

Segmentation of brain tissue from magnetic resonance images

Tina Kapur^{1*}, W. Eric L. Grimson¹, William M. Wells III^{1,2} and Ron Kikinis²

¹Massachusetts Institute of Technology, Artificial Intelligence Laboratory, NE43-750, 545 Technology Square, Cambridge, MA 02139, USA

²Department of Radiology, Harvard Medical School, and Brigham and Women's Hospital, 75 Francis Street, Boston, MA 02115, USA

Abstract

Segmentation of medical imagery is a challenging problem due to the complexity of the images, as well as to the absence of models of the anatomy that fully capture the possible deformations in each structure. The brain is a particularly complex structure, and its segmentation is an important step for many problems, including studies in temporal change detection of morphology, and 3-D visualizations for surgical planning. We present a method for segmentation of brain tissue from magnetic resonance images that is a combination of three existing techniques from the computer vision literature: expectation/maximization segmentation, binary mathematical morphology, and active contour models. Each of these techniques has been customized for the problem of brain tissue segmentation such that the resultant method is more robust than its components. Finally, we present the results of a parallel implementation of this method on IBM's supercomputer Power Visualization System for a database of 20 brain scans each with $256 \times 256 \times 124$ voxels and validate those results against segmentations generated by neuroanatomy experts.

Keywords: deformable models, mathematical morphology, statistical classification, validation

Received December 29, 1995; revised April 2, 1996; accepted April 29, 1996

1. INTRODUCTION

Many current problems in image-guided surgery, therapy evaluation and diagnostic tools strongly benefit from accurate 3-D models of anatomical structures. This implies that automated or semi-automated segmentation methods are of considerable importance to the effective use of medical imagery in clinical and surgical settings. Typically, segmentation involves the isolation of anatomical structures from images obtained using modalities such as computed tomography (CT), X-ray, magnetic resonance imaging (MRI), positron emission tomography (PET), single photon emission computed tomography (SPECT) or ultrasound, with a primary goal of providing accurate representations of key anatomical structures, to be used for: quantitative studies correlating volumes of anatomical structures with pathological or normal development (Shenton *et al.*, 1992; Huppi *et al.*, 1995b; Morocz *et al.*, 1995); or for 3-D visualization of structures for pre- and intra-

operative surgical planning (Pelizzari *et al.*, 1989; Szeliski and Lavalée, 1993; Grimson *et al.*, 1994; Ayache, 1995).

In this work we address the segmentation problem in the context of isolating the brain in MRIs. Figure 1 shows an example segmentation generated by our system from MRI data.

Many issues inherent to medical imagery make segmentation a difficult task. The objects to be segmented from medical imagery are true (rather than approximate) anatomical structures, which are often non-rigid and complex in shape, and exhibit considerable variability from person to person. Moreover, there are no explicit shape models yet available that fully capture the deformations in anatomy. Magnetic resonance images are further complicated due to the limitations in the imaging equipment that lead to a non-linear gain artifact in the images. In addition, the signal is degraded by motion artifacts due to voluntary or involuntary movement of the patient during the scanning process.

In section 2 we present a brief summary of some techniques that have been used for segmentation of medical imagery. In section 3 we present our method for segmentation of brain

*Corresponding author
(e-mail: tkapur@ai.mit.edu)



Figure 1. Segmented brain tissue from MRI using our system.

tissue that combines the strengths of some of the existing technology with some new ideas and that is more robust than its individual components. We also briefly describe the parallel hardware [IBM's supercomputer, the Power Visualization System (PVS)] on which we have implemented our system. In section 4 we discuss issues regarding validation of medical image segmentation, and present a comparison of our results on a database of 20 brains, each of size $256 \times 256 \times 124$ voxels, against segmentations generated by experts. We conclude with a discussion of potential generalizations of our methodology to segment other structures from medical imagery.

2. RELATED WORK

The goal of the present work is to segment MR scans of brain tissue into key anatomical structures. Although there are many rapid scanning techniques in use for MR scanning, here we use gradient-echo images (Westbook and Kaut, 1993). Several previous methods are at least in part applicable to this problem.

Deformable models have been a popular technique for segmentation of medical data, in part because of their ability to encode approximate shape constraints. Systems that use such models typically represent anatomical structures using stacks of deformable contours in 2-D or using actual 3-D deformable surfaces. Example systems include: Cohen and Cohen's 2-D and 3-D Balloons (Cohen and Cohen, 1993), Staib and Dun-

can's Fourier snakes (Staib and Duncan, 1992; Chakraborty *et al.*, 1994), Grzeszczuk and Levin's stochastic deformation method (Grzeszczuk and Levin, 1994), Gindi *et al.*'s 'atlas based snakes' (Chang *et al.*, 1994) and Goble *et al.*'s 'active surfaces' (Gobel *et al.*, 1994).

