

Volumetric Layer Segmentation Using Coupled Surfaces Propagation*

Xiaolan Zeng[†], Lawrence H. Staib[†], Robert T. Schultz[‡] and James S. Duncan[†]
 Department of Electrical Engineering[†] Child Study Center[‡]
 and Diagnostic Radiology Yale University
 Yale University New Haven, CT 06520-8042
 New Haven, CT 06520-8042
 email: xiaolan.zeng@yale.edu

Abstract

The problem of segmenting a volumetric layer of finite thickness is encountered in several important areas within medical image analysis. Key examples include the extraction of the cortical gray matter of the brain and the left ventricle myocardium of the heart. The coupling between the two bounding surfaces of such a layer provides important information that helps to solve the segmentation problem. Here we propose a new approach of coupled surfaces propagation via level set methods, which takes into account coupling as an important constraint. By evolving two embedded surfaces simultaneously, each driven by its own image-derived information while maintaining the coupling, we capture a representation of the two bounding surfaces and achieve automatic segmentation on the layer. Characteristic gray level values, instead of image gradient information alone, are incorporated in deriving the useful image information to drive the surface propagation, which enables our approach to capture the homogeneity inside the layer. The level set implementation offers the advantage of easy initialization, computational efficiency and the ability to capture deep folds of the sulci. As a test example, we apply our approach to unedited 3D Magnetic Resonance(MR) brain images. Our algorithm automatically isolates the brain from non-brain structures and recovers the cortical gray matter.

1 Introduction

Three-dimensional imagery such as MRI and Computed Tomography(CT) is often represented as a set of intensity voxels. In the analysis of such data, we often encounter the problem of segmenting a volumetric layer of finite thickness. For example, in 3D cardiac image analysis, the myocardium of the left ventricle

is a thick structure bounded by the endocardial and epicardial walls. In the analysis of neuroanatomical structures from 3D MR images, the cortex, which is the outermost layer of gray matter in the brain, is bounded by the outer cortical surface (Cerebral Spinal Fluid(CSF)/gray matter interface) and the inner cortical surface (gray/white matter interface). The two surfaces which bound the layer can be viewed as coupled in several ways, ranging from loosely coupled (e.g. the endocardial and the epicardial boundaries) to tight coupling where the thickness of the layer is nearly constant(e.g. the outer and the inner cortical surfaces).

The segmentation of such a volumetric layer must be considered in 3D. Slice by slice segmentation and stacking of the results doesn't use the inter-slice information, hence the result may be either non-optimal or grossly inaccurate. For example, although the volumetric layer of cortical gray matter is constantly about 3.0mm's thick([1]), due to its convoluted nature, an oblique 2D slice that happens to be approximately parallel to a particular sulcus will give the appearance of a much thicker structure. Only by going through the neighboring slices can we get complete information to perform segmentation. Slice by slice manual segmentation of these structures is tedious and labor intensive, hence automatic segmentation is a highly desirable goal. There has been a large amount of work along this line.

1.1 Related work

One form of related effort is region-based method, which exploits homogeneity in images. It primarily depends on the underlying consistency of any relevant feature in the different regions. Following the work of Geman & Geman [7], Markov Random Field (MRF)-based methods have been widely used, which employ energy minimizing techniques to reconstruct a piecewise flat image from the noisy data. In [2] a multispectral voxel classification method is used in conjunction

*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1998.

with connectivity to segment the brain into different tissue types from 3D MR images. A material mixture model[11] is also used for the segmentation problem. Region-based methods typically require further processing to group segmented regions into coherent structure(s).

The typical alternative strategy is boundary finding, of which active contour methods are of special note. They rely mainly on gradient features for segmentation of structures from an image. One of the most generic and popular methods of detecting boundaries is the snakes approach due to Kass et al [8]. One concern regarding this method is that a close initialization has to be provided in order to achieve good final results. A balloon model with a pressure force outward was then introduced in [3] as a way to generalize and solve some of the problems encountered with the above snake method. In [4], 3D deformable surface models using the finite-element method are used to segment 3D images. However, the need to override local smoothness to allow for significant protrusions that a shape may possess (which is highly desirable in order to capture the folds of the sulci, and papillary muscle of the left ventricle) remains a problem.

