

Analyzing Anatomical Structures: Leveraging Multiple Sources of Knowledge

Eric Grimson and Polina Golland

Computer Science and Artificial Intelligence Laboratory,
Massachusetts Institute of Technology,
Cambridge MA 02139

Abstract. Analysis of medical images, especially the extraction of anatomical structures, is a critical component of many medical applications: surgical planning and navigation, and population studies of anatomical shapes for tracking disease progression are two primary examples. We summarize recent trends in segmentation and analysis of shapes, highlighting how different sources of information have been factored into current approaches.

1 Background and Motivation

The application of computer vision techniques to detailed anatomical information from volumetric medical imaging is changing medical practice, in areas ranging from surgery to clinical assessment. Automated reconstruction of precise patient-specific models of anatomic structures from medical images is becoming *de rigueur* for many surgical procedures, disease studies, and clinical evaluation of therapy effectiveness.

For example, neurosurgeons often use navigational aids linked to labeled imagery in order to localize targets. Even the registration to the patient of unprocessed volumetric imagery, such as MRI, is of value to the surgeon. It allows her to see beneath exposed surfaces to localize nearby structures, and to track instruments during minimally invasive procedures, providing faster navigation through narrow openings. But raw imagery is often insufficient, since it is often cluttered and filled with subtle boundaries. By segmenting out distinct anatomical structures, by co-registering functional, biomechanical, or other information to those segmented structures, and by relating all that information to the patient's position, the surgeon's ability to visualize the entire surgical space is greatly enhanced, as is her ability to avoid critical structures while ensuring that targeted tissue is fully excised. This allows surgeons to execute surgeries more quickly, and with less impact on neighboring tissues. Hence, extracting anatomical models from imagery is a key element of emerging surgical techniques. An example is shown in Figure 1.

Understanding healthy development and disease progression also can benefit from computational methods that extract precise models of anatomical substructures from volumetric imagery. For example, in studies of schizophrenia accurately measuring shapes of cortical and subcortical structures and identifying significant differences in shape between diseased and normal populations provides an invaluable tool for

understanding the disease's progression. This is particularly true when shapes of structures also can be tracked over time. In Alzheimer's disease, understanding the relationship between changes in shape and volume of neural structures in correlation with other factors, such as genetic markers or distributions of functional information during specific mental tasks, may provide neuroscientists with computational tools for deepening the understanding of the disease, and its progression. Hence, extracting anatomical models from imagery, especially group statistical properties of those models, is an essential component of emerging neuroscience techniques.

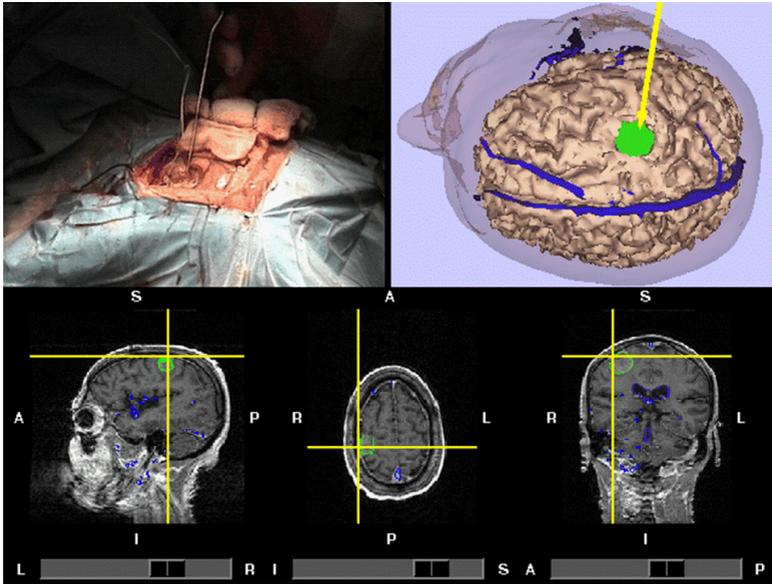


Fig. 1. Example of segmented images in surgical navigation. Segmented structures are overlaid on cross-sectional views of an MRI scan (bottom). The position of a surgical probe is tracked relative to the patient and the segmented scans.

