

Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

HANS LAMECKER, THOMAS LANGE, MARTIN SEEBASS

Segmentation of the Liver using a 3D Statistical Shape Model

Segmentation of the Liver using a 3D Statistical Shape Model

Hans Lamecker*, Thomas Lange†, Martin Seebaß‡

April 8, 2004

Abstract

This paper presents an automatic approach for segmentation of the liver from computer tomography (CT) images based on a 3D statistical shape model. Segmentation of the liver is an important prerequisite in liver surgery planning. One of the major challenges in building a 3D shape model from a training set of segmented instances of an object is the determination of the correspondence between different surfaces. We propose to use a geometric approach that is based on minimizing the distortion of the correspondence mapping between two different surfaces. For the adaption of the shape model to the image data a profile model based on the grey value appearance of the liver and its surrounding tissues in contrast enhanced CT data was developed. The robustness of this method results from a previous nonlinear diffusion filtering of the image data. Special focus is turned to the quantitative evaluation of the segmentation process. Several different error measures are discussed and implemented in a study involving more than 30 livers.

*H. Lamecker is with the department of Scientific Visualization of the Zuse Institute Berlin (ZIB), Takustr. 7, 14195 Berlin, Germany (email: lamecker@zib.de)

†T. Lange is with Robert-Rössle-Klinik, Charité Berlin, Lindenberger Weg 80, 13125 Berlin, Germany (email: thomas.lange@charite.de)

‡M. Seebaß is with Indeed-Visual Concepts GmbH, Ihnestr. 23, 14195 Berlin, Germany (e-mail: seebass@indeed3d.com)

Contents

1	Introduction	3
2	Statistical Shape Modeling	5
2.1	Shape Extraction and Representation	5
2.2	Building the Statistical Model	5
2.3	Establishing Correspondence	6
2.4	Application to image segmentation	8
3	Segmentation of CT data of the liver	9
3.1	Specification of the data	10
3.2	Diffusion filtering	10
3.3	Grey value profile model	11
4	Measures for 3D Shape Comparison	12
4.1	Volume overlap	12
4.2	Symmetric surface distance measures	12
5	Evaluation and Results	13
5.1	Compactness and Completeness	14
5.2	Leave-all-in segmentation	15
5.3	Leave-one-out segmentation	15
6	Discussion and Conclusions	18
A	Optimization method	22

1 Introduction

Pre-operative planning for oncological liver surgery requires the segmentation of the liver tissue. The segmentation is the basis for a 3D anatomical model of the liver, which permits the computation of the resection volume, i.e. the tissue to be removed during surgery [1], [2], [3]. This volume contains the tumor and the part of the liver parenchyma where no blood supply would be provided after the resection. The high variability of the liver shape and vasculature [4] requires individual planning. In the clinical routine the manual segmentation of the liver parenchyma is too time consuming (about one hour per data set). The segmentation time will be even longer with increasing number of slices from improved CT scanners. Therefore methods have to be devised that require few interaction time.

A semi-automatic method, that reduces manual interaction time is presented in [5]. In this approach only a few slices of the original CT stack are interactively segmented using a live-wire approach, and for the intermediate slices a shape based interpolation is performed automatically. Then the interpolated contours of the liver shape are optimized by using the live-wire algorithm with automatically chosen seed points. Manual interaction time is reduced by 60% with this method, but a further reduction is still desirable.

Among the vast variety of automatic image segmentation approaches *deformable models* have proven to be effective and powerful in a large number of medical applications [6]. In many cases simple assumptions on the deformability of the object, like smoothness or elasticity criteria, generate shapes that definitely are not *legal* instances of the object. Including a-priori knowledge of the object to be segmented seems to be sensible to increase robustness. There are only few approaches so far that use 3D deformable models for the segmentation of the liver.

In [3] a globally constrained 3D deformable model proposed in [7] is applied to the segmentation of the liver. This method incorporates prior knowledge about the curvature of one template shape (liver from visible-human data set) in the deformation process. Only few examples have been quantitatively evaluated in this work.

Beichel et al. [8] applied an Active Appearance Model (AAM) [9] to the segmentation of the diaphragm dome which separates the lungs and the heart from the liver. Only the upper part and not the whole liver is segmented. They model the 3D shape of the dome by means of a 2D AAM, where one dimension (z-direction) is represented by a height grey value image above a reference plane.

A 2D and 3D level-set approach is used by Pan et al. [10] to segment the liver. They proposed a novel speed function that controls the front propagation of an implicitly defined surface towards the liver boundary. The idea is to change the speed function dynamically according to the past history of the front. In addition the propagation is constrained by simple a-priori anatomic information (distance between liver and skin).

In our work we will follow the approach of capturing the shape variations of the liver from a set of training data statistically via principal component analysis (PCA) [11]. The deformable model is only allowed to vary within the captured space of variations. This approach seems to be promising for robust and automatic image segmentation, because it provides an efficient parameterization and a compact representation of the object to be segmented, as will be explicated in section 2.

One problem encountered when building such a statistical model of shape is the need to establish correspondences between different shapes in the training set, i.e. to match anatomically equivalent features. The main obstacle is the lack of a suitable shape comparison, due to the difficulty of defining an optimality criterion for a good correspondence. This becomes particularly difficult in 3D. Several approaches have been proposed to face the problem in 3D and to automate the task:

A method for 3D shape correspondence using local geometry and geodesics is described in [12] without its application to statistical shape models. It is evaluated by computing residual surface distances.

In [13] the shapes are represented by their expansion into a series of elliptical harmonics. The PCA is performed on the parametric description in contrast to the point distribution model used in [11]. Only shapes with sphere-like topology can be treated.

Thompson et al. [14] compute a mapping from each shape onto a sphere using a deformable model approach. Correspondence between two spheres, constrained by matched anatomical feature lines, is established by means of a warping algorithm.

Fleute et al. [15] establish correspondence by elastic registration of a template shape with all other shapes based on minimizing Euclidean distance. Closely related is the approach of Frangi et al. [16]. The PCA is performed on the points of the control grids computed from an elastic registration on binary volumetric data.

Brett et al. [17, 18] present two approaches of automatic construction of shape models based on polyhedral shape representation. One minimizes the Euclidean distance between shapes on different levels of detail, obtained by triangle decimation. The related problem of surface folding is relaxed by their second approach of parameterizing surfaces onto common planar domains using harmonic maps. Their method applies only to surfaces homeomorphic to discs and does not guarantee fold-over free mappings however (see [19] for an illustration of this fact).

In [20] the correspondences are defined using an information theoretic minimal description length approach. Here the optimality criterion is the compactness of the model. The optimization of this criterion is very expensive. The method has been applied to 2D-examples and an extension to 3D shapes with sphere-like topology is outlined.

The contributions of this work are as follows. First, we adopt the optimality criterion of *minimal geometric distortion* for shape correspondence, as used by Brett et al. [18]. Instead of harmonic maps we use *convex-combination* maps [21] which are guaranteed to be fold-over free. This idea is combined with a morphing method by Zöckler et al. [22] which is capable of matching shapes of *arbitrary topology*. This approach also allows experts to manually specify corresponding features on two shapes, thereby constraining the correspondence mapping. Note that manual interaction will only be needed in the model building process and not in the segmentation process. To our knowledge this paper is the first one presenting a method that is capable of matching shapes of arbitrary topology fold-over free while minimizing geometric distortion, and applying this method to the generation of statistical shape models.

Second, a statistical shape model of the liver is built from a training set of 43 shapes, extending preliminary work presented by the authors [23]. This is first 3D statistical shape model of the liver applied to the task of image segmentation.