All of the above methods are limited when used to segment the volumetric layer of finite thickness, in the sense that they don't explicitly use the information of how the two bounding surfaces are coupled. There has been some effort made in this direction. Davatzikos et al. [5] introduced the concept of a ribbon for modeling the outer cortex in cross-sectional brain images and proposed an active contour algorithm for determining the spline of such a ribbon. This model is then extended to 3D in [6]. However close initialization and significant human interaction are still needed to force the ribbon into the sulcal folds. It was also reported that the iterative algorithm was fairly computationally demanding, while a multi-scale formulation could be used to decrease the computational load. No processing time was given.

2 Our Approach

The nature of our problem is as follows: the volumetric layer we're trying to recover is bounded by two surfaces. Across each surface, there is a local difference in the gray scale values, while in between the two surfaces there is a homogeneity of certain voxel statistics. The layer is defined completely by its bounding surfaces and the homogeneity in between. Based on these facts, we propose a new approach of coupled surfaces propagation via level set methods, which takes into account the coupling information. By evolving two embedded surfaces simultaneously, each driven by its




Figure 1: Local operator to derive image information

own image-based information while maintaining the coupling, we are able to achieve an automatic, stable and robust segmentation on this special set of volumes. The information used here to drive the surface propagation is based on gray level values instead of image gradient alone, which enables our approach to reward the homogeneity of inside the layer.

2.1 Single surface approach vs. coupled surfaces approach

Because of the limitations of the imaging technique used and volume averaging, it is often observed that in some regions, there's not enough information from the image data to clearly define either the outer bounding surface or the inner bounding surface. When applying a single surface approach, we may very well end up with error in such a region. While using the coupled surfaces approach, information on the partner surface is available through the coupling and improves the performance of the surface finding.

In the case of MR brain images, due to volume averaging, in some regions the boundary between the white matter and gray matter is not well shown, while the CSF appears clearly. The single surface approach may have the inner cortical surface collapse into CSF. However with the coupled surfaces approach, we maintain some minimal distance between the inner cortical surface and CSF, thus preventing the inner cortical surface from going into CSF. There are also places where structures such as eye sockets appear, so that the CSF can't be observed in the image. With the coupled surfaces approach, the information of the white/gray boundary is then used to stop the propagation of the outer cortical surface before it penetrates non-brain structures.

2.2 Image information derivation

Medical images consist of a number of different anatomical regions. The homogeneity of each region can usually be characterized by various voxel statistics in the region. Thus, by using gradient features (information of gray level difference between different tissues) alone, we are actually losing important pieces of information. Here in our approach, instead of using gradient features, we design a local operator which makes use of the gray level information, and

gives a measure of the likelihood of a voxel lying on the boundary between tissue A and tissue B . This model can also be extended to make use of several differently parametric images (such as T1, T2 and PD MR images) or images from different modalities.

We show in Figure 1 the operator we use. At each voxel site s , a small neighborhood around s is drawn. Now given the possible boundary with normal direction $\vec{\theta}$, dividing the neighborhood into parts $R1$ and $R2$, the probability that s lies on the boundary between tissue A and tissue B is:

$$p(\vec{\theta}) = p(R1 \in TissueA) \cdot p(R2 \in TissueB) \quad (1)$$

Given an estimation $\vec{\theta}^*$ of $\vec{\theta}$, we can use $p(\vec{\theta}^*)$ as a measure of the likelihood that s lies on the boundary between tissue A and tissue B .

One way of estimating $\vec{\theta}^*$ is to first generate the vector $P = [p(\vec{\theta}_1), p(\vec{\theta}_2), \dots, p(\vec{\theta}_k)]^T$ where k is the number of possible directions, and $\vec{\theta}^*$ is then the one which corresponds to the biggest element in vector P . In the case when different parameterized images are available, P becomes a matrix with each column being the probability vector obtained from a single parametric image. An estimate of $p(\vec{\theta}^*)$ is then based on the matrix P .

Here we make the assumption of one single parametric image X , in which pixels belonging to tissue A are independently drawn from Gaussian distribution $G(\mu_A, \sigma_A)$, and pixels belonging to tissue B are independently drawn from $G(\mu_B, \sigma_B)$. Then we have

$$p_{AB}(\vec{\theta}) = \prod_{r \in R1} \frac{1}{\sqrt{2\pi}\sigma_A} e^{-\frac{(x_r - \mu_A)^2}{\sigma_A^2}} \cdot \prod_{t \in R2} \frac{1}{\sqrt{2\pi}\sigma_B} e^{-\frac{(x_t - \mu_B)^2}{\sigma_B^2}} \quad (2)$$

In Figure 2, we show examples of our local operator. The local operator was applied to images after we reduced the effects of MR inhomogeneity by correcting using a simple fixed map. The map was determined manually by sampling tissue types throughout the field to decide the average inhomogeneity. Note that more complicated MR image models ([7],[10],[9]) can be used to calculate $p(\vec{\theta})$.

