

Automatic Segmentation of the Pelvic Bones from CT Data Based on a Statistical Shape Model

H. Seim¹, D. Kainmueller¹, M. Heller², H. Lamecker¹, S. Zachow¹, H.-C. Hege¹

¹Medical Planning Group, Zuse-Institute Berlin, Berlin, Germany

²Julius Wolff Institut and Center for Musculoskeletal Surgery Charité-Universitätsmedizin Berlin, Germany

Abstract

We present an algorithm for automatic segmentation of the human pelvic bones from CT datasets that is based on the application of a statistical shape model. The proposed method is divided into three steps: 1) The averaged shape of the pelvis model is initially placed within the CT data using the Generalized Hough Transform, 2) the statistical shape model is then adapted to the image data by a transformation and variation of its shape modes, and 3) a final free-form deformation step based on optimal graph searching is applied to overcome the restrictive character of the statistical shape representation.

We thoroughly evaluated the method on 50 manually segmented CT datasets by performing a leave-one-out study. The Generalized Hough Transform proved to be a reliable method for an automatic initial placement of the shape model within the CT data. Compared to the manual gold standard segmentations, our automatic segmentation approach produced an average surface distance of $1.2 \pm 0.3\text{mm}$ after the adaptation of the statistical shape model, which could be reduced to $0.7 \pm 0.3\text{mm}$ using a final free-form deformation step. Together with an average segmentation time of less than 5 minutes, the results of our study indicate that our method meets the requirements of clinical routine.

Categories and Subject Descriptors (according to ACM CCS): I.4.9 [Image Processing and Computer Vision]: Applications

1. Introduction

1.1. Motivation

In recent years an increasing interest in automatic methods for segmentation of pelvic CT images can be observed. A driving force behind this can be seen in the progression in the field of imaging techniques [FKK00], leading to the availability of an ever increasing amount of high-resolution image data. Such data, which allow for high quality reconstruction of anatomical structures, are the foundation for advanced therapy planning systems [ZZH07].

Whilst manual segmentation of image data by experts is still considered to be the gold standard for defining the anatomical models, reproducible results can typically only be achieved after significant training. Furthermore, this manual process is often time consuming, and thus limits the clinical applicability of new approaches in therapy planning sys-

tems. The manual reconstruction of anatomical models of the pelvic region for radiotherapy planning, for example, can take several hours [PMK04].

In order to make use of new therapy planning systems in both orthopedic and trauma surgery as well as in oncology it is necessary to make individual anatomical features readily available. Therefore it was the goal of this study to develop a framework for fast and accurate automatic segmentation of the pelvis, including left and right hip bone plus the sacrum.

1.2. Related Work

With the aim of providing a tool for fast fully automatic segmentation of pelvic CT images for radiotherapy planning, Haas et al. [HCS*08] developed a coarse to fine approach. They applied a complex combination of low level image analysis techniques (e.g. thresholding, morphological

operations, flood fill, etc.) for segmentation of the pelvic region. The proposed method focuses on the reconstruction of only certain parts of the pelvic bones, the proximal femur and surrounding soft tissue like prostate, bladder and rectum. Though receiving positive feedback from experts, no information is given on the reconstruction quality of the pelvis as a single structure.

To the best of our knowledge, Ehrhardt et al. [EHPP04] were the first to present an approach suitable for automatic three-dimensional reconstruction of the pelvic bones. They employed a male and a female shape atlas reconstructed from the Visible Human datasets. The segmentation was performed by non-linear registration of the anatomical atlas to the patient's image data, taking more than 30 minutes computation time. As a quality measure they introduced the percentage of correctly labeled voxels compared to manual expert segmentations on six patient datasets (95.2–98.5%).

A very similar method was presented in [PKB06]. Here, a non-rigid registration was applied to deform a prototype, in this case a reference dataset containing femur and pelvis represented by few gray values, to patient data. Although tested on clinical CT data, no results of a quantitative analysis have been published and no information is given about initialization and performance of the segmentation process.

