Automatic Segmentation of Mandibles in Low-Dose CT-Data

H. Lamecker[a], S. Zachow[a], A. Wittmers[a], B. Weber[a]

H.-C. Hege[a], B. Elsholtz[b], M. Stiller[b]

[a]*Zuse Institute Berlin, Germany*
[b]*University Hospital Berlin, Charité, Germany*

…

Abstract. Computer assisted planning in cranio-maxillofacial surgery often requires the segmentation and reconstruction of the mandibular bone from CT data. A common imaging modality is cone-beam volumetric tomography, which requires only low doses of radiation yet suffers from small signal to noise ratio and strong artefacts in the presence of metal. This work explores the ability of model-based segmentation using a 3D statistical mandible model for automatic segmentation in such data. Apart from the statistical model, a key ingredient for this method is the deformation strategy for detecting the mandibular bone. Quantitative results support the feasibility of the proposed approach.

Keywords: model-based segmentation, statistical shape model, principal component analysis, correspondence problem, implantology, dental surgery, computer assisted planning

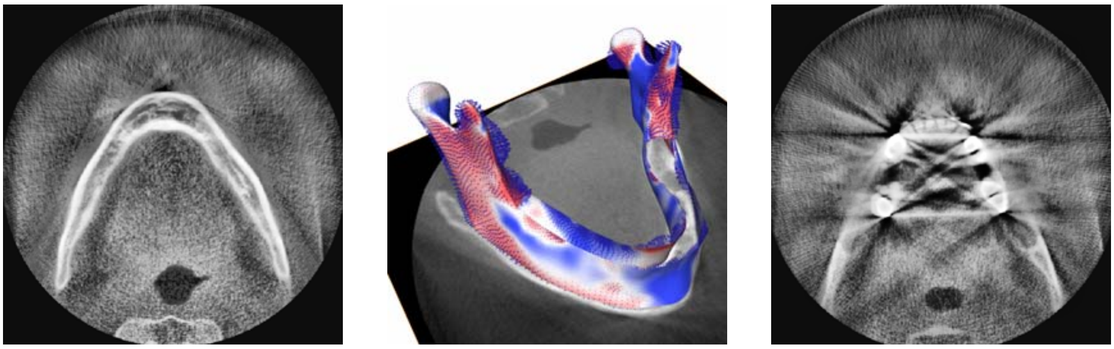

Fig. 1: Left and right: slices from low-dose CBVT, middle: deformable statistical shape model

1. Introduction

Cone-beam volumetric tomography (CBVT) is a widely available technology and a suitable foundation for three-dimensional diagnoses and planning in cranio-maxillofacial surgery [1]. Such scanners can operate with a significantly reduced patient's exposure to radiation compared to conventional CT. As a side effect of the low dose, however, such images are often noisy, and metal artefacts are present (Fig. 1, left and right).

One of the major applications for CBVT is dental imaging. The surgical procedure of placing dental implants requires careful preoperative planning. The surgical plan is guided by prosthetic considerations and anatomical structures, which limit the volume into which implants can be inserted.

The basis for computer aided planning is a segmentation of the mandibular bone in the data. Due to the problems mentioned, this task is difficult to automate with low-level techniques, such as thresholding, region growing or morphological filters. We intend to solve this problem by incorporating a-priori knowledge about the expected anatomical shape of the mandible into a (model-based) segmentation process.

2. Methods

The basic idea of model-based segmentation [2] is to capture the anatomical variability of normally developed mandible shapes in a statistical atlas (training phase) and to match this atlas to a given CT data set via a deformable model approach (segmentation phase).

2.1 Training Phase

The main challenge in performing statistical analysis on a set of 3D shapes lies in the correct identification of anatomically corresponding points on each training surface. Among the many existing approaches for solving this problem, we adopt the method of consistent patch decomposition and parameterization to establish correspondence between different individual surfaces [3]. Each training shape is decomposed into a number of corresponding regions interactively. Each of these regions is then mapped consistently to a common base domain under the constraint of minimizing metric distortion. Concatenating these parameterizations directly yields the desired correspondence map.

