An Integrated Approach for Locating Neuroanatomical Structure from MRI[1]

Lawrence H. Staib*, Amit Chakraborty† and James S. Duncan*
lawrence.staib@yale.edu chakrab@scr.siemens.com james.duncan@yale.edu
* Departments of Electrical Engineering and Diagnostic Radiology, Yale University
333 Cedar Street, New Haven, CT 06520-8042
† Siemens Corporate Research, 755 College Road East, Princeton, NJ 08540

Abstract

The wide availability of high resolution magnetic resonance images (MRI) of the brain has facilitated tremendous progress in neuroscience. Accurate automated segmentation and quantification of neuroanatomical structure from such images is crucial for the advancement of understanding of brain morphology, both in normal variation and in disease. Gradient-based deformable surface finding is a powerful technique for locating structure in three-dimensional images. However, it often suffers from poorly defined edges and noise. This paper proposes a gradient-based deformable surface finding approach that integrates region information. This makes the resulting procedure more robust to noise and improper initialization. In addition, prior shape information may be incorporated. The algorithm uses Gauss's Divergence theorem to find the surface of a homogeneous region-classified area in the image and integrates this with a gray level gradient-based surface finder. Experimental results on synthetic and MR brain images show a significant improvement is achieved as a consequence of the use of this extra information. Further, these improvements are achieved with little increase in computational overhead, an advantage derived from the application of Gauss's Divergence theorem.

1 Introduction

Magnetic resonance imaging (MRI) allows detailed examination of the morphology of the brain at high resolution and *in vivo*. Three dimensional image analysis is important in this domain in order to facilitate the quantitation necessary for better understanding of normal and abnormal structure. In most cases, the analysis requires the precise identification and quantification of structures and abnormalities in the brain in terms of volume, surface area, location and shape. The study of abnormalities and the normal variation of the shape of brain structures is important in characterizing the brain and will likely to lead to an increased understanding of the normal and abnormal morphology. Brain function can be related to morphology by examining subjects with brain disorders and measuring behavioral correlates to morphology in order to establish structure-function relationships. Size differences have been noted in a variety of brain disorders, including, for example, the hippocampus in posttraumatic stress disorder (PTSD) [1], the temporal lobe in learning disabilities [2] and the corpus callosum in normal twins [3]. Shape can be fully characterized by curvature in an invariant way and shape differences have also been found, for example, in the

[1]This work was supported in part by the National Institute of Neurological Disorders and Stroke under Grant NS 35193.

corpus callosum, exhibiting sexual dimorphism [4] and in the cortex, showing shape differences due to atrophy [5].

Such analysis is important not only for structural measurement of anatomy, but for region identification for measurement from functional images, such as positron emission tomography (PET) or single photon emission tomography (SPECT) or functional MRI. These studies even more directly help understanding of brain function.

This paper describes a 3D deformable surface finding methodology that integrates region information and gradient information to find a complete surface that matches the gradient strength in the image and surrounds a homogeneous region. There is a great benefit to consider the entire 3D image set as a whole and analyze it in that way. Often, 3D images are treated as stacks of 2D images, thereby reducing the dimensionality of the problem. While sometimes successful, such methods tend to oversimplify or ignore the 3D properties of the structures under consideration. This is especially important in the brain where most structures bend and turn and do not have a preferred orientation for analysis. The robust identification and measurement of deformable structures such as are found in the brain, is not always achievable by using a single technique that depends on a single image-derived source of information. Thus, there is also a great need for integrated methods that make optimal use of the multiple sources of information.

2 Background

Region-based methods The two principle sources of image derived information that are used by most segmentation methods are region-based and boundary-based. Region-based methods [6, 7, 8, 9, 10, 11, 12, 13] rely on the homogeneity of spatially localized features such as gray level intensity, texture and other local pixel statistics. Homogeneity does not necessarily mean identical pixel values within a particular region, rather it means that the variation within a region is of a smaller extent than that between regions. The advantage of such methods is that they rely directly on the gray level image and thus are less susceptible to noise than methods that use derivative information. Also, if the high frequency information in an image is either missing or is unreliable, the segmented images remain relatively unaffected. Since the typical MR image can be both noisy and have fuzzy boundaries, these features can be very helpful. However, the problem with typical region-based segmentation methods, is that the resulting segmentation depends considerably on the choice of seed points and the region's shape is too dependent on the choice of the actual algorithm used. Also, such methods often result in an over-segmented image. Rule based systems [14] can do better, but are extremely application-specific. Other region-based methods either use probabilistic techniques [10, 15, 16, 17, 11] or use non-linear diffusion methods [12, 13, 18] (see [11] or [12] for the exact mathematical relationship between these methods). These methods perform a type of smoothing that preserves edges. However, isolating objects from the resulting image still requires considerable effort as they also suffer from the problems of poor localization and over-segmentation (related to the problem of choosing the appropriate scale).

Region methods are particularly susceptible to gray level variations over the image. This is particularly relevant for MR images which often suffer from inhomogeneities of the gain field.

The pixel intensities in MR images are often distorted by inhomogeneities due to RF coil field strength. These effects typically result in a low spatial frequency multiplicative corruption of the image data. Such inhomogeneities are particularly troublesome for pixel classification methods because of their reliance on gray level values. Thus, the correction of inhomogeneities is crucial for the application of gray-level-based techniques to MRI. However, methods using gray-level gradients or higher-order statistics are less sensitive to these effects. Intensity calibration based on phantoms has been tried [19] but these methods are limited by the non-linear response of tissue to RF excitation, and gain inhomogeneities due to the interaction of anatomy with the RF coils. Filtering techniques such as homomorphic filtering (a combination of logarithmic transformation and high-pass filtering) or dividing by a low-pass filtered image has been used to correct these distortions [19, 20]. These techniques can be effective but are limited by the variation of intensity between tissue types. If a single tissue type can be identified throughout the image, the distortion can be more accurately measured, assuming that the tissue type should have a homogeneous value wherever it is measured. A number of techniques are based on this idea. Dawant et al. [21] correct for intensity variation by fitting a surface to classified points either from an automated classification operation or from user specified points. They found better correction but sensitivity to operator error using user-specified points. Meyer et al. [22] also classify first to determine the correction factor, correct and then reclassify the image. Wells et al. [23] uses an MRF approach to simultaneously estimate the classification and the distortion field by alternatively iterating their estimation.

A number of techniques for MR brain segmentation rely on voxel classification using region based methods. These techniques assume that one or more feature parameters, such as T1 or T2, of a particular tissue type will have values that cluster in such a way that the different brain tissues can be distinguished from each other. In order to determine these clusters, representative voxels must be manually identified (supervised) or clusters can be automatically determined (unsupervised). Cline et al. [24] use multispectral voxel classification, in conjunction with connectivity, to segment the head into background, facial tissue, brain, cerebrospinal fluid (CSF) and lesions from 3D MR images. This method is limited by the assumption of normality in the probability distributions of the tissues. Gerig et al. [25] use a similar approach. Raya [26] uses multispectral classification in conjunction with a rule base to find brain, CSF and abnormalities from brain MR images. Brummer et al. [27] use classification via histogram thresholding, in conjunction with morphological operations, to detect the brain from MR. However, these region methods all suffer from the problems of field inhomogeneities and poor localization of boundaries. Hence, region based methods are likely to be inadequate by themselves for a reliable segmentation of neuroanatomic structures of interest from MR images of the brain.

Gradient-based methods In contrast to region-based methods, boundary methods primarily use gradient information [28, 29, 30] to locate object boundaries. Such methods have very good localization properties since they rely on the gradient. Gradient methods also are relatively unaffected by changes or inhomogeneities in the gray scale distributions. However, boundary finding in 3D using only local information is not sufficiently robust to the effects of noise, which is enhanced by

differentiation, poor contrast and the presence of other nearby objects.

In addition, some of the methods that are applicable in 2D can no longer be used for 3D images. For example, pixel search methods that follow an optimal path through the two dimensional images cannot naturally be extended to three dimensions because the voxels in a surface have no such ordering. Hough transform methods [31] can be used, but for three dimensional images it is very expensive both in terms of storage and computational costs. To overcome these problems, the use of whole boundary methods has proved successful. These methods avoid the problem of broken or missing edges by imposing a structure to the solution. They augment imperfect image data with shape information provided by a geometric model [32, 33, 34] and if formulated parametrically, they form over-constrained estimates that use a few parameters to describe a large number of points. Such deformable models have been used extensively in medical image analysis [35].

A number of gradient-based methods have been applied to MR brain images. Bomans et al. [36] use a boundary-finding method based on a 3D version of the Marr-Hildreth edge operator [29] to find surfaces the brain and ventricles in MR images. Morphological operators are used to remove small holes and thin connections. Raman et al. [37] track Marr-Hildreth edges from coarse scale to fine scale with the intent of detecting significant edges accurately in MR images of the brain. These methods rely primarily on edges, however, which can be effective in localizing boundaries when there are strong transitions in the intensity but ignore other cues such as homogeneity which can aid in regions where edges are indistinct. Thus, edges by themselves are likely to be insufficient to reliably and accurately segment the brain.

Interactive methods Interactive and semi-automated methods are a compromise between hand tracing and fully automated methods. Kennedy et al. [38, 39] describe a number of semi-automated methods for segmenting MR images of the brain. Hohne and Hanson [40] use mathematical morphology, connected components and thresholding to interactively segment 3D images with feedback from rendered displays. Andreasen et al. [5] uses a combination of manual tracing and thresholding to measure cortical curvature. These techniques rely on simple features and user tracing and decisions to determine the segmentation. User interaction of this sort directly effecting the segmentation leads to variability and inefficiency. It is preferable to reserve interaction for initialization so that the final results are less sensitive to user variability.