The application of a statistical shape model (SSM) for the semi-automatic segmentation of the pelvis was described by Lamecker et al. [LSHD04]. A SSM of the pelvis was trained and evaluated on a set of 23 CT stacks. A gray value based image analysis along surface normals was used to fit the model to the image data resulting in an average segmentation error of 1.8 ± 0.2 mm. The segmentation results depended on the manual initialization of the SSM within the data.

A scheme for creating and validating volumetric statistical shape models of bones is presented in [CES*07]. Taking the male pelvis as an example, a statistical shape model was created from 110 CT datasets. Evaluation with a shape model consisting of 90 training shapes indicated an average surface distance of 1.5 mm in approximating any given pelvis shape using the first 15 eigenmodes. Within the proposed framework the volumetric model was also used for segmentation of new datasets by non-rigid registration, but not evaluated by comparing it to a given ground truth.

1.3. Contribution

This work presents an algorithm for fully automatic segmentation of the pelvic bones from CT data. We suggest a combination of the Generalized Hough Transform, segmentation based on a statistical shape model and a free form segmentation step based on optimal graph searching to combine the potentials of each single technique. An extensive evaluation on 50 CT datasets, all segmented in a fully automatic fashion, revealed a high segmentation quality with only minimal deviations from manual reference segmentations created by

experienced users. To the best of our knowledge our results as presented in Section 5.2 outperform (in terms of error metrics) segmentation results of related segmentation methods for which an evaluation was published (see Section 1.2).

2. Image Data

For our study 50 CT scans were available from a clinical study that aimed to determine the longer-term clinical outcome after unilateral total hip arthroplasty, thus including an implant on either the left or right acetabulum. The data base is composed of half female and half male pelvises. With a voxel size of about $0.9 \times 0.9 \times 5 \text{ mm}^3$ all CT stacks approximately have the same resolution.

The CT datasets were manually labeled by human subjects, well trained in pelvic anatomy and medical image segmentation. These segmentations include the entire pelvis with all three adjoining bones, namely the left and right hip bone as well as the sacrum including the coccyx (tail bone). Each hip bone consists of the ischium, the ilium and the pubis. At the prosthesis side the implant itself was not included in the segmentation. In order to allow for an independent evaluation of the native and the implanted hip bone, all datasets with the implant on the right were mirrored before processing. The analyses were then performed with datasets that all had the native hip joint on the right and the implant on the left side.

3. 3D Statistical Shape Model of the Pelvic Bones

The statistical shape models used in this work (see Section 5) were created from 50 pelvic CT datasets as mentioned in Section 2. Hence the training set meets the recommendation of [CES*07], according to which 40 to 50 training shapes are the minimum to capture intrinsic shape variations of the pelvis. Following the method introduced in [LSHD04], polygonal meshes, created from manual segmentations of each dataset, represent the training set. Based on significant anatomical and geometrical features a consistent subdivision into patches, i.e. regions of the mesh with the topology of a disc, was performed over all training shapes (see Figure 1). Based on these patches, point to point correspondences were established for all training shapes. As a result of this process all training shapes can be represented in a common vector space \mathbb{R}^{3m} , with m being the number of sample points used to discretize the shapes (vertices of the surfaces). After alignment of each training surface to one reference surface, Principle Component Analysis (PCA) on this training set yields a linear model

$$S(\mathbf{b}, T) = T(\bar{\mathbf{v}} + \sum_k b_k \mathbf{p}_k) \quad (1)$$

where $\bar{\mathbf{v}} \in \mathbb{R}^{3m}$ represents the mean shape, $\mathbf{p}_k \in \mathbb{R}^{3m}$ the modes of shape variation (eigenmodes), $b_k \in \mathbb{R}$ the shape weights and T an affine transformation. For each shape mode

k the range of the shape weight b_k is restricted to the minimum and maximum values reached by all training shapes. Any instance of a consistent surface mesh of the pelvic bones comprised in the statistical analysis, i.e. any training shape, may now be represented by such a linear combination. Due to the nature of the applied method the number of shape modes k equals the number of training shapes minus one.

The employed SSMs are non-manifold surface meshes containing two *inner* surfaces separating the sacrum from the adjoining hip bones. The mesh consists of 29619 vertices and 59403 triangles, divided into 21 patches.




