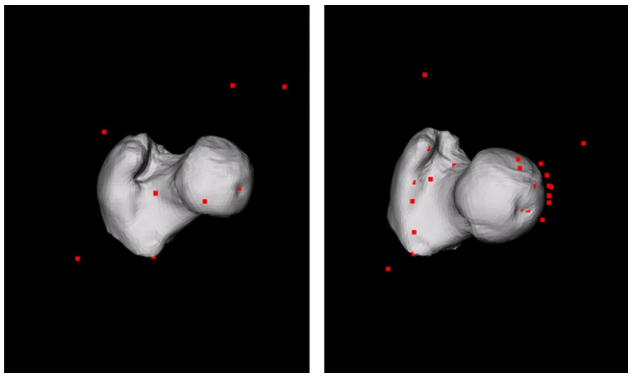

Fig. 7. Owing the screen shots of our method, when the shape of a plastic femur was estimated using surface points digitized using a pointer in our navigation environment. As can be seen the outliers are well eliminated and a stable prediction is realized.

We carried out the experiments on two models, built from different initial training population. The first model was constructed from the entire 30 proximal femurs and the second model was constructed from a subset of 14 proximal femurs, with correspondence optimized across the respective training sets. This helped us to evaluate the effect of training size on the deformation algorithm. Table 1 shows the error results for each of the cadaver bones with different number of digitized surface points using the larger and smaller population. The mean surface error with 3, 27 and 54 selected surface points are tabulated. Fig. 8left shows the surface model of the estimated 3D shape for one of the cadaver bones with color coded distance map to the actual shape. The mean error here with 54 digitized points was 0.85 mm and a median of 0.66 mm using the larger population of 30 proximal femurs. Fig. 8right shows the cumulated statistics across all the cadaver bones. The average of the mean and median errors across the entire set of nine cadaver bones is plotted against the number of digitized points for both the models generated from the smaller and larger population. The average mean surface error with 10 digitized points lies between 2.1 and 2.6 mm and with 54 digitized points the error is 1.7–1.9 mm. The results for predicting the cadaver bones are in the same error range as the leave-one-out experiments. The cadaver experiments helps us conclude that we can indeed estimate quite accurately the 3D shape of an anatomy even with very sparse information.

Table 1
Mean surface errors for nine dry cadaver bones with 3, 27 and 54 selected surface points in the CT-based error scheme

Cadaver bone number	Mean error (mm) w.r.t. number of points					
	Large population			Small population		
	3	27	54	3	27	54
1	2.08	1.90	1.72	2.57	2.02	1.85
2	0.96	0.91	0.85	2.03	1.49	1.23
3	2.44	2.28	2	3.02	2.69	2.50
4	2.55	2.45	2.03	2.92	2.63	2.12
5	2.18	1.99	1.85	1.98	1.87	1.72
6	3.49	3.1	2.54	4.44	3.79	2.65
7	1.73	1.59	1.39	3.15	2.61	2.23
8	2.01	1.87	1.67	1.91	1.75	1.58
9	2.06	2.04	1.83	2.22	2.14	1.64
Average	2.17	2.01	1.76	2.69	2.33	1.95

The errors are tabulated for both the experiments with the larger and smaller population.

4.5. Ultrasound-Initialized deformable bone models

Two different cast proximal femurs were chosen for this study, and CT scans were performed (image resolution was $0.391 \times 0.391 \times 1.0$). Their CT surface models (approximately 5000 vertices) were registered into the anatomy's co-ordinate space using paired point matching and refined using surface matching (Gong et al., 1997) integrated into our in-house optical-tracking navigation system, yielding a registration error of 0.2 mm for this experiment. The registered surface models were considered as "gold" refer-