Vector valued images In MR imaging, voxel values are a function of three tissue parameters: proton density, T1 and T2 relaxation. By using multiecho acquisition sequences, either directly or to compute the MR parameters, vector-valued images can be formed. The different components can be used for multidimensional classification by clustering in this feature space. These clustering methods are limited by their reliance on features of individual positions and ignoring the spatial context. This vector information can also be used to determine boundary features, by calculating vector gradients and using them in the same way as scalar gray level gradients. The gradient of a vector image, $\mathbf{F}$, is the direction and magnitude of greatest change. The formulation of Lee and Cok [41] and Cumani [42] defines the gradient magnitude of a vector image based on the matrix

of first partial derivatives $\mathbf{F}'$. These derivatives can be calculated in the same way as scalar gray-level gradients using smoothed differentiation. The gradient magnitude, in this formulation, is the square root of the largest eigenvalue of $\mathbf{F}'^T\mathbf{F}'$. This matrix can be written as:

$$(\mathbf{F}'^T\mathbf{F}')_{ik} = \sum_{j=1}^{m} \frac{\partial \mathbf{F}}{\partial x_i}\frac{\partial \mathbf{F}}{\partial x_k} \tag{1}$$

The gradient direction is defined by the corresponding eigenvector. This vector gradient represents the direction and magnitude of greatest change in the feature space which directly corresponds with the standard definition of a scalar gradient. Thus, region and gradient based information can be extracted from vector valued MR images and applied in the same way as scalar images.

Prior knowledge An important aspect of user interaction is that the user has the advantage of imposing prior knowledge of the shape of the structure being measured based on their knowledge of neuroanatomy. The advantage of formalizing this information and using it in the segmentation process has begun to be recognized. Staib and Duncan [43, 33] augment the boundary finding process with *a priori* probability information representing the mean shape and the natural variation of the structure to be segmented. Collins et al. [44] segment the brain using an elastic registration to an average brain based on a hierarchical local correlation. The average brain provides strong prior information about the expected image data. Cootes et al. [45] augment a snake-like model with statistics to model structures in medical images in order to locate them. Prior information can act as a strong guide when the intrinsic image information is weak.

Integration Methods

While both the region and boundary methods have their advantages and disadvantages their problems are not necessarily identical. They are not affected in the same way by limitations in the quality of the image. While the presence of noise limits the performance of any image processing algorithm, region-based methods are less affected than gradient-based boundary finding because the gradient is very noise sensitive. Also, if the high frequency information in the image either is missing or is unreliable, boundary finding is more error-prone compared to region-based segmentation. Shape variations, on the other hand, can be better handled using a deformable boundary finding framework when we consider such variations to be generally around an average shape and such information can easily be incorporated as priors [43]. Further, since conventional boundary finding relies on changes in the gray level, rather than their actual values, it is less sensitive to changes in the gray scale distributions, such as MR inhomogeneities. Also, gradient-based methods in general do a better job of edge localization. Given these properties, integrated methods are likely to perform better than either of the methods alone by being able to combine the complementary strengths of these individual methods, as has been pointed out [46, 47].

Unfortunately, however, only a limited amount of previous work has been done seeking to integrate region and boundary information, and primarily for 2D images. Among previous methods, AI-based techniques have been used where production rules are invoked for conflict removal

[46]. In such methods, region growing is done first followed by a binary edge detection step. There are a few disadvantages to this procedure. First, a region classified image is often over-segmented due to the presence of noise. Thus, one needs a validating scheme to distinguish between true and false edges by looking at high gradient, continuity, etc. Also, such schemes have no way of differentiating between multiple regions as it deals with the binary edge map obtained from the region grown image. Further, such methods may suffer from the effects of poor edge localization as is often the case with region based segmentation. Other similar efforts [47, 48] were aimed at integrating region growing with edge-detection rather than finding complex objects.

Probability based approaches [10, 15, 17, 11] typically aim to maximize the *a posteriori* probability of the region classified image given the raw image data by optimization methods like simulated annealing. Integration here is achieved in the local or dense field sense where the edges are used as line processes and the optimization is achieved both over the location of the line processes as well as the pixel classification. Nonlinear diffusion methods [12, 18] achieve a similar sort of integration in a non-probabilistic framework.

Another way of achieving local integration is the reaction-diffusion method [49]. However, the problem of using such local integration methods is that if any one of the processes makes an error (such as a false edge), it is propagated to the final solution. Also, a decision regarding the final object boundary is made by considering the whole space of reaction-diffusion images and choosing one result [49], something that can get very complicated. Finally, the recent work of Zhu et al. [50] has similar motivations as ours, although as yet, the formulation has not been extended 3D images. The algorithm presented here is an extension of our earlier work on integration for boundary finding in 2D images [51].

3 Surface representation

We represent surfaces using a Fourier parameterization [32, 52, 33]. It is a strong model concisely represented in terms of parameters and easily allows the incorporation of prior information, when such information is available. Associated parameter probability distributions can introduce a bias towards an expected range of shapes. It is a natural extension to our 2D boundary parameterization [43]. There are a number of other approaches to three dimensional parametric modeling including generalized cylinders [53], superquadrics [34], hyperquadrics [54] and finite element methods [55, 56, 57].

A surface in three dimensions can be represented by three coordinate functions of two surface parameters as $\mathbf{x}(u,v) = (x(u,v), y(u,v), z(u,v))$, where u and v are the free parameters that vary over the surface. Since there are two free parameters, a function of two parameters is necessary to describe a surface. The Fourier surface representation uses the following basis [32, 33]:

$$\begin{aligned} \phi \quad = \quad & \{1, \cos mu, \sin mu, \cos lv, \sin lv, \cos mu \cos lv, \\ & \sin mu \cos lv, \cos mu \sin lv, \sin mu \cos lv, \ldots \\ & (m = 1, 2, \ldots ;\ l = 1, 2, \ldots)\} \end{aligned} \tag{2}$$

The functions x, y and z are each composed of a weighted sum of the elements of the above basis

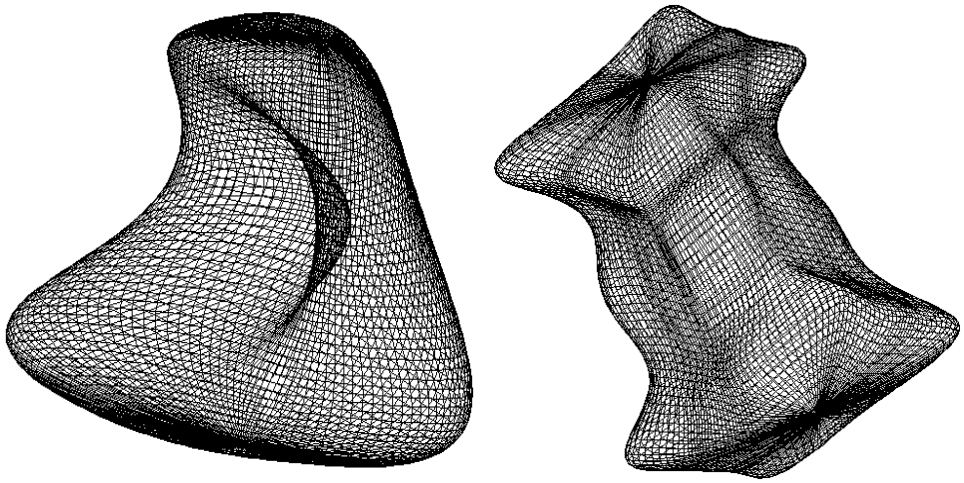

Figure 1: Two closed surface examples using up to four and eight harmonics.

as follows:

$$\begin{aligned} f(u,v) = \sum_{m=0}^{K}\sum_{l=0}^{K} \lambda_{m,l}[\\ & a_{m,l}\cos mu\cos lv + b_{m,l}\sin mu\cos lv + \\ & c_{m,l}\cos mu\sin lv + d_{m,l}\sin mu\sin lv] \end{aligned} \tag{3}$$

where,

$$\lambda_{m,l} \quad = \quad \begin{cases} 1 & for\ m=0,\ l=0 \\ 2 & for\ m>0,\ l=0\ or\ m=0,\ l>0 \\ 4 & for\ m>0,\ l>0 \end{cases}$$

The series is truncated at K, i.e. only a finite number of harmonics are used, in order to limit the search space dimensionality and constrain the space of functions. Taken together, the coefficients form the parameter vector:

$$\vec{p} = (a_x, b_x, c_x, d_x, a_y, b_y, c_y, d_y, a_z, b_z, c_z, d_z) \tag{4}$$

The four basic classes of surfaces in three dimensions are tori (closed tubes), open surfaces (with one edge), tubes (open surfaces with two edges) and closed surfaces. The torus, which is periodic in both the surface variables is formed with the entire basis in Equation 2.

Closed surfaces, suitable for defining regions, are represented using:

$$\phi_{\text{closed}} = \{1, \sin lv, \cos mu\sin lv, \sin mu\sin lv, \ldots\} \tag{5}$$

which forces the functions to be constants at $v = 0, \pi, 2\pi$. This, however, forces the ends together as well. The ends need to be separated by adding a weighted term to each coordinate of the form $sin(v-\pi/2)$ resulting in three more additional shape parameters. Two closed surface examples are shown in Figure 1. It is also possible to represent open surfaces and tubes by this parameterization [33, 52].

4 Region information

Each voxel in the image must be classified into one of a number of regions or classes. Thus, for each voxel, we need to decide or estimate to which class the voxel belongs. There are a variety of approaches to region based segmentation and while there are differences, for our purposes, the performance does not change from one method to the other considerably. The emphasis of this paper is an integrated boundary finding approach given the raw image and the region classified image. The exact method used to get the region classified image is not extremely critical here as long as the output of that method gives reasonable results in part because the integrated method makes the final decision. Any other suitable classification method could be used instead.

