

An Articulated Statistical Shape Model for Accurate Hip Joint Segmentation

Dagmar Kainmueller, Hans Lamecker, Stefan Zachow, Hans-Christian Hege

Abstract—In this paper we propose a framework for fully automatic, robust and accurate segmentation of the human pelvis and proximal femur in CT data. We propose a composite statistical shape model of femur and pelvis with a flexible hip joint, for which we extend the common definition of statistical shape models as well as the common strategy for their adaptation. We do not analyze the joint flexibility statistically, but model it explicitly by rotational parameters describing the bent in a ball-and-socket joint. A leave-one-out evaluation on 50 CT volumes shows that image driven adaptation of our composite shape model robustly produces accurate segmentations of both proximal femur and pelvis. As a second contribution, we evaluate a fine grain multi-object segmentation method based on graph optimization. It relies on accurate initializations of femur and pelvis, which our composite shape model can generate. Simultaneous optimization of both femur and pelvis yields more accurate results than separate optimizations of each structure. Shape model adaptation and graph based optimization are embedded in a fully automatic framework.

I. INTRODUCTION

For patient-specific biomechanical simulations, e.g. of the human lower limb, an accurate reconstruction of the bony anatomy from medical image data is required. This particularly applies to joint regions, as simulation results heavily depend on the anatomy of joints [1]. In CT data, bony tissue usually shows a high intensity contrast to surrounding soft tissues and may be segmented by simple thresholding. However, in joint regions, thresholding is often not sufficient for separating adjacent individual bones. Due to large slice thickness or pathological changes of bones, the joint space may be hard to detect even for human observers.

In this paper, our objective is a fully automatic, robust, accurate and consistent (i.e. non-overlapping) segmentation of both femur and pelvis from CT data. In [2] we achieved robust segmentations of the pelvis alone based on a statistical shape model (SSM) and graph optimization. However, inaccuracies occurred in the acetabulum, as no prior knowledge about the femur was exploited. Segmenting the femoral head without prior knowledge about the pelvis is even more difficult, as shown in Sec. VI. An idea to overcome this problem is to use prior knowledge about the shape and pose relations between femur and pelvis. In [3] we presented a method for simultaneous segmentation of adjacent objects.

D. Kainmueller, S. Zachow and H.C. Hege are with Zuse Institute Berlin (ZIB), Takustr. 7, 14195 Berlin, Germany. H. Lamecker is with INRIA Sophia Antipolis, France. D. Kainmueller is funded by DFG Coll. Research Center SFB760. Thanks to Heiko Seim (ZIB) for providing the shape models of femur and tibia. Thanks to Markus Heller and his group at Charité Julius Wolff Institute, Berlin, Germany, for providing the pelvis model. {kainmueller, lamecker, zachow, hege}@zib.de

This method is capable of producing accurate and consistent segmentations. However, it relies on good initializations of the adjacent objects. To solve this essential initialization task, we propose a composite statistical shape model of femur and pelvis with flexible hip joint.

A. Related Work.

1) *Articulated Shape Models*: Articulated statistical shape models were probably first suggested by Heap and al. [4]. They build a point distribution model as in [5], with the extension that training point sets may represent some subsets of points as polar coordinates with respect to certain pivot points. The resulting point distribution model captures rotational flexibility as far as it is contained in the training data. This is not possible with a purely cartesian point distribution model because of its linearity. Al-Shaher et al. [6] apply an articulated shape model as described in [4] for 2D human shape modelling and suggest a method for adapting their model to 2D landmark point data.

Articulated shape models for medical image analysis were more recently proposed by Klinder et al. [7] and Boisvert et al. [8]. Klinder et al. [7] perform a statistical analysis of rigid transformations between local vertebra coordinate systems to build an articulated spine model. The resulting statistical transformation model then captures spine flexion as contained in the training data. It is applied for rough spine segmentation in CT data, which is refined by an adaptation of multiple vertebra surface models with a feature function that penalizes overlap of adjacent vertebrae. Boisvert et al. [8] suggest a similar approach to build a statistical transformation model for capturing and analyzing the shape variability contained in a training set of scoliotic spines. They apply their model for reconstruction of the 3D spine from 2D landmarks in few radiographs.

2) *Segmentation of Articulated Joints*: Schmid et al. [9] propose a combination of physically-based deformable models and prior shape knowledge (single object SSMs and markov random field modeling of local deformations) for segmentation of femur and hip bones from low resolution MRI datasets. They achieve an average reconstruction error of 1.4 ± 1.1 mm on six datasets. The bone shapes are initialized via manually defined landmarks. No prior knowledge about shape or pose relations between femur and hip bones is exploited. Zoorofi et al. [10] evaluate and compare the suitability of a series of low level techniques to segment the femoral head and acetabulum. For each technique, they identify the percentage of “good” segmentations achieved for a set of 60 ct images based on success- and error measures.

Yin et al. [11] apply multi-object graph cuts for knee-joint bone and cartilage segmentation, starting from rough initial segmentations which they generate by adapting single-object statistical shape models of the knee bones. Based on an evaluation on 17 MR datasets, they report that 94% of the automatic segmentations required further manual border or surface editing on less than 10% of the surface. Liu et al. [12] propose an interactive, graph cuts based framework to separate foot bones in threshold segmentations of CT data. They make use of a weighted graph constructed from a binary volume. They achieve segmentation errors less than 1.5% in terms of Dice’s coefficient.

