Coupling Deformable Models for Multi-object Segmentation

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

Zuse Institute Berlin, Germany

Abstract. For biomechanical simulations, the segmentation of multiple adjacent anatomical structures from medical image data is often required. If adjacent structures are hardly distinguishable in image data, automatic segmentation methods for single structures in general do not yield sufficiently accurate results. To improve segmentation accuracy in these cases, knowledge about adjacent structures must be exploited. Optimal graph searching 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. The problem is how to define suitable correspondences on arbitrary surfaces. In this paper we propose a scheme for constructing a correspondence relation in adjacent regions of two arbitrary surfaces. When applying the thus generated shared displacement directions in segmentation with deformable surfaces, 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 preliminary segmentation results obtained by a graph cut algorithm.

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. 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 distances 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 individual (single) bony structures (e.g. pelvis, femur) by our segmentation framework which is based on statistical shape models (SSM) and free form models (FFM) [1]. Segmentations with SSMs yield good initializations, but lack precision, since previously unknown patient-specific anatomy is generally not contained in the model. With FFMs more precise segmentations can be achieved, but they suffer from a loss of shape

F. Bello, E. Edwards (Eds.): ISBMS 2008, LNCS 5104, pp. 69–78, 2008.

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 are concerned with 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 a proof of concept with very encouraging results.

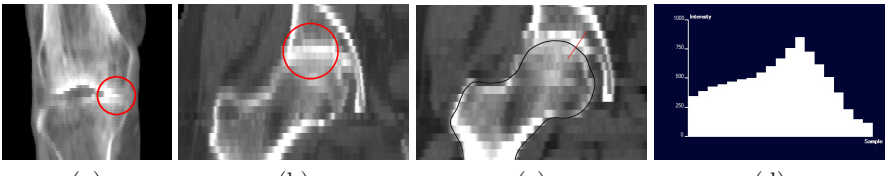

(a) (b) (c) (d)

Fig. 1. (a) CT data of distal femur and proximal tibia, slice distance 2mm, (b) acetabulum and proximal femur, slice distance 4,6mm. Joint space hardly visible in encircled areas. (c) Acetabulum and proximal femur, slice distance 5mm, with surface model cross-section (black) and domain of intensity profile (red). (d) Intensity profile for domain in (c).

2 Related Work

Costa et al. [2] 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. [3] 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. [4] 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. Anyway, 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. [5] 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. However, 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 [6], or if one surface can be obtained by displacing the other along its vertex normals [5]. To the best of our knowledge, there is at present no way for generating shared displacement directions on arbitrary surfaces.

3 Multi-object Segmentation with Graph Cuts

In segmentation with deformable surfaces, *intensity profiles* are commonly used to guide the deformation process. Intensity profiles are intensities sampled f image data along line segments at each vertex of the surface mesh, see Fig. 1c and 1d. Note that the term *intensity profile* or just *profile* may be used as referring to the sampled intensities only, but here 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 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. 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. A smoothing constraint guarantees that new vertex positions on adjacent profiles are at most s sample points away from each other. If sample point i is chosen on a profile as desired new position, sample point i-s or higher must be chosen on adjacent profiles, as illustrated in Fig. 2b. The smaller s is chosen, the

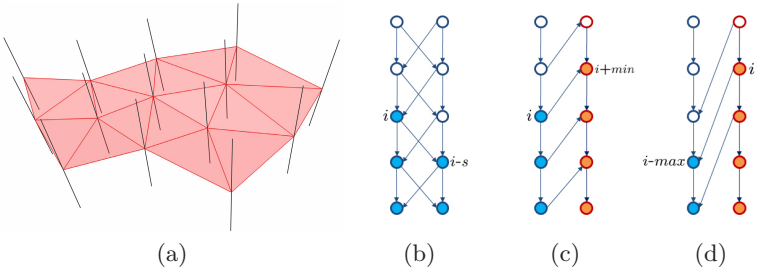

Fig. 2. (a) Triangular surface mesh (red) with profiles (black). (b) Graph edges realizing smoothing constraint. Here, $s = 1$. (c) Minimum distance constraint on shared profile: Sample point i for surface A (blue) entails sample point $i{+}min$ or higher for surface B (red) on duplicated profile. Here, $min = 1$. (d) Maximum distance constraint: Sample point i for surface B entails sample point $i{-}max$ or higher for surface A. Here, $max = 2$.

smoother is the surface resulting from the optimization. Multiple surfaces can be coupled with shared intensity profiles at individual vertices. Shared profiles are duplicated in the graph, so that each surface is equipped with an instance of each shared profile. Minimum and maximum distance constraints guarantee that new vertex positions found on shared profiles by graph optimization are at least min and at most max sample points away from each other, see Fig. 2c and 2d. For more details on graph construction see [6]. 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 [7].