For our purposes, we use a method that has found broad applicability, including in the medical domain, which models the image as a Markov Random Field (MRF) and a Maximum *a posteriori* (MAP) probability approach is used to do the classification [58, 59]. The problem is posed as an objective function optimization, which in this case consists of the *a posteriori* probability of the classified image given the raw data which constitutes the likelihood term, and the prior probability term, which, due to the MRF assumption, is given by a Gibb's distribution.

The task is to determine a segmentation X given the raw image Y and our prior knowledge of X. Thus, the aim is to compute $\arg\max_X \Pr(X|Y)$. We model the region process X by a Markov random field. Due to the Markov property,

$$\Pr(x_i|Y, X_{S/i}) = \Pr(x_i|Y, X_{N_i}) \tag{6}$$

where the subscript S/i represents the whole index set for the image except the ith pixel. The subscript N_i denotes the sites neighboring site i.

Using the Hammersley-Clifford theorem [60], the density of X is given by the Gibbs density of the form,

$$\Pr(X) = \frac{1}{Z} \exp\left\{ - \sum_C V_C(X) \right\} \tag{7}$$

Here, Z is a normalizing constant and the summation is over all cliques C. A clique is a set of points that are neighbors of each other. The clique potentials, V_C, depend only on the pixels that belong to clique C. This MRF model allows us to ensure that the resulting segmentation is smooth, in the sense that neighboring pixels are encouraged to have similar properties. In effect, this helps us to filter out the noise.

Now, given that we use this model, the segmentation can simply be performed as an image estimation process. It can be shown that the MAP objective is equivalent to:

$$x_i = \arg \max_{\{x_i \in l; l=1,2,..,L\}} p(y_i|Y_{N/i}, x_i, X_{N/i}) p(x_i|X_{N/i}) \tag{8}$$

where Y corresponds to the actual image data, X corresponds to the region classified image and l represents the classes in X. The subscript N/i represents the neighborhood of the i^{th} pixel leaving out the i^{th} pixel. A first order neighborhood system having six neighbors (2 neighbors along the three axes) is used. At every iteration, the probability of a particular pixel being classified to different classes is computed and the pixel is assigned to the class that gives the highest probability. The procedure stops when there is no change between iterations.

5 Integrated surface finding objective function

We now define the surface finding objective function. By optimizing this function, we determine the surface parameters which correspond to the structure which matches both the gradient strength in the image and the region homogeneity properties. The development follows along similar lines as our earlier work for 2D images [51].

The input to the problem consists of the actual image I and the region classified image I_s, which is obtained from I after passing it through a region based segmentation step, as discussed above. We assume that the interior of the region enclosed by the boundary that we seek belongs to a single region in I_s. All that this assumption requires is that the intra-region variability should be smaller than the inter-region variability. This assumption can be further relaxed, as was done in the 2D case (see [51] for details). The traditional boundary finding problem does not use the original image directly. Being a gradient based approach, it uses instead the gradient image I_g. As in Staib and Duncan [33, 32], we use the magnitude of the gradient vector at each voxel location. A smooth estimate of I_g can be obtained from I by convolving the input image I with the derivative (taken in the three directions) of a Gaussian kernel and then computing I_g, the magnitude of the above resulting vector image. Alternatively, one can first convolve with a Gaussian to smooth the effects of noise followed by taking a finite difference approximation to the partial derivatives in the three directions and then calculating the magnitude of the gradient vector at each voxel location. Thus, the input to the objective function is the gradient image I_g and the region classified image I_s.

The above surface estimation problem using gradient and region homogeneity information can be posed in the maximum *a posteriori* framework. This is suitable for incorporating *a priori* shape information.

Our aim is to maximize $\Pr(\vec{p}|I_g, I_s)$, where as described in the previous section, $\vec{p}$ is the vector of parameters used to parameterize the contour. First,

$$\Pr(\vec{p}|I_g, I_s) = \frac{\Pr(\vec{p}, I_g, I_s)}{\Pr(I_g, I_s)} \tag{9}$$

$$= \frac{\Pr(I_s|I_g, \vec{p}) \Pr(\vec{p}, I_g)}{\Pr(I_g, I_s)} \tag{10}$$

$$= \frac{\Pr(I_s|I_g, \vec{p}) \Pr(I_g|\vec{p}) \Pr(\vec{p})}{\Pr(I_g, I_s)} \tag{11}$$

Furthermore, ignoring the denominator, which does not change with $\vec{p}$, our aim is to determine (after taking the logarithm),

$$\arg\max_{\vec{p}} \Pr(\vec{p}|I_g, I_s) \equiv \arg\max_{\vec{p}} [\ln \Pr(\vec{p}) + \ln \Pr(I_g|\vec{p}) + \ln \Pr(I_s|I_g, \vec{p})] \tag{12}$$

In Equation 12, we have taken the natural logarithm, which is a monotonically increasing function.

Knowledge of I_g could be used in the calculation of I_s, for example, through the use of line processes [10, 17]. However, if we ignore that information, we are effectively discarding information rather than assuming extra information. Thus, the third term in Equation 12 above is simplified

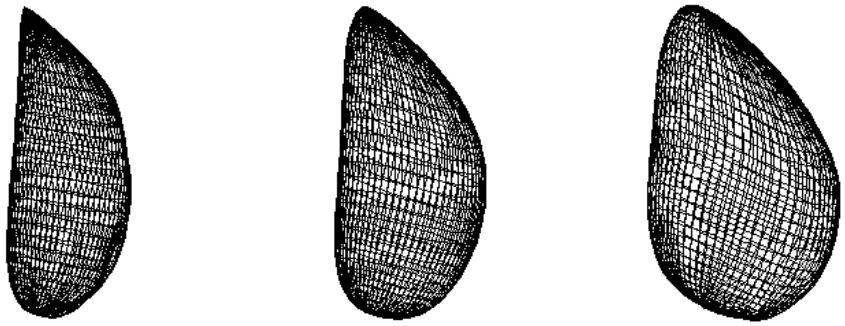

Figure 2: MR example prior. The mean surface (center) is shown with surfaces corresponding to parameters plus and minus one standard deviation. This distribution is used in the example shown in Figures 8 and 9.

using the approximation ignoring the dependence on I_g and we get an objective function that is a sum of three terms, here denoted M (see Chakraborty et al. [51] for additional details):

$$\begin{aligned} \arg\max_{\vec{p}} \; M(\vec{p}, I_g, I_s) \quad &= \quad \arg\max_{\vec{p}} \left[M_{\text{prior}}(\vec{p}) \right. \\ &\quad \left. + M_{\text{gradient}}(I_g, \vec{p}) + M_{\text{region}}(I_s, \vec{p}) \right] \end{aligned} \tag{13}$$

Each of the three terms in the above objective incorporates a different information source.

Prior Term The first term in Equation 13 corresponds to the prior shape term. This prior information is a flexible bias towards more likely shapes. When it is non-uniform, it biases the model towards a particular range of shapes about the mean of the density. The spread in the probability density is due to variability among instances of the object. We use a multivariate Gaussian to model the density. An example density is shown in Figure 2. The middle surface corresponds to the mean parameter values. To the left and right of it are the surfaces corresponding to the mean parameter values plus and minus one standard deviation, respectively.

However, since there might be other objects in the image, we might need an initial estimate of the position to start the optimization process. The information fusion that we present in this case increases the reliability of the surface finding procedure under increased uncertainty in the initial boundary placement and this is borne out by experimental results.

Gradient Term The second term is the gradient likelihood term. It is a measure of the likelihood of the gradient image being the true gradient image corresponding to a particular object boundary. At each point on the surface, the strength of the boundary can be evaluated by the magnitude of the gradient at that particular voxel, given by the gradient image. The likelihood of the image given the boundary parameters can be shown to be proportional to the sum of the magnitude of the gradients at all the points that lie on the surface boundary, given a simplified model assuming that the noise can be approximated by a zero mean Gaussian and that the voxels on the boundary are independent. We can express the above term in the probability expression as the following

area integral (see Staib and Duncan [32, 33] for further details):

$$M_{\text{gradient}}(I_g, \vec{p}) \propto \int\int_{A_{\vec{p}}} I_g(x, y, z) dA \tag{14}$$

where the area element on the surface is given by:

$$dA = |\mathbf{x}_u \times \mathbf{x}_v| du dv \,. \tag{15}$$

Region Term The third term in Equation 13 is responsible for incorporating the region information into the surface finding framework. We expect the bounding surface to surround a homogeneous region. For simplicity, we assume that we have an image where the target object is surrounded by a single background, we assign positive values to the interior of the object and negative values outside. If more than two regions are involved, all pixels of the region that needs to be segmented can be assigned a positive value and the remaining ones negative values, the magnitudes of which reflect how much one expects the target region to be dissociated from the remaining regions. Hence, remote regions are expected to have high negative values, representing a larger penalty for including remote points. This way multiple regions can be handled. Once we have associated positive values with the target object and negative values with points that lie outside, a volume integral that sums up all the points inside the surface is taken. Clearly, this integral would be a maximum when the bounding surface is optimally placed over the object. Thus, the third term in Equation 13 is given by:

$$M_{\text{region}}(I_s, \vec{p}) \propto \int\int\int_{V_{\vec{p}}} I_s(x, y, z) dV \tag{16}$$

Hence, we finally have:

$$\begin{aligned} \arg\max_{\vec{p}} M(\vec{p}, I_g, I_s) &= \max_{\vec{p}} [M_{\text{prior}}(\vec{p}) \\ &\quad + M_{\text{gradient}}(I_g, \vec{p}) + M_{\text{region}}(I_s, \vec{p})] \\ &\equiv \max_{\vec{p}} \left[M_{\text{prior}}(\vec{p}) + K_1 \int\int_{A_{\vec{p}}} I_g(x, y, z) dA \right. \\ &\quad \left. + K_2 \int\int\int_{V_{\vec{p}}} I_s(x, y, z) dV \right] \end{aligned} \tag{17}$$

where K_1 and K_2 are the weighting constants which signify the relative importance of the two terms in the above equation.