B. Contribution

As a methodological contribution, we extend the definition of SSMs [5] by rotation parameters describing the flexion in the hip joint. In contrast to related work on Articulated Shape Models [4], [6], [7], [8], we model the flexion *explicitly* and not via a statistical analysis of joint posture or articulation in a training set. As we consider the joint posture not to be a patient specific anatomical property, we see no need to analyze it statistically. The hip joint can be approximated by an ideal ball-and-socket joint, so its behaviour is known and can be modeled explicitly without analyzing a training set. We perform statistical analysis on a set of aligned training shapes to capture just the *shape variation*, undisturbed by *joint posture variation*. We extend the common approach for adapting SSMs to image data [5] to cope with the additional parameters, and we show how to generate such a model.

Our second contribution is a detailed leave-one-out evaluation on 50 CT datasets. Particularly for the femoral head, our composite model outperforms a single femur SSM. A framework that combines the adaptation of the composite model with successive graph based multi object optimization [3] solves our initial objective, as it produces robust, accurate and consistent segmentations of both femur and pelvis, without any manual interaction.

The paper is organized as follows: We describe the SSM extension in Sec. II and III. Sec. IV describes the image data this work is based on and its usage for building the articulated shape model of femur and pelvis. In Sec. V we describe our fully automatic segmentation framework. Evaluation results are given in Sec. VI and discussed in Sec. VII. Sec. VIII concludes and debates ideas on future work.

II. AN SSM WITH ARTICULATED BALL-AND-SOCKET JOINT

Our work is motivated by the idea that *joint flexibility* is not the same as *shape variation*. This implies that joint flexibility should not be modelled implicitly via the modes of shape variation of an SSM, but explicitly with separate parameters. In this Section, we describe a composite model that consists of two objects P and Q and a rotation center c of a ball-and-socket joint in which the objects are linked. The modes of shape variation cover the whole object compound. We first give the common definition of SSMs, and then extend it to capture articulated joint flexibility by introducing

additional degrees of freedom. Then we describe how to generate such a model from a training set of unaligned shapes.

A. Model Definition

A statistical shape model (SSM) is commonly [5] computed by performing principal component analysis (PCA) on a training set of aligned shapes $\mathbf{s}_i \in \mathbb{R}^{3m}$, $i = 1 \dots n$, where each \mathbf{s}_i is defined by a set of m points in space. PCA on the training set of shapes yields a linear model of the form

$$S(\mathbf{b}) = \bar{\mathbf{s}} + \sum_{k=1}^{n-1} b_k \mathbf{p}_k \quad (1)$$

where $\bar{\mathbf{s}} \in \mathbb{R}^{3m}$ represents the mean shape vector, $\mathbf{p}_k \in \mathbb{R}^{3m}$ the modes of shape variation (eigenmodes), $b_k \in \mathbb{R}$ the shape weights, and $\mathbf{b} = (b_1, \dots, b_{n-1})^T$ the shape weight vector. Any consistent shape comprised in the statistical analysis, i.e. any training shape, can now be represented by an instance $S(\mathbf{b})$ of the SSM. We denote the position of point/vertex $j \in \{1 \dots m\}$ in an instance $S(\mathbf{b})$ (i.e. its $(3j - 2)$ th, $(3j - 1)$ th and $3j$ th component) as $\vec{v}_j(\mathbf{b})$. A transformation $T \in \mathbb{R}^{4m \times 4m}$ is commonly included in the model to describe transformed shapes. T is composed of an affine transformation $t \in \mathbb{R}^{4 \times 4}$ (usually rigid with scale) and zero matrices $\mathbf{0} \in \mathbb{R}^{4 \times 4}$,

$$T = \begin{pmatrix} t & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & t & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & t \end{pmatrix} \quad (2)$$

Assuming (without a change of notation) that vertex positions $\vec{v}_j(\mathbf{b})$ are represented in homogeneous coordinates $\in \mathbb{R}^4$, the model becomes $S(\mathbf{b}, T) = T(S(\mathbf{b}))$.

Assume vertex number c with position $\vec{v}_c(\mathbf{b})$ represents the joint center of a composite model $S(\mathbf{b})$, all vertices $\vec{v}_j(\mathbf{b})$ with $j < c$ represent the first object (P), and all vertices with $j > c$ represent the second object (Q). Then we extend the model by a matrix $R \in \mathbb{R}^{4m \times 4m}$ that describes the relative transformation of object Q with respect to object P. R is composed of a matrix $r \in \mathbb{R}^{4 \times 4}$ that defines a rotation around $\vec{v}_c(\mathbf{b})$, an identity matrix \mathbf{E}_{4c} and zero matrices,

$$R = \begin{pmatrix} \mathbf{E}_{4c} & \mathbf{0}_{4c, 4(m-c)} \\ \mathbf{0}_{4(m-c), 4c} & \begin{pmatrix} r & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & r & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & r \end{pmatrix} \end{pmatrix}$$

$$r = \text{Transl}(\vec{v}_c(\mathbf{b})) \circ \text{Rot}(\alpha, \beta, \gamma) \circ \text{Transl}(-\vec{v}_c(\mathbf{b})) \quad (3)$$

with the parameters α, β, γ describing the rotation in terms of Euler angles. Note that R depends on the weight vector \mathbf{b} via the rotation center $\vec{v}_c(\mathbf{b})$ as well as on the Euler angles α, β, γ . In the following, we denote the Euler angles as $\theta :=$

(α, β, γ) . The model then becomes

$$S(\mathbf{b}, T, \theta) = (T \circ R(\theta, \vec{v}_c(\mathbf{b}))) (S(\mathbf{b})). \quad (4)$$

We call such a model *Joint Statistical Shape Model* (JSSM).