Figure 1: Mean shape \bar{v} of the statistical shape model created from 50 pelvises. The different shades indicate the patch structure used for correspondence finding.

4. Segmentation Framework

With the focus on fully automatic segmentation our method follows a global to local approach divided into three phases: *pose initialization* to estimate position, orientation and scale of the pelvis in arbitrary image data, *statistical shape adaptation* to provide a good local initialization for the final *free-form segmentation* phase, which is used to overcome the restrictive character of statistical shape models and allows for a segmentation of shapes not captured by the SSM. A detailed description of each step is given in the following sections.

4.1. Pose Initialization by 3D Generalized Hough Transform

The Hough Transform is a well known gradient based, global method for robust detection of parametric objects in image data. This makes it an ideal tool for spatial initializations of more local methods, like our segmentation steps. An overview is given in [IK88]. For *real world objects*, however, it is often impossible to find a parametric description. Ballard [Bal81] proposed a method to overcome this limitation called the Generalized Hough Transform (GHT). The pose initialization of the pelvis from CT Data presented in our work closely follows the GHT based approach for 3D object detection introduced by Khoshelham [Kho07].

Instead of an analytic shape description a triangulated surface mesh is used as template shape. At each vertex c_i of

the template shape, the surface normal \vec{n}_i is extracted and its orientation angles ϕ_i and ψ_i are computed (see Figure 2). Additionally a 2-D look-up table (or R-table) is generated, that stores a vector \vec{v}_i for every c_i , indexed by ϕ_i and ψ_i . The vector \vec{v}_i describes the position of c_i relative to a pre-defined reference point r in object space, e.g. the center of gravity of the template shape.




Figure 2: A surface point's relative position \vec{v}_i stored in the R-table indexed by its normal vector's orientation angles.

To recognize an object, for every edge voxel at position e in image space a look-up in the R-table is performed depending on the orientation angles of the image gradient at e . Now, each vertex c_i of the template shape with similar normal orientation, i.e. the respective entry position in the R-table, votes for a position $p_i = e + v_i$ in image space. The maximum entry in the thus generated accumulator corresponds to the reference position with the most votes, i.e. the position where the reference point r is most probably located. By repeating this process with adjusted R-tables for a rotated and scaled template shape, an affine Transformation T^0 can be extracted to initialize the SSM segmentation.

The robustness of the GHT strongly depends on the quality of the template shape. The statistical shape model allows for an extraction of the mean shape \bar{v} , promising to give the *most general* representation of the pelvis' shape, and thus being the best possible template shape. Another disadvantage of the method is its performance. GHT can be considered very slow, since it is a brute-force method. We chose to make two assumptions on the data to reduce the computational cost without losing generality: 1) Limiting the range of scale and rotation (H_{scale}, H_{rot}), since the patient can be considered z -axis aligned, and 2) Using only bone gradients, for which we can roughly estimate a characteristic intensity threshold (H_{thresh}).

4.2. Adaptation of the Statistical Shape Model

By the use of an SSM, prior knowledge about the typical shape of the pelvis is incorporated into the segmentation process to constrain it where the image information is not reliable. Any segmentation resulting from SSM adaptation is an instance of the shape model and therefore should have a more or less plausible shape.

Segmentation using the SSM (1) is the task of finding the set of position and shape parameters such that $S(b, T)$ approximates the (unknown) shape $R \in \mathbb{R}^{3m}$ to be segmented as good as possible. We denote the resulting segmentation $R^* = S(b^*, T^*)$. However, the location and shape of R is only encoded implicitly in the image data $I: \mathbb{R}^3 \rightarrow \mathbb{R}$. Therefore the computation of R^* proceeds iteratively. Let $R^i = S(b^i, T^i)$ denote the segmentation in iteration i :