In our experience, snake-like models serve best in an interactive setting because of the need for manual initialization of the starting positions and the adjusting of parameters for acceptable performance. Fourier snakes obviate the need for manual initialization by constructing shape models from sample data and are attractive for smooth surfaces like that of the corpus callosum, but would require a prohibitive amount of computation if applied to surfaces as complex as that of the brain. Atlas-based snakes also remain to be demonstrated on a surface other than that of the smooth putamen. It is also worthwhile to note that the efficacy of snakes as an interactive tool depends greatly on the presence of a powerful graphical user interface that allows the user easily, naturally, and efficiently to control the different forces acting on different parts of the snake.

Other systems have used binary morphological operators such as erosion and dilation to incorporate topological information into segmentation algorithms (Gerig *et al.*, 1991; Brummer *et al.*, 1993; Sandor and Leahy, 1994). Though binary morphology provides a simple and efficient way for incorporating distance and neighborhood information in segmentation, it requires a prior binarization of the image into object and background regions, which is not necessarily a trivial thresholding.

In addition, statistical methods, such as maximum likelihood, Bayesian decision theory, and principal component analysis have been used for intensity-based as well as location based segmentation of medical imagery (Collins *et al.*, 1992; Cootes *et al.*, 1994; Wells *et al.*, 1994).

Our approach is to combine aspects of each of these three classes of techniques into a hybrid system that is an improvement on each of the individual methods.

3. OUR METHOD

In our method for segmentation of brain tissue from magnetic resonance images, we combine methods that individually exploit grey level, topological and spatial information in the images. The specific techniques we use are: expectation/maximization (EM) segmentation for an intensity-based correction and classification of the data, binary morphology and connectivity for incorporation of relative topological information, and balloon-based deformable contours for addition of spatial information to the segmentation process.

We begin the discussion with a description of the input data to our algorithm, followed by a description of the model




Figure 2. An annotated gradient echo MR slice (air, csf and cranium are dark in these images).

of the brain tissue that is implicit in our algorithm, and then present details of each of the three steps in our algorithm. We conclude the section with a brief discussion of a parallel implementation of the algorithm that we are currently using at the Surgical Planning Laboratory of Brigham and Women's hospital to routinely segment brains for clinical and research purposes.

3.1. Description of input data

We use MR scans of the head as input to our algorithm. We are currently working with gradient echo images acquired using a General Electric Signa 1.5 Tesla clinical MR imager. The voxel size is approximately $1 \times 1 \times 1.5$ mm and there are $256 \times 256 \times 124$ voxels per data set [see Shenton *et al.* (1992) for details].

Tissue classes visible in such MRI scans include white and grey matter, cerebrospinal fluid (csf), meninges (the protective membranes surrounding the brain), skull, muscle, fat, skin or air (see Figure 2). Pathology introduces the additional classes of edema, tumor, hemorrhage or other abnormality.

3.2. Model for the brain

Our first task was to construct a model for the brain that would guide our segmentation process. This model is implicit in our algorithms and was constructed based on expert opinion from the Surgical Planning Laboratory, Brigham and Women's Hospital, and Harvard Medical School.

We represent the brain as the largest region consisting of white and grey matter, located approximately in the center of the head, and surrounded by csf and meninges. If our segmenter could clearly identify white and grey matter, without falsely including other tissue, and if it could do so while clearly isolating this material for surrounding tissue, that would be sufficient. Unfortunately, several other tissue types can cause the white and grey matter to be connected to other structures. For example, the meninges are surrounded by the cranium. Blood vessels and nerves connect the cranium to the brain tissue. In particular, the connectors are, in order of increasing thickness: bridging veins from the cerebral cortex to dura, and from dura to the skull, the second cranial nerve, or optic nerve, the vertebral arteries around foramen magnum, and the external carotid artery in the region of the temporalis muscle. Thus, we need some way of removing these connecting structures so that we can isolate the white and grey matter. This is compounded by the fact that there is some natural overlap in the intensity distributions of brain versus non-brain structures. Additional overlap in the intensity distributions is typically introduced due to limitations of the imaging process, and noise may be introduced due to movement of the patient during the acquisition of the scan.

This model of the brain strongly suggests the use of intensity distributions as well as absolute and relative spatial arrangement of the various structures in the head to aid the segmentation process. In fact, several groups have approached

this task with techniques that indicate a similar underlying model for the brain tissue (Brummer *et al.*, 1993; Sandor and Leahy, 1994).

3.3. Details of method

Using the model described above, we divide the segmentation of the brain into three steps. We will illustrate the details of each of these steps using two example MRI scans. Although the algorithm operates in 3-D, 2-D cross-sections are shown for convenience. In the first example, the first two steps of our method suffice to generate a segmentation, while the second example additionally requires the third step.

