Multi-object Segmentation with Coupled Deformable Models

Dagmar Kainmueller, Hans Lamecker, Stefan Zachow

Zuse Institute Berlin,
Takustr. 7, 14195 Berlin
⟨`kainmueller@zib.de`⟩

Abstract

For biomechanical simulations, the segmentation of multiple adjacent anatomical structures from medical image data is often required. If adjacent structures are barely distinguishable in image data, in general automatic segmentation methods for single structures do not yield sufficiently accurate results. To improve segmentation accuracy in these cases, knowledge about adjacent structures must be exploited. Optimal graph searching (graph cuts) based on deformable surface models allows for a simultaneous segmentation of multiple adjacent objects. However, this method requires a correspondence relation between vertices of adjacent surface meshes. Line segments, each containing two corresponding vertices, may then serve as shared displacement directions in the segmentation process. In this paper we propose a scheme for constructing a correspondence relation in adjacent regions of two arbitrary surfaces. This correspondence relation implies shared displacement directions that we apply for segmentation with deformable surfaces. Here, overlap of the surfaces is guaranteed not to occur. We show correspondence relations for regions on a femoral head and acetabulum and other adjacent structures, as well as an evaluation of segmentation results on 50 ct images of the hip joint.

1 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 location of joints Heller et al. [2001]. In CT data, bony tissue can usually be reconstructed by simple thresholding. However, in joint regions, thresholding is often not sufficient for separating adjacent individual bones from each other. Due to large slice thickness or pathological changes of the bones, the joint space may be hard to detect even for a human observer. Fig. 1a and 1b show exemplary situations.

We can achieve good initializations for bony structures (e.g. pelvis, femur) by our segmentation framework which is based on initialization with the Generalized Hough Trans-

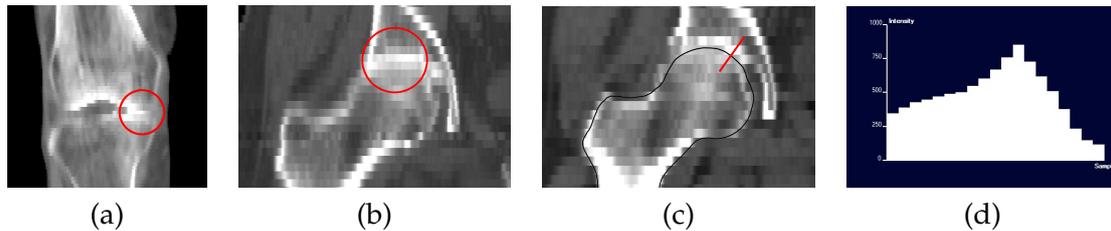(a)   (b)   (c)   (d)

Figure 1: (a) CT data of distal femur and proximal tibia, slice thickness 2mm, (b) acetabulum and proximal femur, slice thickness 4,6mm. The joint space is barely visible in encircled areas. (c) Acetabulum and proximal femur, slice thickness 5mm, with surface model cross-section (black) and domain of intensity profile (line segment). (d) Intensity profile for domain in (c).

form (GHT) and the adaptation of statistical shape models (SSM) Seim et al. [2008]. Segmentations with SSMs, as originaly proposed by Cootes et al. Cootes et al. [1995], yield very good initializations, but lack precision, since previously unknown patient-specific anatomy is generally not contained in the model. With Graph based optimization on Deformable Models, more precise segmentations can be achieved, but they suffer from a loss of shape knowledge, causing inaccurate interpolations where the object to be segmented cannot be distinguished from adjacent structures in image data. Furthermore, the lack of image information may generally lead to overlapping segmentation results when adjacent structures are segmented separately.

A basic idea for improving segmentation results and simultaneously solving the overlap problem is to segment multiple adjacent objects at the same time and incorporate some knowledge about their spatial relationship. The problem we address in this paper is how to establish a suitable coupling of two arbitrary adjacent deformable surface models (triangular meshes), assuming that a good initialization of the two models is given. The contribution of this work is a construction scheme for shared displacement directions for two arbitrary surfaces, that is, line segments along which vertices of both surfaces can be displaced in a deformable surface segmentation framework. We used the coupling realized by these shared displacement directions for fine grain multi-object segmentation based on graph cuts. This paper presents an evaluation on 50 ct images of the pelvic bones and proximal femur. Sufficiently accurate initial segmentations were achieved with an articulated, composite statistical shape model of the pelvis and proximal femur, as presented in Kainmueller et al. [Sept. 2009].