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 to a manual expert segmentation. For two such well-initialized surface models, we identify a potential overlap region by growing each surface by a user-specified profile length, see Fig. 3a and 3b. Our concern is to couple the surfaces with shared profiles wherever single profiles would reach the overlap region.

4.1 Properties of Shared Intensity Profiles

Which conditions must hold on shared profiles to generally prevent any overlap after surface deformation? Note one necessary condition posed on new vertex positions of coupled surface meshes A and B: 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, i.e. $\|x^*_A - x\| < \|x^*_B - x\|$ (if we disregard that surfaces can completely swap sides along shared profiles). See Fig. 3c. In the following we assume that this condition holds.

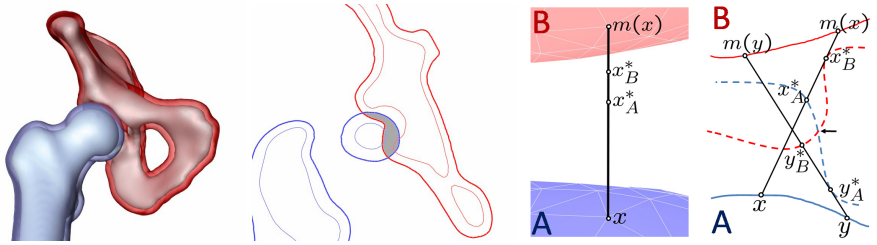

(a) (b) (c) (d)

Fig. 3. (a) Proximal femur and ilium with transparent isosurfaces at same distance. (b) 2D cross-section. Thin lines: surface contours. Thick lines: isosurface contours at the same distance. Grey region: potential overlap region. (c) Necessary condition: x_A^* lies closer to x than x_B^*. (d) Contours A (blue) and B (red). Black lines: Intersecting connections of mapped points. Dotted lines: Exemplary deformed contours with intersection (black arrow).

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. 3d. Now we examine this situation in 3D. For this purpose we consider a bijective mapping m of piecewise *continuous* regions R_A and R_B on surfaces A and B: 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.

Overlap after surface deformation cannot occur if the mapping m satisfies what we call *non-intersection* condition, i.e. 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 $[x, m(x)]$ and $[y, m(y)]$. 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 a sufficient criterion 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 as all points with the same distance to A and B. Compute the normals n on M, wherever M is smooth.
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 a sphere that touches the surface in more than one point, but does 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. 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 mapped regions are not empty: At least all normals on M where the distance function $d : M \rightarrow \mathbb{R}, x \mapsto d(x, A)$ has a local minimum connect corresponding points. (Note that for all x $d(x, A) = d(x, B)$.) Such local minima exist, as we are dealing with closed surfaces. The linear connections of the respective corresponding points also do not contain any point of the skeleton of M.

As a proof, we first examine properties of the normals on M at local minima of d: The connection of any point x on M to a closest point a_x on A is normal on A. (Imagine a sphere with radius $d(x)$ around x. This sphere touches A in a_x.) This applies to closest points b_x on B, accordingly. Local minima of the function d are located at points m on M where isosurfaces of the distance transforms of A and B *touch*. Let a_m and b_m be closest points to m on A and B, respectively. As isosurfaces with value $d(m)$ touch in m and the distances from a_m to m and b_m to m are both $d(m)$, spheres around a_m and b_m with radius $d(m)$ touch in m. This implies that the connections between a_m and m as well as b_m and m both are perpendicular to M. The direct connection between a_m and b_m therefore contains m and is perpendicular to A, B and M.

Now we show that the connection g between a_m and b_m does not meet the skeleton of M: Assume it did. Let p be the closest point to m on g that lies on the skeleton. Assume without loss of generality that p lies closer to a_m than to b_m. Then there is a circle around p that touches M in multiple points, but does not intersect with M. Let r denote the radius of that circle R. r must be smaller or equal to the distance $\|p - m\|$, otherwise m would lie *inside* this circle. Let n be one of the points that R touches, $n \neq m$. Then the distance between n and a_m is shorter than the distance between m and a_m. This is a contradiction to m being the closest point to a_m on M. Hence, there is no point on g that lies on the skeleton of M.