B. Model Generation

Given a set of (unaligned) training shapes $\mathbf{s}_i \in \mathbb{R}^{3m}$, $i \in \{1 \dots n\}$ with corresponding vertices, each containing vertices $\mathbf{p}_i \in \mathbb{R}^{3(c-1)}$ of an object P and vertices $\mathbf{q}_i \in \mathbb{R}^{3(m-c)}$ of an object Q as well as a joint center $\vec{v}_i \in \mathbb{R}^3$. Commonly, an affine transformation T_i (rigid or rigid with scale) is computed for each training vector that optimally aligns it to a reference vector (e.g. the first training vector), $T_i = \operatorname{argmin}_T \|T(\mathbf{s}_i) - \mathbf{s}_1\|^2$. This would not make sense in our case, as the relative transformations of the objects Q with respect to “their” objects P do not necessarily correspond (i.e. the individual objects Q are generally not all flexed “in the same way” in the joint center). Instead, as a first step of alignment, we align the vertices \mathbf{p}_i of object P only. Therefore we compute a transformation $T_i \in \mathbb{R}^{4m \times 4m}$ for each training vector such that just its components \mathbf{p}_i are aligned with reference components \mathbf{p}_1 , namely $T_i = \operatorname{argmin}_T \|T(\mathbf{p}_i) - \mathbf{p}_1\|^2$. This results in new training vectors $T_i(\mathbf{s}_i)$. A transformation $R_i \in \mathbb{R}^{4m \times 4m}$ is then computed for each $T_i(\mathbf{s}_i)$ that aligns the orientation of its components \mathbf{q}_i to the orientation of reference components \mathbf{q}_1 . Therefore we compute a transformation (rigid with scale) $A_i = \operatorname{argmin}_A \|A(T_i(\mathbf{q}_i)) - \mathbf{q}_1\|^2$. We omit the scale factor from A_i , so that the transformed object Q is scaled with the same factor as the respective object P, but its orientation corresponds to reference components \mathbf{q}_1 . This results in a rigid transformation B_i . We add a translation C_i such that the joint center $T_i(\vec{v}_i)$ is a fixed point of $R_i := C_i \circ B_i$, i.e. $R_i(T_i(\vec{v}_i)) = T_i(\vec{v}_i)$. Without change of notation we assume that T_i and R_i are replenished with identity and zero matrices as in Eq. 2 and 3. Now we perform PCA over the set of training vectors $(R_i \circ T_i)(\mathbf{s}_i)$ to generate the mean shape $\bar{\mathbf{s}}$ and $n - 1$ shape vectors \mathbf{p}_k . Note that the alignment of the training shapes is optimal (in terms of the mean distance of corresponding vertices) for the components of object P, but not necessarily for Q or for the object compound. In effect, the model we generate is hierarchic in the sense that firstly object P is chosen for alignment to a reference, and subsequently object Q is rotated in a corresponding way.

III. ADAPTING A JSSM TO A TARGET SHAPE

In this Section, we propose an approach for fitting a JSSM to a target shape. This is essential for image driven iterative model adaptation, where the model is fitted to a temporary mesh in each step. For a description of adaptation strategies for single shape models, see e.g. [5]. Transformation and shape parameters are usually adapted separately. Our JSSM fitting approach copes with the additional degrees of freedom in the articular joint in both transformation and shape adaptation. For an instance $S(\mathbf{b}, E_{Am}, 0)$ of a shape model, we denote object P’s vertices (i.e. the first $4(c-1)$ components) as $\mathbf{p}(\mathbf{b})$, object Q’s vertices as $\mathbf{q}(\mathbf{b})$ and the joint center as

$\vec{v}_c(\mathbf{b})$. Vertices of a model $S(\mathbf{b}, T, \theta)$ with transformations are denoted as $\mathbf{p}(\mathbf{b}, T, \theta)$ or alike. As for the target shape $\mathbf{s}^* \in \mathbb{R}^{4m}$, we denote object P’s vertices as \mathbf{p}^* , object Q’s vertices as \mathbf{q}^* , and joint center as \vec{v}_c^* .