Third, the model is applied to the automatic segmentation of 33 CT data sets and the results are evaluated in detail. In order to use a statistical shape model in a segmentation procedure, a method for adapting the model to the data is needed. A common way is to model grey value profiles along surface normals statistically from the training data sets [11] and deform the shape according to the best-matching profile. For the CT data we are dealing with, we propose a different approach, where we use a *fixed model* for the profiles. This becomes possible after filtering the data with a nonlinear isotropic diffusion filter, as will be explained in detail in section 3.

The paper is organized as follows: In section 2 we will review the notions of statistical shape modeling and elaborate on our approach of establishing correspondence between shapes. Section 3

presents details of the adaption strategy specific to the segmentation of the liver from CT data. In section 4 we will focus on quantitative evaluation criteria that our results will be based on. Many different criteria are used in the literature and we will discuss some of them in detail, in order to allow comparisons with other methods. Section 5 presents the results obtained by applying our statistical shape model and the grey value model in a segmentation procedure.

2 Statistical Shape Modeling

2.1 Shape Extraction and Representation

Building a statistical shape model requires the setup of a training set, which in this work is based on 43 CT data sets acquired from a helical CT scanner. The training data comprise livers of 25 women and 18 men. Each data set was manually segmented by labeling the liver tissue. Surface reconstruction from the labeled images yields triangulated surfaces with a stepped structure due to the anisotropy of the voxels (slice thickness of 5 mm, resolution of 1.4 mm in axial slices). To obtain a superior representation of the true anatomical 3D-shape the labelings are interpolated by inserting intermediate slices between each two consecutive slices of the original image stack, such that the resulting voxel size is nearly isotropic. The method of variational implicit function interpolation [24] is well suited for this task. For reasons of efficiency all surfaces are simplified [25] by reducing the number of triangles, obtaining meshes with about 25000 triangles and 12500 nodes. All these steps are performed with the help of the visualization and modeling software Amira [26].

2.2 Building the Statistical Model

A statistical model is built from a training set of shapes \mathbf{v}_i ($i = 1, \dots, N$). Each shape \mathbf{v}_i is given by M points sampled on its surface (thus $\mathbf{v}_i \in \mathbb{R}^{3M}$). Using standard principal component analysis [27] each shape vector can be expressed using a linear model of the form

$$\mathbf{v}_i = \bar{\mathbf{v}} + \mathbf{P}\mathbf{b}_i = \bar{\mathbf{v}} + \sum_k \mathbf{p}^k b_i^k \quad (1)$$

where $\bar{\mathbf{v}}$ is the mean shape vector and $\mathbf{P} = \{\mathbf{p}^k\}$ the matrix of eigenvectors of the covariance matrix. The corresponding eigenvalues $\{\lambda^k\}$ describe the amount of variance in the direction of the eigenvectors. The shape parameters $\mathbf{b} = \{b^k\}$ control the modes of variation.

In order to obtain a correct statistical model all M points on each surface (a) must correspond in an anatomical meaningful way, *and* (b) their coordinates must be given relative to a common frame of reference (alignment). This is crucial, since incorrect correspondences can either introduce too much variation or lead to illegal instances of the model.

Note that - in general - these two goals can be accomplished independently of one another. Often, however, an initial alignment precedes the computation of corresponding points (e.g. when corresponding points are supposed to be closest points). For the generation of a model of the liver, we will employ a method, that does *not* rely on an initial alignment (see section 2.3). Once correspondence has been established, the alignment will be computed in the following way:

Without loss of generality, let \mathbf{v}_1 define the reference coordinate system, which all other shapes \mathbf{v}'_j ($j > 1$) will be aligned to. Let $\mathbf{x}_{1,k}$ denote the coordinates of the shape vector \mathbf{v}_1 , and $\mathbf{x}'_{j,k}$ the




Figure 1: A homeomorphism f between the shapes \mathcal{S}_1 and \mathcal{S}_2 is computed by mapping each patch of the two shapes onto a disc using the shape preserving mappings ϕ_1 and ϕ_2 respectively. The boundaries are mapped to the unit circle by fixing all branch points according to their average arc-length on both surfaces to be matched. All edges on the boundaries are mapped according to their arc-length on the original surfaces. The resulting mapping for one patch is given by $f = \phi_2^{-1} \circ \phi_1$.

coordinates of \mathbf{v}'_j respectively. Then we compute a rigid transformation $\mathbf{T}_{j,\min}$ that minimizes the sum of squared distances between corresponding pairs $\mathbf{x}_{1,k}$ and $\mathbf{x}'_{j,k}$ ($k = 1, \dots, M$), for all $j > 1$:

$$\mathbf{T}_{j,\min} = \arg \min_{\mathbf{T}} \sum_{k=1}^M \|\mathbf{x}_{1,k} - \mathbf{T}\mathbf{x}'_{j,k}\|^2$$

This minimization problem can be solved by singular value decomposition [28]. The coordinates of the final shape vector \mathbf{v}_j will then be $\mathbf{x}_{j,k} = \mathbf{T}_{j,\min}\mathbf{x}'_{j,k}$.

This procedure introduces some bias towards the reference system \mathbf{v}_1 , which can be reduced by iterative refinement (see e.g. [16]). However we have seen in our experiments that the influence of this defect is rather small.

2.3 Establishing Correspondence

The task of establishing correspondence between two topologically equivalent 3D shapes $\mathcal{S}_1, \mathcal{S}_2$, represented as triangular meshes, is to compute a homeomorphic mapping

$$f : \mathcal{S}_1 \rightarrow \mathcal{S}_2 ,$$

possibly under additional constraints (e.g. specifically given correspondences of points or lines of anatomical or geometrical significance). Depending on these constraints, there may be many solutions to this problem. Hence there is a need to impose further restrictions, that define optimal correspondences.

Since we are dealing with anatomical shapes, we propose that f is required to introduce as little distortion as possible, i.e. preserve the metric structure of \mathcal{S}_1 by keeping angles and lengths approximately fixed.

Instead of computing f directly we define a canonical base domain D to which each of the shapes \mathcal{S}_1 and \mathcal{S}_2 is mapped. D must be homeomorphic to \mathcal{S}_1 or \mathcal{S}_2 .

$$\begin{array}{ccc} \mathcal{S}_1 & \xrightarrow{f_2^{-1} \circ f_1} & \mathcal{S}_2 \\ & \searrow f_1 & \swarrow f_2 \\ & D & \end{array} \quad (2)$$

Hence we are left with the problem of finding a mapping $f : \mathcal{S}_i \rightarrow D$ for each shape i , consistent with the constraints given.

Surfaces homeomorphic to discs. Let us consider the easiest case, where \mathcal{S}_1 and D are homeomorphic to a disc. Moreover let $D \subset \mathbb{R}^2$ be a convex region. As an additional constraint, we require the boundary of \mathcal{S}_1 to be mapped smoothly to the boundary of D .

One possible candidate for a map f that minimizes metric distortion could be piecewise linear (PL) harmonic maps [29]. However these maps may introduce fold-overs, as was shown in [19]. An alternative, which is guaranteed to be fold-over free, are so-called *convex-combination maps* [21].