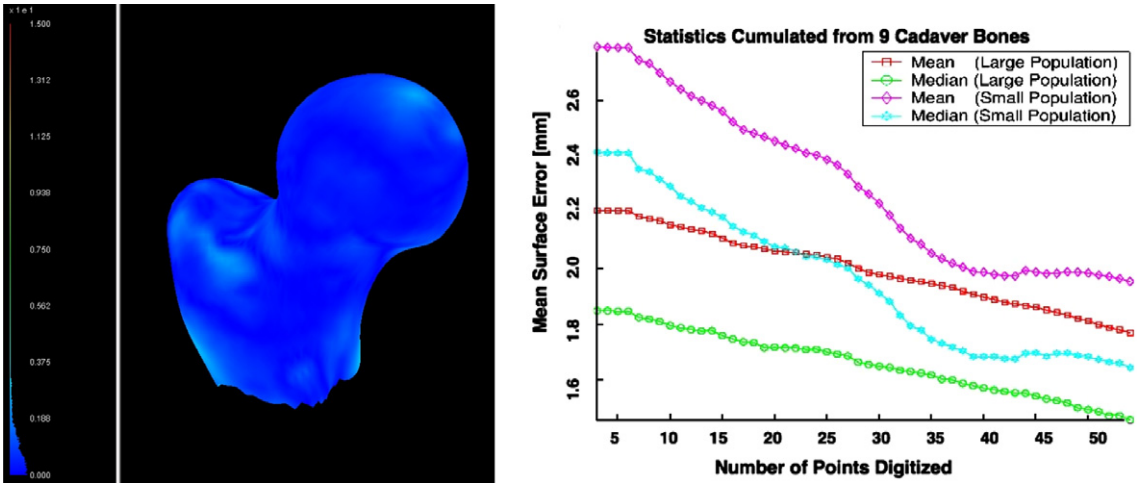

Fig. 8. Left: The surface model that was estimated for one of the cadaver bones with color coded distance map to the actual shape. The mean error with 54 digitized points was 0.85 mm and a median of 0.66 mm. Right: Error statistics cumulated across all the cadaver bones. The average mean and median errors are plotted against the number of digitized points for both the models generated from smaller and larger population.

ences, used for error measurements (computed with Mesh (Aspert et al., 2002)) of the predicted bone models.

To initialize as well as provide surface points for our bone deformation method a classical ultrasound system (Kontron Sigma 330™, Basel, Switzerland) with a 7.5 MHz linear array transducer, was used for the experiments. The system provides conventional B-Mode imaging. Calibration was achievable in less than 5 min, using a minimum of ultrasound images, and has a high reported accuracy (Kowal et al., 2003). The bone contours used in our experiment were automatically segmented from the image planes of the calibrated tracked 2D B-mode probe, thereby yielding a cloud of points in the co-ordinate space of the anatomy. Our automatic segmentation approach requires an average of 0.8 s of computation for each ultrasound image frame and has a mean accuracy of 0.42 mm (Kowal et al., 2001).

The cloud of segmented ultrasound points were provided as input to the deformation algorithm. The result was a predicted model in the anatomy's co-ordinate space. To obtain stable estimates of the errors we performed a series of five trials per bone, the results of which are tabulated in Table 2. The results in this table show the mean and median surface errors for the predicted shapes with respect to the "gold" references for each bone, using 24–26 digitized surface points. They also include averaged mean and median surface errors for each scenario, to help gauge the repeatability of each experiment. Fig. 9 shows for each bone one case of ultrasound based shape prediction, with predicted shape overlaid to its respective "gold" reference. From the results we can see that ultrasound imaging could be used along with our deformation algorithm to estimate an appropriate 3D model of the anatomy.