A. Transformation Adaptation

Given a shape instance $S(\mathbf{b}, E_{Am}, 0)$, we determine a new transformation T_{new} and rotation θ_{new} , while the shape parameters remain constant. We compute T_{new} by alignment of only object P’s points: $T_{new} = \operatorname{argmin}_T \|\mathbf{p}(\mathbf{b}, T, 0) - \mathbf{p}^*\|^2$. Then we compute θ_{new} by rotational alignment of $\mathbf{q}(\mathbf{b})$ to inversely transformed target points around the joint center $\vec{v}_c(\mathbf{b})$: $\theta_{new} = \operatorname{argmin}_\theta \|\mathbf{q}(\mathbf{b}, E_{Am}, \theta) - T_{new}^{-1}(\mathbf{q}^*)\|^2$. Note that global transformation and local rotation are optimized separately here. Thus the combination of parameters (T_{new}, θ_{new}) does not necessarily minimize the distance of the adapted model to the target shape.

B. Shape Adaptation

Given a model instance $S(\mathbf{b}, T, \theta)$. While transformation and rotation angles stay constant here, we want to determine adapted shape parameters \mathbf{b}_{new} . Therefore we compute the shape parameters that optimally fit the untransformed model vertices to the inversely transformed target vertices:

$$\mathbf{b}_{new} = \operatorname{argmin}_{\mathbf{b}} \|S(\mathbf{b}, E_{Am}, 0) - \operatorname{Inv}(\mathbf{p}^*, \vec{v}_c^*, \mathbf{q}^*)\|^2 \quad (5)$$

$$\text{with } \operatorname{Inv} = (R^{-1}(\theta, T^{-1}(\vec{v}_c^*)) \circ T^{-1}) \quad (6)$$

Note that the shape adaptation also changes the rotation matrix R of the shape model: $R_{new} = R(\theta, \vec{v}_c(\mathbf{b}_{new}))$. Also the inverse rotation matrix $R^{-1}(\theta, T^{-1}(\vec{v}_c^*))$ is generally not the inverse of the rotation matrix before adaption, $R(\theta, \vec{v}_c(\mathbf{b}))$. In effect, \vec{b}_{new} does not necessarily minimize the distance of $S(\mathbf{b}, T, \theta)$ to the *untransformed* target vertices $(\mathbf{p}^*, \vec{v}_c^*, \mathbf{q}^*)$.

IV. A JSSM OF PELVIS AND FEMUR WITH FLEXIBLE HIP JOINT

In this Section we describe the generation of an application specific model. The training shapes are reconstructed from tomographic image data. The individual bent of the hip joint has to be eliminated from each training shape before performing PCA on the training set of shapes. As for the image data, datasets containing both the complete pelvis and femur were not available for this study. This is a common situation, as in most clinical cases only one object of interest is recorded, which then contains only part of the adjacent objects. It is not desirable to include in the model only the part which is captured by all datasets. This is because we want to use our shape model to segment new, unseen image data that may capture a wider part of this object.

A. Image Data

For model generation we dispose of 50 pelvis CT datasets (set P). With a voxel size of about $0.9 \times 0.9 \times 5$ mm all datasets approximately have the same resolution. They all display the patient specific pelvis, the femoral head, and various amounts of the femoral body. At maximum,

about half the femoral body is displayed. The pelvis, the right proximal femur, and the center of the right hip joint were manually labelled by a human observer (gold standard labels). For the pelvis, three labels were defined, one for each hip bone and one for the sacrum. For more details on set P, see [2]. Additionally we make use of 30 femur CT datasets (set F) with a slice resolution of about 0.5×0.5 mm and slice distances of 0.5 to 1.5 mm. All datasets display the whole femur, but only part of the pelvis. The right femur was manually labelled by human observers.

B. Training Shapes

From set P, triangular pelvis surfaces with corresponding mesh topologies were established from the gold standard labels as described in [2]. Among other anatomically relevant structures on the pelvis surfaces, the acetabulum was identified manually and forms a separate patch. Analogously, triangular femur surfaces with corresponding mesh topologies were established from the labels of set F. Here, the femoral head was identified manually and forms a separate patch. From the femur surfaces, we generated an SSM of the complete femur. This model was then fitted to the partial femur labels contained in set P. Thus we get approximate patient specific femur surfaces for P with corresponding mesh topologies. Their distal ends are purely extrapolated by the model. However, this way we can cope with the varying amount of the femoral body displayed in the CT datasets of P. Now, together with the manually marked joint centers, we have 50 unaligned training vectors $\mathbf{s}_1, \dots, \mathbf{s}_n$ composed of 29619 pelvis vertices, one joint center and 8461 femur vertices. We align the training vectors as proposed in Sec. II-B, and perform PCA on the resulting vectors.

V. SEGMENTATION FRAMEWORK

Our fully automatic framework for pelvis and femur segmentation in CT data is composed of the following steps: Initialization, image driven JSSM or SSM adaptation, and graph based optimization of single or coupled objects. This Section describes all steps as well as the overall segmentation algorithm with various parametrizations as used in our experiments (see Sec. VI).

A. Initialization

The pose initialization of the pelvis in CT data as contained in our framework closely follows a global approach for 3D object detection introduced by Khoshelham [13]. It is based on the Generalized Hough Transform (GHT). For a detailed description, see [2]. When using the single femur SSM, as for pose initialization in (pelvis-) CT data, we can reuse the transformation that results from pelvis initialization, as the femur SSM is located in an anatomically consistent way with respect to the pelvis SSM.