Without loss of generality, let $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{x}_{n+1}, \dots, \mathbf{x}_N\}$ be the coordinates of the vertices of the surface triangulation \mathcal{S}_1 . Assume that $\mathbf{x}_1, \dots, \mathbf{x}_n$ are internal nodes, while $\mathbf{x}_{n+1}, \dots, \mathbf{x}_N$ lie on the boundary. Furthermore let $\mathbf{u}_i \in D \subset \mathbb{R}^2$ ($i = 1, \dots, N$) be the corresponding coordinates in the parameter domain. The mapping

$$\phi : X \rightarrow D, \mathbf{x}_i \mapsto \mathbf{u}_i = \phi(\mathbf{x}_i)$$

is called a *convex combination map*, if

$$\mathbf{u}_i = \sum_{j=1}^N \lambda_{ij} \mathbf{u}_j \quad (3)$$

for all internal nodes $i = 1, \dots, n$, and

$$\sum_{j=1}^N \lambda_{ij} = 1 \quad , \quad \lambda_{ij} \begin{cases} = 0 & (i, j) \notin E \\ > 0 & (i, j) \in E \end{cases} \quad (4)$$

where E is the set of all edges of the triangulation. It was shown by Floater [21] that ϕ does not produce any fold-overs in the planar triangulation, and every node \mathbf{u}_i lies in the convex hull of ∂D . Computing ϕ amounts to solving the linear system (3) under the constraints of (4). Floater also showed that there exists a unique solution to this problem. Since (3) is a sparse linear system of equations, it can be solved efficiently.

The question remains on how to choose convex the combination weights λ_{ij} . The most obvious choice would be to set $\lambda_{ij} = 1/n_i$, where n_i is the number of neighbors of vertex \mathbf{x}_i . This mapping is called the *barycentric mapping*, since every vertex in the parameter domain will lie in the barycenter of its neighbors. Yet we are interested in producing a mapping, that preserves the metric structure of the surface \mathcal{S}_1 as good as possible.

We use here the idea of locally approximating the *geodesic polar map* [30], first presented by Floater [21]. Thereby λ_{ij} are computed as convex weights in the planar domain, such that the angles and lengths of the original triangulation are preserved as good as possible (see Figure 2). The resulting map is called *shape preserving map*.

Arbitrary surfaces. To establish correspondence between two arbitrary surfaces, we employ a method presented by Zöckler et al. [22]. An arbitrary surface is thereby divided into patches, each homeomorphic to a disc. This decomposition needs not only to be topologically equivalent on all surfaces of the training set but also to represent similar semantic regions in order to get a meaningful correspondence. Each corresponding pair of patches is mapped onto a common planar convex




Figure 2: Approximating the geodesic polar mapping around the star-like neighborhood of a node of the surface triangulation: nodes \mathbf{x}_{j_k} are mapped to \mathbf{p}_k such that $\|\mathbf{p}_k - \mathbf{p}\| = \|\mathbf{x}_{j_k} - \mathbf{x}_i\|$ and $\theta'_k = 2\pi \theta_k / \sum_{l=1}^{n_i} \theta_l$.

base domain. We choose the unit disc (radius 1, centered at the origin). To achieve continuity across patch borders the patch boundaries are mapped to the unit circle according to their average arc-length on the two surfaces to be matched (see figure 1). Branch points (vertices, that belong to more than one boundary) are thus fixed on the circle, while all boundary edges are mapped according to the arc-length on their original surface. If we repeat the procedure for each pair of patches, we obtain a mapping from each vertex of \mathcal{S}_1 onto the surface of \mathcal{S}_2 . To create the statistical model of the liver we computed the correspondence from one arbitrarily selected surface to all other shapes. The reference surface and hence the final liver model consist of about 12000 points.

The patch decomposition of the liver surfaces is chosen to be along lines of high curvature (feature lines), since these lines represent anatomical features that can be found on each liver. We divide the surfaces into four patches: 1. the lower border of the left lobe, 2. the lower border of the right lobe plus the caudate lobe, 3. the lower border of the quadratic lobe and 4. the whole upper part of the liver surface (see figure 3). The feature lines are drawn on the surface by manually selecting points where two or more patch boundaries would meet (branch points). Some intermediate points along the feature lines are added manually. These points are then automatically connected by computing geodesic shortest paths between them. In almost all cases we use a metric that favors paths along lines of high curvature. Only in one region we use a pure distance measure (see 3 right). Following this procedure the user has to specify no more than 10 landmarks per surface, resulting in 4 patches, divided by 6 patch boundaries and 4 branch points.

Note that since the detection of anatomical features, as described above, is independent of the location and orientation of the surface in space, the whole process of establishing correspondences does not rely on any initial alignment of the surfaces.

2.4 Application to image segmentation

The segmentation strategy based on a statistical shape model does not essentially differ from the one used in [11] apart from the fact that we use a different iteration scheme. The core adjustment procedure of the segmentation strategy consists of two steps:

- First a displacement vector $\Delta \mathbf{x}_i$ for each point \mathbf{x}_i on the model surface is computed from an analysis of the grey value profile along the surface normal. In [11] $\Delta \mathbf{x}_i$ is computed from




Figure 3: Decomposition of the liver surface into 4 patches seen from caudal (left) and posterior (right) direction: lower border of left lobe (LL), lower border of right plus caudate lobe (LR+CL), quadratic lobe (QL) and upper part of the liver surface (U). The 6 patch boundaries are defined along lines of high curvature (feature lines) and meet at 4 branch points (black dots). Only between the upper and the lower right part no distinct curvature exists. In this case the closest adjacent feature lines are connected by the geodesic shortest path.

statistically modeling the profiles in the training set and matching the current profile to this model. We will use a fixed model to compute $\Delta \mathbf{x}_i$, that will be explained in detail in section 3.3.

- Next a weighted least squares approximation between the shape model $\mathcal{S}(\mathbf{b}, T) = T(\bar{\mathbf{v}} + P\mathbf{b})$ and the displaced surface points $\mathbf{x}_i + \Delta \mathbf{x}_i$ is computed. The approximation is either performed with respect to the position parameters T (position adjustment) or with respect to the shape parameters \mathbf{b} (shape adjustment). To prevent extreme shapes the allowed range for the shape parameters is restricted to the minimal and maximal values as derived from the training data.

The segmentation process is initialized by positioning the mean shape $\bar{\mathbf{v}}$ into the image data. In our application this can automatically be done by a translation in the cranial-caudal direction, such that the cranial borders of the bounding boxes from the CT data and the mean shape match.

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 the three most important modes are fitted to the data. For each subsequent level three more modes are added until the maximum number of modes is reached.

3 Segmentation of CT data of the liver

The approach of model based segmentation as described above is general except for the computation of the displacement field $\Delta \mathbf{x}_i$. This depends on the type of data to be segmented. In the pre-operative planning for hepatic surgery, CT is the most common imaging modality. We will therefore describe a method for the computation of the displacements $\Delta \mathbf{x}_i$, that is suited for this type of data.

3.1 Specification of the data

The data that were examined during the model building process were taken from the PACS of the radiological department. It served as the basis for the assumptions of the grey value model as explained in section 3.3. Thus we will describe our observations here in more detail.

Most CT data were acquired after contrast-agent injection in the portal venous phase on different helical CT scanners with varying protocols and included healthy subjects and patients with lesions. Most of the liver parenchyma is clearly defined in the CT data. Nevertheless there are some regions where a distinction is difficult:

- Because of the movement of the heart and partial volume effects the boundary between the heart and the liver often cannot be specified.
- All intra-hepatic vessels were considered to be part of the liver up to the point where they enter the liver. The exact location of the entry point often is difficult to determine.
- In some slices the vena cava is surrounded by liver parenchyma. In order to avoid sharp and small shape features in the model, we considered it part of the liver in these slices. Ultimately these vessels could be modeled separately.