1. $R^0 := S(0, T^0)$. The computation of T^0 is described in Section 4.1.
2. Compute a displacement vector field $\Delta R^i \in \mathbb{R}^{3m}$ defined on the current segmentation R^i , i.e. a vector $\Delta r_k \in \mathbb{R}^3$ is assigned to each vertex $k \in \mathbb{N}$ of the surface R^i . It describes the desired deformation of the model towards the (unknown) surface R in the underlying image data I .
3. Project the displacements onto the SSM by solving the optimization problem $(b^{i+1}, T^{i+1}) = \operatorname{argmin}_{b, T} |(R^i + \Delta R^i) - S(b, T)|^2$.
4. Update $i \leftarrow i + 1$ and return to step (2) if convergence has not been achieved, i.e. if $|S(b^i, T^i) - S(b^{i+1}, T^{i+1})| > 3m \cdot \epsilon$; else return $b^* = b^i$ and $T^* = T^i$.

The resulting segmentation R^* generally does not equal the optimal solution $\operatorname{argmin}_{b, T} |R - S(b, T)|^2$. It is not clear how well R^* approximates R . This depends on the quality of the shape model, the nature of the iterative approach and on the computation of ΔR^i .

The displacement vector field ΔR^i is computed on the basis of the image data I . Therefore, an intensity model for the encoding of R in I must be established. A cost function c_k is computed for each vertex k of the surface by analyzing a 1D intensity profile at vertex position x_k along the (unit) surface normal u . This profile is sampled equidistantly over a length L at the set of points: $P_k = \{x_k^n := x_k + [(i-1)/(N_p-1) - 1/2] \cdot L \cdot u \text{ with } n = 1, \dots, N_p\}$. The cost function $c_k: P_k \rightarrow \mathbb{R}_0^+$ assigns some cost to each sampling point:

$$c_k(x_k^n) = \frac{-g_{min}}{|\nabla I(x_i)|} \text{ if (3), else } c_k(x_k^n) = 1 \quad (2)$$

$$I(x_k^n) \in [t1, t2] \text{ and } |\nabla I(x_k^n)| < -g_{min} \quad (3)$$

In (3), $t1$ and $t2$ define an intensity window of interest and g_{min} defines a threshold for gradient magnitude. The position of the minimum of the cost function c_k yields the displacement vector $\Delta r_k = x_k^* - x$, with $x_k^* = \operatorname{argmin}_{x_k^n} c_k(x_k^n)$, $x_k^n \in P$, $i \in \{1, \dots, N_p\}$, at vertex k of the surface.

4.3. Free-form Segmentation

As mentioned in Section 4.2, segmentation with an SSM reduces the search space to the shape model space, which is appropriate for producing a robust segmentation result. Anyway, in general, new, unknown shapes are not captured by the SSM model and therefore cannot be reproduced by a weighted combination of eigenvalues. To overcome this

limitation, we apply a free form deformation of the surface model, using the SSM segmentation result as initialization.

Adding displacement vectors individually to each vertex of the surface model is not feasible in terms of robustness. However, optimal graph searching algorithms allow for a global minimization of the sum of costs for each vertex displacement $\sum_k c_k(x_k^*)$, while respecting hard constraints on shape preservation. The hard constraints are realized by means of graph edges that connect sample points on profiles in such a way that a non empty minimum closed set in the graph defines the optimal surfaces.

The shape preservation constraint (also known as smoothing constraint) guarantees that new vertex positions on adjacent profiles are no more than s sample points away from each other. If sample point x_k^n is chosen on a profile P_k as desired new position, sample point x_l^m with $m > n - s$ must be chosen on adjacent profiles P_l . The smaller s is chosen, the more alike are the initial surface and the surface resulting from the optimization. For more details on graph construction see [LWCS06].

The minimum closed set problem can be transformed to a minimum s-t-cut problem, which is solved in polynomial time by maximum flow algorithms. An experimental comparison of various maximum flow algorithms can be found in [BK04].

4.4. Overall Segmentation Algorithm

The segmentation algorithm consists of a series of steps combining the methods presented in the sections above (see Table 1). The following methods are applied: pose initialization (GHT, see Section 4.1), shape model fitting (SSM, see Section 4.2), and optimization via graph cuts (OPT, see Section 4.3). For each step the parameters for its particular method are listed in the second column labeled *Details*.