2.3 Level set method

Level set methods ([15],[13],[12]) are powerful numerical techniques for analyzing and computing interface motion. The essential idea here is to first represent the surface (in our case) of interest as a front $\gamma(t)$. Now given a moving closed front $\gamma(t)$, the idea





Figure 2: (a)-(d) Axial and (e)-(h) coronal slices from 3D brain images. (a),(e): original images; (b),(f): result from gradient operator; (c),(g): result from our local operator $P_{BC}(\vec{\theta})$, B =gray matter, C =white matter; (d),(h): $P_{AB}(\vec{\theta})$, A =CSF, B =gray matter.

is to produce an Eulerian formulation for the motion of this surface propagating along its normal direction with speed F , where F can be a function of the surface characteristics (such as the curvature, normal direction etc.) and the image characteristics (e.g. the gray level, or gradient etc.) This is done by embedding the propagating interface as the zero level set of a higher dimensional function Ψ defined by $\Psi(x, t) = d$, where d is the signed distance from position x to $\gamma(t)$. The equation of evolution of Ψ , inside which our surface is embedded as the zero level set is given by:

$$\Psi_t + F |\nabla \Psi| = 0 \quad (3)$$

The major advantages of using this method over other active contour strategies include the following:


 Figure 3: Functions g and h used in speed term design

first, the evolving function $\Psi(x, t)$ remains a function. However the propagating hypersurface $\gamma(t)$ may change topology, break, merge and form sharp corners as the function Ψ evolves. Second, the intrinsic geometric properties of the front may be easily determined from the level function Ψ . For example, at any point of the front the normal vector is given by $\vec{n} = \nabla \Psi$, and the curvature is easily obtained from the divergence of the gradient of the unit normal vector to front, that is, $K = \nabla \cdot \frac{\nabla \Psi}{|\nabla \Psi|}$.

2.4 Coupled surfaces propagation, speed term design

In solving the problem of segmenting the volumetric layer bounded by two surfaces, we consider two moving interfaces describing the outer bounding surface and the inner bounding surface respectively. Starting from either inside the inner bounding surface or outside the outer bounding surface, with an offset in between, the interfaces propagate along the normal direction stopping at the desired place, while maintaining the distance between them.


Embedding each surface as the zero level set in its own level function, we have two equations:

$$\Psi_{in_t} + F_{in} |\nabla \Psi_{in}| = 0 \quad (4)$$

$$\Psi_{out_t} + F_{out} |\nabla \Psi_{out}| = 0 \quad (5)$$

where F_{in} and F_{out} are functions of the surface normal direction, image-derived information and distance between the two surfaces. The coupling is embedded in the designing of F_{in} and F_{out} . The basic ideas are that at places where the distance between the two surfaces is within the normal range, the two surfaces propagate according to the image-based information; where the distance between the two surfaces is out of the normal range, then the distance imposes a constraint on the propagation of the surfaces.

With the level set implementation we have an easy and natural way to evaluate the distance between the two surfaces. Recall that the value of the level function of a front at any point is simply the distance from this point to the current front, which as in [15], is


 Figure 4: Narrow band implementation: dynamic construction of the narrow band and update of the level function ψ within the narrow band are performed in the neighborhood of the current surface.

calculated as the shortest distance from this point to all the points on the front. In our case of the two moving surfaces, for any point on the inner moving surface, the distance to the outer moving surface is the value Ψ_{out} at this point, and vice versa for the point on the outer moving surface. Hence, we write

$$F_{in} = g(p_{BC}(\vec{\theta}^*))h(\Psi_{out}) \quad (6)$$

$$F_{out} = g(p_{AB}(\vec{\theta}^*))h(\Psi_{in}) \quad (7)$$

where g and h are the functions as shown in Figure 3, and A, B, C denote CSF, gray matter and white matter respectively.

Function g maps larger probability to slower speed, i.e., as the probability gets larger, g tends to zero, while as the probability gets near zero, g tends to a constant. Function h penalizes the distance off the normal range, as the distance goes out of normal range, h goes to zero. Thus, each surface moves with constant speed along the normal direction, slows down when either the image-based information becomes strong or the distance to the other surface moves away from the normal range, and finally stops when the image-derived information is strong enough or the distance to the other surfaces is out of the normal range. The speed term only has meaning on the front, i.e. zero level set, it is then extended from the zero level set to the whole image grid as in [15].