2 Related Work

Costa et al. Costa et al. [2007] employ a non-overlapping constraint for coupled segmentation of prostate and bladder with deformable models. They propose a force that drives two models apart if intersections occur in the segmentation process. This method principally allows for free form deformations of the models while coping with overlap. If present, a statistical shape model can be enforced on one of the models. If one structure is better distinguishable from the background than the other, an asymmetric non-overlap force can be applied. This approach yields promising results in prostate and bladder segmentation. However, displacements are not found *simultaneously* for both objects: Apart from the non-overlap force

which only exists if an overlap has already occurred, one object does not take into account any knowledge about the presence of the other object in displacement computation.

Tsai et al. Tsai et al. [2003] build composite statistical shape models by applying principal component analysis to a training set of implicit (signed distance function) representations of multiple objects. They apply such models to the segmentation of subcortical brain structures and male lower abdominal structures (prostate gland, rectum, obturator muscles). Babalola et al. Babalola et al. [2007] build a composite active appearance model on the basis of explicit (surface-mesh) representations of multiple subcortical brain structures. They apply this model to obtain a good initialization of brain structure models to accurately segment the caudate in a single object segmentation framework. Composite statistical shape models yield a tight coupling of the deformations of multiple objects. Ideally, no overlap between adjacent objects should be possible in model space. Nevertheless, neither of the two approaches allows for a fine grain free form multi-object segmentation, as model deformation is bound to the respective shape space.

Li et al. Li et al. [2005] solve the overlap problem with optimal graph searching in a deformable model segmentation framework. They apply their method to the segmentation of bone and cartilage in 3D MRI of human ankles. Vertex normals of the bone surface are used as shared displacement directions for bone and cartilage surfaces. Thus, one direction may be used to search for two object boundaries, thereby allowing for a completely simultaneous segmentation process. Methods involving shared displacement directions have been described for surfaces on which corresponding vertices are easily found. This holds for height field or cylindrical surfaces in regular grids Li et al. [2006], or if one surface can be obtained by displacing the other along its vertex normals Li et al. [2005].

Furthermore, Yin et al. Yin et al. [2008] apply multi-object graph cuts for knee-joint bone and cartilage segmentation. Similar to the method we propose in this paper, they couple adjacent surfaces with shared displacement directions. However, they do not deal with the problem that overlapping segmentation results might occur if the shared displacement directions do not fulfil certain conditions. Recently, Liu et al. Liu et al. [2009] have proposed a method for segmenting the inner and outer wall of airway trees and lung vascular trees. Their approach, which also uses the multi-object graph cuts method proposed by Li et al. Li et al. [2005], is similar to ours as it makes use of medial axes of an inital segmentation to define displacement directions with an appropriate length while avoiding self-intersections. However, the displacement directions are specifically designed for tree-like structures with multiple walls. They cannot be directly used for arbitrary adjacent surface meshes.

To the best of our knowledge, there is at present no method that generates shared displacement directions on arbitrary adjacent surfaces that are guaranteed to generally prevent overlapping segmentation results.

3 Multi-object segmentation with graph cuts

In segmentation with deformable surfaces, *intensity profiles* at each vertex of the surface mesh are commonly used to guide the deformation process. An intensity profile is a set of intensity samples taken from image data along a line segment at a vertex of the surface mesh, see Fig. 1c and 1d. Note that the term *intensity profile* or just *profile* may be used to refer to the sampled intensities only, yet we use it to refer to the domain of the sampling, i.e. a line segment, as well. Fig. 2a shows exemplary profiles on a triangular surface mesh. Profiles are commonly

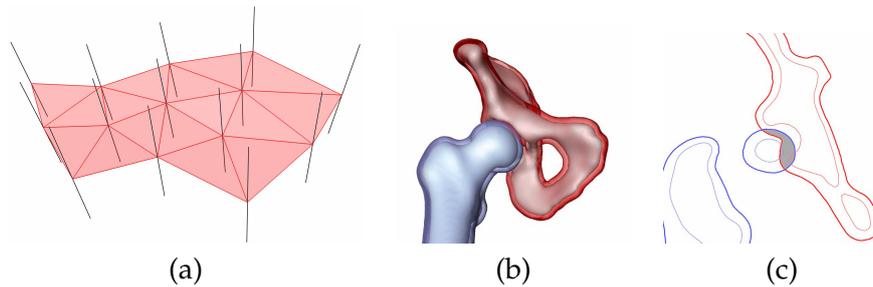(a)   (b)   (c)