Table 2
Mean and median surface errors from their actual surfaces in five different trials for the two cast femur bones, using ultrasound to acquire surface points

Trial number	Bone 1 error (mm)		Bone 2 error (mm)	
	Mean	Median	Mean	Median
1	3.88	3.56	2.94	2.12
2	4.37	4.28	5.30	4.98
3	6.88	6.34	3.79	3.48
4	4.75	4.51	4.57	4.54
5	3.08	2.53	3.12	2.84
Average	4.59	4.35	3.95	3.59

5. Discussion

In this paper we have presented a novel anatomical shape deformation technique to predict the three-dimensional model of a given anatomy using statistical shape models. The proposed shape deformation is especially attractive in the scenario of sparse set of surface points, and can also seamlessly handle small and large sets of digitized points, which is an innovative concept. We have shown that we can robustly estimate a realistic patient-specific three-dimensional model of a given anatomy. Our formulation of the problem as a Mahanalobis weighted least squares error minimization, and the novel solution scheme by linear system solving, enable us to rapidly generate 3D models for visualization. The running times of the algorithm on a Pentium 4 machine with 512 MB of RAM, running SuSe Linux 9.0, are in the range of 1–2 s. Our formulation also enables incorporating the complete set of eigenvectors.

The ρ parameter helps us to relax the probability term to get a better estimate as more points are digitized. The effect of the ρ parameter is not significantly noticed in the case when the population size is small. This is because the error gets stabilized and uniform after the first few points are digitized and there is not much information that could be extracted by adding additional points in this case. Hence the ρ factor seems not to contribute much as was observed in the proximal femur model with a population size of only 14 members. On the contrary in the study using the hippocampus population (Rajamani et al., 2004c) the effect of the ρ parameter was significantly visible and it contributes in a

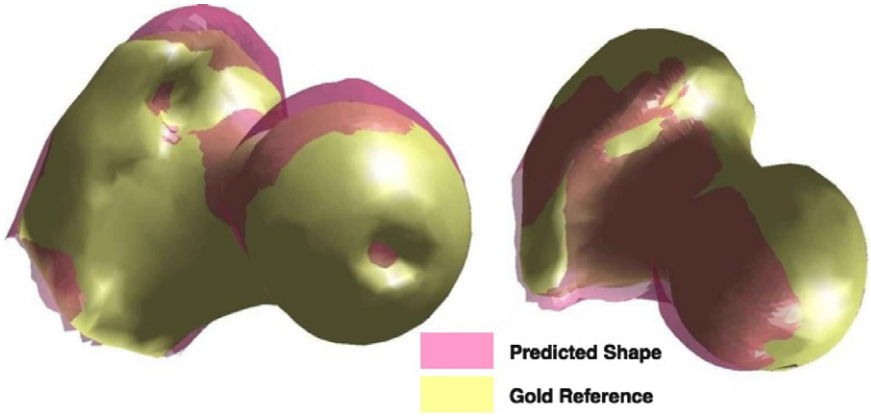

Fig. 9. Ultrasound-based prediction: Predicted models overlaid onto "gold" references. Bone 1 (left): 3.08 mm mean error and Bone 2 (right): 2.94 mm mean error.

significant way to decrease the error and achieve better convergence.

Rectangular or box filters are quite easy to incorporate for outlier rejection, but they do not smartly handle the outliers and are quite rigid. Outlier resistance based on Gaussian function seems to be an option but they do not fall off sharply enough. The novel usage of M-estimators based weighting enables for a smart estimation mechanism that is robust to outliers. Direct incorporation of the M-estimator into the function definition makes the minimization problem non-linear and not easy to solve. Hence we incorporated the M-estimating weight function into the fitting function. This gives us all the advantages of the M-estimator and also we do not lose the linearity of the problem. Normalization of the shapes prior to modeling and normalizing the digitized data makes our method size invariant. Our experiments on the cast proximal femur data show that our method can robustly reject outliers.

We also see that there is indeed a dependency of the deformation algorithm on the size of the training population. The larger population as expected estimates the shape better in most of the cases. In seven out of the nine bones the larger population has better estimation properties. In two of the cases where the smaller population seems to outperform, the difference between the two is in the order of 0.1 mm. It has to be emphasized that our population is still not large enough to capture all the possible variability of the shape. It could very well be the case that for a much larger population the shape variability is better captured and estimates even these two bones better than the smaller population. It is to be expected, therefore, that an increase in population size would increase the accuracy of the method.