Step	Details
1. GHT	Image resampling: $5x5x5mm^3$, $H_{scale} := [0.8, 1.2]$, $H_{rot} := \pm 10^\circ$, $H_{thresh} := 100HU$
2. SSM	Profile length $L := 50$ mm, sampling points $N_p := 50$, $[t1, t2] := [120, 320]$, $g_{min} := 50/mm$, $\epsilon := 0.1$ mm, all shape modes adapted
3. SSM	$L := 20$ mm
4. OPT	$L := 10$ mm, $N_p := 20$, Shape preservation $s := 2$

Table 1: Overall segmentation algorithm. Parameters remain the same if not stated otherwise.

5. Results

In a leave-one-out study all of the 50 CT stacks available for this study were segmented in a fully automatic fashion applying the algorithm presented in Section 4.4. For each data set to be segmented a shape model was generated excluding the respective pelvis from the training set. Thus we ensure the independence of training and testing data. The complete segmentation of one CT stack, i.e. data import, pose initialization and segmentation, took approximately 4 : 20 minutes on a 64-bit desktop PC (2,66GHz Core, 8GB RAM).

5.1. Evaluation Method

After each step (pose initialization, statistical shape model adaptation and free-form deformation) the resulting surface meshes were converted to labeled voxel representations with the same resolution as the respective original image dataset. The subdivision of the statistical shape allowed for a separate conversion of each sub-structure, i.e. creating one voxel label for the right hip bone, the left hip bone and the sacrum. With automatic and manual segmentations having the same structure, a quantitative evaluation on single labels as well as any combination of these labels is possible.

For evaluation, we used the following metrics: signed and absolute relative volume difference (SVD, AVD), volumetric overlap error (OE), average symmetric surface distance (AD), average symmetric roots mean square surface distance (RMS) and maximum surface distance (MD), each as described in [vGHS07]. For computation of these metrics we applied ANN, a library for computing approximate nearest neighbors (available at: <http://www.cs.umd.edu/~mount/ANN/>).

5.2. Evaluation Results

The summarized results are given in Tables 2, 3 and 4. For each error metric, the average (bold) and standard deviation over all 50 test cases is shown. Results for the whole pelvis (All), the right hip bone (RHB), the left hip bone (LHB) and the sacrum (S) are listed in separate rows. Table 2 shows the results after GHT initialization of the pelvis, Table 3 refers to the results after SSM adaptation and Table 4 lists the final results after free form segmentation. Additionally the result statistics are presented as box-and-whisker diagrams in Figure 3, where the median (black diamond), the interquartile range, i.e. 50% of the data (boxes), and an outlier-free min-max range are visualized for all metrics and all phases.

Final segmentation results for the complete pelvis exhibit an average AD of $0.7 \pm 0.3mm$, an average RMS of $1.9 \pm 0.6mm$ and an average MD of $16.5 \pm 5mm$. The average volumetric errors are $-0.5 \pm 3.9%$ for the SVD, $2.9 \pm 2.6%$ for the AVD and $14.4 \pm 3.4%$ for the OE. Figure 4 shows three segmentation results that are exemplary for a good, an average and a bad case in terms of their AD.

GHT	SVD	AVD	OE	AD	RMS	MD
	[%]	[%]	[%]	[mm]	[mm]	[mm]
All	0.8	12.1	49.0	3.8	5.4	24.1
std	16.2	10.7	8.3	1.3	1.8	6.5
RHB	1.0	12.6	50.6	3.4	4.7	18.6
std	16.8	10.9	8.7	1.2	1.6	4.3
LHB	1.7	14.5	53.5	3.8	5.3	20.3
std	19.2	12.5	14.0	2.2	2.7	7.1
S	0.7	12.3	48.0	4.5	5.8	19.7
std	16.4	10.8	13.6	2.1	2.4	5.9

Table 2: Evaluation results (initialization phase).