Figure 2: (a) Triangular surface mesh with profiles in normal direction. (b) Proximal femur and ilium. Transparent: isosurfaces in distance transforms. (c) 2D cross-section. Thin lines: surface contours. Thick: isosurface contours. Grey: potential overlap region.

defined to run along vertex normals, but other directions may be chosen as well. On each profile, a cost function is derived from image data for a number of equidistant sampling points. The minimum cost sample point on a profile may serve as a desired (locally optimal) position for the respective vertex.

However, graph cut algorithms allow for a *global* optimization of the sum of costs for each vertex displacement while respecting hard constraints on the distance of multiple objects and on single object smoothness. Multiple surfaces can be coupled with *shared intensity profiles* at individual vertices. Minimum and maximum distance constraints can be enforced on the new vertex positions that the graph optimization finds for each surface on each shared profile. For more details on graph construction see Li et al. [2006]. The optimization problem can be solved in polynomial time, as discussed in Boykov and Kolmogorov [2004].

4   Non-overlapping surface deformations

Before shared profiles can be constructed, surface models of adjacent structures must be initialized to have a reasonable spatial relation to each other. For bony structures of the lower limb we can achieve initializations in CT data with a maximum surface distance of about 1cm from a manual expert segmentation. For two such well-initialized surface models, we identify a potential overlap region by dilating each surface by a user-defined amount l (Fig. 2b and 2c). Our concern is to couple the surfaces with shared profiles wherever single profiles of length $2 \cdot l$ would reach the overlap region.

4.1   Properties of shared intensity profiles

In the following we give the conditions which must hold on shared profiles to generally prevent any overlap after surface deformation. One necessary condition is imposed on new vertex positions of coupled surface meshes A and B:

  1. Let x be a vertex on A and $m(x)$ the corresponding vertex on B. Then the new position x_A^* of x must lie closer to A than the new position x_B^* of $m(x)$ in profile direction.

In the following we assume that this condition holds, i.e. $\|x_A^* - x\| < \|x_B^* - x\|$, see also Fig. 3a. We disregard the fact that surfaces can completely swap sides along shared profiles.

In 2D, when establishing shared profiles between deformable contours A and B, intersecting profiles can cause overlap of the deformed contours, as illustrated in Fig. 3b. Now we

examine this situation in 3D. For this purpose we consider a bijective mapping m of piece-wise *continuous* regions R_A and R_B on surfaces A and $B \subset \mathbb{R}^3$: Shared profiles on triangular meshes can be seen as a finite set of line segments connecting corresponding triangle vertices x and $m(x)$. Hence they partially define or are embedded in such a mapping. However, during surface deformation, not only each vertex, but *each point* on R_A and R_B is displaced along a line segment that leads to its corresponding point. Profiles may not intersect with each other, while other line segments that connect mapped point pairs do. Thus we are interested in properties of mappings m and not only in properties of shared profiles.

An overlap after surface deformation cannot occur if the mapping m satisfies what we call *non-intersection* condition, i.e:

2. No two connections of two mapped point pairs intersect.

For a proof, let R_A^* and R_B^* be the deformed regions. R_A^* is the image of a function f_A that maps each point x on R_A to a point x_A^* on the line segment $\{x + \lambda \cdot (m(x) - x)|0 \leq \lambda \leq 1\}$. Each x_A^* is defined by an individual $\lambda_A(x) \in [0, 1]$. Likewise R_B^* is the image of f_B that maps each x on R_A to an $x_B^* = x + \lambda_B(x) \cdot (m(x) - x)$, with the additional constraint that $\lambda_A(x) < \lambda_B(x)$ for all $x \in R_A$. If the deformed regions intersect, we have $x, y \in R_A$ with $f_A(x) = f_B(x)$, i.e. $x + \lambda_A(x) \cdot (m(x) - x) = y + \lambda_B(y) \cdot (m(y) - y)$. The constraint $\lambda_A(x) < \lambda_B(x)$ implies $x \neq y$. This implies an intersection of the line segments $\{x + \lambda \cdot (m(x) - x)|0 \leq \lambda \leq 1\}$ and $\{y + \lambda \cdot (m(y) - y)|0 \leq \lambda \leq 1\}$, i.e. the line segments that connect the mapped point pairs $(x, m(x))$ and $(y, m(y))$, respectively. In reverse, if no two connections of two mapped point pairs intersect, the deformed regions do not intersect either. Note that the non-intersection condition also implies the continuity of m for homeomorphic regions R_A and R_B, yet continuity alone is not sufficient for preventing overlap.