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. Concerning the reverse direction of the surface distance, no general assertions can be made.

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, see Fig. 5. 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. 4c and 4d). This region is then displaced onto the surfaces of each femur and ilium in vertex normal direction. These displaced patches are merged into the respective surface mesh by removing the original triangles in this region and connecting the boundary of the remaining mesh to the boundary of the patch. The resulting surfaces are shown in Fig. 4e and 4f.

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 meshes' vertices are defined, as shown in Fig. 4g and 4h. 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

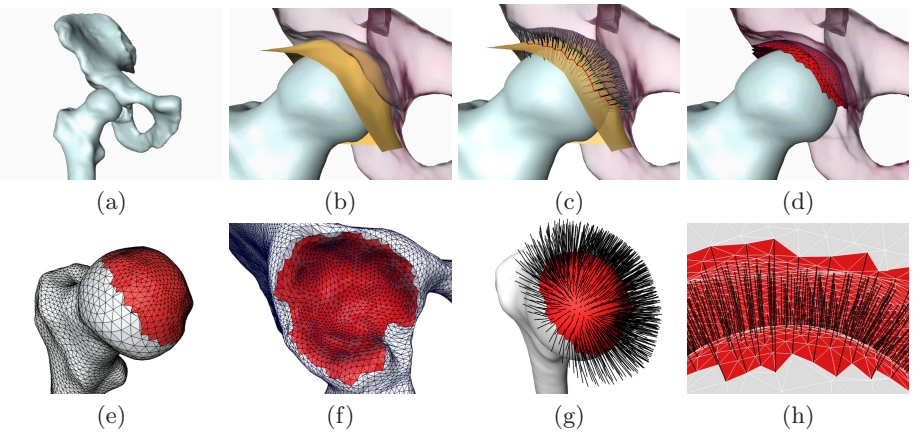

Fig. 4. Construction of shared profiles. (a) Proximal femur and right ilium. (b) Mid surface (yellow) in region of interest. (c) Normals on mid surface entering both femur and illium. (d) Extracted region on mid surface (red). (e,f) femoral head and acetabulum with integrated displaced region (red). (g,h) Vectors coupling femoral head and acetabulum (black).

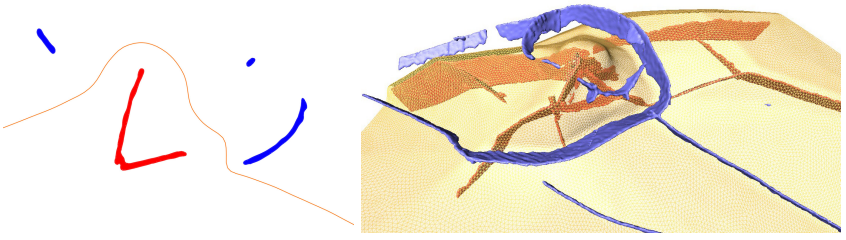

(a) (b)

Fig. 5. 2D cross-section (a) and 3D view (b) of mid surface (yellow), upper skeleton (blue) and lower skeleton (red). For a better impression of the skeleton, the mid surface shown here is not restricted to a region of interest.

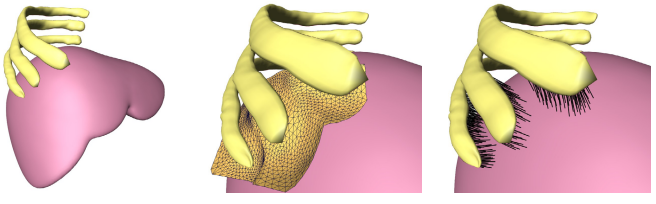

(a) (b) (c)

Fig. 6. (a) Surface models of liver (red) and nearby ribs (yellow). (b) Mid surface in region of interest. (c) Coupling vectors between liver and ribs.

a user-specified maximal length. As another example, Fig. 6 shows a liver model and a model of surrounding ribs which are connected. Three regions with shared profiles are identified here.

5 Results

In a first investigation we computed segmentations of a femoral head and ilium and a distal femur and proximal tibia in CT data with 5mm slice distance. Single surface models were initialized individually with SSM fitting as described in [1]. Fig. 7a shows resulting initializations. Note the overlap of the initialized models. Shared profiles (Fig. 7b) were established as proposed in Section 4.3. A graph was then constructed as in [6], containing these shared profiles as well as traditional single profiles [1] in non-adjacent regions of the surfaces. As a cost function on the profiles we used thresholding in the image data, see Fig. 7d. Costs are low at sample points where intensities change from above to below the threshold, and high everywhere else. We applied the graph cut algorithm proposed in [7] to find optimal surfaces. Fig. 7c shows the resulting segmentations. The overlap was resolved and a reasonable interpolation was found in image regions where intensities lie above the threshold due to low image quality.