SSM	SVD	AVD	OE	AD	RMS	MD
	[%]	[%]	[%]	[mm]	[mm]	[mm]
All	-3.9	5.0	21.8	1.2	2.2	16.1
std	5.3	4.2	3.6	0.3	0.6	5.4
RHB	-5.1	5.7	20.8	0.9	1.5	9.2
std	4.5	3.6	2.7	0.2	0.2	1.9
LHB	-4.6	6.5	24.4	1.0	1.9	11.4
std	6.4	4.5	4.1	0.2	0.3	2.0
S	-0.8	8.0	25.9	1.8	3.0	15.6
std	11.5	8.3	6.0	0.7	1.2	5.7

Table 3: Evaluation results (active shape adaptation).

OPT	SVD	AVD	OE	AD	RMS	MD
	[%]	[%]	[%]	[mm]	[mm]	[mm]
All	-0.5	2.9	14.4	0.7	1.9	16.5
std	3.9	2.6	3.4	0.3	0.6	5.0
RHB	-3.3	3.9	11.9	0.4	1.1	9.2
std	3.6	2.9	2.7	0.1	0.3	2.0
LHB	-1.2	3.7	15.6	0.6	1.5	10.8
std	5.0	3.4	3.8	0.2	0.3	2.4
S	4.8	8.4	21.1	1.4	2.8	16.0
std	10.7	8.2	6.3	0.7	1.3	5.4

Table 4: Evaluation results (free-form phase).

Segmentation qualities of the three anatomical structures, namely left hip bone, right hip bone and sacrum, exhibit significant differences. The right hip bone, with an average AD of $0.4 \pm 0.1mm$ for the final results, performs slightly better than its left counterpart reaching $0.6 \pm 0.2mm$ (see Table 4). Both hip bones and the whole pelvis reach an average AD smaller or equal to the in-plane resolution of the CT data. Compared to the hip bones, the average AD for the final sacrum segmentations is relatively large with $1.4 \pm 0.7mm$. Note as well that the respective average SVD is largely positive with $4.8 \pm 10.7%$ as compared to the hip bones, which means that the sacrum is over-segmented in most cases.

In all 50 cases the mean pelvic model was initialized correctly within the data using the GHT. The largest average distance after pose initialization of only $6.3mm$ supports this statement. The evolution of the error metrics over the three segmentation phases (see Figures 3 and 5) shows that error metrics are improved in each phase, with the only exception that the average MD does not change significantly from the result of the SSM phase to the final segmentation result.




Figure 3: Surface distance error metrics after initialization, SSM adaptation and free-form step.




Figure 4: From left to right, transversal, coronal and 3D-view of a good case (top, AD 0.5mm), an average case (middle, AD 0.7mm) and difficult case (bottom, AD 1.2mm). The outlines identify the manual (dark) and automatic segmentations (light). The 3D views show the automatic segmentation, with the segmentation error encoded by intensity.

5.3. Discussion

The GHT proved to be a robust method to estimate the pose and size of the pelvis in CT datasets. The fact, that a single mean template shape is sufficient to reliably initialize the

pelvis despite its varying anatomy, can be attributed to the uniqueness of the shape within the human anatomy and its distinctive image features.

The automatic segmentation still has some deficiencies.

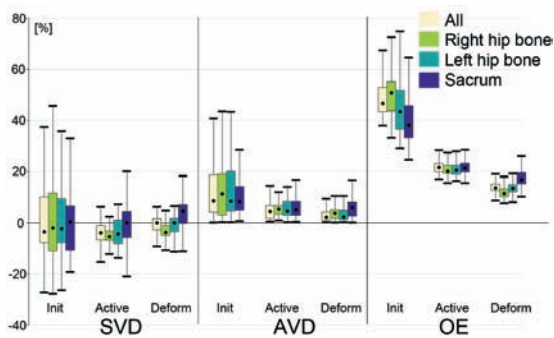


Figure 5: Volumetric error metrics after initialization, SSM adaptation and free-form step.




Figure 6: Three difficult cases of the native acetabulum. In contrast to the gold standard (dark contour) the automatic segmentation leaks into the femoral head (light contour).

Two main reasons for the difference in segmentation quality of right and left hip bone have been identified (see Figure 4): 1) The automatic segmentation is likely to be pulled towards the outer contour of the implant due to the image features, whereas the manual segmentation left out the implant. This results in a larger error at the implanted acetabulum. 2) Another problem in the acetabulum region, concerning pathological and healthy side, is a leaking of the segmentation into the femoral head, due to the lack of image features (see Figure 6). However, the shape preserving constraint of the free-form step confines leaking.