To enhance surgical practice, and to extend neuroscientific understanding of diseases, we need models of shapes, and thus we need segmentation algorithms to extract structures from multi-modal images. These methods have evolved in capability, as increasingly sophisticated models of anatomy and of the image formation process have been embedded within them. Early methods simply relied intensity thresholds to separate structures. These methods fail when confronted with bias fields in the imagery, so approaches were developed that either model the field as a smooth function (e.g., a polynomial) or employ non-parametric techniques, (e.g., Expectation Maximization (EM)), to account for image acquisition artifacts and to create statistical models of tissue response. Because individual voxel responses can be consistent with several tissue types, especially in the presence of noise, Markov Random Field (MRF) methods have been added to impose local continuity. They capture the notion that nearby labels for tissue type should influence the labeling of a particular voxel, since tissue is generally locally continuous. Atlas-based methods

capture more global constraints, especially the expected position and spatial extent of structures, by providing spatial templates in which to interpret voxel intensities, either by defining prior probabilities on tissue type as a function of position, or by defining initial positions for boundaries of structures, which are then warped to best match the intensity data.

Structural measurements in morphological studies have seen a parallel refinement. Early methods simply recorded volumes of structures, measurements that have improved in specificity as segmentation methods enabled increasingly detailed and accurate delineation of the structures of interest. However, volume alone is often insufficient to detect changes in a structure or to differentiate different populations. Often a difference in structure between two distinct populations is localized in specific components of that structure; or the progression of a disease will result in variations in a particular component of a structure.

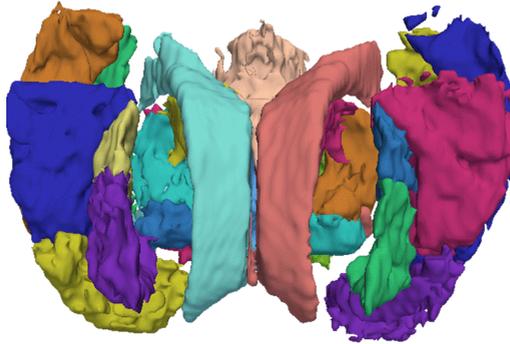


Fig. 2. An example segmentation from a neurosurgical case. Thirty-one different brain structures have been segmented using MRF-based voxel classification methods. The goals of the proposed project include improving on such segmentation by adding shape constraints. While these segmentations are a good start, there are residual errors near subtle boundaries that would be improved by adding shape information.

Thus, we believe that a natural next stage in the evolution of segmentation methods and associated measurements of structures is the incorporation of shape information into the process. To do this, we need to learn the typical shapes of structures and their statistical variation across normal populations. We believe that this information can be extracted from sets of training images, and that the resulting statistical distributions on shape can be used to guide segmentation of new images. Moreover, if we can capture shape distributions, we can use them as a comparative tool: to identify differences in shape between populations or to track changes in shape with time.

We have suggested that capturing statistics of shape can aid in segmenting new scans. One clear application of this tool is in surgical planning and guidance. By providing the surgeon with a detailed reconstruction of the anatomy, she can plan access to the operative site, and by registering the model to the actual position of the patient, she can track the position of surgical instruments relative to structures in the

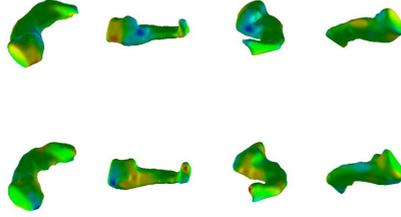


Fig. 3. Visualization of the shape differences in the right hippocampus between the schizophrenia group and the normal controls captured by a discriminative model learned from example images. Using the learned classifier based on the two groups, one can examine the amount of deformation, from blue (moving inwards) to red (moving outwards) needed to shift a normal example towards the schizophrenia group, and vice versa.

reconstructed model. Figure 2 shows an example segmentation of a neuroimage into 31 structures derived using atlas-based segmentation, followed by a MRF label regularization.

Shape based methods can provide more than just better segmentation, however. The shape analysis methods also may improve the quantitative assessment of anatomical development during disease progression. In particular, we can create shape-based classifiers, using machine learning tools applied to our shape distributions, to distinguish between different populations. Such classifiers not only provide a mechanism for assigning new instances to the appropriate class, they can be used to determine what parts of the shape change in significant ways between the two classes, as illustrated in Figure 3 for a study of hippocampal shape in schizophrenia. The colormap indicates the amount of deformation required to make the normal hippocampus shown here be likely to come from the schizophrenia group with respect to the trained classifier function. Such visualization techniques allow one to isolate shape changes to key elements of a structure.