4.2 Mapping of non-empty surface regions

The following scheme establishes a mapping of topologically equivalent regions on smooth, closed surfaces A and B that satisfies the non-intersection condition:

1. Compute the mid surface M (all points with same distance to A and B), and its normals n where possible.

2. Define finite length vectors v on M as follows: Scale normals n and inverse normals $-n$ to some length exceeding the maximum distance of any two points x on A and y on B. Then trim the scaled normals at the skeleton of M if they reach it. (The skeleton of a surface is the set of points that are centers of spheres that touch the surface in more than one point, but do not cut it.) The resulting vectors then do not intersect with each other.

3. Iteratively map points on A and B cut by the same vector v. Start with vectors on local minima of the signed euclidean distance function $d : M \rightarrow \mathbb{R}, x \mapsto d(x, A)$ on M. (Note that $d(x, A) = d(x, B)$ for all x.) Grow the regions considered on M as long as corresponding regions on A and B have the same topologies.

This results in an intersection free mapping of regions. The mapping is not empty: At least all vectors defined on local minima of d connect corresponding points. Such local minima exist, as we are dealing with closed surfaces. The linear connections of the respective corresponding points do not contain any point of the skeleton of M. For a proof, see Kainmueller et al. [2008].
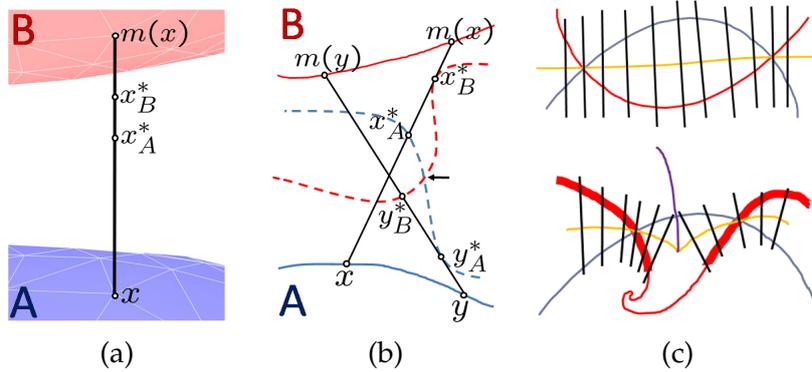
Figure 3: (a) Necessary condition: x_A^* lies closer to x than x_B^*. (b) Contours A, B. Black lines: Connections of mapped points intersect. Dotted: Exemplary deformed contours, intersecting at black arrow. (c) Exemplary situations in 2D. Red, blue: contours to be coupled. Yellow: mid contour. Purple: Skeleton of mid contour. Black: shared displacement directions. Top: Overlapping regions can be entirely mapped. Bottom: Overlapping regions cannot be entirely mapped. The thin part of the red contour remains unmapped

The mapped regions are guaranteed not to overlap after surface deformation along the shared profiles. There might, however, be regions where an initial overlap cannot be resolved, because some overlapping regions cannot be mapped. An overlap in unmapped regions is not resolved with our method. For exemplary situations, see Fig. 3c. Appropriate initial segmentations are required to avoid situations as in Fig. 3c (bottom). If e.g. the initial surfaces do not overlap at all, the whole surfaces are guaranteed not to overlap after surface deformation along the shared profiles.

4.3 Generating shared intensity profiles

Based on the scheme for constructing an intersection-free mapping as proposed in Section 4.2, we realize a construction algorithm for shared profiles on pairs of adjacent *triangular* surfaces A and B. In the process, we modify the connectivity of parts of the surface meshes. The one-sided surface distance of the modified surface to the original surface is always zero. No general assertions can be made regarding the reverse direction of the surface distance.