The comparatively large error metrics for the sacrum are caused by an over-segmentation of the lower lumbar vertebra due to the low resolution and the lack of meaningful image information in this area (see Figure 7) also indicated by the sacrum's positive average SVD value, as listed in Table 4. Since the sacrum is the smallest of the three structures in terms of surface area and volume these errors have a rather small impact on the overall result.

The stagnation of the maximum distance values from the SSM adaptation to the free-form step may be caused by the chosen profile length $L = 10\text{mm}$ for the free form step. This idea is supported by the fact that the MD values after SSM adaptation are in a range of 10 to 20mm, which is more than double the length of a profile on one side of a surface vertex.




Figure 7: Distance of automatic to reference segmentation reveals largest errors at the proximal end of the sacrum.

6. Conclusion and Future Work

An algorithm was presented for fully automatic segmentation of the pelvic bones from CT data, combining the potentials of the Generalized Hough Transform, segmentation based on a statistical shape model and a free form segmentation step based on optimal graph searching. An extensive evaluation of the proposed algorithm on 50 CT datasets revealed a high segmentation quality. Deviations from manual reference segmentations created by experienced users are below the in-plane resolution of the CT data.

Although providing very good spatial initializations for the SSM adaptation, a further exploitation of the GHT is recommended for future studies. At the moment only the affine transformation of the SSM is initialized. The use of different template shapes, each representing one of the principal modes of variation, may be used to transfer the shape weights of the template shape to the statistical shape model. By doing so, we expect a further reduction of the time of convergence for the adaptation of the statistical shape model.

With the adaptation of the statistical shape model we did outperform the results from past studies applying an SSM of the pelvis in terms of accuracy when used for segmentation [LSHD04] or the representation of shapes that are not contained in the model [CES*07]. According to the error metrics this step provides a fairly good initialization for the last free-form deformation step. However, an improvement, i.e. an extension of the training set and creation of different shape classes (e.g. male/female), may further improve the outcome of this step. This requires a sufficiently large set of training shapes for each shape class.

After local initialization by a SSM a free-form deformation was applied for the final segmentation of the bony surface. Although performing very well in most areas of the pelvis, we could identify regions where the proposed method still needs improvement, namely the acetabula and the sacrum. Incorporating the lower lumbar vertebra into the SSM may easily reduce the error for the sacrum. However, the acetabula need more attention and are already subject to current work, where multiple surfaces are coupled with shared intensity profiles for segmentation (e.g. femoral head and acetabulum) [KLZH08]. This approach is likely to further improve the automatic segmentation results especially in challenging situations such as low-resolution image data or strong cartilage wear (see Figure 6).

Future studies should include an evaluation of the inter-user variability for multiple manual segmentations. This would help compare our error metrics of the automatic segmentation to those yielded by human subjects and gain additional knowledge about the ground truth. We expect, at least for the hip bones, an inter-user variance similar to the results presented in this work.

Acknowledgements

Heiko Seim is supported by the European Commission under the FP6 IST Project DeSSOS (027252). Dagmar Kainmüller is supported by DFG Collaborative Research Center SFB 760. Hans Lamecker is funded by the German Research Center MATHEON in Berlin. We would like to thank Prof. Dr. Carsten Perka and Dr. Jörg Schröder (Center for Musculoskeletal Surgery, Charité - Universitätsmedizin Berlin) for providing the CT data sets, Alexander Wurl and Philippe Moewis (Julius Wolff Institut and Center for Musculoskeletal Surgery, Charité - Universitätsmedizin Berlin) for preparing the data.