3.2 Diffusion filtering

In a preprocessing step the CT data are smoothed with a nonlinear isotropic diffusion filter to reduce noise but preserve the liver boundary. An overview of nonlinear diffusion filtering is given in [31]. Let $f(x)$ be the image then a filtered image $u(x, t)$ is calculated by solving the partial differential equation

$$\partial_t u = \operatorname{div} (g(|\nabla u_\sigma|^2) \nabla u) \quad \text{with} \quad u(x, 0) = f(x). \quad (5)$$

The diffusivity function

$$g(s) := \begin{cases} 1 & (s \leq 0) \\ 1 - \exp\left(\frac{-3.315}{(s/\lambda)^4}\right) & (s > 0) \end{cases} \quad (6)$$

is decreasing with the magnitude of the edge detector $|\nabla u_\sigma|^2$. u_σ is a smoothed version of u which is obtained by convolving u with a Gaussian of standard deviation σ . Equation 5 can efficiently be solved by using the additive operator splitting scheme suggested in [32].

For this work we implemented 2D diffusion filtering and applied the filter slice by slice to the image data, using the values $\sigma = 3.0$, $\lambda = 3.5$ and $t = 100$. The effect of the diffusion filter on a CT slice is illustrated in figure 4. It can be observed, that it produces piecewise almost constant regions that are separated by sharp edges. The extraction of a grey value profile normal to the liver surface from the filtered CT data leads to profiles shown in figure 4 in the middle. The liver boundary is located in the middle of the profile at sample point 25. The inside of the liver is the region to the left of this point. The profile extracted from the diffusion filtered data is almost constant inside the liver, while the edge at the boundary is preserved nicely.

The filtered image could be the basis for a semi-automatic segmentation due to the features described above. Yet experiments with a region-growing approach showed that in most cases manual corrections would be necessary.




Figure 4: Effect of nonlinear diffusion filtering. Left: original unfiltered CT data. Middle: Unfiltered (dashed line) and diffusion-filtered (solid line) profiles. Right: diffusion filtered CT data.

3.3 Grey value profile model

The diffusion filtering leads to very homogeneous grey values within the liver tissue, while the application of the contrast agent leads to significant differences of the grey values between the liver and the neighboring tissues. This allows an explicit modeling of profiles normal to the model surface. These profiles are sampled from the data by tri-linear interpolation. Their length is 25 mm in both directions, 50 sample points are taken.

The average grey value \bar{g} is the only parameter for the profile model, that has to be determined for each liver individually. It can be estimated easily by a few mouse clicks in the data, due to the homogeneity of the grey values. Any grey value that lies within some interval Δg around \bar{g} is said to be a typical grey value for this specific liver. Δg was chosen constant for all livers.

In order to compute the displacement vectors $\Delta \mathbf{x}_i$ at the points \mathbf{x}_i on the model surface we have to distinguish three cases (let $I(\mathbf{x})$ be the grey value at point \mathbf{x}):

1. $|I(\mathbf{x}_i) - \bar{g}| \leq \Delta g$ and $|I(\mathbf{x}) - \bar{g}| \leq \Delta g$ for all \mathbf{x} on the inner side of the profile:
It is very likely that \mathbf{x}_i lies inside the liver and we move outward until we find a grey value at point \mathbf{x} that is not typical for the liver. The displacement $\Delta \mathbf{x}_i$ is then given by $\mathbf{x} - \mathbf{x}_i$.
2. $|I(\mathbf{x}_i) - \bar{g}| \leq \Delta g$ and there exists an \mathbf{x} on the inner side of the profile where $|I(\mathbf{x}) - \bar{g}| > \Delta g$:
Two different reasons for this situation are possible: Either the profile intersects an intra-hepatic vessel or lesion inside the liver. Or \mathbf{x} lies in a tissue that has a similar grey value like the liver and the profile intersects the tissue boundary on the inner side of the profile. In some cases for example the right kidney has a similar contrast enhancement like the liver. In those cases no decision whether \mathbf{x} lies inside or outside the liver can be made and the weight at point \mathbf{x}_i is assigned 0 in the least squares approximation.
3. $|I(\mathbf{x}_i) - \bar{g}| > \Delta g$:
It is likely that \mathbf{x}_i lies outside the liver in this case. We then move inward until we find a typical liver grey value. The displacement $\Delta \mathbf{x}_i$ is given by $\mathbf{x} - \mathbf{x}_i$. However, in some cases \mathbf{x}_i might lie within intra-hepatic vessels or lesions inside the liver. Usually these cases are statistically insignificant as long as those regions do not occur close to the liver boundary.

In cases where the profile intersects the surface more than once or leaves the bounding box of the CT data, the profile will be restricted to a suitable length. If a surface point does not lie inside the bounding box of the CT at all, this point is assigned weight 0 in the weighted least squares approximation. These cases are checked at each adjustment step of the segmentation process.

4 Measures for 3D Shape Comparison

In order to quantify the segmentation accuracy, suitable measures for comparing 3D shapes have to be devised. Different measures are used in the literature which often makes it difficult to compare the results. In this section we give precise definitions for evaluation criteria, that we believe are useful and intuitive for the user (in our case the surgeon). We distinguish between volume and surface based measures.

4.1 Volume overlap

This measure is interesting in particular for the application we have in mind, since the surgeons must not remove too much tissue in order to sustain the vital liver functions. We define the relative (symmetric) volume difference as in [5] to be

$$v(\mathcal{S}, \mathcal{S}') = 1 - \frac{|V_{\mathcal{S}} \cap V_{\mathcal{S}'}|}{1/2 \cdot (|V_{\mathcal{S}}| + |V_{\mathcal{S}'}|)}. \quad (7)$$

To evaluate the volume difference, we sample the interior of the surface on a regular grid via a scan conversion. We use the grid spacings of the original CT data. Nevertheless some error is introduced here, since the scan directions are somewhat arbitrary.

4.2 Symmetric surface distance measures

Given two surfaces \mathcal{S} and \mathcal{S}' we define the distance $d(\mathbf{x}, \mathcal{S}')$ between a point \mathbf{x} on a surface \mathcal{S} and the surface \mathcal{S}' as:

$$d(\mathbf{x}, \mathcal{S}') = \min_{\mathbf{x}' \in \mathcal{S}'} \|\mathbf{x} - \mathbf{x}'\|_2 \quad (8)$$

where $\|\cdot\|_2$ denotes the Euclidean norm. The following distance measures are chosen to be symmetric when exchanging \mathcal{S} with \mathcal{S}' . This is important if we do not want to lose important information, as exemplified in Figure 5.

Mean Distance:

$$d_{\text{mean}}(\mathcal{S}, \mathcal{S}') = \frac{1}{|\mathcal{S}| + |\mathcal{S}'|} \left(\int_{\mathbf{x} \in \mathcal{S}} d(\mathbf{x}, \mathcal{S}') d\mathcal{S} + \int_{\mathbf{x} \in \mathcal{S}'} d(\mathbf{x}, \mathcal{S}) d\mathcal{S} \right) \quad (9)$$

where $|\mathcal{S}|$ denotes the area of the surface.