Fig. 4 shows the construction pipeline for an exemplary femur and ilium. First, the mid surface M between the objects is computed as the zero level of the objects' distance transforms, subtracted from each other (Fig. 4b). M is discretized and triangulated to form a mesh of high regularity. The skeleton of M is approximated by uniformly displacing M along its surface normals in both directions and identifying self intersection points. Then we identify the region on M where its vertex normals, scaled to a user-specified maximal length, enter both femur and ilium, without reaching the skeleton of M first (Fig. 4b and 4c). This region is displaced onto each femur and ilium in vertex normal direction. Each displaced patch is merged into the respective surface mesh by removing all original triangles on the mesh which are surrounded by the patch boundary, and connecting the boundaries of remaining the mesh and the patch. Fig. 4d and 4e show the resulting surfaces.

As a result, we obtain a bijective mapping of the displaced (continuous) patches that satisfies the non-intersection condition. Thereby shared profiles between the modified surface
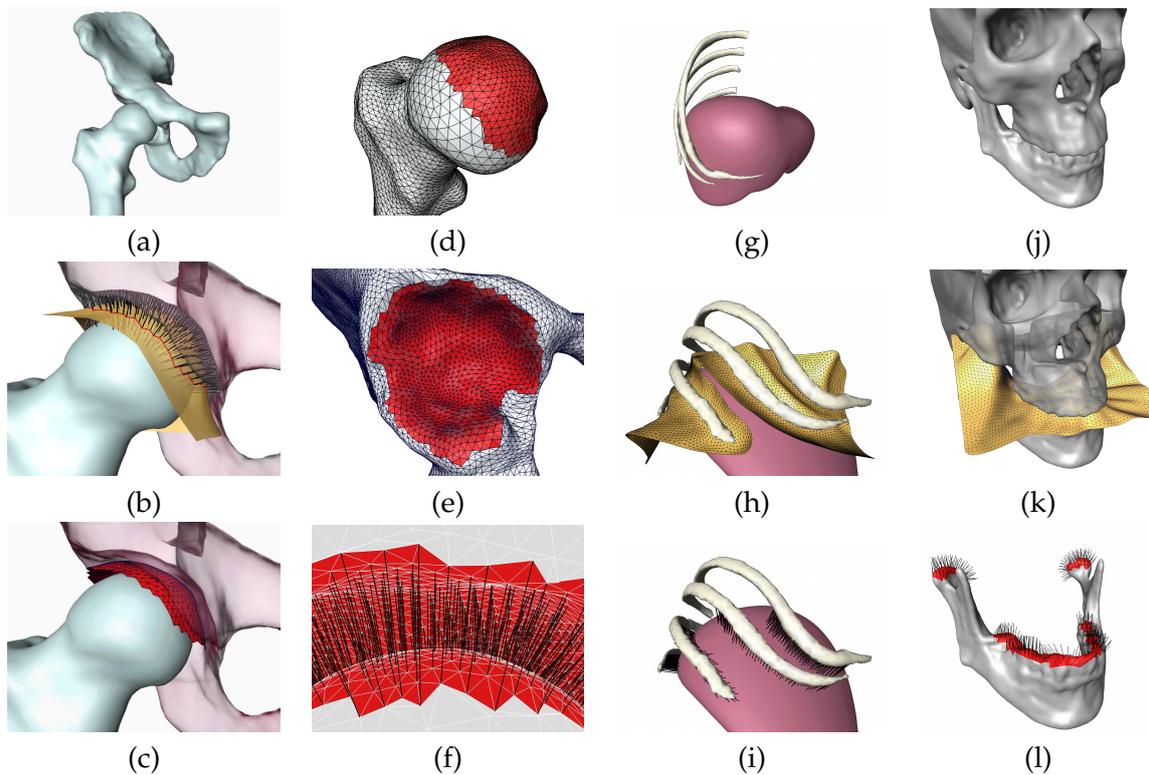
Figure 4: Construction of shared profiles. (a) Surface models of proximal femur and right ilium. (b) Mid surface (yellow) in region of interest, with normals entering both femur and illium. (c) Extracted region on mid surface (red). (d,e) Surface meshes with integrated displaced region (red). (f) Vectors coupling femoral head and acetabulum. (g) Surface models of liver and nearby ribs. (h) Mid surface in region of interest. (i) Coupling vectors between liver and ribs. (j) Surface models of skull and mandibular bone. (k) Mid surface. (l) Coupling vectors between mandibular bone and skull.

meshes' vertices are defined, as shown in Fig. 4f and 4g. We let the shared profiles reach *into* the surfaces until they meet the skeleton of M, or the inner skeleton of the respective surface, or they reach a user-specified maximal length. Fig. 4g-4i show a liver model and a model of surrounding ribs which are connected. Three regions with shared profiles are identified here. As another example, Fig. 4j-4l shows models of a skull and a mandibular bone with shared profiles. Multiple regions are coupled here, as well.