Due to the level set formulation, we have a notion of the inside and outside of the current moving front, and this information is embedded in the normal direction \vec{n} . It actually helps to reduce the feasible space of possible $\vec{\theta}$ s, or provides an estimate of $\vec{\theta}$, thus obtaining a better result.

2.5 Narrow band implementation

The algorithm is implemented using a narrow band method, which modifies the level set method so that it only affects the points close to the current propagating fronts (see Figure 4). Based on the fact that any point




Figure 5: (a): Original synthetic image of a ribbon; (b): initialization for the algorithm; (c): an intermediate stage of propagation; (d): final ribbon captured with the constraint of nearly constant width imposed.

b in the narrow band of the current surface should be within some neighborhood of a certain point a on the surface, the narrow band is constructed dynamically in the neighborhood of each point on the current surface by including points that lie less than some given distance (bandwidth) away from that particular point. Also since a point b in the narrow band can be within the neighborhoods of several points a_1, \dots, a_k on the current surface, we update the value of the level function Ψ at b to be $sign(\Psi) \min_{i=1, \dots, k} |dist(b, a_i)|$. Thus, for a surface represented using N points, the construction of its narrow band and the update of ψ in the narrow band is an $O(N)$ calculation. The bandwidth is chosen such that the inner surface lies within the narrow band of the outer surface and vice versa for the outer surface, to ensure that the distance-based correspondence between the coupled surfaces falls out automatically. Thus, at each time step, the current position of the propagating coupled surfaces and the surrounding narrow bands are estimated, and the whole process repeats until the speed terms for both the inner and outer surfaces reach a zero value everywhere.

3 Experimental Result

For ease of illustration, we first show results employing a 2D synthetic image. As shown in Figure 5, the original image contains a closed ribbon (i.e. a layer in 2D). Other than at the folds, the ribbon is of nearly constant width. We see from the result that the algorithm nicely utilizes the constraint of nearly constant width on the folds. This property is highly desirable in capturing the sulci on the outer cortical surface from the MR images of the brain.

More substantively we have also tested the algorithm on 14 sets of 3D MR brain images to segment the brain from surrounding non-brain matter (e.g. skull) and the cortical gray matter from the overlaying CSF and underlying white matter. Subjects were imaged with a GE 1.5 T scanner, using a 3D SPGR volume acquisition. Contiguous, 1.2 mm thick sagittal images were acquired covering the entire brain with the




Figure 6: Propagation of the outer bounding surface(left) and inner bounding surface(right); (a) initialization; (b),(c) intermediate steps; (d) final result.

following parameters: TR=24, TE=5, flip angle=45, matrix=192x256, NEX=2, FOV=30 cm, 124 images. This yields images with essentially isotropic voxels (1.17 x 1.17 mm in plane, 1.2 mm thick). The images were preprocessed to reduce the effects of MR bias field inhomogeneity using a simple standard non-linear map, which was also a step before the expert manual tracing described below.

The algorithm works in the following fashion. We start by initializing several pairs of concentric spheres inside the white matter in the 3D brain volume, with an offset between the outer spheres and the corresponding inner spheres. The outer spheres then grow out and merge to capture the outer cortical surface, while the inner spheres grow out to capture the inner cortical surface. The propagation is driven by the image-derived information, constrained by the requirement of maintaining the coupling between the outer and inner surfaces. The captured outer cortical surface defines the outer boundary of the brain,




Figure 7: Three views of the outer(left) and inner(right) cortical surfaces. (a): sagittal view; (b): axial view; (c): coronal view.

brain volume TP rate(%)	brain volume FP rate(%)	cortex TP rate(%) (axial)	cortex TP rate(%) (coronal)
94.3	3.2	86.9	86.6
93.4	3.8	87.1	85.9
95.3	4.0	87.0	85.6
95.4	3.3	86.2	85.7
94.6	5.1	83.7	82.6
97.2	5.1	88.2	86.8
95.3	3.9	87.3	86.2
95.5	3.7	86.6	85.3
95.6	3.8	88.6	85.6
95.0	3.4	87.9	88.8
94.0	2.0	89.7	89.5
94.5	4.8	87.0	84.4
95.3	4.0	84.7	85.4
97.2	3.9	87.0	87.8

Table 1: Comparison of the whole brain volume and cortical gray matter from 3D algorithm vs. from expert slice by slice manual tracing on 14 3D MR brain images. TP: true positive; FP: false positive.