Root Mean Square Distance:

$$d_{\text{rms}}(\mathcal{S}, \mathcal{S}') = \sqrt{\frac{1}{|\mathcal{S}| + |\mathcal{S}'|} \left(\int_{\mathbf{x} \in \mathcal{S}} d(\mathbf{x}, \mathcal{S}')^2 d\mathcal{S} + \int_{\mathbf{x} \in \mathcal{S}'} d(\mathbf{x}, \mathcal{S})^2 d\mathcal{S} \right)} \quad (10)$$

Area of deviation: The idea is to have a measure that is not as local as the Hausdorff distance $\max_{\mathbf{x} \in \mathcal{S}} d(\mathbf{x}, \mathcal{S}')$ and not as global as the mean distance and allows intuitive imagination of shape




Figure 5: Deficiency of the one-sided surface distance: The maximal distance from \mathcal{S} to \mathcal{S}' in the region shown is much smaller than the maximal distance from \mathcal{S}' to \mathcal{S} .

difference. The relative surface area on which the deviations are larger than some threshold t presents such a compromise:

$$d_r(\mathcal{S}, \mathcal{S}', t) = \frac{1}{|\mathcal{S}| + |\mathcal{S}'|} \left(\int_{\mathbf{x} \in \mathcal{S}} \Theta(d(\mathbf{x}, \mathcal{S}') - t) d\mathcal{S} + \int_{\mathbf{x} \in \mathcal{S}'} \Theta(d(\mathbf{x}, \mathcal{S}) - t) d\mathcal{S} \right) \quad (11)$$

where Θ is the Heaviside-Theta function.

5 Evaluation and Results

Using the methods described above we will focus on the following questions:

- How *complete* is our model, i.e. how well does it describe arbitrary shapes that are not explicitly incorporated into the model?
- Is our grey value profile model well suited?
- Using the shape and the profile model, which accuracy can be achieved when segmenting CT data sets, that are not contained in the model?

For the experiments conducted below, all data from the model building process were considered, except for some cases where the assumptions for our grey value model (see section 3.3) failed.

Since the data were taken from different scanners using different protocols 10 out of 43 data sets had to be excluded from the analysis for basically two reasons. In five cases the data were acquired in a different phase of contrast enhancement, that would not be applied for liver surgery planning. In the other five cases there were larger bright or dark peripheral regions within the liver tissue, which mostly were due to lesions. These problems could be solved by an improved grey-value model,

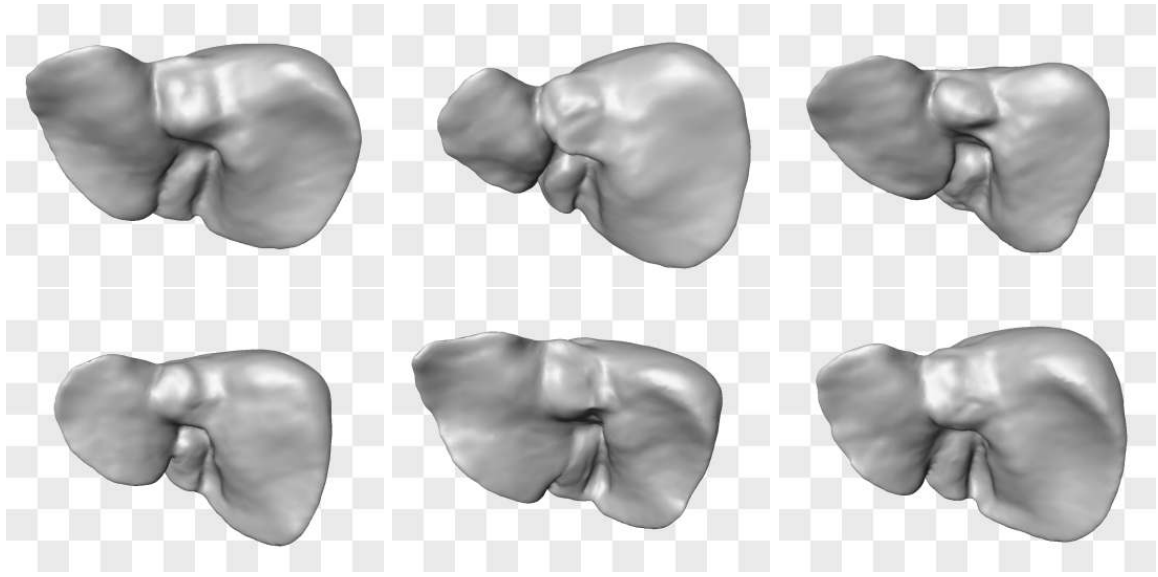


Figure 6: Variability of a statistical model of the liver shape made from 43 training data sets: in the left column the eigenmode with the largest variance λ_1 is varied between $\pm 3\sqrt{\lambda_1}$, in the second and third column the modes with the second and third largest variance are shown respectively.

where parts of the liver surface that border specific organs (e.g. heart, lung, ribs, kidney, vena cava) are treated separately. This is subject to future work. Hence all experiments were performed on a pool of 33 data sets, where the manual segmentations are referred to as the gold-standard (they will be called reference surfaces).

5.1 Compactness and Completeness

The compactness of a statistical shape model is its ability to describe the variability of a shape using as few modes as possible. Figure 7 shows the cumulative relative variance of a model built from 43 shapes. 21 modes already describe 95% of the total variance. The shape variability is displayed in Figure 6. The first three eigenmodes are varied within three standard deviations.

The completeness of a statistical model is its generalization ability or its ability to describe shapes that have not been incorporated in the training set. We test the completeness of our model in a leave-one-out test: for each of our 43 liver shapes, we match a model built from all other 42 shapes to the particular liver shape. This is achieved by minimizing the symmetric RMS distance (10) between the model \mathcal{S} and the reference surface \mathcal{S}' over the shape parameters $\mathbf{b} = \{b^k\}$ and the transformation parameters of a rigid transformation \mathbf{T} :

$$\min_{\mathbf{b}, \mathbf{T}} \{d_{\text{RMS}}(\mathcal{S}(\mathbf{b}, \mathbf{T}), \mathcal{S}')\} \quad (12)$$

The results of the shape optimization for the pool of 33 data sets are displayed in the middle curve of Figure 11. On average the shapes can be represented with a mean distance error of $1.9(\pm 0.3)$ mm. The average value for the surface area on which the deviation is larger than 5 mm is $4.9(\pm 2.5)$ % (see row 3 of Table 1). We describe in the appendix how the minimization (12) is performed.

For a complete model, we also should expect the number of modes making up for more than 95% of the total variance to converge against a fixed number, when increasing the number of data in the training set. We call this fixed number the dimensionality of the 95% projection space of




Figure 7: Cumulative relative variance of the statistical liver model built from 43 training shapes.

the model. Although in Figure 8 convergence cannot be observed, a promising trend can be seen. Convergence can only be tested if more data are included in the training set.

5.2 Leave-all-in segmentation

To analyze the performance of our grey value profile model, we segment all training data with a shape model that incorporates all training shapes (see Figure 9 for the different steps of the iterative segmentation procedure). Therefore we expect the results only to depend on the profile modeling and the overall segmentation strategy, but not the statistical shape model itself, since all shapes are represented in the model. The result of this experiment can be seen in the lower curve in Figure 11. On average the mean distance of the segmented surface to the reference surface is $1.1(\pm 0.3)$ mm. The average value for the surface area on which the deviation is larger than 5 mm is $1.8(\pm 1.8)$ % (see row 2 of Table 1).

Another test of the grey value profile modeling is to initialize the model with the position and shape parameters of the manual segmentation and start the segmentation process. With an ideal profile model the surface should not change in this case. However, on average the resulting segmentation differs by a mean surface distance of $0.9(\pm 0.4)$ mm and the average value for the surface area on which the deviation is larger than 5 mm is $1.6(\pm 1.7)$ % (see row 1 of Table 1).