5 Results

We computed segmentations of the femoral head and acetabulum in 50 CT images with 5mm slice thickness. Manually defined gold standard segmentations of the pelvis and the proximal femur were availible for this data. Surface models were initialized with a composite SSM as described in Kainmueller et al. [Sept. 2009].

Shared profiles were established as proposed in Section 4.3. A graph was then constructed as in Li et al. [2006], containing these shared profiles as well as traditional single profiles in non-adjacent regions of the surfaces, as in Seim et al. [2008]. We used intensity

(a)   (b)   (c)
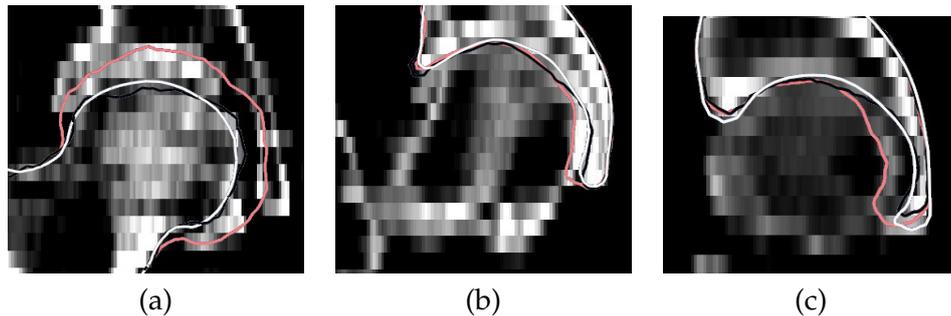
Figure 5: Black: coupled optimization. White: Gold standard. (a) Red: Single femur optimization. Coupled optimization prevents mis-adaptation of the femoral head to features inside the pelvis. (b,c) Red: Single pelvis optimization, 20mm profiles. Minimum distance constraints prevent mis-adaptations in coupled optimitzation.

	Initialization		Single OPT		Coupled OPT	
	ad	**md**	**ad**	**md**	**ad**	**md**
F	**1.7** (0.6)	**4.5** (1.5)	**1.7** (1.0)	**5.8** (2.1)	**1.1** (0.5)	**4.3** (1.4)
A	**1.4** (0.5)	**5.4** (1.9)	**0.9** (0.5)	**4.7** (2.4)	**0.8** (0.5)	**4.7** (2.2)

Table 1: Initialization, single, and coupled optimization results for femoral head (F) and acetabulum (A). All entries in mm. Bold: Average error metrics on 50 CT images. In brackets: Standard deviation.

and gradient thresholding and in the image data to derive a cost function on the profiles. Costs are low at sample points where intensities change more than a gradient threshold from bright to dark within a certain intensity window. Costs are high everywhere else. The exact definition of the cost function can be found in Kainmueller et al. [Sept. 2009]. We applied the graph cut algorithm proposed in Boykov and Kolmogorov [2004] to find optimal surfaces. For comparison, we optimized both surfaces individually (using the respective part of the same graph), without any coupling by shared profiles.

The focus of our evaluation was to determine how accurate the automatic segmentations are in the area of the right hip joint. 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 assymmetric measures. Table 1 lists the average error metrics for the femoral head and acetabulum. Results are given for the initialized surfaces, the single-object optimizations, and the multi-object optimizations on coupled surfaces.

6   Discussion

We proposed a method for coupling arbitrary adjacent surfaces by means of shared profiles. Two surfaces are guaranteed not to overlap if deformed along the shared profiles. Optimal graph searching on the graph obtained from the coupled surfaces produced very accurate

segmentations of the femoral head and acetabulum in 50 ct images of the hip joint.

With the composite SSM, situations as in 3c (bottom), i.e. the problem of unmapped, overlapping regions did not occur for our 50 test cases. However, in case that the initial composite SSM segmentations show such an overlap (which is unlikely, but possible), it is not guaranteed that the overlapping regions can be completely mapped.