Figure 8: Axial(top) and coronal(bottom) slices from 3D images. (a): original image; (b): cortical gray matter from manual tracing; (c): cortical gray matter from our 3D algorithm.

automatically segmenting brain from non-brain matter. The cortical gray matter is captured between the inner and outer cortical surfaces. Figure 6 shows the growing process of the outer and inner spheres.

The initialization only requires the user to specify several pairs of concentric spheres (more specifically, only the center of the concentric spheres and the radius of the inner sphere for each pair), which can be done within minutes. Once the starting spheres are planted, there's no further user interaction. In our implementation, the normal range of the distance between the outer and inner surfaces is set to be 3.0mm with an allowed variation of -20% to 30%. It should be emphasized that neither the number nor the placement of the spheres during the initialization affect the accuracy or reproducibility of the final result (due to the ability of the level set method to handle topology change), only the computational efficiency of the algorithm.

Expert manual tracing is done in the following steps: first isolate the brain from non-brain structures, then define the CSF/gray matter interface and finally define the gray/white matter interface. All these are done slice by slice. For each data set, the first two steps alone take 3-4 hours, and the third step is many times more time consuming (usually 1-2 days). The algorithm described in this paper runs in about 1.5 hours on a R10000 SGI Indigo2 machine. We are currently working on the optimization of our codes, which should help to reduce the processing time more.

A quantitative comparison of our algorithm and the expert manual tracing are provided in the table. In




Figure 9: 3D volume rendering of the cortical gray matter from our 3D algorithm with oblique cutting planes. The convoluted thin bright ribbons are the cortical gray matter captured on the cutting plane, the darker parts are the rendering of the out-of-plane gray matter.

the first two columns, we compare the whole brain volume captured from our 3D algorithm (referred to as V_a) and that from the expert tracing (referred to as V_e). We denote the overlap between V_a and V_e as V_{ae} , and the part that is in V_a but not in V_e as $V_{ae'}$. The true positive rate is then defined to be the ratio of the size of V_{ae} to the size of V_e , while the false positive rate is defined to be the ratio of the size of $V_{ae'}$ to the size of V_e . The true positive rate for the whole brain ranges from 93.4% to 97.2% in the sample of 14 images. Only part of the brain stem and cerebellum are recovered, which contributes to most part of the error in the true positive rate. This occurs because the constraint of nearly constant thickness is not as well satisfied in these areas as in the cerebral hemispheres. The missing part of the brain stem and cerebellum can be easily picked up by hand editing within minutes. The 2%-5.1% false positive rate is in part a function of the variability in expert tracing performance, and in part due to the over inclusiveness of the algorithm. The over inclusiveness is due to the voxels adjacent to brain with similar gray level values. From the 3D views shown in Figure 7, we see our algorithm well defines the outer and inner cortical surfaces.

The results of cortical gray matter recovered from our algorithm is compared to the expert results on one axial and one coronal slices to provide the true positive rate in the table. The true positive rate here is defined in the same fashion as for the whole brain volume. The true positive rate varies from 82.6%-89.7%. The

study of repeated manual tracings by experts shows a reproducibility rate (true positive rate when compare results from expert repetition) of 90%-95%. Comparatively, our algorithm performs reasonably well. Figure 8 shows the comparison of cortical gray matter from our algorithm and expert result on several slices. In Figure 9, we show the 3D volume rendering of the cortical gray matter, with oblique cutting planes.

Whole brain volume and volume of its constituent parts, white and gray matter, are typically the measurements of interest for comparison among different subjects in studies of neuroanatomy. Thus, a second way to analyze the utility of our algorithm is to compute reliability statistics on volume measurement. Reliability is best assessed by the intraclass correlation (ICC) ([14]), because it is sensitive to both the rank order associations and mean differences produced by the two measurement procedures (expert and algorithm). Although the tracer is expert in defining neuroanatomical features, hand tracing inevitably produces some error, and thus it is important to determine how large the algorithm's error is relative to the expert's own rate of error. Repeat measurement of whole brain volume on 15 subjects by two independent expert tracers involved virtually no error (ICC = .997; average volumes differed by 0.6%). Whole brain volume measured by the algorithm compared to the expert tracing was also excellent (ICC = .990; average volumes differed by 1.35%). For gray/white matter segmentation, a single expert tracer was able to reproduce gray matter volumes across 7 subjects with an average of 3.1% error (ICC = .972). Our algorithm, on the other hand, averaged 4.4% error compared to the expert tracings (ICC = .923). Thus, for both whole brain volume and gray matter volume, the coupled surfaces algorithm produced measurements that were highly similar to expert tracings.