5.3 Leave-one-out segmentation

Since both the shape optimization and the profile-model test yield satisfactory results, we are now interested to see how the actual segmentation on arbitrary CT data sets performs. We carried out a leave-one-out test on all 33 data sets each time segmenting it with a model made from 20, 30 and 42 shapes. The resulting mean distance errors are plotted in Figure 10.

The best result is achieved with the model made from 42 shapes: the average mean surface distance is $2.3(\pm 0.3)$ mm and the average value for the surface area on which the deviation is larger than 5 mm is $9.0(\pm 3.4)$ %. More overall error measures are given in rows 4-6 in Table 1. We observe that the segmentation improves the more degrees of freedom (shape modes) we apply. One example




Figure 8: Dimensionality of the model: the number of modes needed to account for at least 95% of the total variance is plotted against the number of training shapes contained in each model.

	d_{mean} [mm]	d_{RMS} [mm]	$d_r(3 \text{ mm})$ [%]	$d_r(4 \text{ mm})$ [%]	$d_r(5 \text{ mm})$ [%]	v [%]
leave-all-in optimal guess (43)	0.9 ± 0.4	1.3 ± 0.5	4.5 ± 4.2	2.4 ± 2.6	1.6 ± 1.7	3.3 ± 1.8
leave-all-in segmentation (43)	1.1 ± 0.3	1.5 ± 0.5	5.6 ± 4.7	3.0 ± 2.7	1.8 ± 1.8	3.8 ± 1.7
shape model optimization (42)	1.9 ± 0.3	2.5 ± 0.3	19.3 ± 5.9	9.8 ± 4.2	4.9 ± 2.5	5.9 ± 1.8
leave-one-out segmentation (42)	2.3 ± 0.3	3.1 ± 0.5	26.3 ± 6.2	15.3 ± 4.8	9.0 ± 3.4	7.0 ± 1.8
leave-one-out segmentation (30)	2.6 ± 0.4	3.4 ± 0.5	31.7 ± 6.6	20.1 ± 5.7	12.5 ± 4.4	7.8 ± 2.1
leave-one-out segmentation (20)	3.0 ± 0.5	3.9 ± 0.7	38.3 ± 8.2	25.9 ± 7.6	17.1 ± 6.5	8.8 ± 2.4

Table 1: Results of the leave-all-in segmentation (row 1: starting from optimal guess, row 2: standard initialization), the leave-one-out shape model optimization (row 3) and the leave-one-out segmentation for three different models (row 4-6). The number in brackets indicates how many training shapes are incorporated in the particular model. 33 data sets were segmented automatically. The different error measures are defined in section 4.




Figure 9: Steps of the segmentation process for a data set contained in the statistical liver model. At each of the six steps the model at its current position and shape is shown as a 3D surface. Below each surface one axial slice with the intersection line of the surface is displayed. Step 1: initial positioning of the mean model in the CT data such that the upper borders of the bounding boxes match. Step 2: result of the initial optimization of the position parameters. Step 3-5: results of the combined position and shape optimization with increasing number of shape modes (3, 9, 21). Step 6: final result (42 shape modes).

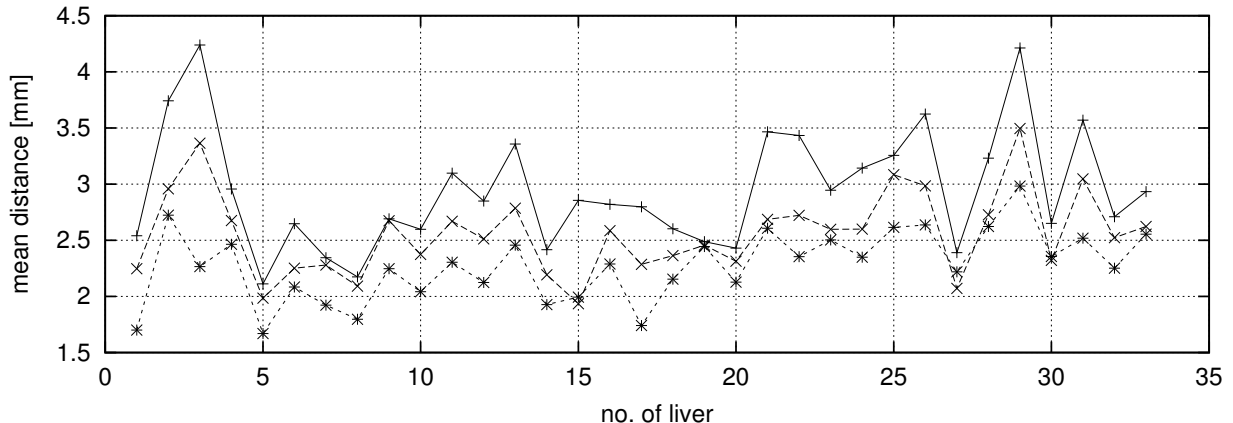


Figure 10: Leave-one-out segmentation results: the upper curve shows the mean surface distance of the automatic segmentation with a model built from 20 training shapes to the corresponding manual segmentation. The middle and lower curve show the same results but with models built from 30 and 42 training shapes respectively. The graphs show an improvement of the segmentation results with increasing number of degrees of freedom for most of the data sets.

of a resulting segmentation with a mean surface distance of $1.9(\pm 1.6)$ mm is shown in Figure 12 in 3 axial and 3 coronal slices. The relative surface area on which the deviation is larger than 5 mm is 5.6% in this example.

6 Discussion and Conclusions

We present an automatic method for the segmentation of the liver from CT data. It is based on a 3D statistical shape model built from 43 training data sets. A geometric approach for the solution of the correspondence problem in 3D is used. This algorithm is efficient and extends easily to arbitrary topologies. With an appropriate pre-processing (diffusion filtering) of the CT data to be segmented a robust adaption of the model to the data is achieved. Over a set of 33 test data sets we achieve an accuracy of $2.3(\pm 0.3)$ mm mean surface distance between the automatic results of our method and manual segmentations.

As mentioned in the introduction the following methods tackle the problem of automatic 3D liver segmentation. Beichel et al. [8] segment only the upper part of the liver, and it is not obvious how to extend their method to the whole liver parenchyma. Two other works by Pan et al. [10] and Soler et al. [3] apply their methods to rather small samples of 5 individual scans each. A volume overlap of 94 to 96% in the 2D case and 87 to 94% in the 3D case is reported by Pan et al. Soler et al. evaluate the one-sided mean surface distance between the automatic result and the manual segmentation. They achieve an accuracy of 2 mm. We consider symmetric measures more appropriate (refer to Figure 5).

The accuracy of our results is in a similar range (see Table 1). However several problems arise when quantitatively comparing segmentation results:

- Image qualities and resolutions differ between studies.
- Missing gold-standard: manual segmentations vary among different experts and are even not exactly reproducible by the same expert.




Figure 11: Comparison between the optimal surface distance to be achieved by using a model made from 42 training shapes (middle curve) versus a leave-one-out segmentation with the same model (upper curve). The differences should ideally be zero, hence they indicate the quality of the segmentation strategy, particularly the profile adaption method. As a reference the lower curve shows the results from the leave-all-in segmentation.




Figure 12: Example of a leave-one-out (dark contours) versus the leave-all-in (bright line) segmentation result. Three axial and three coronal slices with intersection lines of the 3D surfaces are shown. The mean surface distance errors of the leave-one-out and the leave-all-in segmentation are 1.9 mm and 0.8 mm, respectively. The deviation is larger than 5 mm on 5.6% and 0% of the surface area, respectively.

- Various similarity measures and discretizations are used to compare the results with the manual segmentation.