References

- [Bal81] BALLARD D. H.: Generalizing the Hough Transform to Detect Arbitrary Shapes. *Pattern Recognit* 13, 2 (1981), 111–122.
- [BK04] BOYKOV Y. Y., KOLMOGOROV V.: An Experimental Comparison of Min-Cut/Max-Flow Algorithms for Energy Minimization in Vision. *IEEE Trans Pattern Anal Mach Intell* 26, 9 (September 2004), 1124–1137.
- [CES*07] CHINTALAPANI G., ELLINGSEN L. M., SADOWSKY O., PRINCE J. L., TAYLOR R. H.: Statistical atlases of bone anatomy: construction, iterative improvement and validation. *Med Image Comput Comput Assist Interv Int Conf Med Image Comput Comput Assist Interv* 10, Pt 1 (2007), 499–506.
- [EHPP04] EHRHARDT J., HANDELS H., PLÖTZ W., PÖPPL S. J.: Atlas-based recognition of anatomical structures and landmarks and the automatic computation of orthopedic parameters. *Methods Inf Med* 43, 4 (2004), 391–397.
- [FKK00] FUCHS T., KACHELRIESS M., KALENDER W. A.: Technical advances in multi-slice spiral CT. *Eur J Radiol* 36, 2 (Nov 2000), 69–73.
- [HCS*08] HAAS B., CORADI T., SCHOLZ M., KUNZ P., HUBER M., OPPITZ U., ANDRÉ L., LENGKEEK V., HUYSKENS D., VAN ESCH A., REDDICK R.: Automatic segmentation of thoracic and pelvic CT images for radiotherapy planning using implicit anatomic knowledge and organ-specific segmentation strategies. *Phys Med Biol* 53, 6 (Mar 2008), 1751–1771.
- [IK88] ILLINGWORTH J., KITTLER J.: A survey of the Hough transform. *Comput Vision Graph Image Process* 44, 1 (1988), 87–116.
- [Kho07] KHOSHELHAM K.: Extending Generalized Hough Transform to Detect 3D Objects in Laser Range Data. In *ISPRS Workshop on Laser Scanning, Proceedings, LS 2007* (Espoo, Finland, September 2007), pp. 206–210.
- [KLZH08] KAINMUELLER D., LAMECKER H., ZACHOW S., HEGE H.-C.: Coupling Deformable Models for Multi-object Segmentation. In *Biomedical Simulation, Proceedings of the 4th Int. Symposium, ISBMS* (London, UK, July 2008), Bello F., Edwards E., (Eds.), vol. 5104 of *Lecture Notes in Computer Science*, Springer, pp. 69–78.
- [LSHD04] LAMECKER H., SEEBASS M., HEGE H.-C., DEUFLHARD P.: A 3D Statistical Shape Model of the Pelvic Bone for Segmentation. In *Proceedings of SPIE - Volume 5370 Medical Imaging 2004: Image Processing* (May 2004), Fitzpatrick J., Sonka M., (Eds.), pp. 1341–1351.
- [LWCS06] LI K., WU X., CHEN D. Z., SONKA M.: Optimal Surface Segmentation in Volumetric Images-A Graph-Theoretic Approach. *IEEE Trans Pattern Anal Mach Intell* 28, 1 (2006), 119–134.
- [PKB06] PETERSSON J., KNUTSSON H., BORGA M.: Automatic Hip Bone Segmentation Using Non-Rigid Registration. In *Pattern Recognition, Proceedings of the 18th Int. Conference, ICPR* (Hongkong, China, August 2006), IEEE Computer Society, pp. 946–949.
- [PMK04] PEKAR V., MCNUTT T. R., KAUS M. R.: Automated model-based organ delineation for radiotherapy planning in prostatic region. *Int J Radiat Oncol Biol Phys* 60, 3 (November 2004), 973–980.
- [vGHS07] VAN GINNEKEN B., HEIMANN T., STYNER M.: 3D Segmentation in the Clinic: A Grand Challenge. In *3D Segmentation in the Clinic: A Grand Challenge, MICCAI Workshop, Proceedings* (2007), Heimann T., Styner M., van Ginneken B., (Eds.), pp. 7–15.
- [ZZH07] ZACHOW S., ZILSKE M., HEGE H.-C.: 3D Reconstruction of Individual Anatomy from Medical Image Data: Segmentation and Geometry Processing. In *25. ANSYS Conference & CADFEM Users' Meeting* (Dresden, 2007). Proc. CD 2.12.15.