As a result of this process, all training shapes can be represented in a common vector space of dimension 3n, where n is the number of sample points used to discretize the shapes (vertices of the surfaces). Principal component analysis (PCA) on this set of vectors provides a compact representation of the variability within the training set, resulting in a linear model (average shape v plus the main modes of variation P within the sample): $S(b,T) = T(v + Pb)$. The variables (degrees of freedom) of this model are the weights b of the eigenmodes of the covariance matrix (shape weights) and a linear transformation T.

3.2 Segmentation Phase

Segmentation using a statistical model can be formulated as a registration problem. Let R denote the surface of the shape in the data to be segmented. Then we must solve the minimization problem:

$$(b^*, T^*) = \arg\min_{b,T} \| R - S(b,T) \|^2$$

The final segmentation is an approximation to R given by $R^* = S(b^*, T^*)$. In the segmentation setting, the location and shape of R is encoded implicitly in the image data I to be segmented. Therefore, one must derive a model for this encoding. This model certainly depends on the image data I and of the shape to be segmented. It can be incorporated in the segmentation in the following way: let $R^i = S(b^i, T^i)$ denote some segmentation at time i. Furthermore let ΔR^i be a vector field defined on the current segmentation R^i describing the desired deformation as implied by the underlying model. We proceed by solving iteratively the minimization problems

$$(b^{i+1}, T^{i+1}) = \arg\min_{b,T} \| (R^i + \Delta R^i) - S(b,T) \|^2$$

until some suitable stopping condition is met. Thus the segmentation strategy becomes an iterative procedure of deforming the shape model to the image data I. The minimization is performed separately either with respect to the position parameters T (position adjustment) or with respect to the shape parameters b (shape adjustment). This amounts to solving a linear system, if the discrete deformation field ΔR^i is defined on the vertices of the mesh. In order to prevent extreme shapes from emerging the allowed range for the shape parameters is constrained, e.g. to the minimal and maximal values as derived from the training data.

The segmentation process is initialized by positioning the shape model into the image data I. This can be done automatically, e.g. by placing the shape model in the centre of the bounding box of the image data. At each iteration of the segmentation process, several position adjustments are performed until no further significant improvement is achieved. Then a single shape adjustment is applied. This process is repeated until convergence. To improve the robustness a multilevel strategy has been implemented where the number of shape modes in the weighted least squares approximation is increased successively. At the first level only a few important modes are fitted to the data. For each subsequent level additional modes are considered until the maximum number of modes is reached.

3.2 Deformation Model

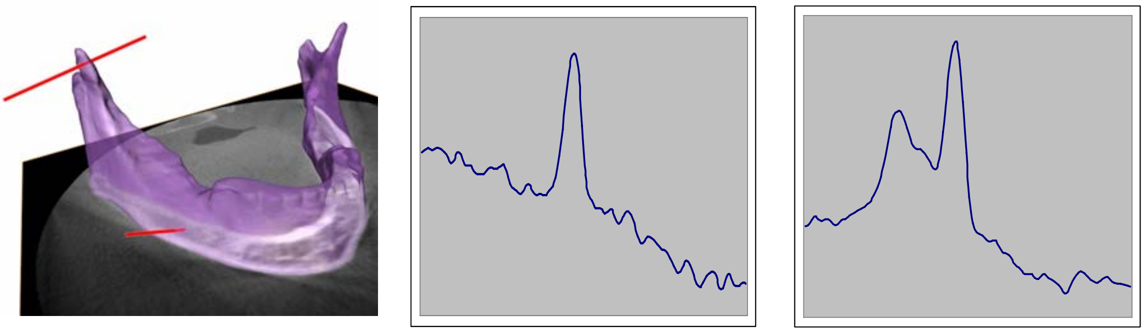

Fig. 2: Design of deformation model. Left: two profiles ($L = 4$cm) normal to the surface, middle: profile in upper mandibular region (single peak), right: profile in lower mandibular region (twin-peak).