B. Image Driven JSSM Adaptation

Segmentation using the JSSM (4) is the task of finding transformation and shape parameters such that $S(\mathbf{b}, T, \theta)$ approximates the unknown target shape $R^* \in \mathbb{R}^{3m}$ as

good as possible. The computation proceeds iteratively. Let $R^i = S(\mathbf{b}^i, T^i, \theta^i)$ denote the segmentation in iteration i : A displacement vector field $\Delta R^i \in \mathbb{R}^{3m}$ is computed that assigns a vector $\Delta r_j \in \mathbb{R}$ to each vertex $j \in \mathbb{N}$ of R^i . It describes the desired deformation of the model towards R^* in the underlying image data I . Then, transformation and rotation parameters (T, θ) and shape parameters \mathbf{b} of the JSSM are alternately adapted to the target shape $(R^i + \Delta R^i)$ as described in Sec. III. For details on single object SSM adaptation, see [2]. The following paragraph explains how the displacement vector field ΔR^i is generated.

C. Image Analysis

The displacement vector field ΔR for surface R is computed by analyzing 1D intensity profiles in the image data $I : \mathbb{R}^3 \rightarrow \mathbb{R}$: For each vertex j of R , I is sampled over a length L along the surface normal u_j at vertex position v_j . A cost function $c_j : P_j \rightarrow \mathbb{R}_0^+$ is computed on the set of sampling points $P_j = \{v_j^n := v_j + (\frac{n-1}{N_j-1} - 0.5) \cdot L \cdot u_j | n = 1, \dots, N_j\}$. The displacement vector at vertex j is then defined as $\Delta r_j = v_j^* - v_j$, with $v_j^* = \operatorname{argmin}_{v_j^n} c_j(v_j^n)$. In the following we drop the index j for clarity. Our framework employs three cost functions: A “standard” cost function cS , a “conservative” cost function cC and an “ignore” cost function cI . We define $cS(v^n) =$

$$\begin{cases} (2i+1) \left(\frac{-g}{dI(v^n)} + 2 \frac{|m-n|}{N} \right) & \text{if (8)} \\ 7 \left(\frac{-g}{dI(v^n)} + 2 \frac{|m-n|}{N} \right) & \text{if (9)} \\ 30 + 2 \frac{|m-n|}{N} & \text{else, with conditions} \end{cases} \quad (7)$$

$$I(v^n) \in [t+iw, t+(i+1)w] \wedge dI(v^n) < -g, i=0, 1, 2 \quad (8)$$

$$I(v^n) \in [t, t+3w] \text{ and } dI(v^n) < -0.5 \cdot g \quad (9)$$

Here, t defines an intensity threshold, w an intensity window width, and g a threshold for gradient magnitude. $dI(v^n)$ denotes the directional derivative of I along u , and $m \in \{1, \frac{N+1}{2}, N\}$ is a “preferred sample index”, i.e. sample points closer to v_m get slightly lower costs. We define cI as $cI(v^n) = 0$ if $n = \frac{N+1}{2}$, and $cI(v^n) = 10$ otherwise. cC is defined as $cC(v^n) = cS(v^n)$ if there are two distinct intervals of sample points with $cS(v^n) < 30$, and at least one sample point inbetween with $dI(v^n) > g_{min}$, and $cC(v^n) = cI(v^n)$ otherwise. Thus, features are ignored if only one downward gradient (from bone to background) is found along the profile. This makes sense for the femoral head: A downward gradient is often not present at the femur surface, but only inside the pelvis. Note that we determined all cost functions in a heuristic manner. None of the parameters were automatically learned from the training images. The cost functions are specifically designed to segment bones in CT data, especially in joint regions.

D. Graph Based Multi-object Optimization

In JSSM adaptation, the minimum cost sample point on a profile serves as a desired (locally optimal) new position for the respective vertex. However, graph cut algorithms allow

Step on Model	Details	
1. GHT	all	as in [2]
2. SMA	all	Profile length $L := 50\text{mm}$, sampling points $N := 50$, intensity parameters $t := 120$, $w := 200$, gradient threshold $g := 50/\text{mm}$. Adapt only transform. parameters.
	F	Cost function cI or cS with $m = \frac{N+1}{2}$ for femoral head, cS with $m = N$ otherwise.
	P	cS with $m = \frac{N+1}{2}$ for acetabulum, $m = N$ otherwise.
	C	cS with $m = \frac{N+1}{2}$ for acetabulum and femoral head, $m = N$ otherwise.
3. SMA	all	$L := 20\text{mm}$. Alternating update of transformation parameters and the first 30 shape parameters, (b_1, \dots, b_{30}) .
4. SMA	all	$L := 10\text{mm}$. Consider all shape parameters \vec{b} in update.
5. COU	C	$L := 10\text{mm}$ or 20mm .
6. OPT	all	$L := 10\text{mm}$ or 20mm , $N_p := 20$ or 40 . Shape preserving 1mm .
	F	cC or c_k^S with $m = \frac{N+1}{2}$ for femoral head, or cC only for coupled femur patches resulting from COU .
	C	cC for femoral head. Min/max distance constraint $0\text{mm}/8\text{mm}$.