4 Summary and Future Directions

In this paper, we presented a new approach to the problem of segmenting a volumetric layer of finite thickness. Two bounding surfaces and the homogeneous nature of the gray level values in between the surfaces define such a layer. By evolving two embedded surfaces simultaneously, each driven by its own image-derived information while maintaining the coupling, a final representation of the bounding surfaces and automatic segmentation of the layer are achieved. A local operator based on gray level values is designed to capture the information to drive surface finding. The use of such gray level-based information instead of the image gradient, gives our algorithm the ability to capture the homogeneity of tissue, and in

this way it improves the performance. The level set implementation offers the advantage of easy initialization, ability to handle complex geometry and topology change, as well as the ready evaluation of several characteristics of the layer, such as the surface curvature and the distance map between the two bounding surfaces.

Curvature maps for both the outer and inner bounding surfaces can be calculated using equation in 2.3. The information of high curvature point can then be used to guide 3D tracing and mapping. Also, the thickness map between the outer and inner surfaces can be readily evaluated through Ψ_{in} and Ψ_{out} , which can be used to detect pathology.

Future directions for this work include the following: finer design of the local feature operator to better model the volume averaging effect, better capturing the homogeneity of the volume, the calculation of the curvature of the bounding surfaces, and the calculation of the thickness map to help in abnormality diagnosis.

Acknowledgment

This work was supported in part by NSF grant IRI-9530768, and NIH grants NINDS R01 NS35 193-01 and NICHD 5 P01 HDIDC 35482. The authors would like to thank James Rambo and Lawrence Win for their help with visualizing and manual tracing of the MR data.

References

- [1] S.M.Blinkov and I.I.Glezer. *The Human Brain In Figures and Tables. A Quantitative Handbook*. Chapter X, pp182. Basic Books,Inc., Publishers, Plenum Press, 1968.
- [2] H.E.Cline, W.E.Lorensen, R. Kikinis, and F.Jolesz. Three-dimensional segmentation of MR images of the head using probability and connectivity. *Journal of Computer Assisted Tomography*, 14(6):1037-1045, Nov./Dec. 1990.
- [3] L.D.Cohen. On active contour models and balloons. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 53(2):211-218, Mar, 1991.
- [4] L.D.Cohen and I.Cohen. Finite-Element methods for active contour models and balloons for 2-D and 3-D images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1131-1147, Nov, 1993.
- [5] C.A.Davatzikos and J.Prince. An active contour model for mapping the cortex. *IEEE Transactions on Medical Imaging*, 14(1)65-80, March, 1995.
- [6] C.Davatzikos and R.N.Bryan. Using a deformable surface model to obtain a shape representation of cortex. *IEEE Transactions on Medical Imaging*, 15(6)785-795, 1996.
- [7] D. Geman and S. Geman. Stochastic relaxation, Gibbs distribution and Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721-741, 1984.
- [8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:312-331, 1988.
- [9] S.Lakshmanan and H.Derin. Simultaneous Parameter Estimation and Segmentation of Gibbs Random Fields Using Simulated Annealing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):799-810, 1989.
- [10] R.Leahy, T.Hebert and R.Lee. Applications of Markov Random Fields in medical imaging. *Information Processing in Medical Imaging*. pp:1-14. Wiley-Liss Inc, 1991.
- [11] Z.Liang, R.F.Jaszczak and R.E.Coleman. Parameter estimation of finite mixtures using the EM algorithm and Information criteria with application to medical image processing. *IEEE Transactions on Nuclear Science*, Vol.39, No.4, pp1126-1133, 1992.
- [12] R.Malladi, R.Kimmel, D.Adalsteinsson, G.Sapiro, V.Caselles and J.A. Sethian. A geometric approach to segmentation and analysis of 3D medical images. *Proceedings of MMBIA*, 1996.
- [13] R.Malladi, J.A.Sethian and B.C.Vemuri. Shape modeling with front propagation: a level set approach, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):158-174, Feb, 1995.
- [14] R.T.Schultz and A.Chakraborty. Magnetic resonance image analysis. In E. Bigler (Ed.), *Handbook of Human Brain Function: Neuroimaging*, pp:9-51. New York: Plenum Press, 1996.
- [15] J.A.Sethian. *Level set methods:evolving interfaces in geometry, fluid mechanics, computer vision and materials science*. Cambridge University Press, 1996.