The design of the deformation model ΔR^i for segmenting mandibular bone from low-dose CT data is based on the analysis of one dimensional profiles along the normal of the shape model's surface. Typically, such profiles show two distinct peaks that indicate the transition from the callus and the marrow of the bone (see Fig. 2). In the upper mandibular regions the bone is rather thin and often exhibits only a single strong peak. Thus, the deformation vector ΔR^i at a given vertex is determined as follows:

- sample a profile normal to the current surface R^i of length L (such that half of the profile is located on the inside and on the outside of the surface)
- apply smoothing on the profile, e.g. by gauss or median filtering
- detect the two major peaks
- if the magnitude of the peaks differs significantly, discard smaller of the two
- move outwards from the right peak to the point of inflection

Fig. 1 shows a deformation vector field ΔR^i for a given intermediate segmentation R^i. The (signed) magnitude $|\Delta R^i|$ is color-coded on the surface, arrows indicate the deformation vectors. The teeth region is not considered in this analysis, because of metal artefacts inherent in this region.

3. Results

3.1 Statistical Shape Model

Up to now, the training set of shapes is composed of 13 individual mandible shapes, reconstructed interactively from conventional CT data (Fig. 4). Each surface is decomposed into 8 regions in a symmetric manner (previously described in [4]): lower corpus mandibulae, teeth region, ramus mandibulae and the caput mandibulae (Fig. 3). The teeth region (two patches) is not considered in the statistical analysis because, in general, teeth differ in their number and the topology of the reconstructed geometry (contact vs. non-contact) from patient to patient. Moreover, the deformation model for the segmentation is not designed to include this region.

Fig. 3: Patch decomposition of mandibular bone for the construction of the statistical shape model.

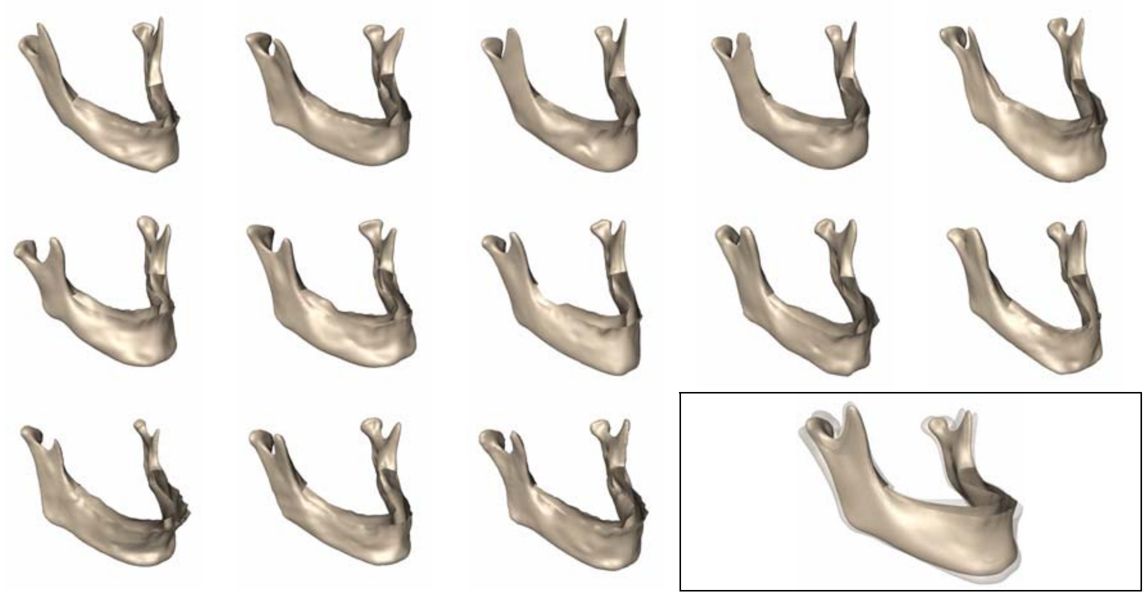

Fig. 4: Training set of mandibles (without teeth region), lower right: average shape plus first variation.

3.3 Surface Reconstruction of the Gold-Standard

For the evaluation of the segmentation process 15 data sets from a NewTom DVT scanner were available. These were segmented interactively by anatomical experts and serve as the gold standard for the evaluation of the proposed methods. For comparison, the statistical model is directly fitted to the surface R reconstructed from the segmentation of each of 15 data sets by solving the minimization problem

$$(b^*, T^*) = \arg\min_{b,T} d\{R, S(b,T)\}^2$$

Here $d(S, S')$ denotes the (squared) surface distance between S and S'. For details of solving this optimization problem refer to [3]. Results are shown in Fig. 5.