Volume to Area Integral Of the last two terms in Equation 17, one is an area integral and the other is a volume integral. In general, computing an area integral is much less expensive compared to a volume integral ($\mathcal{O}(N^2)$ versus $\mathcal{O}(N^3)$, where N is the diameter of the object). Thus, we can save a lot of computation, especially when we carry out an iterative optimization procedure, if we convert the volume integral to an area integral. An area integral must already be computed in any case because the second term which is present in the original surface finding method already

involves the computation of an area integral. Thus, the order of the computational complexity is not increased. The above conversion can be done using Gauss' divergence theorem [61] as follows. First, construct the functions,

$$\begin{aligned}
\mathbf{F}_x(x,y,z) &= \frac{1}{3}\int_0^x I_s(\alpha,y,z)d\alpha \\
\mathbf{F}_y(x,y,z) &= \frac{1}{3}\int_0^y I_s(x,\beta,z)d\beta \\
\mathbf{F}_z(x,y,z) &= \frac{1}{3}\int_0^z I_s(x,y,\gamma)d\gamma
\end{aligned} \tag{18}$$

so that,

$$\frac{\partial \mathbf{F}_x}{\partial x} + \frac{\partial \mathbf{F}_y}{\partial y} + \frac{\partial \mathbf{F}_z}{\partial z} = \nabla \cdot \mathbf{F} = I_s \tag{19}$$

where $\mathbf{F} = (\mathbf{F}_x, \mathbf{F}_y, \mathbf{F}_z)$.

The definition of $\mathbf{F}$ is done in such a way that the C^1 continuity requirement in the statement of the above theorem is met. Given these definitions, we have,

$$\begin{aligned}
\int\int\int_{V_{\vec{p}}} I_s(x,y,z)dV &= \int\int_{A_{\vec{p}}} \mathbf{F}\cdot dA \\
&= \int\int_{A_{\vec{p}}} \mathbf{F}\cdot(\mathbf{x}_u \times \mathbf{x}_v)dudv \\
&= \int\int_{A_{\vec{p}}} [\mathbf{F}_x(y_u z_v - z_u y_v) + \mathbf{F}_y(z_u x_v - x_u z_v) \\
&\quad + \mathbf{F}_z(x_u y_v - y_u x_v)]dudv
\end{aligned} \tag{20}$$

We can also see that:

$$\begin{aligned}
\int\int\int_{V_{\vec{p}}} I_s(x,y,z)dV &= 3\int\int_{A_{\vec{p}}} \mathbf{F}_x(y_u z_v - z_u y_v)dudv \\
&= 3\int\int_{A_{\vec{p}}} \mathbf{F}_y(z_u x_v - x_u z_v)dudv \\
&= 3\int\int_{A_{\vec{p}}} \mathbf{F}_z(x_u y_v - y_u x_v)dudv
\end{aligned} \tag{21}$$

Substituting Equation 20 into Equation 17 we finally get,

$$\begin{aligned}
\max_{\vec{p}}\ M(I_g, I_s, \vec{p}) &= \max_{\vec{p}}\,[M_{\text{prior}}(\vec{p}) \\
&\quad + K_1\int\int_{A_{\vec{p}}} I_g(x,y,z)dA \\
&\quad + K_2\int\int_{A_{\vec{p}}} \mathbf{F}\cdot dA]
\end{aligned} \tag{22}$$

The calculation of $\mathbf{F}$ is done only once at the start of the optimization process. These calculations merely involve summing up the values of the voxels in the image I_s. Further, the derivatives, which we need during the optimization process, are the values of the image I_s itself. Thus, the use of the additional region information hardly introduces any extra computational burden to the deformable surface finding process.

In the above, we have presented a 3D gradient based surface finding procedure that introduces a matching term that incorporates information that is derived from region based segmentation. Further, the use of Gauss's divergence theorem allows us to reduce the whole objective calculation to computing surface integrals only, rather than both surface and volume integrals.

6 Evaluation and Optimization

The objective function in Equation 22, can be evaluated by numerical integration. The gradient of the objective is necessary for optimization. The surface integrals require differentiation of the area element on the surface $\mathcal{A}$ given by Equation 15. The derivative of the objective is given by:

$$\begin{aligned}\frac{\partial M}{\partial p_x} &= \frac{\partial M_{prior}(\vec{p})}{\partial p_x} \\ &\quad + K_1 \int\int_{A_{\vec{p}}} \left[I_g(x,y,z) \frac{\partial}{\partial p_x} |\mathbf{x}_u \times \mathbf{x}_v| \right. \\ &\quad \left. + \frac{\partial I_g(x,y,z)}{\partial x} \frac{\partial x(\vec{p},u,v)}{\partial p_x} |\mathbf{x}_u \times \mathbf{x}_v| \right] du dv \\ &\quad + 3K_2 \int\int_{A_{\vec{p}}} I_s(x,y,z)(y_u z_v - z_u y_v) \frac{\partial x(\vec{p},u,v)}{\partial p_x} du dv \end{aligned} \tag{23}$$

and similarly for y and z. This expression can also be evaluated by numerical integration. Expressions such as $\frac{\partial I_g(x,y,z)}{\partial x}$ can be obtained using discrete derivative calculation. Other expressions like $\frac{\partial x(\vec{p},u,v)}{\partial p_x}$ and $\mathbf{x}_u$ and $\mathbf{x}_v$ can be obtained analytically from Equations 2 and 3. The derivatives of the prior terms can be determined by analytical differentiation, as in the 2D case [43, 33].

Optimization is achieved using the conjugate gradient method, which is a local gradient optimization method. For surface finding, even local maximization involves a lot of computation. Thus, to avoid even further computational burden, global optimization methods were not considered at the cost, however, of not being able to guarantee global convergence. Through the use of prior information, the method is likely to be initialized close to the actual location, thus making global optimization methods less of a necessity.

7 Results

Experiments were carried out both with synthetic and MR brain images to verify the performance of the above mentioned method. The experiments were run on a Sun Sparcstation 10 with an average execution time of 20 minutes. In order to evaluate the performance quantitatively, we need a method to calculate the error between two surfaces expressed parametrically. The error is defined as the average distance between each point on the estimated surface and the closest point on the true surface [33]. That is, the error between surfaces S and $\hat{S}$ is defined as:

$$e(S,\hat{S}) = \frac{\int_{(u,v)\in\hat{S}} \min_{(u',v')\in S} |S(u',v') - \hat{S}(u,v)| dA}{\int_{(u,v)\in\hat{S}} dA} \tag{24}$$

This can be computed discretely by first taking a distance transform of a binary volume representing the true surface [33]. The result is then correlated with the binary volume representing the estimated surface, which gives the minimum distance between the estimated and the true surface. The result is then normalized by the area of the estimated surface.

We first used a synthetic example to evaluate the algorithm developed. This is useful because for this case we have exact knowledge of the true surface boundary. Comparisons of the integrated method were done against the traditional gradient based surface finding approach.

Figure 3 shows a simple synthetic example of a closed surface with added Gaussian noise. The signal to noise ratio (SNR) is defined here as the ratio of gray-level contrast to the standard deviation of the added noise and for this image it was 1.6. The initial surface was roughly placed on the target object. The combined method performed distinctly better, noticeable especially at the bottom of the left and right slices. The surface finder diverges under these noise conditions when using gradient information alone, while it converges appropriately for the integrated method.

Figure 4 shows a comparison of the two methods under increasing noise conditions. The measured error is shown here as a function of the noise level imposed on the image for both methods. The vertical axis corresponds to the SNR of the noise used. The y-axis gives a measure of the distance between the estimated surface and the true one using Equation 24. The integrated method clearly performs better under high noise conditions. This test demonstrates the robustness of the integrated method to noise.

Performance was also evaluated with respect to initialization. Figure 5 shows the performance when the vertical shift was varied from the true position, keeping the initialization for the other parameters fixed. This shows at larger capture region for the integrated method. In other words, the integrated method succeeds in converging to the desired target object much further away than in the gradient-only case. Thus, the the region within which the initialization must be in order to converge is larger for the integrated algorithm.

In Figures 6 and 7, we use the algorithm on a three dimensional magnetic resonance human brain image. The target object is the right thalamus which is only subtly distinguished in terms of gray level from the surrounding structures and without strong gradients at the margins. Neither source of information is strong. The gradient is not sufficient by itself to cause a proper solution and the optimization does not converge. However, with the integration of region information, a good delineation is found due to the combination of features.

In Figures 8 and 9, we demonstrate the performance on a 3D MRI to determine the head of the right caudate nucleus. The target borders both the brighter white matter and the darker CSF. Here again, the gradient information is not enough by itself and the surface is not found. However, by using the integrated method with both region and gradient information, the proper boundary is found.