As a natural extension we have explored the use of ultrasound imaging for non-invasive intra-operative surface points digitization. Ultrasound effectively solves the problems posed by limited surgical access and is an ideal way to acquire points from otherwise inaccesible regions. We have seen above that ultrasound initialized deformable bone models in our experimental conditions can provide a stable and repeatable prediction for bone visualization.

The results that we have achieved in our different experiments show great promise for the potential of our method to be applicable in clinical settings. We are confident that our method can be utilized already for clinical visualization applications, where we can provide 3D models from very limited sets of digitized points. We also consider that we are on the right track towards the ultimate use of our shape estimation procedure for surgical guidance. Current error ranges for shapes estimated from few landmark points are not quite in the range of surgical usability, which is typically in the area of 1.5 mm average error (Livyatan et al., 2003). This accuracy could be obtained by using more densely sampled clouds of digitized points, which would increase the amount of information used for model fitting.

The proposed technology brings a variety of advantages to orthopaedic and other surgical procedures, such as improved accuracy and safety, often reduced radiation exposure, and improved surgical reality through 3D visualization. In particular navigation based on shape deformation opens the door to more minimally invasive approaches. Future work will include the expansion of our training set of bone shapes to build more comprehensive shape models. We will also target other anatomies like the distal femur, the entire femur and the spine. Additionally, the order and the location of the surface points that are provided to the shape deformation algorithm might have an influence on the overall estimation result. Evaluation and quantification of this dependence could help in identifying the best ordered set of landmarks and surface points for a given application. Regarding the use of ultrasound for shape estimation, we are also working on improving the methods for bone detection to improve the accuracy.

Acknowledgements

We thank Sarang Joshi from University of North Carolina at Chapel Hill, USA for insightful discussions about modeling and shape prediction. The CT datasets were provided by Frank Langlotz from MEM Research Center, Bern Switzerland and the segmentation and MSS preprocessing of the CT data were done by Gabor Zsemlye from ETH Zurich, Switzerland. The MDL tools were provided by Rhodri H. Davies and Chris Taylor from the University of Manchester, UK. This work was supported by the Swiss National Science Foundation project Computer Aided and Image Guided Medical Interventions (NCCR CO-ME).

References

Aspert, N., Santa-Cruz, D., Ebrahimi, T., 2002. Mesh: Measuring errors between surfaces using the Hausdorff distance. In: Proceedings of the IEEE International Conference on Multimedia and Expo, vol. I, pp. 705–708.

Bookstein, F.L., 1986. Size and shape spaces for landmark data in two dimensions. Statistical Science 1, 181–242.

Brechbühler, C., Gerig, G., Kübler, O., 1995. Parametrization of closed surfaces for 3-D shape description. Computer Vision and Image Understanding 61 (2), 154–170.

Chan, C.S., Edwards, P.J., Hawkes, D.J., 2003. Integration of ultrasound-based registration with statistical shape models for computer-assisted orthopaedic surgery. In: Proceedings of SPIE Medical Imaging, pp. 414–424.

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

Cootes, T.F., Edwards, G.J., Taylor, C.J., 1998. Active appearance models. In: Proceedings of European Conference on Computer Vision, pp. 484–498.

Cootes, T.F., Hill, A., Taylor, C.J., Haslam, J., 1994. The use of active shape models for locating structures in medical images. Image and Vision Computing 12 (6), 355–366.

Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J., 2002a. 3D statistical shape models using direct optimisation of description length. In: Proceedings of European Conference on Computer Vision, pp. 3–20.

Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J., 2002b. A minimum description length approach to statistical shape modeling. IEEE Transaction on Medical Imaging 21 (5), 525–537.

Dryden, I., Mardia, K., 1998. Statistical Shape Analysis. John Wiley and Sons.