TABLE I

VARIANTS OF THE OVERALL SEGMENTATION ALGORITHM. FROM TOP TO BOTTOM, PARAMETERS DO NOT CHANGE IF NOT NOTED OTHERWISE.

for a *global* optimization of the sum of costs for each vertex displacement while respecting hard constraints on surface smoothness and optionally on the distance between multiple surfaces. For the latter, multiple surfaces must be coupled with *shared intensity profiles* at individual vertices. For more details on graph construction see [14]. In [3], we proposed a construction algorithm for shared intensity profiles on pairs of adjacent triangular surface meshes P and Q that yields a bijective mapping between *coupled patches* on both meshes. In the process, the connectivity of parts of the original meshes is modified. For more details, see [3].

E. Overall Segmentation Algorithm

The segmentation algorithm consists of a series of steps combining the methods presented above, see Table I. The following methods are applied: pose initialization (GHT), SSM or JSSM adaptation (SMA), coupling with shared displacement directions (COU) and optimization via graph cuts (OPT). For each step the particular parameters are listed in the *Details* column. Details that are the same for all models are labelled as **all** in the second column, details for a single femur model as **F**, for a single pelvis model as **P**, and for a combined or coupled model with **C**.

VI. RESULTS

We segmented the 50 pelvis CT datasets of set P (see Sec. IV) with the different variants of the fully automatic algorithm as defined in Table I. For the segmentation of dataset i , we removed the respective training surface s_i from the pelvis SSM or JSSM. The complete segmentation of one CT stack took about 4 : 20 to 6 : 00 minutes on a desktop PC (2,66GHz Core, 8GB RAM), depending on the shape model (femur, pelvis or composite model) and the type of optimization.

	SFM1		SFM2		SPM		JSSM	
	ad	md	ad	md	ad	md	ad	md
F	3.1	8.3	2.2	5.7	not in model		1.7	4.5
	1.4	3.7	1.0	2.8			0.6	1.5
A	not in model		not in model		1.5	5.5	1.4	5.4
					0.5	2.0	0.5	1.9

TABLE II

SSM AND JSSM ADAPTATION RESULTS FOR FEMORAL HEAD (F) AND ACETABULUM (A). ALL ENTRIES IN MM. BOLD: AVERAGE MEASURE OVER 50 DATASETS. BELOW: STANDARD DEVIATION.

A. Evaluation Method.

The main focus of our evaluation is to determine how accurate the automatic segmentations are in the area of the right hip joint. We particularly focus on the accuracy of JSSM adaptations as compared to single SSM adaptations, and analogously on the accuracy of coupled optimization as compared to single object optimizations. As accuracy measures we computed the average surface distance (ad) and maximum surface distance (md) of the right acetabulum and femoral head from the respective gold standard surface. Each structure is defined as a set of vertex numbers on the gold standard meshes. Thus we could compute the distance from these vertices to the respective automatic segmentation. Note that **ad** and **md** are asymmetric measures.

A second focus of the evaluation is to measure the overall accuracy of the pelvis reconstruction. This way our results can be compared with previous and related work. We performed an evaluation of left and right hip bone, sacrum, and the complete pelvis as in [2]. We compute the 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 [15].

B. Results for Shape Model Adaptation.

We carried out two experiments with the single femur SSM, namely SFM1 with cost function cS for the femoral head, and SFM2 with cost function cI for the femoral head. Experiments SPM and JSSM were performed with the single pelvis SSM and the JSSM, respectively, both with cost function cS . Table II lists the average error metrics for the femoral head and acetabulum. For experiment JSSM, Table IV gives the average error metrics for the pelvic bones.

C. Results for Graph Based Optimization.

For the femur, we carried out four single object optimization experiments, all of which starting from the femur results of JSSM, namely SF1 with 10mm profiles and cost function cS for the femoral head, SF2 with 10mm profiles and cost function cC for the femoral head, SFC1 and SFC2 with 10mm profiles and 20mm profiles, respectively, and cost function cC for the coupled femur patches. For the pelvis, two single object optimization experiments were carried out, both with cost function cS and 20mm profiles, namely SP1, starting from the results of PSSM, and SP2, starting from the pelvis results of JSSM.

	SF1		SF2		SFC1		SFC2		SP1		SP2		C1		C2	
	ad	md	ad	md	ad	md	ad	md	ad	md	ad	md	ad	md	ad	md
F	1.5	5.0	1.4	4.5	1.3	4.4	1.7	5.8	not in model	not in model	1.2	4.2	1.1	4.3	0.7	2.1
	0.7	2.1	0.6	1.8	0.6	1.8	1.0	2.1			0.5	1.6	0.5	1.4		
A	not in model	not in model	not in model	not in model	not in model	not in model	not in model	not in model	0.9	4.8	0.9	4.7	0.9	4.7	0.8	4.7
									0.5	2.4	0.5	2.4	0.5	2.1	0.5	2.2

TABLE III

SINGLE AND COUPLED OPTIMIZATION RESULTS FOR FEMORAL HEAD (F) AND ACETABULUM (A). ALL ENTRIES IN MM. BOLD: AVERAGE ON 50 DATASETS. BELOW: STANDARD DEVIATION.