3.3.1. Gain correction

Background on EM segmentation. Traditional intensity based segmentation relies on elements of the same tissue type having MR intensities that are clustered around a mean characteristic value, and on each cluster being well separated from other tissue clusters. Ideally, this means that segmentation is a straightforward matter of choosing the right thresholds to separate each cluster, but in practice this is not as easy as the ideal case would suggest (Brummer *et al.*, 1993). Most MR scanners have inhomogeneities in the imaging equipment, which give rise to a smoothly varying, non-linear gain field. While the human visual system easily compensates for this field, the gain can perturb the intensity distributions, causing them to overlap significantly and thus lead to substantial misclassification in traditional intensity-based classification methods (e.g. Cline *et al.*, 1990). Figure 3 shows histograms of the intensity distribution in the head in an MR scan, and of the intensity distribution in segmented brain from the same scan. A visual comparison of the brain histogram with the complete head histogram shows that the intensities for the brain neither form distinct peaks in the histogram, nor are clearly separated from the rest of the structures in the head in any other way, which are the two criteria on which intensity-based segmentation relies.

One way to deal with this problem is to use a statistical classification scheme, in which one solves both for the assignment of segmentation labels to voxels and for the gain field artifact at each voxel. A simple version of this approach for dual-echo MR data (Gerig *et al.*, 1991) has previously appeared. A more general method was introduced by Wells *et al.* (1994), whose method uses knowledge of tissue properties and RF coil inhomogeneity to correct for the spatial distortion in the signal. The method, EM segmentation (or expectation/maximization segmentation), produces simultaneous classification of tissue classes and estimation of the gain field due to inhomogeneities in the RF coil in MRI. The following is a description of the details of the method presented in Wells *et al.* (1994, 1995,




Figure 3. Intensity distributions for the head and brain in MRI.

1996).

Intra- and inter-scan MRI intensity inhomogeneities can be modeled with a spatially varying factor called the *gain field*

that multiplies the intensity data. Log transformation of the data allows the artifact to be modeled as an additive bias field.

EM segmentation is based on the following observations.

- When the bias field is known, classification of voxels based on their signal intensity is straightforward. After the bias field is subtracted from the signal, conventional statistical intensity classifiers work well for the determination of tissue classes.
- If, on the other hand, the tissue class of each voxel is known, then it is feasible to estimate the bias field—by subtracting the appropriate tissue class mean from the observed intensity at each voxel. A low-pass filter may be used to reduce the noise level in the resulting estimate of the bias field.

Of course, individually, the above observations are not very useful. While each of the important unknowns (tissue class and intensity correction) can be determined from the other, at the outset neither one is known. EM segmentation works by alternately using partial information about one of the unknowns to improve the estimate of the other, until good estimates are obtained for both unknowns.

Statistical models are used for the following aspects of the domain.

- The intensity properties of each tissue are modeled by probability densities conditioned on tissue class and the local value of the bias field.
- The overall probability of each tissue in the volume is modeled by a spatially stationary prior probability model.
- A probability density is used to model the entire bias field. This model is used to enforce the constraint that the intensity inhomogeneities of MR scanners are spatially smooth.

The above modeling is used to construct an estimator of the unknown bias field. The bias field estimate is obtained by solving a non-linear optimization problem that depends on the scan data. When formulated in this way, the bias field estimation problem is a natural candidate for solution by the EM algorithm. The EM algorithm is often used in estimation problems that would be easy, if the values of some discrete ‘hidden’ variables were known. In this application the hidden variables are the tissue classes of the voxels.

The EM algorithm is iterative, and alternates between two steps, the ‘E step’, where the hidden variables are estimated (based on the most recent estimate of the ‘ordinary’ variables) and the ‘M step’, where the ordinary variables are estimated (based on the most recent estimate of the ‘hidden’ variables). In this application, the E step corresponds to estimating the

tissue probabilities at each voxel based on the most recent bias field estimate, while the M step corresponds to estimating the bias field, based on the most recent tissue probability estimates.

In practice, EM segmentation converges in ~ 7 –15 iterations. Convergence here means that the location and value of the maximum bias is within a predefined tolerance of values in the previous iteration. The segmentation method may be started on either of the E or M steps, and initial estimates will be needed for one of the unknowns, bias field or tissue probabilities. The algorithm is usually started on the E step, with the bias field initialized to zero at all voxels. Training data—sample intensities from each tissue type occurring—are used to construct the probability densities for different tissue classes which are used in the E step. This training process is done once for each different type of acquisition (such as gradient echo or T1–T2 pairs) and need not be repeated for each scan that is segmented.

This method has been used successfully to segment > 1000 scans in a multiple sclerosis study with training data collected from only one data set (Wells *et al.*, 1995). There is no additional manual intervention after the initial training data, and this method is one of the most aggressively tested intensity based classifiers.

EM segmentation in context of our method. Initially, we use EM segmentation to correct for the gain introduced in the data by the imaging process. We use a single channel, non-parametric, multi-class implementation of the segmenter that is described in Wells *et al.* (1994, 1995). Training points are used for white matter, grey matter, csf and skin, and therefore the resultant tissue classifications correspond to these four classes. The output of this stage is a set of classified voxels.

Figures 4 and 5 show the results of EM segmentation on our two examples. The top left image in each figure is the input given to the segmenter, and the rest of the images (top to bottom, left to right) show the tissue classifications generated by the EM segmenter in successive iterations.