In the future it would be interesting to compare different methods on the same data pool with equal similarity measures.

For our method we were interested to analyze the source of inaccuracies and how we can improve the segmentation results. We have therefore performed experiments to test the main building blocks of our method: the *statistical shape model* and the *adaption process* based on the grey value profile model.

The expectation of a high variability of liver shape has been confirmed in our compactness and completeness evaluation. The dimensionality of the 95% projection space does not converge within the training set of 43 shapes, i.e., the model is not yet complete. This is also indicated by the results of the shape model optimization to each reference shape and the fact that even modes with small variance increase the ability to detect shape features. However there might be an upper limit for the number of modes above which the robustness of the segmentation process decreases. But this limit has not been reached yet: the comparison of the segmentation results from training sets of 20, 30, and 42 livers shows significant improvements with increasing training set, for the mean values (Table 1) as well as for the majority of the individual segmentations (Figure 10).

Apart from increasing the training data, the shape modeling can be improved by other means: in the future we intend to analyze how the point correspondences influence the performance of the segmentation. The objective of minimizing local distortion of the surfaces is only approximatively met by our method. As an improvement we intend to relax the surface triangulation after the initial parameterization by moving the triangle nodes on the surface, constraining only those nodes that represent distinct features (landmarks, nodes along lines of high curvature, etc.).

During the segmentation process lateral displacements are not desirable, since they do not modify the shape. Hence a second relaxation process at each adjustment step of the algorithm might further improve the results.

We are also interested to compare our method of establishing correspondences with other methods, e.g. by Davies et al. [20], Kelemen et al. [13] or Frangi et al. [16] especially with respect to its influence on the segmentation process.

The adaption process of the model to the image data is robust, due to the preprocessing of each image slice with the 2D diffusion filter and the multilevel iteration scheme. We tested this by doing *leave-all-in* segmentations, i.e. applying the segmentation procedure to data where the shape is incorporated in the statistical model.

In a first experiment we initialized the model with the position and shape of the manual segmentation to analyze the error due to the profile model only. The mean surface distance of the resulting segmentation to the reference surface was 0.9 mm on average over all 33 data sets. On the major part of each surface the profile model performs very well: the distance is even smaller than the average value and comparable to the differences between multiple manual segmentations of the same data set. Larger deviations occur at sharp edges and at the point where the vena cava enters the liver. At that point the vena cava has been defined to be a part of the liver in the manual segmentations, which is not reflected by our profile model. As a general trend, the surface distance is larger for data sets with a low average grey value of the liver, i.e. a low contrast between the liver and its

surroundings.

In the second experiment the model was initialized in the same way as in the leave-one-out test. The mean surface distance was 1.1 mm on average, which is almost as accurate as the result from the first experiment. This shows that not only the profile model works well but also the multilevel iteration scheme is a robust strategy.

For a complete model, we expect the segmentation results to show a similar accuracy as the *leave-all-in* tests (compare bright contours in Figure 12), so that for the large majority of cases no interactive post-processing of the automatic liver segmentation would be necessary.

In future work we want to

- improve the profile model by marking parts of the liver surface and applying specific profile models on those parts. By doing this we expect to capture some of the exceptional cases which had to be excluded in this study.
- add separate models of the vena cava, the right kidney, and eventually parts of the heart to the shape model. We expect to improve the segmentation accuracy in the regions where the liver bounds to these organs.
- develop a profile model for other imaging modalities as well. Due to the increasing use of MR scanners this modality is of particular interest. Note, that the shape model can be reused in this case, only the profile model has to be adapted.
- implement an anisotropic edge enhancing 3D filter. Although the isotropic 2D filtering works well, smoothing along edges and across slices might further improve the results.

Within the framework of a statistical segmentation algorithm, of course not all data might be segmented automatically in a correct manner. Yet the goal remains to get correct results in as many cases as possible. Our results suggest that an approach based on a statistical shape model is well suited to achieve this goal.

Acknowledgment

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) project "Intraoperative Navigation" 201879, by the DFG collaborative research project "Hyperthermia: Clinical Aspects and Methodology" (SFB 273) and by the DFG Research Center "Mathematics for key technologies" (FZT 86) in Berlin.




Figure 13: Computation of the gradient of the symmetric root-mean-square distance. Left: from \mathcal{S} to \mathcal{S}' . Right: from \mathcal{S}' to \mathcal{S}

A Optimization method

In this appendix, we explain how we match a statistical model of shape $\mathcal{S}(\mathbf{b}, \mathbf{T})$ to a given surface \mathcal{S}' by minimizing the symmetric RMS-distance (10):

$$\min_{\mathbf{b}, \mathbf{T}} \{d_{\text{RMS}}(\mathcal{S}(\mathbf{b}, \mathbf{T}), \mathcal{S}')\}$$

For the discrete case of triangulated surfaces the integrals in (10) are evaluated only at the sample points, i.e. the triangle vertices on the surface. With dense enough sampling we can therefore approximate the surface area $|\mathcal{S}|$ by the number of sample points M on the surface. The discrete RMS-distance (10) is then given by

$$d_{\text{ams}}(\mathcal{S}, \mathcal{S}') = \sqrt{\frac{1}{M + M'} \left(\sum_{\substack{j=1 \\ \mathbf{x}_j \in \mathcal{S}}}^M d(\mathbf{x}_j, \mathcal{S}')^2 + \sum_{\substack{j=1 \\ \mathbf{x}_j \in \mathcal{S}'}}^{M'} d(\mathbf{x}_j, \mathcal{S})^2 \right)} \quad (13)$$

To speed up the calculations we use an octree data structure.

We choose an iterative optimization strategy, where at each iteration we first compute optimal rigid transformation parameters \mathbf{T} by applying the ICP-algorithm [33]. We then perform the optimization with respect to the shape parameters \mathbf{b} . This procedure is repeated until convergence.

By separating the optimization (12) for the shape and transformation parameters respectively, we can speed up the calculations, since we will now be able to compute an explicit gradient of (13) with respect to the shape parameters \mathbf{b} . Hence we can use a quasi-Newton optimization algorithm [34], that is more efficient than methods that do not make use of explicit gradients. The gradient of (13) with respect to \mathbf{b} can be computed as follows:

Let $\mathbf{d}(\mathbf{x}_j, \mathcal{S}')$ be the vector from point \mathbf{x}_j on surface \mathcal{S} to the closest point \mathbf{r}_j on surface \mathcal{S}' , (see Figure 13, left):

$$\mathbf{d}(\mathbf{x}_j, \mathcal{S}') = \mathbf{r}_j - \mathbf{x}_j \quad , \quad |\mathbf{d}(\mathbf{x}_j, \mathcal{S}')| = d(\mathbf{x}_j, \mathcal{S}')$$

Since the surface \mathcal{S} is given by the linear model (1) we have

$$\mathbf{x}_j = \bar{\mathbf{v}}_j + \sum_k \mathbf{p}_j^k b^k$$

where $\bar{\mathbf{v}}_j \in \mathbb{R}^3$ denotes the coordinates of the j -th sample point of the mean surface (analogously $\mathbf{p}_j^k \in \mathbb{R}^3$). We can now compute

$$\frac{\partial}{\partial b^k} |\mathbf{d}(\mathbf{x}_j, \mathcal{S}')|^2 = 2 \mathbf{d}(\mathbf{x}_j, \mathcal{S}') \cdot \mathbf{p}_j^k$$

How do we deal with the second term of (13) under the square-root? Here the points \mathbf{x}_j lie on \mathcal{S}' and thus we cannot treat it in the same manner as described above. The solution is to invert the vectors $\mathbf{d}(\mathbf{x}_j, \mathcal{S})$, $\mathbf{x}_j \in \mathcal{S}'$ in the following way (see Figure 13, right):

For each $\mathbf{x}_j \in \mathcal{S}'$ we find the sample point $\mathbf{q}_j \in \mathcal{S}$ that is closest to \mathbf{r}_j . We then replace the vector $\mathbf{d}(\mathbf{x}_j, \mathcal{S})$ by $\mathbf{d}(\mathbf{q}_j, \mathbf{x}_j)$. In this way we get a reasonably well approximation of the symmetric RMS distance (13).