2 Classes of Segmentation Methods

Segmentation is a central component in many medical image analysis applications. It provides a fundamental basis for surgical planning, surgical navigation, analysis of disease progression, and therapy evaluation in essence by extracting geometrically accurate, patient-specific reconstructions of anatomical structures, segmentation provides a foundation for understanding structural correlates with disease and for visualizing and planning interventions. Unfortunately, segmentation is inherently under-constrained: the image intensity alone is often insufficient to delineate anatomical structures. Human operators employ their knowledge of anatomy when they perform segmentation; thus incorporating such knowledge into the automatic methods has improved the quality of segmentation results. However, the problem of accurate automatic segmentation of neuroimages is not yet fully solved, and we believe that building computational mechanisms for modeling anatomical shape and its variability will significantly improve segmentation accuracy.

To understand the role of shape in segmentation, we find it useful to group image segmentation methods into three broad classes: voxel classification, atlas-based segmentation, and boundary localization. The output of the first class of algorithms is a volume in which each voxel is labeled with an associated tissue class. The second group infers segmentation through registering a previously segmented image (atlas) to a new input image. Algorithms in the third class are based on deformable contours that are evolved to explain the novel image.

2.1 Voxel Classification

The original methods in this class used intensity information at a specific voxel to decide on the tissue type enclosed by that voxel [1,2,3,4]. These methods measured intensity responses for different tissue classes from training data (such as voxels selected by a user based on uniqueness of tissue type) and derived a set of optimal thresholds for assigning a tissue class to each voxel. This approach has been refined by our group and other researchers to deal with non-linear gain (bias) artifacts by simultaneously solving for a smooth spatially variant gain field and the identification of tissues at voxels [5,6,7]. To handle noise and uncertainty in the intensity information, the next group of algorithms in this class incorporated local geometric information by employing Markov Random Field (MRF) models [8,9]. The idea is to use information about tissue likelihoods at neighboring voxels to influence the assignment of a tissue label at a specific voxel. The original formulation has been refined by several groups, including ours, to include local estimation of individual and pairwise priors from a set of scans previously segmented by an expert and registered together to form an atlas. This probabilistic atlas is then registered to a new scan to transform the likelihoods into appropriate locations [10,11].

Algorithms in this class provide excellent quality segmentation for gross voxel assignment into three large classes: white matter, gray matter, and cerebrospinal fluid (CSF). The results can be used for analysis of overall (total volume) tissue changes in an individual or a population, but since the methods operate on an implicit representation of the anatomy as a discrete map of tissue labels, they do not produce an explicit representation of the structure boundaries. Instead, an additional step of extracting the boundary surfaces is required if we are to utilize the results of voxel classification in visualization for surgical planning and navigation, or population studies focusing on anatomical shape. Moreover, explicit shape models that involve global descriptors do not easily fit within this intrinsically local framework. Recent work focuses on more detailed segmentation into individual subcortical structures and cortical parcellation by constructing location-specific Markov priors for every structure [10,12] effectively reducing the representation of the anatomy to the level of individual voxels and voxel neighborhoods.

2.2 Atlas-Based Segmentation

The methods in this class seek to induce a segmentation of a new scan by deforming a given segmented atlas image to the novel grayscale image and by applying the estimated transformation, or warp, to the label map of the atlas. The atlas generally includes at least one scan that has been carefully annotated, perhaps along with

likelihoods used to describe the variance seen across a population. If the atlas can be successfully matched to a novel scan, then all information present in the atlas, such as the tissue labels, is known with respect to the new image. The registration process is often reduced to an optimization of a weighted sum of two terms: the quality of the image fit, i.e., how well the atlas grayscale image matches the novel scan, and a penalty term on the amount of deformation between the atlas and the image. The deformations required to adequately match the atlas to novel scans are typically sufficiently high-dimensional, requiring a substantial amount of regularization. Examples of regularization methods include representing the deformation field as a sum of basis vectors, which allows for a coarse to fine solution [13,14], or representing the large deformation as a concatenation of a set of small deformations that more efficiently fit into the optimization framework [15] or employing a cascade of similarity groups beginning with rigid transformations, and subsequently allowing more flexible warps, such as piece-wise affine [16,17] or elastic [18,19]. Tissue deformation models can also assist in regularizing the deformation field when dealing with anatomy containing very different structures such as bone, muscle, and CSF [20,21]. Depending on the image properties, the quality of the fit is evaluated using image intensity, its gradient, texture measures, or discrete landmark points. A comprehensive survey of relevant registration techniques can be found in [22].