For the femur, coupled optimization performs significantly better than single optimization (ad 1.1 versus 1.7mm). In fact, single optimization does not improve the initial segmentation at all, as the femur adapts to image features inside the pelvis, see Fig. 5a. This is prevented by the distance constraints in coupled optimization. As for the pelvis, coupled optimization is slightly 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. 5b and 5c.

Acknowledgements

References

K.O. Babalola, V. Petrovic, T.F. Cootes, C.J. Taylor, C.J. Twining, T.G. Williams, and A Mills. Automatic Segmentation of the Caudate Nuclei using Active Appearance Models. In T. Heimann, M. Styner, and B. van Ginneken, editors, *3D Segmentation in the Clinic: A Grand Challenge*, pages 57–64, 2007.

Y. Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 26(9):1124–1137, September 2004.

T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active Shape Models - Their Training and Application. *Comput. Vis. Image Underst.*, 61(1):38–59, 1995. ISSN 1077-3142. doi: http://dx.doi.org/10.1006/cviu.1995.1004.

María Jimena Costa, Hervé Delingette, Sébastien Novellas, and Nicholas Ayache. Automatic segmentation of bladder and prostate using coupled 3d deformable models. In Nicholas Ayache, Sébastien Ourselin, and Anthony J. Maeder, editors, *MICCAI (1)*, volume 4791 of *Lecture Notes in Computer Science*, pages 252–260. Springer, 2007. ISBN 978-3-540-75756-6.

MO. Heller, G. Bergmann, G. Deuretzbacher, L. Dürselen, M. Pohl, L. Claes, NP. Haas, and GN. Duda. Musculo-skeletal loading conditions at the hip during walking and stair climbing. *J Biomech*, 34:883–93, 2001.

Dagmar Kainmueller, Hans Lamecker, Stefan Zachow, and Hans-Christian Hege. Coupling Deformable Models for Multi-object Segmentation. In F. Bello and E. Edwards, editors, *ISBMS*, volume 5104 of *Lecture Notes in Computer Science*, pages 69–78. Springer, 2008.

Dagmar Kainmueller, Hans Lamecker, Stefan Zachow, and Hans-Christian Hege. Coupling Deformable Models for Multi-object Segmentation. In *accepted at the IEEE Conference on Engineering in Medicine and Biology (EMBC)*, Minnesota, Minneapolis, USA, Sept. 2009.

Kang Li, Steven Millington, Xiaodong Wu, Danny Z. Chen, and Milan Sonka. Simultaneous segmentation of multiple closed surfaces using optimal graph searching. In Gary E. Christensen and Milan Sonka, editors, *IPMI*, volume 3565 of *Lecture Notes in Computer Science*, pages 406–417. Springer, 2005. ISBN 3-540-26545-7.

Kang Li, Xiaodong Wu, Danny Z. Chen, and Milan Sonka. Optimal surface segmentation in volumetric images-a graph-theoretic approach. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28 (1):119–134, 2006. ISSN 0162-8828. doi: http://dx.doi.org/10.1109/TPAMI.2006.19.

Xiaomin Liu, Danny Z. Chen, Xiaodong Wu, and Milan Sonka. Optimal graph search based image segmentation for objects with complex topologies. volume 7259, page 725915. SPIE, 2009. doi: 10.1117/12.811704. URL http://link.aip.org/link/?PSI/7259/725915/1.

Heiko Seim, Dagmar Kainmueller, Markus Heller, Hans Lamecker, Stefan Zachow, and Hans-Christian 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)*, pages 93–100, Delft, Netherlands, 2008.

Andy Tsai, William M. Wells III, Clare Tempany, W. Eric L. Grimson, and Alan S. Willsky. Coupled multi-shape model and mutual information for medical image segmentation. In Christopher J. Taylor and J. Alison Noble, editors, *IPMI*, volume 2732 of *Lecture Notes in Computer Science*, pages 185–197. Springer, 2003. ISBN 3-540-40560-7.

Yin Yin, Xiangmin Zhang, and Milan Sonka. Optimal multi-object multi-surface graph search segmentation: Full-joint cartilage delineation in 3d. In Stephen McKenna and Jesse Hoey, editors, *MIUA 2008*, pages 104–108, 2008.