3.2 Segmentation Strategy and Results

Initially the shape model is placed in the centre of the bounding box of the CT data. The number of sampled points along each profile is kept fixed at all times and the segmentation process is divided into two phases. In the first phase, only the lower corpus mandibulae is matched to the CT data, i.e. the deformation vectors of all other patches equalled 0, L was set to 5 cm. This yields a good initialization for the second phase, where the ramus mandibulae is also segmented, and L is reduced to 2 cm - resulting in a more accurate deformation of the model and at the same time preventing profiles in the upper mandibular region from (falsely) detecting the maxilla. The accuracy of the segmentation is measured by computing the surface distance between the automatic result with the gold standard. Results are shown in Fig. 5.

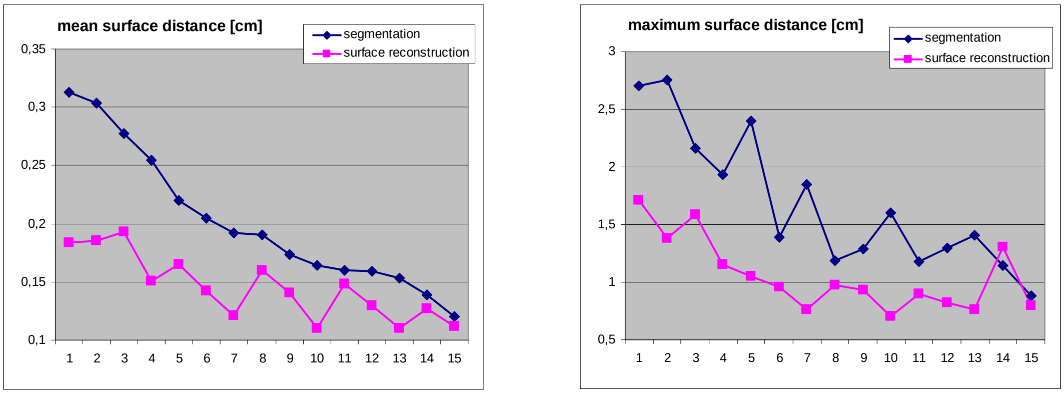

Fig. 5: Mean and maximum surface distances between gold standard and results of segmentation and surface reconstruction for each of the 15 NewTom data sets.

3. Discussion and Conclusions

The results of the direct surface reconstruction indicate the best possible reconstruction to be achieved with the available statistical shape model. The deviations between the segmentation and the surface reconstruction are due to shortcomings in the segmentation process, i.e. inaccuracies in the deformation model and/or the segmentation strategy. Improvements are subject to future work.

However the results show that statistical models of 3D shapes offer a promising approach for automating the segmentation of volumetric data. Incorporating a-priori knowledge about the object and the data seems a feasible way to segment image data in the presence of high noise-to-signal ratios or artefacts - as is the case in low-dose CT data. The usage of a statistical 3D shape model of the mandible is a suitable approach due to the mandible’s characteristic shape and well defined topology.

While the proposed method may require some manual interaction for post-processing, we expect that with increasing number of samples in the training set, interaction will be required only in a minority of cases. Since the model-based approach yields a good initialization for such interaction, this may even be avoided completely by means of locally elastic deformations, thus providing a fully automated segmentation method.

References

[1] A. Schramm, M. Rücker, N. Sakkas et al.: The Use of cone beam CT in cranio-maxillofacial surgery, in Proc. CARS (H. Lemke et al., eds), Int Congress Series (1281), Elsevier, 2005, pp. 1200-1204.

[2] T. Cootes, A. Hill, C. Taylor, J. Haslam: The Use of Active Shape Models for Locating Structures in Medical Images, Image and Vision Computing, 12 (6), pp. 355-366, 1994.

[3] H. Lamecker, T. Lange, M. Seebaß: Segmentation of the Liver using a 3D Statistical Shape Model, Technical Report, ZIB PaperWeb 04-09, 2004.

[4] S. Zachow, H. Lamecker, B. Elsholtz, M. Stiller: Reconstruction of Mandibular Dysplasia using a Statistical 3D-Shape Model, in Proc. CARS (H. Lemke et al, eds.), Int. Congress Series (1281), Elsevier, 2005, pp. 1238-1243.