Since this step classifies voxels purely on the basis of their signal intensities, and there is a natural overlap between intensity distributions of the brain and non-brain structures, the misclassifications (of surrounding tissue, skin and scalp as white or grey matter) that we see in the results were predictable. The next step aims to reduce some of this misclassification by using neighborhood and connectivity information.

3.3.2. Removal of thin connectors using morphological operations

One way to deal with the potential misclassification is to separate thin connections between similar structures, then use connectivity to isolate major structures with similar responses.




Figure 4. Top to bottom, left to right: input image and tissue classification generated by successive iterations of the EM segmenter (white matter is brightest, grey matter is medium grey, and csf and air are black). Bottom right is the final segmentation.




Figure 5. Top to bottom, left to right: input image and tissue classification generated by successive iterations of the EM segmenter (white matter is brightest, grey matter is medium grey, and csf and air are black). Bottom right is the final segmentation.



Figure 6. Examples of erosion and dilation on a binary brain image.

Background on mathematical morphology. Image morphology provides a way to incorporate neighborhood and distance information into algorithms (see Serra, 1982; Haralick *et al.*, 1987 for detailed treatment of morphological operators). The basic idea in mathematical morphology is to convolve an image with a given mask (known as the structuring element) and to binarize the result of the convolution using a given function. Choice of convolution mask and binarization function depend on the particular morphological operator being used.

Binary morphology has been used in several segmentation systems, and we provide here functional descriptions of morphological elements as applicable in our work.

- **Erosion:** an erosion operation on an image I containing labels 0 and 1, with a structuring element S , changes the value of pixel i in I from 1 to 0, if the result of convolving S with I , centered at i , is less than some predetermined value. We have set this value to be the area of S , which is basically the number of pixels that are 1 in the structuring element itself. The structuring element (also known as the *erosion kernel*) determines the details of how a particular erosion thins boundaries.
- **Dilation:** dual to erosion, a dilation operation on an image I containing labels 0 and 1, with a structuring element S , changes the value of pixel i in I from 0 to 1, if the result of convolving S with I , centered at i , is more than some predetermined value. We have set this value to be zero. The structuring element (also known as the *dilation kernel*) determines the details of how a particular dilation grows boundaries in an image.
- **Conditional dilation:** a conditional dilation is a dilation operation with the added condition that only pixels that are 1 in a second binary image, I_c , (the image on which the dilation is conditioned), will be changed to 1 by the dilation process. It is equivalent to masking the results

of the dilation by the image I_c .

- **Opening:** an opening operation consists of an erosion followed by a dilation with the same structuring element.
- **Closing:** a closing operation consists of a dilation followed by an erosion with the same structuring element.

As an example, Figure 6 shows (from left to right) a binarized MR cross-section, erosion of the MR image with a circular structuring element of radius 3, conditional dilation of the largest connected component in the eroded image with a circular structuring element of radius 4. Since the dilation is conditioned on the original image, no boundaries are expanded in this process.

Mathematical morphology in context of our method. This step uses morphological operations (in 3-D) to incorporate neighborhood information into the tissue-labeled image obtained from EM segmentation. The strategy is to use morphological operators to ‘shave off’ the misclassified nerve fibers and muscles connecting the brain tissue to the cranium, and then use connectivity to find the largest connected component of white and grey matter in the image. Similar methods have been used by others (e.g. Hohne and Hanson, 1992). Specifically, the sequence of operations performed is as follows.

- Perform an erosion operation on the input with a spherical (in real space which implies elliptical in image space due to the anisotropy of the voxels) structuring element with radius corresponding to the thickness of the connectors between brain and the cranium (determined empirically, and held constant over scans), so that it eliminates connections from the brain to any misclassified non-brain structure.
- Find the largest 3-D connected component with tissue labels corresponding to the brain.

- Dilate the brain component obtained in the previous step by a structuring element slightly larger in size than the one used in the erosion, conditioned on the brain labels in the input image. Since the dilation is conditioned on the original image, no boundaries are expanded in this process. This corresponds approximately to restoring the boundaries of the brain component that were distorted in the erosion step.

The result of this stage is an improved segmentation of the tissue types, which incorporates topological information into the results of the pure intensity classification. Figure 7 illustrates the results of this step on the first of our two example scans, and is representative of the case in which the EM segmentation step combined with the morphology step achieves an isolation of the brain.

Occasionally (in $\sim 10\%$ of the over 100 scans we have segmented), due to the variation in the size of the connecting elements from the brain tissue to the cranium, the empirically determined radius of the erosion kernel does not adequately model the width of the connectors between the brain and non-brain structures, and therefore the brain tissue is not isolated at the end of this step. This case is illustrated in the results of the morphological operations on our second example (Figure 8). Such scenarios are currently detected by manual inspection, and lead to the use of the third step of our algorithm which uses manually initialized deformable models to annihilate connections between the brain and spurious structures. More automatic methods for detecting these cases are currently under development.