References

- [1] G. Glombitza, W. Lamade, A. M. Demeris, M. R. Göpfert, and A. Mayer, “Virtual planning of liver resections: image processing, visualization and volumetric evaluation,” *International Journal of Medical Informatics*, vol. 53, no. 2-3, pp. 225–237, 1999.
- [2] D. Selle, W. Spindler, B. Preim, and H.-O. Peitgen, *Mathematics Unlimited – 2001 and Beyond*, ser. Springer’s Special Book for the World Mathematical Year 2000, 2000, pp. 1039–1059.
- [3] L. Soler, H. Delingette, G. Malandain, J. Montagnat, N. Ayache, C. Koehl, O. Dourthe, B. Malassagne, M. Smith, D. Mutter, and J. Marescaux, “Fully automatic anatomical, pathological, and functional segmentation from ct scans for hepatic surgery,” *Computer Aided Surgery*, vol. 6, no. 3, pp. 131–142, 2001.
- [4] J. Fasel, D. Selle, and C. E. et al., “Segmental anatomy of the liver: Poor correlation with ct,” *Radiology*, vol. 206, pp. 151–156, 1998.
- [5] A. Schenk, G. Prause, and H.-O. Peitgen, “Efficient semiautomatic segmentation of 3d objects in medical images,” in *MICCAI 2001*, ser. LNCS, vol. 1935. Springer, 2001, pp. 186–195.
- [6] T. McInerney and D. Terzopoulos, “Deformable models in medical image analysis: a survey,” *Medical Image Analysis*, vol. 1, no. 2, pp. 91–108, 1996.
- [7] J. Montagnat and H. Delingette, “Volumetric medical images segmentation using shape constrained deformable models,” in *CVRMed*, 1997, pp. 13–22.
- [8] R. Beichel, S. Mitchell, E. Sorantin, F. Leberl, A. Goshtasby, and M. Sonka, “Shape- and appearance-based segmentation of volumetric medical images,” in *International Conference on Image Processing*, vol. 2, 2001, pp. 589–592.
- [9] T. Cootes, G. J. Edwards, and C. J. Taylor, “Active appearance models,” in *European Conference on Computer Vision*, H. Burkhardt and B. Neumann, Eds., vol. 2. Springer, 1998, pp. 484–498.

- [10] S. Pan and B. M. Dawant, "Automatic 3d segmentation of the liver from abdominal ct images: a level-set approach," in *SPIE 2001*, vol. 4322, 2001, pp. 128–138.
- [11] T. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "Use of active shape models for locating structures in medical images," *Image and Vision Computing*, vol. 12, pp. 355–366, 1994.
- [12] Y. Wang, B. S. Peterson, and L. H. Staib, "Shape-based 3d surface correspondence using geodesics and local geometry," in *CVPR 2000*, vol. 2, 2000, pp. 644–651.
- [13] A. Kelemen, G. Székely, and G. Gerig, "Three-dimensional model-based segmentation of brain mri," *IEEE Trans. on Medical Imaging*, vol. 18, no. 10, pp. 828–839, 1999.
- [14] P. M. Thompson and A. W. Toga, "Detection, visualization and animation of abnormal anatomic structure with a deformable probabilistic brain atlas based on random vector field transformations," *Medical Image Analysis*, vol. 1, no. 4, pp. 271–294, 1996.
- [15] M. Fleute, S. Lavalée, and R. Julliard, "Incorporating a statistically based shape model into a system for computed-assisted anterior cruciate ligament surgery," *Medical Image Analysis*, vol. 3, no. 3, pp. 209–222, 1999.
- [16] A. F. Frangi, D. Rückert, J. A. Schnabel, and W. J. Niessen, "Construction of multiple-object three-dimensional statistical shape models: Application to cardiac modelling," *IEEE Transactions on Medical Imaging*, vol. 21, no. 9, pp. 1151–1166, 2002.
- [17] A. D. Brett and C. J. Taylor, "A method of automated landmark generation for automated 3d pdm construction," *Imag. Vis. Comp.*, vol. 18, no. 9, pp. 739–748, 2000.
- [18] ———, "Automated construction of 3d shape models using harmonic maps," in *Medical Image Understanding and Analysis*, 2000, pp. 175–178.
- [19] T. Duchamp, A. Certain, T. DeRose, and W. Stuetzle, "Hierarchical computation of PL harmonic embeddings," University of Washington, Tech. Rep., 1997.
- [20] R. H. Davies, C. Twining, T. Cootes, J. C. Waterton, and C. J. Taylor, "A minimum description length approach to statistical shape modelling," *IEEE Transactions on Medical Imaging*, vol. 21, no. 5, pp. 525–537, May 2002.
- [21] M. S. Floater, "Parameterization and smooth approximation of surface triangulations," *Computer Aided Geometric Design*, vol. 14, no. 3, pp. 231–250, 1997.
- [22] M. Zöckler, D. Stalling, and H.-C. Hege, "Fast and intuitive generation of geometric shape transitions," *The Visual Computer*, vol. 16, no. 5, pp. 241–253, 2000. [Online]. Available: <http://visinfo.zib.de/EVlib/Show?EVL-2000-36>
- [23] H. Lamecker, T. Lange, and M. Seebaß, "A statistical shape model for the liver," in *MICCAI 2002*, 2002, pp. 422–427.
- [24] G. Turk and J. F. O'Brien, "Shape transformation using variational implicit functions," in *SIGGRAPH 99*, 1999, pp. 335–342.
- [25] M. Garland and P. S. Heckbert, "Simplification using quadratic error metrics," in *SIGGRAPH 97*, 1997, pp. 209–216.

- [26] D. Stalling, M. Westerhoff, and H.-C. Hege, "Amira - an object oriented system for visual data analysis," in *Visualization Handbook*, C. R. Johnson and C. D. Hansen, Eds. Academic Press, to appear. [Online]. Available: www.amiravis.com
- [27] A. Basilevsky, *Statistical Factor Analysis and Related Methods: Theory and Applications*. Wiley-Interscience, 1994.
- [28] A. D. McLachlan, "Least squares fitting of two structures," *J. Mol. Biol.*, vol. 128, pp. 74–79, 1979.
- [29] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle, "Multiresolution analysis of arbitrary meshes," in *SIGGRAPH*, vol. 29, 1995, pp. 173–182.
- [30] B. O'NEILL, *Elementary differential geometry*. Acadaemic press, 1966.
- [31] J. Weickert, "A review of nonlinear diffusion filtering," in *Scale-Space Theories in Computer Vision*, 1997, pp. 3–28.
- [32] J. Weickert, B. M. ter Haar Romeny, and M. A. Viergever, "Efficient and reliable schemes for nonlinear diffusion filtering," *IEEE Trans. Image Proc.*, vol. 7, no. 3, pp. 398–410, 1998.
- [33] J. P. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. of Pattern Anal. and Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.
- [34] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C – The Art of Scientific Computing*. Cambridge University Press, 1986.