	JSSM				C2			
	OE	AD	RMS	MD	OE	AD	RMS	MD
	[%]	[mm]	[mm]	[mm]	[%]	[mm]	[mm]	[mm]
All	19.9 (3.0)	1.0 (0.3)	2.0 (0.6)	15.4 (4.5)	12.7 (3.1)	0.6 (0.2)	1.6 (0.6)	15.7 (5.0)
RHB	18.0 (2.3)	0.7 (0.1)	1.3 (0.2)	8.5 (2.1)	9.7 (2.0)	0.3 (0.1)	0.8 (0.2)	7.6 (1.9)
LHB	22.8 (3.8)	1.0 (3.8)	1.8 (0.4)	11.4 (2.2)	14.8 (3.7)	0.6 (0.2)	1.5 (0.4)	11.4 (2.4)
S	23.5 (5.4)	1.6 (0.6)	2.7 (1.2)	14.6 (4.9)	18.5 (5.9)	1.1 (0.6)	2.4 (1.2)	14.9 (5.4)

TABLE IV

EVALUATION RESULTS ON PELVIS: JSSM ADAPTATION AND COUPLED OPTIMIZATION. OVERALL PELVIS (ALL), RIGHT HIP BONE (RHB), LEFT HIP BONE (LHB), AND SACRUM (H). EACH MEASURE WITH STANDARD DEVIATION IN BRACKETS.

Two multi object optimization experiments were performed, namely C1 with 10mm profiles and C2 with 20mm profiles. Both started from the coupled femur and pelvis results of JSSM and applied cost function cC for the coupled femur patches. No multi object optimization experiment was performed starting from single SSM results.

Table III lists the average error metrics for the femoral head and acetabulum. For C2, Table IV lists the average error metrics for the pelvic bones. Note that the metrics ad and AD, and also md and MD, cannot be directly compared, as their computation methods differ fundamentally. However, as a reference example, we computed both ad and AD for the overall pelvis surfaces from C2. The averages are 0.8 ± 0.2 mm and $AD 0.6 \pm 0.2$ mm.

VII. DISCUSSION

A. Shape Model Adaptation.

Concerning the femur, the JSSM yields robust segmentations of the femoral head (ad 1.7 ± 0.6 mm). We could not achieve such an accuracy with the single femur SSM, which yielded the best result when ignoring the femoral head in the adaptation process (ad 2.2 ± 1.0 mm). As the femoral head is then completely extrapolated by the model, oversegmentations as well as undersegmentations occur, see Fig. 1(a) and 1(b). When treating the femoral head of the single femur SSM with the standard cost function (as for the JSSM!), it adapts to “false” image features in many cases, as high gradients are commonly present inside the pelvis where the bone density decreases, and also on the pelvis surface opposite the acetabulum, see Fig. 1(c) and 1(d).

As for the pelvis, the single pelvis SSM already yields robust segmentations of the acetabulum (ad 1.5 ± 0.5 mm). The JSSM performs slightly, but not significantly, better (ad 1.4 ± 0.5 mm). The accuracies of one individual JSSM result and the respective single-SSM result differ at maximum 0.3 mm (ad). No major differences are visible. However, we

consider the average results as a hint that in the JSSM the acetabulum is deterred from features inside the femur.

B. Optimization.

For the femur, coupled optimization is significantly better than single optimization (ad 1.1 versus 1.4mm). It is also more robust in the sense that longer intensity profiles can be used without the risk of the femoral head’s adapting to image features inside the pelvis. Coupled optimization with 20mm intensity profiles produces the overall best result for the femoral head. Single optimization with 20mm intensity profiles produces large oversegmentations, the femoral head adapting to features inside the pelvis. Therefore we used 10mm intensity profiles for most single femur optimization experiments. The best single optimization result was reached with 10mm intensity profiles and the conservative strategy for the femoral head. However, in some cases 10mm intensity profiles do not reach any image features, see Fig. 2(a). Furthermore even if features are reached, the conservative strategy often ignores them, see Fig. 2(b). For coupled optimization however, 20mm profiles could be used, as the femur mesh is prevented from adapting to the inside-pelvis features by the pelvis mesh. Apart from the constraints imposed by coupled optimization, another advantage in experiments C1 and C2 is that the coupled patches are defined on the femoral head and can be treated separately, in our case with the conservative strategy. We performed SFC1 and SFC2 for direct comparison to C1 and C2.

As for the pelvis, coupled optimization is not significantly better than single optimization (ad 0.8 versus 0.9mm). Although most individual results are similar, coupled optimization outperforms single optimization in a few cases due to the distance constraints imposed by the shared intensity profiles, see Fig. 2(c-e).

VIII. CONCLUSION

We proposed a fully automatic framework for simultaneous segmentation of the pelvis and proximal femur from CT




Fig. 1. Black: JSSM. White: Gold standard. (a,b) Red/grey: SFM2. Ignoring the femoral head causes (a) under- and (b) oversegmentations. (c,d) Red/grey: SFM1. Femoral head adapts to features (c) inside and (d) on the pelvis surface.