3.3.3. Refinement using deformable contour models

Background on snakes and balloons. Deformable models are a popular component of medical image segmentation systems due to their ability to encode approximate shape constraints. For segmentation purposes, anatomical structures can be modeled using stacks of deformable contours in 2-D or using actual 3-D deformable surfaces. Also known as active contour models, these provide a method for minimizing an objective function to obtain a contour of interest, especially if an approximate location of the contour is available. We start with some background on two earlier deformable contour models: snakes (Witkin *et al.*, 1988) and balloons (Cohen, 1991) in order to familiarize the reader with the essential components of such models.

A deformable contour is a planar curve which has an initial position and an objective function associated with it. A special class of deformable contours called snakes was introduced by Witkin *et al.* (1988) in which the initial position is specified interactively by the user and the objective function is

referred to as the energy of the snake. By analogy to physical systems, the snake slithers to minimize its energy over time. This energy of the snake (E_{snake}) is expressed as a sum of two components: its internal energy ($E_{internal}$) and its external energy ($E_{external}$)

$$E_{snake} = E_{internal} + E_{external}. \quad (1)$$

The internal energy term imposes a piecewise smoothness constraint on the snake by preferring low first and second derivatives along the contour:

$$E_{internal} = \int_s (w_1(s) ||v'(s)||^2 + w_2(s) ||v''(s)||^2) ds, \quad (2)$$

where s is arc length, derivatives are with respect to s , and $v(s)$ stands for the ordered pair $(x(s), y(s))$, which denotes a point along the contour. The choice of w_1 and w_2 reflects the penalty associated with first and second derivatives along the contour respectively. w_1 is also known as the coefficient of elasticity, and w_2 as the coefficient of rigidity for the snake.

The external energy term in Equation (1) is responsible for attracting the snake to interesting features in the image. The exact expression for $E_{external}$ would depend on the characteristics of the features of interest. For example, if points of high brightness gradient magnitude in the image are interesting for a particular image, then the expression for the external energy can be:

$$E_{external} = -||\nabla I(v(s))||^2.$$

Finding a local minima for E_{snake} from Equation (1) corresponds to solving the following Euler–Lagrange equation for v :

$$-(w_1 v')' + (w_2 v'')'' + \nabla E_{external}(v) = 0.$$

with boundary conditions specifying if the snake is a closed contour, or the derivatives are discontinuous at the end points. This equation is then written in matrix form as

$$Av = F,$$

where

$$F(v) = -\nabla E_{external}. \quad (3)$$

Here A is a pentadiagonal banded matrix that represents the ‘stiffness’ properties of the snake, v is the position vector of the snake, and F is gradient of the external energy of the snake, or the external force acting on it.

The evolution equation:

$$\frac{dv}{dt} - Av = F,$$

is solved to obtain the v that is closest to the initial position. As $\frac{dv}{dt}$ tends to zero, we get a solution to the system $Av = F$.




Figure 7. Top to bottom, left to right: EM segmentation from Figure 5, binarized image, eroded image, largest connected component in eroded image, dilated connected component, conditionally dilated connected component.



Figure 8. Top to bottom, left to right: EM segmentation from Figure 4, binarized image, eroded image, largest connected component in eroded image, dilated connected component, conditionally dilated connected component. Note that the slice shown here is a cross-section of a full 3-D data which happens not to show the connectors between brain and non-brain structures.

Formulating this evolution problem using finite differences with time step τ , we obtain a system of the form (Cohen, 1991):

$$(I + \tau A)v^t = v^{t-1} + \tau F(v^{t-1}),$$

where v^t denotes the position vector of the snake at time t , and I is the identity matrix. The system is considered to have reached equilibrium when the difference between v^t and v^{t-1} is below some threshold.

The 2-D balloon model extends the snake energy to include a ‘balloon’ force, which can either be an inflation force, or a deflation force. The external force F in Equation (3) is changed to

$$F = k_1 \mathbf{n}(s) + k \frac{\nabla E_{\text{external}}}{\|\nabla E_{\text{external}}\|}, \quad (4)$$

where $\mathbf{n}(s)$ is a unit vector normal to the contour at point $v(s)$, and $|k_1|$ is the amplitude of this normal force. Changing the sign of k_1 makes this normal force inflating or deflating. Comparing Equation (4) with Equation (3) shows that the balloon model uses only the direction of the gradient of the external force, while snakes employ the magnitude of this force. This normalization of the external force, along with careful selection of the magnitude of k , constrains the motion of each point on the contour to at most one pixel per iteration. This avoids instabilities due to time discretization, and is discussed in detail in Cohen (1991).

Contrasting the Balloon and Snake models, we note that incorporation of the normal force in the balloon model allows the initial position of the contour to be further from the intended final position, while still enabling convergence. As well, in the balloon model, the initial position can lie either inside or outside the intended contour, while the snake model requires the initial position to surround the intended contour if regions of zero image force lie between the initial and intended contours, since snakes can only collapse to a point in the absence of image forces.