Fleute, M., Lavallée, S., 1998. Building a complete surface model from sparse data using statistical shape models. In: Lecture Notes in Computer Science, Proceedings of Medical Image Computing and Computer Assisted Interventions, vol. 1496, pp. 879–887.

Gong, J., Bächler, R., Sati, M., Nolte, L.-P., 1997. Restricted surface matching: a new approach to registration in computer assisted surgery. In: Proceedings of CVRMed-MRCAS, pp. 597–605.

Heitz, G., Rohlfing, T., Maurer, C., 2005. Statistical shape model generation using nonrigid deformation of a template mesh. In: Proceedings of SPIE Medical Imaging, vol. 5747, pp. 1411–1421.

Jolliffe, I., 1986. Principal Component Analysis. Springer.

Kelemen, A., Székely, G., Gerig, G., 1999. Elastic model-based segmentation of 3D neuroradiological data sets. IEEE Transactions on Medical Imaging 18 (October), 828–839.

Kendall, D., 1989. A survey of the statistical theory of shape. Statistical Science 4 (2), 87–120.

Kotcheff, A., Taylor, C., 1998. Automatic construction of eigenshape models by direct optimization. Medical Image Analysis 2, 303–314.

Kowal, J., Amstutz, C., Caversaccio, M., Nolte, L., 2003. On the development and comparative evaluation of an ultrasound B-mode probe calibration unit. Computer Aided Surgery.

Kowal, J., Amstutz, C., Nolte, L., 2001. On B-mode ultrasound based registration for computer assisted orthopaedic surgery. In: Proceedings of the 6th International Symposium on CAOS.

Livyatan, H., Yaniv, Z., Joskowicz, L., 2003. Gradient-based 2-D/3-D rigid registration of fluoroscopic X-ray to CT. IEEE Transactions on Medical Imaging 22 (11), 1395–1406.

Rajamani, K., González Ballester, M., Nolte, L., Styner, M., 2005. A novel and stable approach to anatomical structure morphing for intraoperative 3d visualization. In: Proceedings of SPIE Medical Imaging, vol. 5744, pp. 718–725.

Rajamani, K.T., Joshi, S., Styner, M., 2004a. Bone model morphing for enhanced surgical visualization. In: Proceedings of IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 1255–1258.

Rajamani, K.T., Nolte, L., Styner, M., 2004b. Bone morphing with statistical shape models for enhanced visualization. In: Proceedings of SPIE Medical Imaging, pp. 122–130.

Rajamani, K.T., Nolte, L., Styner, M., 2004c. A novel approach to anatomical structure morphing for intraoperative visualization. In: Proceedings of Medical Image Computing and Computer Aided Interventions, pp. 478–485.

Small, C., 1996. The Statistical Theory of Shape. Springer.

Stindel, E., Briard, J., Merloz, P., Plaweski, S., Dubrana, F., Lefevre, C., Troccaz, J., 2002. Bone morphing: 3D morphological data for total knee arthroplasty. Computer Aided Surgery 7 (3), 156–168.

Styner, M., Brechbühler, C., Székely, G., Gerig, G., 2000. Parametric estimate of intensity inhomogeneities applied to MRI. IEEE Transactions on Medical Imaging 19 (3), 153–165.

Styner, M.A., Rajamani, K.T., Nolte, L.P., Zsemlye, G., Szekely, G., Taylor, C.J., Davies, R.H., 2003. Evaluation of 3D correspondence methods for model building. In: Proceedings of Information Processing in Medical Imaging, pp. 63–75.

Turk, M., Pentland, A., 1991. Eigenfaces for recognition. Journal of Cognitive Neuroscience 3 (1), 71–86.

Wang, Y., Bradley, S.P., Staib, L.H., 2000. Shape-based 3D surface correspondence using geodesics and local geometry. In: Proceedings of Computer Vision and Pattern Recognition, pp. 644–651.