The algorithms in this class have also been used to map tissue class priors onto a new scan as an initialization step in voxel classification methods that further refine the initial segmentation by evoking MRF-based segmentation [10,11]. Like the voxel classification algorithms, these methods operate on the implicit voxel-based representation of anatomy and thus take little advantage of information on anatomical shape variability. They typically rely on the intensity information and the inherent geometric properties of the mapping (continuity, smoothness, etc.) to guide registration. Active Appearance Models [23,24] are a notable exception, as they achieve registration of a template to a novel image based on the shape variability model. However, we group this particular method with the boundary detection algorithms, as it manipulates an explicit boundary representation both for shape modeling and for registration.

2.3 Boundary Detection

Rather than label each voxel by tissue type, the methods in this group search for boundaries between different tissue types. Standard methods use deformable contours, or "snakes" [25]: they evolve a contour (or a surface in 3D), typically by shrinking or expanding it in proportion to its local curvature, until it reaches a strong intensity boundary. Robust variants that use regularization to reduce the sensitivity to noise include balloons [26], t-snakes [27], and pedal snakes [28].

Boundary evolution techniques explicitly manipulate the object boundary surface and therefore lend themselves naturally to shape-based extensions. The Active Shape Models method [29,30] extracts a set of corresponding points on the outline of every training example and employs Principal Component Analysis (PCA) to build a linear model of variation, typically keeping only a few principal components. The resulting probabilistic model is then used to constrain the space of deformations of the "representative" example (mean shape) when matching it to a new image. In addition

to the variance of point positions, Active Appearance Models [23,24] include the intensity distributions over the object. The model fitting stage searches over the eigen-coefficients with the goal of matching the grayscale values of the model and the image. This framework has also been extended to multi-shape segmentation, where the statistical model captures shape of several structures, as well as their relative position and orientation, helping to overcome poor local intensity contrast [31]. Variants of this approach have been demonstrated that manipulate parametric representations of the boundary -- the Fourier descriptors in 2D [32, 33] and the spherical harmonics in 3D [34] -- by transferring the problem of modeling and template matching into the space of the representation coefficients instead of the boundary points.

The level-set implementation of curve evolution [35,36] overcomes several shortcomings of the traditional snakes. It is more robust to initialization, allows topology changes, and is more stable numerically as it operates on the volumetric grid, thus eliminating the need for re-parameterization of the snake. Similarly to the original snakes, the level-set implementation can be augmented to include prior shape information in steering the evolving boundary towards the "typical" shape. Examples of application-specific shape constraints include coupled evolution of two surfaces for cortical segmentation [37], or topology preservation of the resulting surface through local constraints [38,39]. We demonstrated a level-set counterpart of the Active Shape Models by applying PCA to the signed distance transforms of the training examples and by introducing a term into the evolution dynamics that "pulls" the evolving distance transform towards the most likely shape under the probability distribution obtained through PCA [40,41,42]. A multi-shape segmentation based on this principle has recently been demonstrated by our and other groups [43,44].

3 Next Steps in Segmentation

To summarize, segmentation methods have seen a series of developments, each stage incorporating additional information. Thus, early methods simply relied on intensity thresholds to separate structures. These methods fail when confronted with bias fields in the imagery, so approaches were developed that either model the field as a smooth function or employ non-parametric techniques to account for image acquisition artifacts and to create statistical models of tissue response. Because individual voxel responses can be consistent with several tissue types, especially in the presence of noise, Markov Random Field (MRF) methods have been added to impose local continuity. They capture the notion that nearby labels for tissue type should influence the labeling of a particular voxel, since tissue is generally locally continuous. Atlas-based methods capture more global constraints, especially the expected position and spatial extent of structures, by providing spatial templates in which to interpret voxel intensities, either by defining prior probabilities on tissue type as a function of position, or by defining initial positions for boundaries of structures, which are then warped to best match the intensity data. Finally, template-based segmentation has served as a basis for most methods that incorporate prior shape information into segmentation. Both the original deformable contour methods and the more recent

level-set implementations have been augmented with probabilistic shape models that modify the evolution of the boundary.

Although considerable progress has been made in creating sophisticated segmentation algorithms and in analyzing populations of segmented structures to capture statistical models of variation, important challenges remain. Key among these are as follows:

- methods that robustly leverage shape information to enhance segmentation;
- a demonstration that shape information significantly improves the quality of segmentation;
- a demonstration that such improved segmentation provides an added value in surgical planning and navigation;
- methods that capture statistical variations in distributions of shape information;
- a demonstration that these statistical models can be used to create classifiers that distinguish between diseased and normal populations;
- a demonstration that classifiers based on shape differences are more accurate than simple volumetric measures;
- methods that use these classifiers to visualize differences in shapes of populations; and
- using the results of the statistical analysis to find correlations between disease states and changes in shapes of structures.