Snakes and balloons have been typically used with creative external forces to segment various anatomical structures. Using snakes to track 2-D contours in 3-D, coupling of the snakes to region information in the image, using simulated annealing to find the global minimum of the snake energy, and imposing shape priors are notable extensions to the original snakes.

Snakes and balloons in context of our method. The third step in our segmentation algorithm is the use of deformable contour models to refine the result of the brain tissue estimate obtained using morphological operations, since morphology is occasionally unable to remove all connectors from the brain to the cranium (as shown in the example in Figure 8). The

intuition behind this step is to incorporate substantial spatial information into the segmentation process via manual intervention.

For our deformable model, we started with the classical snakes (Witkin *et al.*, 1988), applied the refinements discussed in Cohen (1991) and added a radial force to the contours. The direction of this radial force, instead of being determined a priori [as is the case in the direction of the normal force in Cohen (1991)], is determined by the results of the previous steps. Specifically, voxels that are classified as brain in the previous steps exert a radial force directed towards the center of the contour, and the voxels classified as non-brain exert a force radially away from the center of the contour. From an implementation point of view, instead of using a fixed predetermined direction for the balloon force by selecting a sign for the factor k_1 in Equation (4) which is independent of the image characteristics, we define a signed balloon force direction vector, B , with one entry per voxel of the input data. The sign at the i th position of vector B indicates whether the voxel i exerts a force along the inward or outward normal to the evolving contour. This vector B is determined using the brain estimate obtained after applying the morphological operations as discussed in the previous section. If voxel i is classified as brain tissue at the end of the morphological routines, then $B[i]$ gets a positive sign, so that the voxel pushes the contour in the direction of the normal towards the boundary of the brain estimate, otherwise $B[i]$ gets a negative sign. The external force now becomes:

$$F = kB(s)\mathbf{n}(s),$$

where $B(s)$ is the direction of the balloon force exerted by the image at the point $v(s)$ of the contour, and $\mathbf{n}(s)$ is the normal to the local tangent vector to the contour at s . The unsigned constant k is used to determine the magnitude of the balloon force. If $B(s)$ is positive, motion is along $\mathbf{n}(s)$, which corresponds to an inflating force applied to the contour at s . Negative $B(s)$ results in a deflating force applied to the contour at $v(s)$.

This deformable contour model is employed in our system as follows. If the brain tissue has not been isolated by the end of the morphological operations, an expert manually initializes the true boundary of the brain tissue in a few carefully chosen slices, and sets the parameters for the ‘internal’ and ‘external’ forces of the deformable contour (i.e. the coefficients of elasticity, rigidity and the weight associated with the image data). The number and accuracy requirements of the initial contours, as well as the choice of the contour parameters, depend on the performance of the previous two steps on a particular data set. Once initialized, these contours are propagated slice-by-slice through the entire volume until the

Table 1. Sample running times for segmentation steps.

Algorithm	Time per iteration (s)	Number of iterations in average case
EM segmentation	30	10
Erosion/dilation	10 (kernel radius 3)	1
3-D connectivity	30	1

brain volume is isolated^a. This step often requires as little intervention as the user zeroing in on the spurious connecting element (this is done by inspection, but casual observation indicates that these occur in a few select places, which suggests possibilities for incorporating strong spatial priors into the segmentation process, and thereby reducing the need for manual intervention) and drawing the correct brain boundary through it in a couple of slices, adjusting the contour parameters so that the weight of the image forces in the neighborhood of the offending connector is very low, and allowing the contour to be propagated to the adjacent few slices (since a connector is usually only a few slices in thickness).

As an alternative to using deformable contour models as the final step, morphological operations with specialized kernels (larger size, different connectivity) were also explored as a means for eliminating the undesirable connecting elements. In that approach, manual intervention was required in the form of choosing the optimal kernel. Manual intervention appears necessary if acceptable segmentation is to be achieved in 100% of the cases. Deformable models were our editing tool of choice because they provide an elegant framework in which various degrees of manual intervention—by way of positioning the initial contour, and determining the weight on the image and smoothness forces—can be incorporated with ease.

We have presented a method for segmentation of brain tissue from MR images, which combines intensity, topological, and spatial information into the process. Next, we briefly discuss the implementation of this method on a supercomputer, IBM Power Visualization System.

3.4. Implementation

Implementation of the EM segmentation, morphological operations, and connectivity has been done on an IBM PVS. In this section we present a brief description of the architecture of the PVS (IBM, 1993), along with sample running times for each algorithm.

The PVS is a platform for scientific visualization, imag-

ing, and graphics applications. It is composed of a parallel processor server, a video controller, and a disk array. The server at the Surgical Planning Lab, Brigham and Women's Hospital contains 32 processors, 500 Mbyte of shared memory, HiPPI channels, all interconnected by a 1280 Mbyte/s backplane. Each processor consists of a 40 MHz Intel i860 microprocessor, 16 Mbyte of private memory, and a buffered interface to the high-speed backplane. The server attaches to an IBM RS6000, which serves as the support processor, and provides access to standard networks. The video controller takes images over a 100 Mbyte/s HiPPI channel from the server and displays them at a resolution of 1280×1024 . The disk array subsystem provides data storage for 20 Gbyte, and communicates with the server at sustained data transfer rate of 55 Mbyte/s using a HiPPI channel.