Fig. 2. Black: coupled optimization with 20mm profiles (C2). White: Gold standard. (a,b) Red/grey: Single femur optimization (SF2). (a) 10mm profiles too short. (b) Conservative strategy ignores features. (c-e) Red/grey: Single pelvis optimization, 20mm profiles (SP2). (c,d) Minimum and (e) maximum distance constraints prevent mis-adaptations in coupled optimization.

data. The framework is based on a composite statistical shape model of pelvis and femur with a flexible hip joint. We gave a general definition of such a model, and showed how it can be adapted to target shapes. We evaluated the framework on 50 CT datasets. The composite statistical shape model yields more accurate segmentation results in the hip joint region than single statistical shape models. Furthermore, the composite shape model segmentation results are suitable as initializations for fine grain multi object segmentation based on a method for coupling adjacent meshes and graph cuts. We showed that coupled graph based optimization robustly leads to very accurate segmentations of the hip joint region. The average mean surface distance from gold standard segmentations is 0.8mm for the acetabulum, and 1.1mm for the femoral head. As future work we are to establish a composite model with more than two objects and flexibility in other types of joints than a ball and socket joint, e.g. hinge joints.

REFERENCES

- [1] M. Heller, G. Bergmann, G. Deuretzbacher, L. Dürselen, M. Pohl, L. Claes, N. Haas, and G. Duda, "Musculo-skeletal loading conditions at the hip during walking and stair climbing," *J Biomech*, vol. 34, pp. 883–93, 2001.
- [2] H. Seim, D. Kainmueller, M. Heller, H. Lamecker, S. Zachow, and H.-C. Hege, "Automatic segmentation of the pelvic bones from ct data based on a statistical shape model," in *Eurographics Workshop on Visual Computing for Biomedicine (VCBM)*, Delft, Netherlands, 2008, pp. 93–100.
- [3] D. Kainmueller, H. Lamecker, S. Zachow, and H.-C. Hege, "Coupling Deformable Models for Multi-object Segmentation," in *ISBMS*, ser. LNCS, F. Bello and E. Edwards, Eds., vol. 5104. Springer, 2008, pp. 69–78.
- [4] T. Heap and D. Hogg, "Extending the point distribution model using polar coordinates," *Image Vision Comput.*, vol. 14, no. 8, pp. 589–599, 1996.
- [5] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active Shape Models - Their Training and Application," *Comput. Vis. Image Underst.*, vol. 61, no. 1, pp. 38–59, 1995.
- [6] A. A. Al-Shaher and E. R. Hancock, "Modelling human shape with articulated shape mixtures," in *SSPR/SPR*, ser. Lecture Notes in Computer Science, A. L. N. Fred, T. Caelli, R. P. W. Duin, A. C. Campilho, and D. de Ridder, Eds., vol. 3138. Springer, 2004, pp. 304–314.
- [7] T. Klinder, R. Wolz, C. Lorenz, A. Franz, and J. Ostermann, "Spine segmentation using articulated shape models," in *MICCAI (1)*, ser. LNCS, D. N. Metaxas, L. Axel, G. Fichtinger, and G. Székely, Eds., vol. 5241. Springer, 2008, pp. 227–234.
- [8] J. Boisvert, F. Chieriet, X. Pennec, H. Labelle, and N. Ayache, "Geometric variability of the scoliotic spine using statistics on articulated shape models," *IEEE Trans. Med. Imaging*, vol. 27, no. 4, pp. 557–568, 2008.
- [9] J. Schmid and N. Magnenat-Thalmann, "MRI Bone Segmentation Using Deformable Models and Shape Priors," in *MICCAI (1)*, ser. LNCS, D. N. Metaxas, L. Axel, G. Fichtinger, and G. Székely, Eds., vol. 5241. Springer, 2008, pp. 119–126.
- [10] R. A. Zoroofi, Y. Sato, T. Sasama, T. Nishii, N. Sugano, K. Yonenobu, H. Yoshikawa, T. Ochi, and S. Tamura, "Automated segmentation of acetabulum and femoral head from 3-d ct images," *IEEE TITB*, vol. 7, no. 4, pp. 329–343, 2003.
- [11] Y. Yin, X. Zhang, and M. Sonka, "Optimal multi-object multi-surface graph search segmentation: Full-joint cartilage delineation in 3d," in *MIUA 2008*, S. McKenna and J. Hoey, Eds., 2008, pp. 104–108.
- [12] L. Liu, D. Raber, D. Nopachai, P. Commean, D. Sinacore, F. Prior, R. Pless, and T. Ju, "Interactive separation of segmented bones in ct volumes using graph cut," in *MICCAI (1)*, ser. LNCS, D. N. Metaxas, L. Axel, G. Fichtinger, and G. Székely, Eds., vol. 5241. Springer, 2008, pp. 296–304.
- [13] K. Khoshelham, "Extending generalized hough transform to detect 3d objects in laser range data," in *Proceedings of the ISPRS Workshop Laser Scanning 2007 and SilviLaser 2007*, ser. IAPRS, P. Rönholm, H. Hyypää, and J. Hyypää, Eds., vol. XXXVI, 2007, pp. 206–211.
- [14] K. Li, X. Wu, D. Z. Chen, and M. Sonka, "Optimal surface segmentation in volumetric images—a graph-theoretic approach," *IEEE TPAMI*, vol. 28, no. 1, pp. 119–134, 2006.
- [15] B. van Ginneken, T. Heimann, and M. Styner, "3D Segmentation in the Clinic: A Grand Challenge," in *3D Segmentation in the Clinic: A Grand Challenge*, T. Heimann, M. Styner, and B. van Ginneken, Eds., 2007, pp. 7–15.