References

1. Cline, H., Lorensen, W., Kikinis, R., Jolesz, F., Three-dimensional segmentation of MR images of the head using probability and connectivity. *Journal of Computer Assisted Tomography*, 14(6):1037-1045, 1990.
2. Kohn, M.I., Tanna, N.K., Hermann, G.T., Resnick, S.M., Mozley, P.D., Gur, R.E., Alavi, A., Zimmerman, R.A., Gur, R.C., Analysis of brain and cerebrospinal fluid volumes with MR imaging. Part I. Methods, reliability, and validation. *Radiology*, 178:115-122, 1991.
3. Lim, L.O., Pfefferbaum, A., Segmentation of MR brain images into cerebrospinal fluid spaces, white and gray matter. *J. Comput. Assisted Tomography*. 13:588 1989.
4. Vannier, M.W., Pilgram, T.K., Speidal, C.M., Neumann, L.R., Rickman, D.L., Schertz, L.D., Validation of magnetic resonance imaging (MRI) multispectral tissue classification. *Computer Medical Imaging and Graphics*, 15:217- 223, 1991.
5. Brechbuhler, C., Gerig, G., Szekely, G., Compensation of spatial inhomogeneity in MRI based on a multi-valued image model and a parametric bias estimate. In: *Proc. VBC'96*, LNCS 1131:141-146, Springer, 1996.
6. Wells, W., Grimson, W.E.L., Kikinis, R., Jolesz, F., Statistical intensity correlation and segmentation of MRI Data. *Visualization in Biomedical Computing*, 1994
7. Van Leemput, K., Maes, F., Vanermeulen, D., Suetens, P., Automated model-based bias field correction of MR images of the brain, *IEEE TMI*, 18(10): 885-895, 1999.
8. Kapur, T., Grimson, W.E.L., Wells, W.M., Kikinis, R., Enhanced spatial priors for segmentation of magnetic resonance imagery. *Proceedings MICCAI* 457-468, 1998.
9. Leahy, R., Hebert, T., Lee, R., Applications of Markov random field models in medical imaging. In *Proceedings of IMPI'89*, 1-14, 1989.

10. Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A.M., Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain, *Neuron*, 33, 2002.
11. Pohl, K.M., Wells, W.M., Guimond, A., Kasai, K., Shenton, M.E., Kikinis, R., Grimson, W.E.L., Warfield, S.K., Incorporating non-rigid registration into expectation maximization algorithm to segment MR images, Fifth MICCAI, Tokyo, Japan, pp. 564-572, 2002.
12. Pohl, K.M., Bouix, S., Kikinis, R., Grimson, W.E.L., Anatomical guided segmentation with non-stationary tissue class distributions in an Expectation-Maximization Framework, 2004 IEEE International Symposium on Biomedical Imaging, Arlington, VA, 2004, 81-84.
13. Christensen, G., Rabbitt, R., Miller, M., Deformable templates using large deformation kinematics. *IEEE Trans. on Image Processing* 5(10):1435-1447, 1996.
14. Christensen, G., Consistent linear-elastic transformations for image matching. IPMI 1999.
15. Christensen, G.E., He, J., Consistent nonlinear elastic image registration, IEEE MMBIA, Kauai, Hawaii, Dec. 2001, pp. 37-43.
16. Collins, D.L., Holmes, C.J., Peters, T.M., Evans, A.C., Automatic 3D model-based neuroanatomical segmentation, *Human Brain Mapping*, 3(3):190-208, 1995.
17. Grenander, G., Miller, M., Representations of knowledge in complex systems. *Journal of the Royal Statistical Society B*, 56:249-603, 1993.
18. Bajcsy, R., Kovacic, S., Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing*, 46:1-21, 1989.
19. Warfield, S., Robatino, A., Dengler, J., Jolesz, F., Kikinis, R. Nonlinear registration and template driven segmentation. in brain warping. Ed. Arthur W. Toga, (Progressive Publishing Alternatives) Ch.4:67-84, 1998.
20. Edwards, P., Hill, D., Little, J., Hawkes, D., Deformation for image-guided interventions using a three component tissue model. *Medical Image Analysis*, 2(4): 355-367, 1998.
21. Little, J., Hill, D., Hawkes, D., Deformations Incorporating Rigid Structures. *Mathematical Methods in Biomedical Image Analysis*, 1996.
22. Maintz, J., Viergever, M., A survey of medical image registration. *Medical Image Analysis*, 2(1):1-36, 1998.
23. Cootes, T.F., Edwards, G.J., Taylor, C.J., Active appearance models. In *Proceedings of the European Conference on Computer Vision*. 2:484-498, 1998.
24. Cootes, T.F., Edwards, G.J., Taylor, C.J., Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 23(6):681-685, 2001.
25. Kass, M., Witkin, A., Terzopoulos, D., Snakes: Active contour models. *IJCV* 1:321-331, 1988.
26. Cohen, L., On active contour models and balloons. *CVGIP: IU*, 53(2):211-218, 1991.
27. McInerney, T., Terzopoulos, D., Medical image segmentation using topologically adaptable surfaces. *Conf. Computer Vision, Virtual Reality, and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery (CVRMed-MRCAS)*, 23-32, 1997.
28. Guo, Y., Vemuri, B., Hybrid geometric active models for shape recovery in medical images. *Int'l Conf. Inf. Proc. in Med. Imaging*, pages 112-125, Springer-Verlag, 1999.
29. Cootes, T., Taylor, C., Cooper, D., Graham, J., Training models of shape from sets of examples. *Proceedings British Machine Vision Conference*, 9-18, Springer-Verlag, 1992.
30. Cootes, T., Taylor, C., Cooper, D., Graham, J., Active shape models - their training and application. *Computer Vision and Image Understanding*, 1995.
31. Cootes, T., Beeston, C., Edwards, G., Taylor, C., Unified framework for atlas matching using active appearance models. *Information Processing in Medical Imaging* 1999.