Table 3.3.3 shows sample running times for different steps of the segmentation process on a $256 \times 256 \times 124$ data set. The intent of this table is simply to give the reader a rough idea of how long each of these steps takes on this parallel machine, and to show that these running times make the process realistic for use in actual surgical planning tasks.

The deformable contour code runs on standard Sun/SGI workstations, and it takes 2–3 seconds for a typical contour to converge from an initial position on one slice. Since there is considerable variation in the time it takes a user to initialize contours, as well as the number of slices that the contours are propagated to, the total time of this step varies from case to case.

4. THE VALIDATION PROBLEM AND RESULTS

While several systems for segmentation of medical data are currently in use in various research laboratories and hospitals, the issue of correct validation is often ignored. The import of this statement is not that these systems perform unsatisfactorily, rather it is an illustration of the fact that the medical image processing community is reaching the consensus that validating segmentations of medical data is a hard, and perhaps even a poorly defined, problem.

^aWe think of this step as operational in '2.5-D' since the estimate of the brain is obtained by processing the data as a 3-D volume, and smoothness of the brain boundary between adjacent slices is imposed by the propagation of contours between adjacent slices.

Table 2. Strengths and weaknesses of validation methods.

#	Segmentation task	Strengths	Weaknesses
1	Brain tissue (Brummer <i>et al.</i> , 1993)	Simple	Subjective
2	White and grey matter (Gerig <i>et al.</i> , 1991; Kikinis <i>et al.</i> , 1992; Wells <i>et al.</i> , 1994)	Less subjective	15% variation in manual segmentations and labor intensive
3	White and grey matter (Wells <i>et al.</i> , 1994), cardiac walls and corpus callosum (Chakraborty <i>et al.</i> , 1994)	Scientific approach	Not sufficient in all applications
4	Cardiac motion (Shi <i>et al.</i> , 1994)	Reliable	Invasive
5	Brain tissue (Gobel <i>et al.</i> , 1994)	Reliable	Labor intensive

Each row summarizes one of five validation schemes: (1) visual inspection, (2) comparison with manual segmentation, (3) testing on synthetic data, (4) use of fiducials on subjects, (5) use of fiducials and/or cadavers.

4.1. The validation problem

A validation method can be thought of as a combination of two components. One component is the notion of a ‘ground truth’ against which the results of an algorithm are to be judged. The second component is a measure for establishing the deviation of the results from this ground truth. For their second component, most validation schemes use standard statistical methods of finding means, modes, variances, standard deviations, or root mean squared errors. The first component requires developing the notion of a ground truth for a segmentation algorithm and that is where one tends to run into difficulty. The problem is not that there is no ground truth for medical data, for certainly there is; the problem is that the ground truth is not typically available to the segmentation systems in any form that they can readily use. For instance, in the case of segmentation of brain tissue from MR images, there is indeed a true boundary of the brain tissue for each patient, but we do not know what it is. Approximations to the true boundary can be obtained in the form of manual segmentation by experts of neuroanatomy, but experiments have shown that there can be up to 15% variations in classifications generated by different experts (Kikinis *et al.*, 1992; Wells *et al.*, 1994). We could test our method on synthetic images, but it is difficult to generate synthetic images that capture the complexity and deformability of the human head. We could test our method on a phantom brain, but we are not aware of the existence of one. We could place fiducials or markers randomly in the patient in known areas of the brain, and use that as a subsampled ground truth to measure the classifications against, but that is too invasive to be practical. A reasonable alternative would be to use cadaver brains for marker based validation, but we are not aware of many systems that actually do this. One way of working around the ground truth issue would be to note that segmentation is usually a pre-processing step for other applications, and instead to measure the suitability of various segmentation algorithms to particular applications based on

the performance of the applications that build upon them. Of course, this ‘lazy evaluation’ would only be useful in applications where the validation problem is easier, for example in the case of registration of data from different modalities where external markers can be used (Mellor, 1995) to facilitate validation.

In the rest of this section we briefly summarize the methods that have been typically used for validation of segmentation algorithms, their strengths and weaknesses, and then present our results based on the validation scheme that we have chosen.

4.2. Validation methods

The following are representative of schemes used for validation of medical segmentation results.

- Method 1: visual inspection
- Method 2: comparison with manual segmentation
- Method 3: testing on synthetic data
- Method 4: use of fiducials on patients
- Method 5: use of fiducials and/or cadavers

Table 4 presents for each of these methods a summary of the particular segmentation applications for which it been used, as well as its strengths and weaknesses.

4.3. Our validation scheme and results

We use a combination of methods 1 and 2, i.e. visual inspection by radiologists and comparison with manual segmentation for validating our results.