8 Conclusions

We have presented in this paper an integrated method for surface finding in MR brain images using both region and gradient information. As the examples show, the integrated approach is

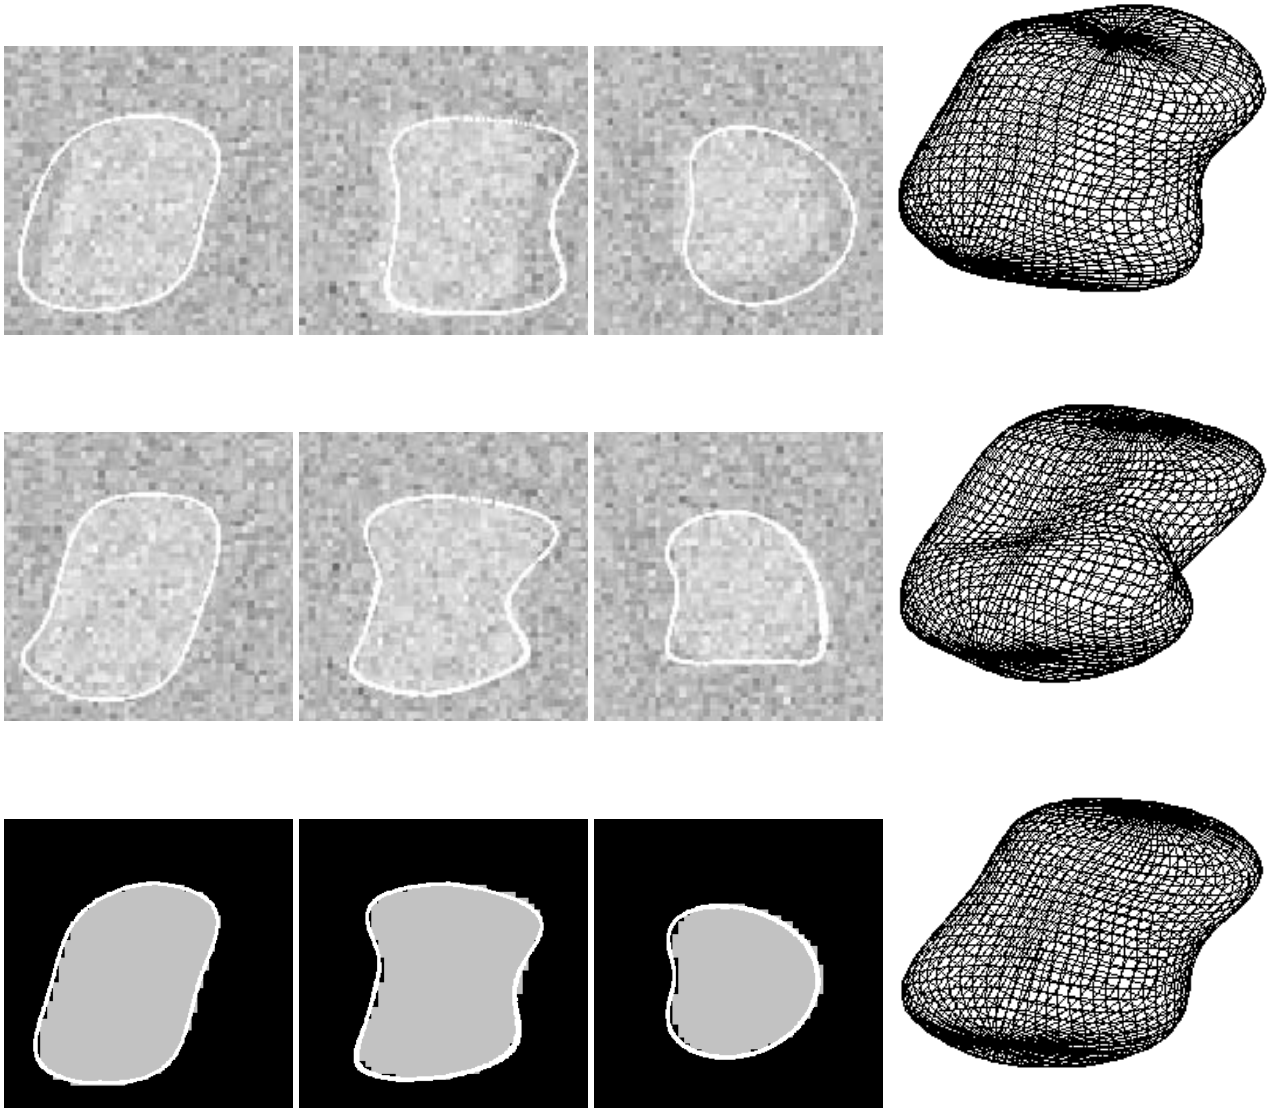

Figure 3: Surface finding for a noisy synthetic image with and without region information. (a) Top: Three perpendicular slices through the 3D noisy image ($48 \times 48 \times 48$) are shown with the initial surface and the wireframe. (b) Middle: The same slices through the same 3D image are shown with the surface obtained using only the gradient information, and the corresponding wireframe. (c) Bottom: The surface obtained using both gradient and region information from the noisy image, is shown here with the same slices through the noise-free image (for comparison), along with the corresponding wireframe. The solution found is visually indistinguishable from the true surface.

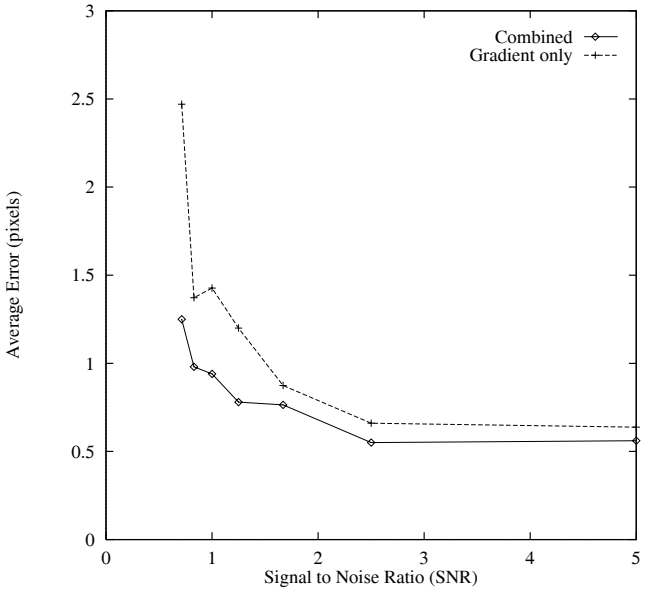

Figure 4: Noise performance of the surface finder with and without region information. The combined method has a lower average error for this example, especially at low SNR.

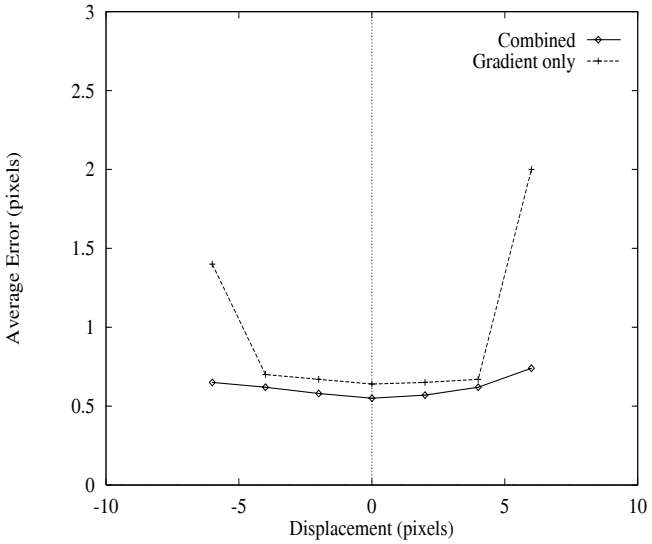

Figure 5: Performance of the surface finder with and without region information under different starting positions. This was varied by shifting the initialization vertically. Clearly, the combined method is superior.

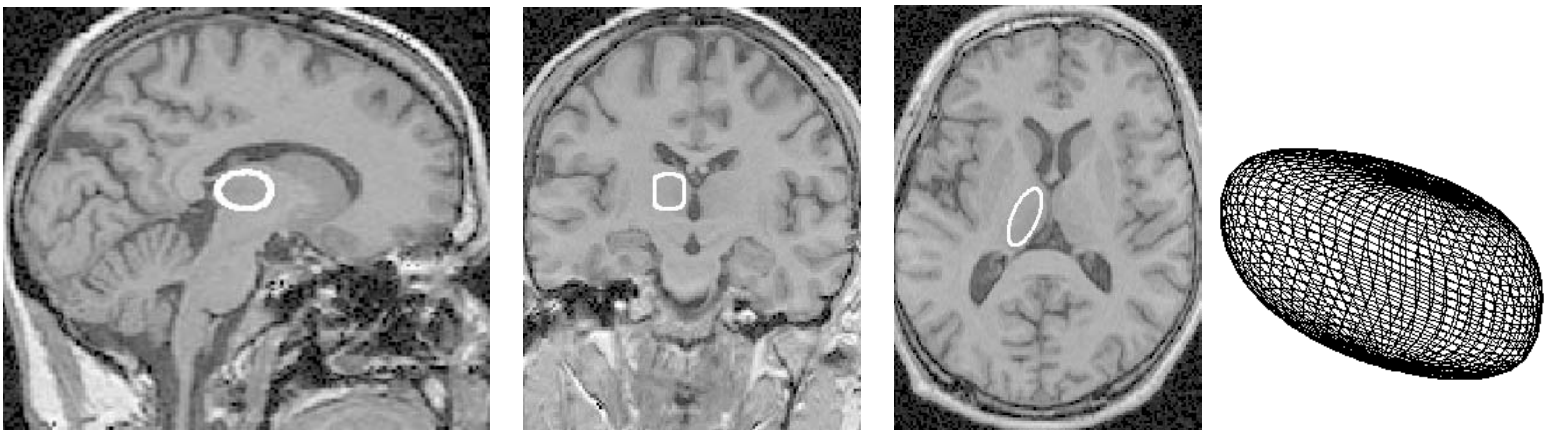

Figure 6: Initialization for surface finding for the right thalamus in an MR human brain gradient echo image, shown in Figure 7. Three perpendicular slices through the 3D image shown with the initial surface and the wireframe.

more robust to both increased amounts of noise as well as increasingly displaced initialization of the initial boundary. Thus, there is an improvement over the conventional gradient based boundary finding. It is important to note that this improvement in performance is achieved without significantly increasing the computational burden, due to the appropriate application of Gauss's divergence theorem. Application of this method on MR brain images results in noticeable improvement as shown.