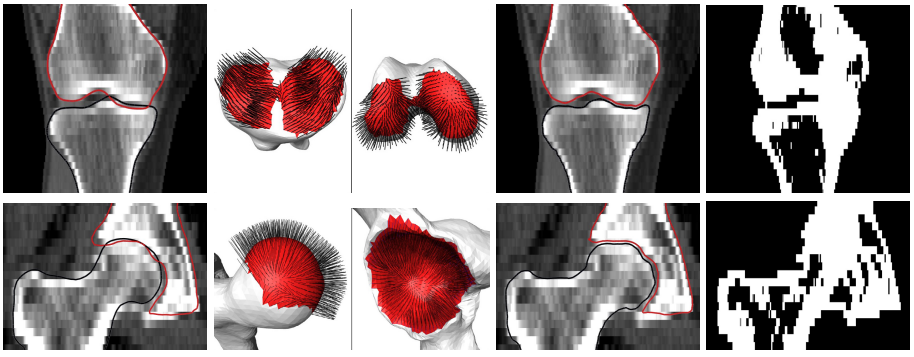

Fig. 7. First row: Distal femur and proximal tibia. Second row: Acetabulum and proximal femur. (a) Model initialization with statistical shape models. (b) Shared profiles. (c) Graph cuts optimization. (d) Threshold information used for model fitting.

6 Discussion

We proposed a method for coupling arbitrary adjacent surfaces by means of shared profiles. The shared profiles are constructed in a way that they satisfy a non-intersection condition. This property guarantees that two surfaces do not overlap if deformed along the shared profiles. We applied the coupling method to femur, tibia and ilium models which were initialized in exemplary CT data by SSM fitting. Optimal graph searching on the graph obtained from the coupled surfaces produced encouraging results.

The next task will be to evaluate segmentation results quantitatively on a set of 3D CT datasets by comparing them to manual expert segmentations. For this purpose, as a first step, variations among segmentations performed by different experts must be determined. This is of particular importance here as we are dealing with segmentations of adjacent objects which often are hardly distinguishable in image data, even for human observers. Until now, simple thresholding was performed for establishing a cost function. We are working on a more elaborate cost function design that considers more information (e.g. the gradient) contained in image data. We also plan to establish an automatic method for rough initialization of pelvis and femur models in CT datasets as starting positions for SSM segmentation. Until now, this first initialization requires manual interaction.

Acknowledgements

Dagmar Kainmueller is funded by DFG Collaborative Research Center SFB760. Hans Lamecker is funded by DFG Research Center Matheon in Berlin. Thanks to Heiko Seim (ZIB) for providing the shape models of femur and tibia. Thanks to Charité Center for Musculoskeletal Surgery for providing the pelvis model.

References

1. Kainmueller, D., Lange, T., Lamecker, H.: Shape Constrained Automatic Segmentation of the Liver based on a Heuristic Intensity Model. In: 3D Segmentation in the Clinic: A Grand Challenge, pp. 109–116 (2007)
2. Costa, M.J., Delingette, H., Novellas, S., Ayache, N.: Automatic segmentation of bladder and prostate using coupled 3d deformable models. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 252–260. Springer, Heidelberg (2007)
3. Tsai, A., Wells, W., Tempany, C., Grimson, W.E.L., Willsky, A.S.: Coupled multi-shape model and mutual information for medical image segmentation. In: Taylor, C.J., Noble, J.A. (eds.) IPMI 2003. LNCS, vol. 2732, pp. 185–197. Springer, Heidelberg (2003)
4. Babalola, K.O., Petrovic, V., Cootes, T., Taylor, C., Twining, C., Williams, T., Mills, A.: Automatic Segmentation of the Caudate Nuclei using Active Appearance Models. In: 3D Segmentation in the Clinic: A Grand Challenge, pp. 57–64 (2007)
5. Li, K., Millington, S., Wu, X., Chen, D.Z., Sonka, M.: Simultaneous segmentation of multiple closed surfaces using optimal graph searching. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565, pp. 406–417. Springer, Heidelberg (2005)
6. 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), 119–134 (2006)
7. Boykov, Y.Y., Kolmogorov, V.: 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 (2004)