32. Wang, Y., Staib, L., Boundary finding with correspondence using statistical shape models. CVPR 1998.
33. Szekely, G., Kelemen, A., Brechbuler, C., Gerig, G., Segmentation of 2D and 3D objects from MRI volume data using constrained elastic deformations of flexible fourier contour and surface models. *Medical Image Analysis*, 1(1):19-34, 1996.
34. Kelemen, A., Szekely, G., Gerig, G., Three-dimensional model-based segmentation. In *Proceedings of IEEE International Workshop on Model Based 3D Image Analysis*, Bombay, India, 87-96, 1998.
35. Caselles, V., Kimmel, R., Sapiro, G., Geodesic active contours. *IJCV*, 22(1):61-79, 1997.
36. Kichenassamy, A., Kumar, A., Olver, P., Tannenbaum, A., Yezzi, A., Gradient flows and geometric active contour models. In *Proc. IEEE ICCV*, 810-815, 1995.
37. Zeng, X., Staib, L.H., Schultz, R.T., Duncan, J.S., Segmentation and measurement of the cortex from 3D MR images using coupled surfaces propagation. *IEEE TMI*, 18(10), 1999.
38. Han, X., Xu, C., Prince, J.L., A topology preserving level set method for geometric deformable models, *IEEE Trans. PAMI*, 25(6):755-768, 2003.
39. Ségonne, F., Grimson, E., Fischl, B., A Genetic algorithm for the topology correction of cortical surfaces, *IPMI 2005*
40. Leventon, M., Grimson, W.E.L., Faugeras, O., Statistical shape influence in geodesic active contours. *CVPR*, 2000.
41. Leventon, M., Faugeras, O., Grimson, W.E.L., Wells, W.M., Level set based segmentation with intensity and curvature priors. *MMBIA*, 2000.
42. Tsai, A., Yezzi, A., Willsky, A.S., A curve evolution approach to smoothing and segmentation using the Mumford-Shah functional. *CVPR 2000*: 1119-1124
43. Tsai, A., Wells, W., Tempany, C., Grimson, E., Willsky, A., Coupled Multi-shape model and mutual information for medical image segmentation, In *Proceedings IPMI, LNCS 2732*:185-197, 2003.
44. Yang, J., Staib, L.H., Duncan, J.S., Neighbor-constrained segmentation with 3d deformable models. In *Proceedings of Information Processing in Medical Imaging (IPMI'03)*, LNCS 2732:198-209, 2003.
45. Pohl, K.M., Warfield, S.K., Kikinis, R., Grimson, W.E.L., Wells, W.M., Coupling statistical segmentation and PCA shape modeling, *MICCAI, LNCS 3216*, Springer 2004, 151-159.