References

[1] J. D. Bremner, P. Randall, E. Vermetten, L. Staib, R. A. Bronen, C. Mazure, S. Capelli, G. McCarthy, R. B. Innis, and D. S. Charney. MRI-based measurement of hippocampal volume in posttraumatic stress disorder related to childhood physical and sexual abuse – A preliminary report. *Biol. Psychiatry*, 41:23–32, 1997.

[2] R. T. Schultz, N. K. Cho, L. H. Staib, L. E. Kier, J. M. Fletcher, S. E. Shaywitz, D. P. Shankweiler, J. C. Gore, J. S. Duncan, and B. A. Shaywitz. Brain morphology in normal and dyslexic children: The influence of sex and age. *Annals of Neurology*, 35(6):732–742, June 1994.

[3] R. T. Schultz and L. H. Staib. Corpus callosum morphology in twins: Evidence for heritability and sex differences in structure-function relationships. *J. Inter. Neuropsychological Soc.*, 1:179, 1995.

[4] L. S. Allen, M. F. Richey, et al. Sex differences in the corpus callosum of the living human being. *J. Neuroscience*, 11:933–942, 1991.

[5] N. C. Andreasen, G. Harris, T. Cizadlo, S. Arndt, and D. S. O'Leary. Techniques for measuring sulcal/gyral patterns in the brain as visualized through magnetic resonance scanning: BRAINPLOT. *Proc. Natl. Acad. Sci. USA*, 90:93–97, January 1994.

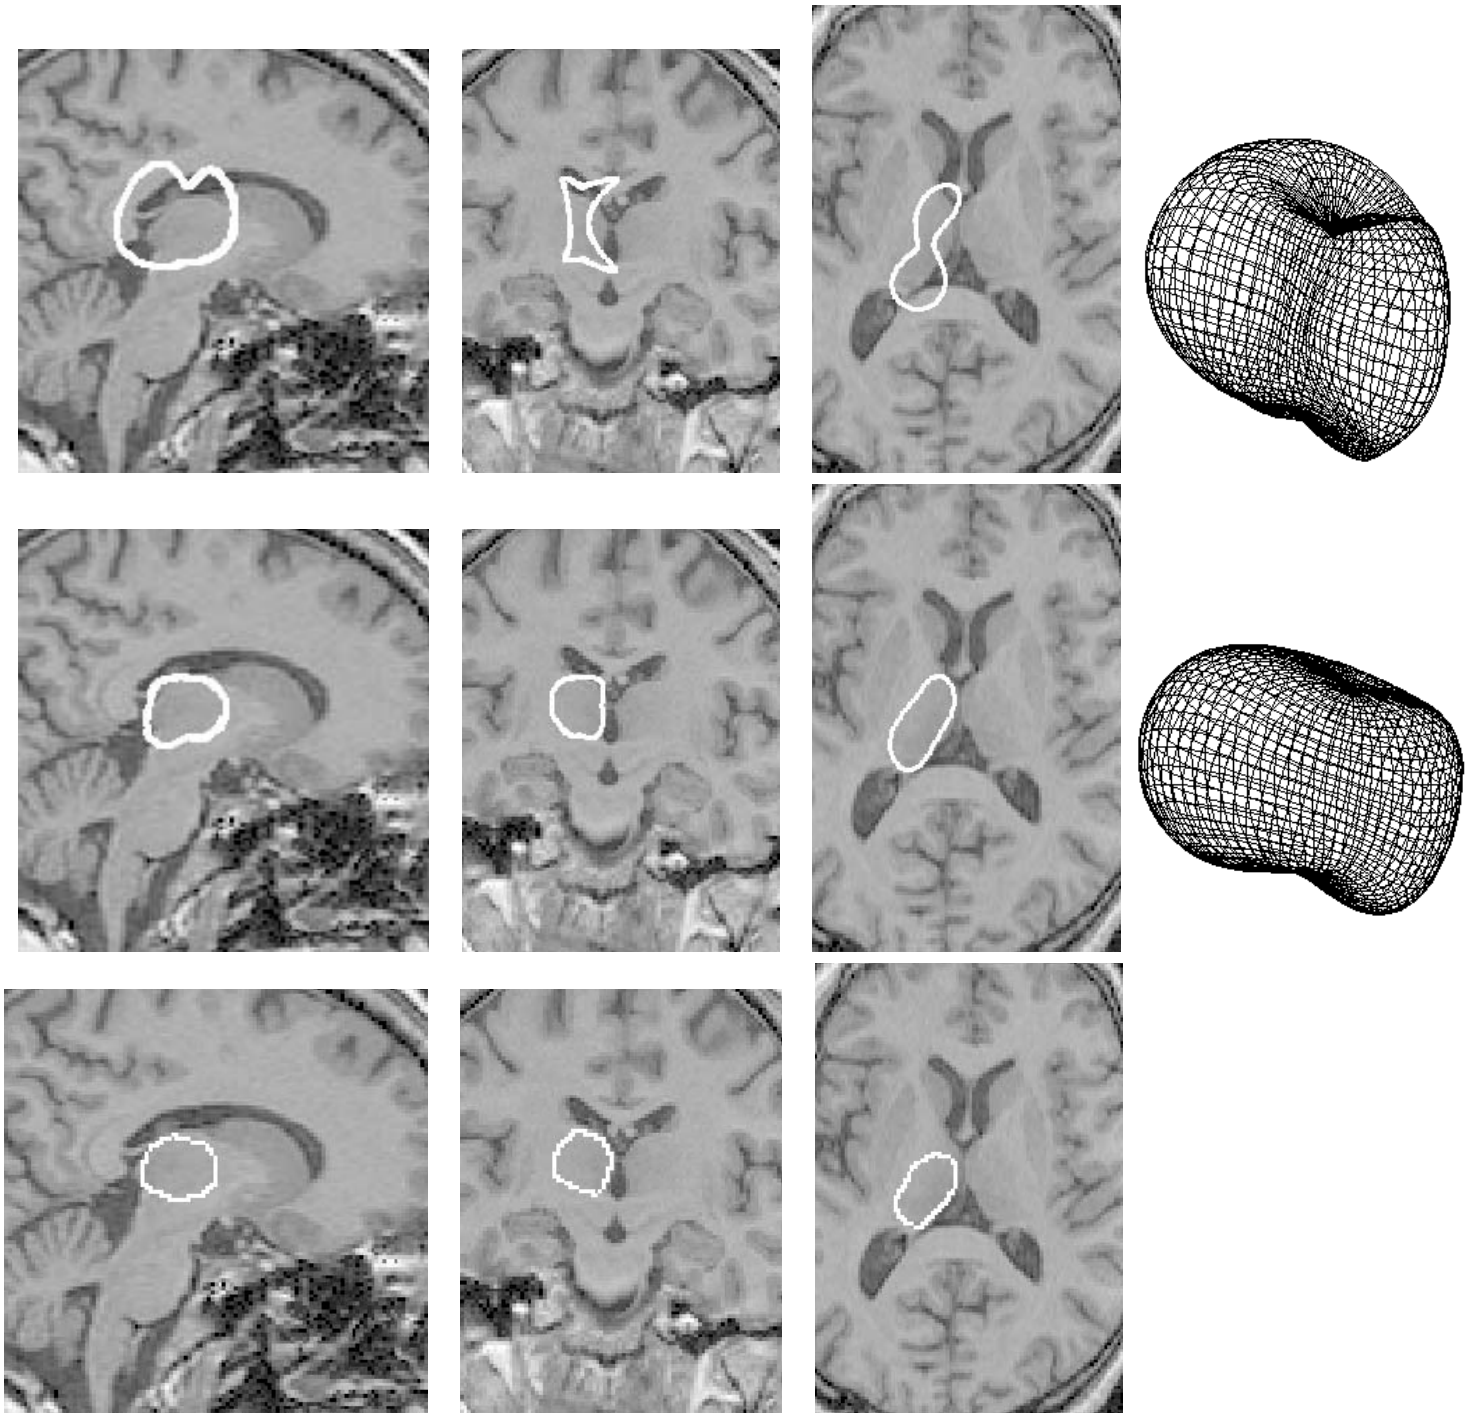

Figure 7: Results of surface finding for the right thalamus in an MR human brain gradient echo image with and without region information. (a) Top: Three perpendicular slices through the 3D image (1.2mm^3 voxels) are shown with the poor surface obtained using only the gradient information and the corresponding wireframe. (b) Middle: The same slices are shown with the thalamus surface obtained using both the gradient and the region information and the wireframe. (c) Bottom: Manual delineation of the same structure shows agreement with the results of the combined method.

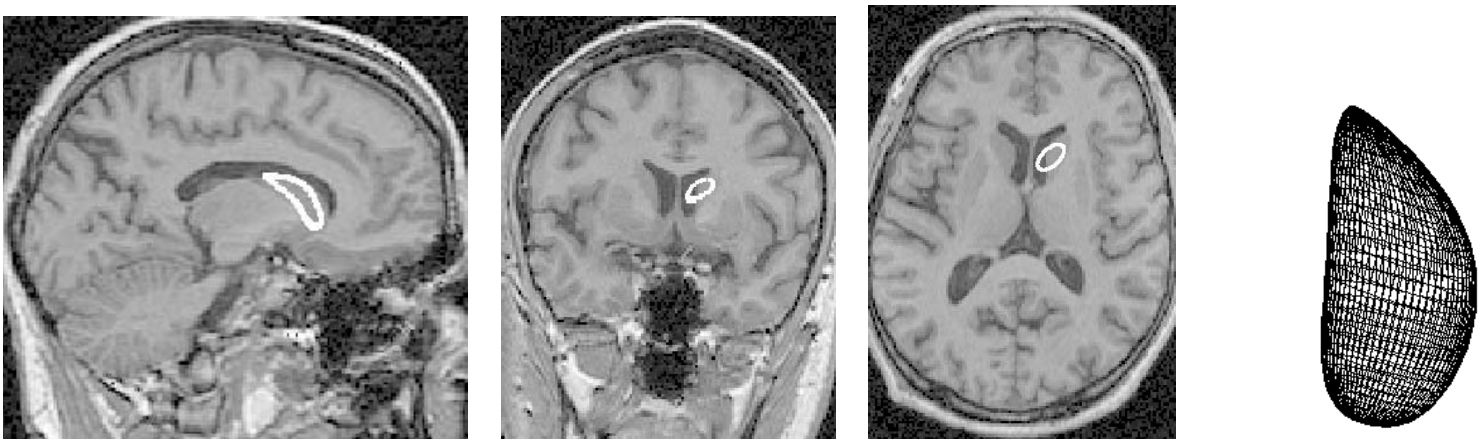

Figure 8: Initialization for surface finding of the head of the left caudate nucleus in an MR human brain gradient echo image, shown in Figure 9. Three perpendicular slices through the 3D image ($1.2mm^3$ voxels) are shown with the initial surface and the wireframe.

[6] T. Taxt, P. J. Flynn, and A. K. Jain. Segmentation of document images. *IEEE Trans. Pattern Anal. Machine Intell.*, 11:1322–1329, 1989.

[7] M. D. Levine and A. M. Nazif. Dynamic measurement of computer generated image segmentation. *IEEE Trans. System, Man and Cybernetics*, 7:155–164, 1985.

[8] J. Kittler and J. Illingworth. On threshold selection using clustering criteria. *IEEE Trans. System, Man and Cybernetics*, 15:652–655, 1985.

[9] R. Kohler. A segmentation based on thresholding. *Comp. Vision Graphics Image Proc.*, 15:319–338, 1981.

[10] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Machine Intell.*, 6(6):721–741, November 1984.

[11] D. Geiger and A. Yuille. A common framework for image segmentation. *Int. J. Computer Vision*, 6:227–243, 1991.

[12] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intell.*, 12(7):629–639, July 1990.

[13] B. ter Haar Romeny. *Geometry Driven Diffusion in Computer Vision*. Kluwer, 1994.

[14] P. J. Burt, T. H. Hong, and A. Rosenfeld. Segmentation and estimation of region properties through co-operative hierarchical computation. *IEEE Trans. System, Man and Cybernetics*, 11:802–809, 1981.

[15] D. Mumford and J. Shah. Boundary detection by minimizing functionals. *Proc. Comp. Vision Pattern Recog.*, page 22, 1985.

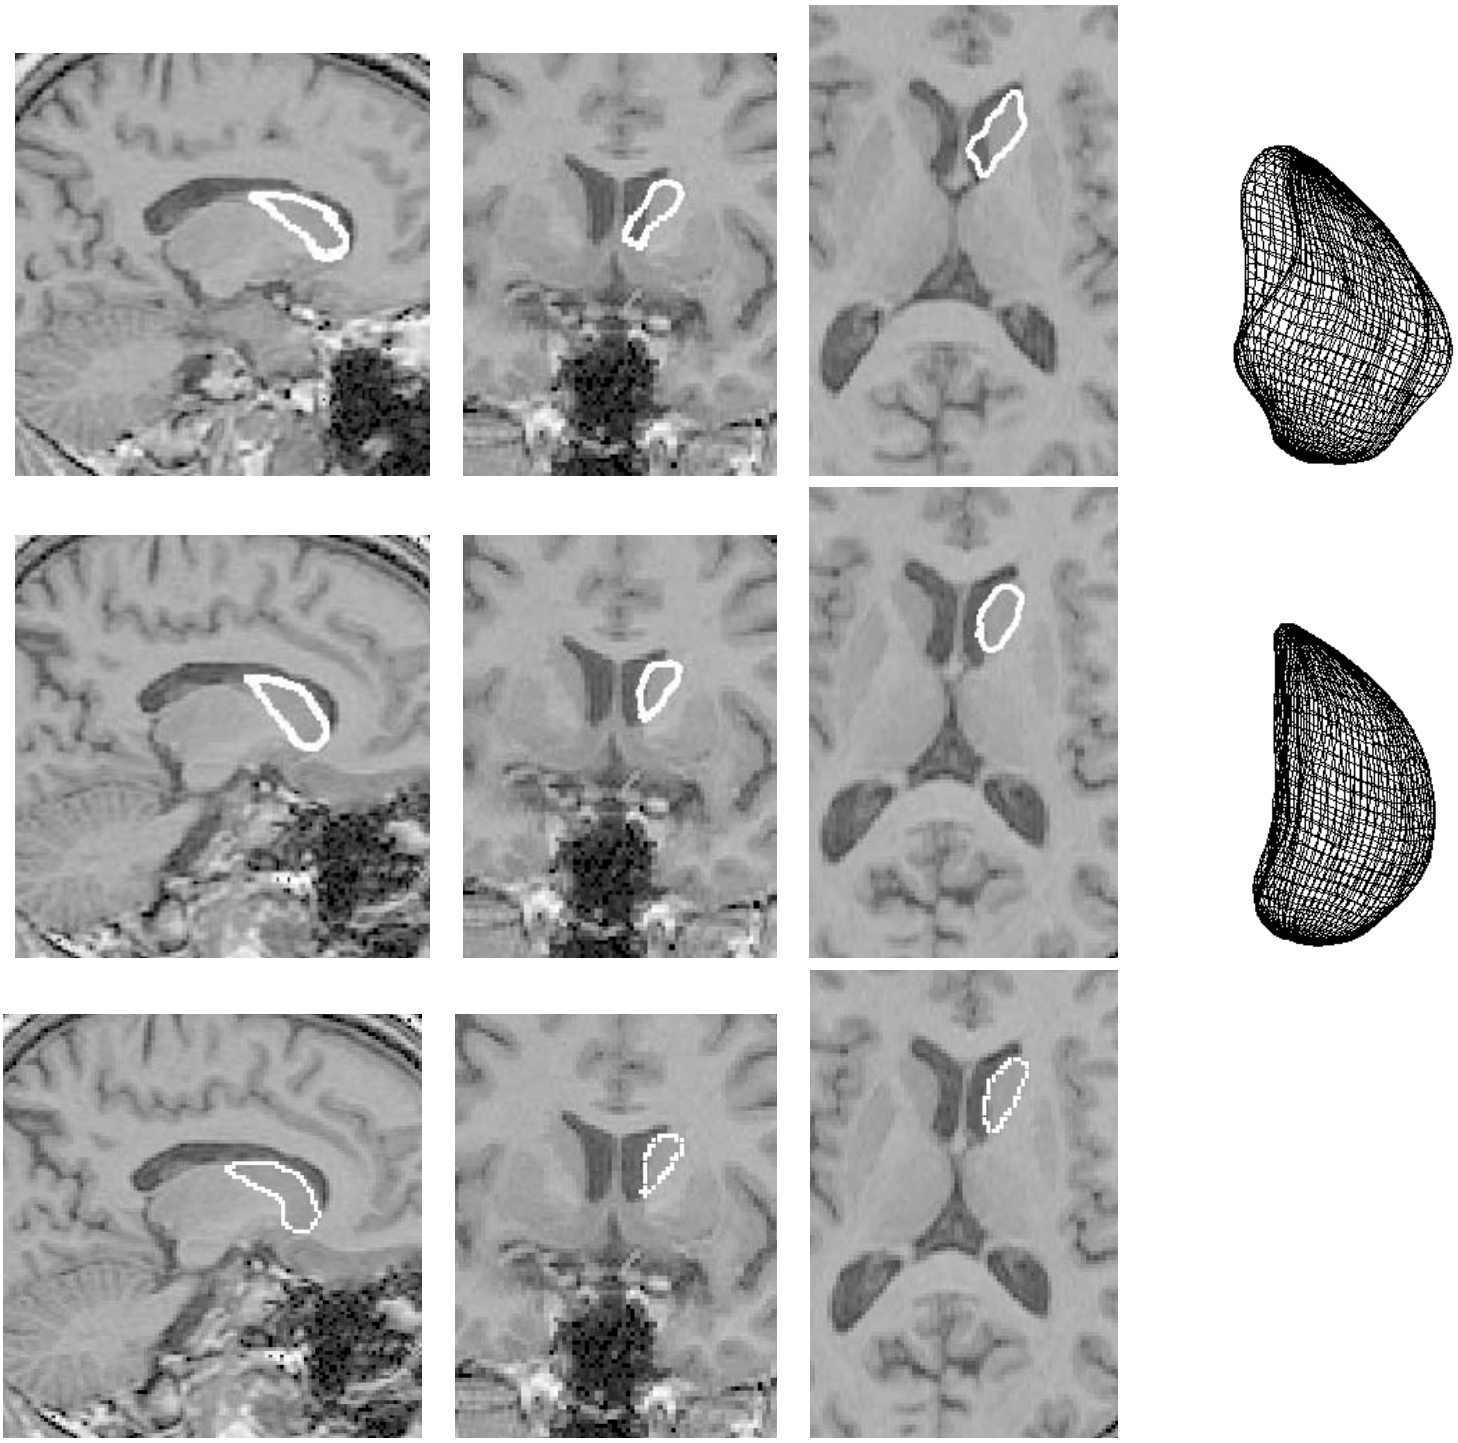

Figure 9: Results of surface finding for the head of the left caudate nucleus in an MR human brain gradient echo image with and without region information. (a) Top: Three perpendicular slices through the 3D image ($1.2mm^3$ voxels) are shown with the poor surface obtained using only the gradient information and the corresponding wireframe. (b) Middle: The final surface of the head of the caudate nucleus obtained using both the gradient and the region information is shown. While smoother, due to the surface parametrization, it substantially agrees with the manual delineation. (c) Bottom: Manual delineation of the same structure.

[16] Y. G. Leclerc. Constructing simple stable descriptions for image partitioning. *Int. J. Computer Vision*, 3:73–102, 1989.

[17] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, Cambridge, MA, 1987.

[18] L. Alvarez, P. L. Lions, and J. M. Morel. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal of Numerical Analysis*, 29:845–866, 1992.

[19] L. Axel, J. Constantini, and J. Listerud. Intensity correction in surface-coil MR imaging. *AJR*, 148:418–420, 1987.

[20] K. O. Lim and A. Pfefferbaum. Segmentation of MR brain images into cerebrospinal fluid spaces, white and gray matter. *J. Comp. Assisted Tomogr.*, 13(4):588–593, Jul./Aug. 1989.

[21] B. M. Dawant, A. P. Zijdenbos, and R. A. Margolin. Correction of intensity variations in MR images for computer-aided tissue classification. *IEEE Trans. Med. Imaging*, 12(4):770–781, December 1993.

[22] C. R. Meyer, P. H. Bland, and J. Pipe. Retrospective correction of intensity inhomogeneities in MRI. *IEEE Trans. Med. Imaging*, 14(1):36–41, March 1995.

[23] W. M. Wells, W. Grimson, R. Kikinis, and F. A. Jolesz. Statistical intensity correction and segmentation of MRI data. In R. A. Robb, editor, *Visualization Biomed. Comp. 1994, Proc. SPIE 2359*, pages 148–159, 1994.

[24] H. E. Cline, W. E. Lorensen, R. Kikinis, and F. Jolesz. Three-dimensional segmentation of MR images of the head using probability and connectivity. *J. Comp. Assisted Tomogr.*, 14(6):1037–1045, Nov./Dec. 1990.

[25] G. Gerig, J. Martin, R. Kikinis, O. Kubler, M. Shenton, and F. A. Jolesz. Automating segmentation of dual-echo MR head data. In A. Colchester and D. Hawkes, editors, *Information Proc. Med. Imaging*, pages 175–185. Springer-Verlag, Berlin, 1991.

[26] S. P. Raya. Low-level segmentation of 3-D magnetic resonance brain images - A rule-based system. *IEEE Trans. Med. Imaging*, 9(3):327–337, September 1990.

[27] M. E. Brummer, R. M. Mersereau, R. L. Eisner, and R. Lewine. Automatic detection of brain contours in MRI data sets. In A. Colchester and D. Hawkes, editors, *Information Proc. Med. Imaging*, pages 188–204. Springer-Verlag, Berlin, 1991.

[28] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Machine Intell.*, 8(6):679–698, November 1986.

[29] D. Marr and E. Hildreth. Theory of edge detection. *Proc. Roy. Soc. London B*, 207:187–217, 1980.

[30] A. Rosenfeld and A. Kak. *Digital Picture Processing, 2nd ed.* Academic Press, New York, 2nd edition, 1982.

[31] D. H. Ballard and C. M. Brown. *Computer Vision*. Prentice-Hall, Englewood Cliffs, 1982.

[32] L. H. Staib. *Parametrically Deformable Contour Models for Image Analysis*. PhD thesis, Yale University, New Haven, CT, 1990.

[33] L. H. Staib and J. S. Duncan. Model-based deformable surface finding for medical images. *IEEE Trans. Med. Imaging*, 15(5):720–731, 1996.

[34] F. Solina and R. Bajcsy. Recovery of parametric models from range images: The case for superquadrics with global deformations. *IEEE Trans. Pattern Anal. Machine Intell.*, 12(2):131–147, February 1990.

[35] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis. In *Proc. Workshop Math. Meth. Biomed. Image Anal.*, pages 171–180, June 1996.

[36] M. Bomans, K. Hohne, U. Tiede, and M. Riemer. 3-D segmentation of MR images of the head for 3-D display. *IEEE Trans. Med. Imaging*, 9(2):177–183, June 1990.

[37] S. V. Raman, S. Sarkar, and K. L. Boyer. Tissue boundary refinement in magnetic resonance images using contour-based scale space matching. *IEEE Trans. Med. Imaging*, 10(2):109–121, June 1991.

[38] D. N. Kennedy, P. A. Filipek, and V. S. Caviness. Anatomic segmentation and volumetric calculations in nuclear magnetic resonance imaging. *IEEE Trans. Med. Imaging*, 8(1):1–7, March 1989.

[39] P. A. Filipek, D. N. Kennedy, V. S. Caviness, et al. Magnetic resonance imaging-based brain morphometry: Development and application to normal subjects. *Ann. Neurol.*, 25(1):61–67, January 1989.

[40] K. H. Hohne and W. A. Hanson. Interactive 3D segmentation of MRI and CT volumes using morphological operations. *J. Comp. Assisted Tomogr.*, 16(2):285–294, 1992.

[41] H. Lee and D. R. Cok. Detecting boundaries in a vector field. *IEEE Trans. Signal Process.*, 39(5):1181–1194, May 1991.

[42] A. Cumani. Edge detection in multispectral images. *CVGIP: Graphical Models Image Process.*, 53(1):40–51, January 1991.

[43] L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(11):1061–1075, November 1992.

[44] D. Collins, T. Peters, W. Dai, and A. Evans. Model based segmentation of individual brain structures from MRI data. In R. A. Robb, editor, *Visualization Biomed. Comp. 1992, Proc. SPIE 1808*, pages 10–23, 1992.

[45] T. Cootes, A. Hill, C. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. In H. H. Barrett and A. F. Gmitro, editors, *Information Proc. Med. Imaging*, pages 33–47. LNCS 687, Springer-Verlag, Berlin, 1993.

[46] T. Pavlidis and Y. Liow. Integrating region growing and edge detection. *IEEE Trans. Pattern Anal. Machine Intell.*, 12:225–233, 1990.

[47] C. C. Chu and J. K. Agarwal. The integration of image segmentation maps using region and edge information. *IEEE Trans. Pattern Anal. Machine Intell.*, 15:1241–1252, 1993.

[48] J. F. Haddon and J. F. Boyce. Image segmentation by unifying region and boundary information. *IEEE Trans. Pattern Anal. Machine Intell.*, 12:929–948, 1990.

[49] H. Tek and B. B. Kimia. Image segmentation by reaction diffusion bubbles. *Proc. Fifth Int. Conf. Comp. Vision*, pages 156–162, 1995.

[50] S. C. Zhu, T. S. Lee, and A. L. Yuille. Region competition: Unifying snakes, region growing and Bayes/MDL for multi-band image segmentation. *Proc. Int. Conf. Comp. Vision*, pages 416–423, 1995.

[51] A. Chakraborty, L. H. Staib, and J. S. Duncan. Deformable boundary finding in medical images by integrating gradient and region information. *IEEE Trans. Med. Imaging*, 15(6):859–870, 1996.

[52] L. H. Staib and J. S. Duncan. Deformable Fourier models for surface finding in 3D images. In R. A. Robb, editor, *Visualization Biomed. Comp. 1992, Proc. SPIE 1808*, pages 90–104, 1992.

[53] K. Rao and R. Nevatia. Computing volume descriptions from sparse 3-D data. *Int. J. Computer Vision*, 2(1):33–50, 1988.

[54] S. Kumar, S. Han, D. Goldgof, and K. Bowyer. On recovering hyperquadrics from range data. *IEEE Trans. Pattern Anal. Machine Intell.*, 17:1079–1083, 1995.

[55] I. Cohen, L. Cohen, and N. Ayache. Using deformable surfaces to segment 3-D images and infer differential structures. *Comp. Vision Graphics Image Proc.*, 56(2):242–263, 1992.

[56] L. D. Cohen and I. Cohen. Finite element methods for active contour models and balloons for 2D and 3D images. *IEEE Trans. Pattern Anal. Machine Intell.*, 15(11):1131–1147, November 1993.

[57] L. D. Cohen, E. Bardinet, and N. Ayache. Surface reconstruction using active contour models. *Proc. SPIE Conference on Geometric Methods in Computer Vision*, 1993.

[58] B. S. Manjunath T. Simchony and R. Chellappa. Stochastic and deterministic networks for texture segmentation. *IEEE Trans. Acoustics Speech Signal Proc.*, 38:1039–1049, 1990.

[59] B. S. Manjunath and R. Chellappa. Unsupervised texture segmentation using markov random field models. *IEEE Trans. Pattern Anal. Machine Intell.*, 13:478–482, 1991.

[60] J. Besag. On the statistical analysis of dirty pictures. *J. Royal Stat. Soc.*, 48:259–302, 1986.

[61] P. Baxandall and H. Liebeck. *Vector Calculus*. Oxford University Press, 1986.