Recently our segmenter has gained usage in three volumetric studies underway at Harvard Medical School and Brigham and Women’s Hospital: a schizophrenia study that examines the white matter and grey matter content in the temporal lobes (Shenton *et al.*, 1992), a study of the change in brain composition in newborns (Huppi *et al.*, 1995a, b), and a study of effects of altitude change on brain volume (Morocz *et al.*, 1995).

Table 3. Difference in classification between our method and manual segmentation.

Case no	Number of different voxels	Percentage difference
1	120 463	1.72
2	129 648	1.85
3	153 631	2.19
4	124 335	1.77
5	150 297	2.14
6	133 743	1.91
7	192 502	2.75
8	221 291	3.16
9	188 528	2.69
10	171 776	2.45
11	176 724	2.52
12	144 990	2.07
13	136 753	1.95
14	184 274	2.63
15	142 237	2.03
16	258 752	3.69
17	194 908	2.78
18	197 774	2.82
19	101 356	1.45
20	103 202	1.19

The middle column shows the number of pixels that were classified differently by our 3-class (white matter, grey matter, csf) segmentation method as compared with manual segmentation for each of 20 cases of size $256 \times 256 \times 124$ voxels.

Quantitative results for the comparison of our segmentation against expert segmentations are given below. All presented results are for gradient echo data sets of size $256 \times 256 \times 124$ voxels.

4.3.1. Quantitative results

Table 4.3 shows a simple comparison of the number and percentage of pixels that are classified differently by our segmenter as compared with expert segmentation, based on a 3-class segmentation of the data into white matter, grey matter and csf. Note that while the total number of voxels classified differently in the two cases appears large, compared with the total volume in each data set, only a small fraction (typically 1–2%) are actually different.

To test the extent to which this difference in classification affects the overall segmentation, we ran a second test. Tables 4.3 and 4.3.2 present the result of a test in which the edges of the brain tissue generated by our segmentations and the edges generated by the manual segmentations are compared. This comparison is performed in two steps: first by finding the mean distance (city-block) from each edge in our segmentation to edges in the manual segmentation (which basically penalizes false positives) as shown in Table 4.3, and then by

Table 4. Measuring false positives.

Case no	Mean d	% at $d = 0$	% at $d = 1$	% at $d = 2$	% at $d > 2$
1	1.10	19.84	69.05	7.65	3.47
2	1.23	8.19	73.49	12.67	5.65
3	0.84	31.49	61.27	5.10	2.13
4	1.07	20.91	66.45	7.29	5.35
5	1.11	11.00	75.19	9.27	4.54
6	0.85	28.20	65.20	4.82	1.78
7	0.76	38.58	55.20	4.42	1.80
8	0.81	34.74	57.45	5.04	2.78
9	0.52	57.18	39.12	2.29	1.41
10	1.21	22.34	63.32	6.78	7.56
11	0.77	42.97	47.00	5.14	4.89
12	2.22	15.27	61.86	8.13	14.74
13	1.67	17.96	58.14	8.43	15.47
14	1.36	13.12	68.58	8.84	9.47
15	1.59	13.63	63.66	10.33	12.38
16	6.44	37.79	31.49	4.58	26.14
17	1.29	42.25	40.81	5.32	11.63
18	1.26	41.16	40.10	5.94	12.80
19	1.15	16.53	71.84	8.29	3.35
20	1.54	30.51	64.55	2.19	2.04

The second column shows the mean distance d (in pixels) between a pixel that falls on the boundary of the brain in the manual segmentation and the nearest brain-boundary pixel in our segmentation. The rest of the columns show the percentage of brain-boundary pixels in the manual segmentation at different distances from the nearest brain-boundary pixels in our segmentation.

finding the mean distance from an edge in the manual segmentation to edges in our segmentation (which penalizes false negatives) as shown in Table 4.3.2. We also include some other statistics on the edges, such as the percentage of edges in our segmentation that coincide with (or are one and two pixels away from) edges in the manual segmentation. We repeat this test to compute the number of edges in the manual segmentation that coincide with, or are within one or two pixels of the edges of our segmentation. The idea behind these tests is to convey the ‘goodness of fit’ between the brain boundaries produced by our segmentation and ones produced manually. As can be seen by the numbers in these tables, we are usually within 2–3% of manual segmentations, which is considered acceptable performance in the applications with which we are dealing (Kikinis *et al.*, 1992).

Although we did not test the distribution of false positives and false negatives in a detailed manner, we observed from sampling of the data that these errors appeared to be randomly uniformly distributed, rather than aggregated in particular places.

We illustrate our results in Figure 9 using a representative slice from an MR scan. The top left image shows the edges of brain in our segmentation overlaid on the input image (in




Figure 9. Top left, input image overlaid with brain edges from our segmentation; top right, input image overlaid with brain edges from manual segmentation; bottom left, our segmentation overlaid with brain edges from manual segmentation; bottom right, manual segmentation overlaid with brain edges from our segmentation.

white), and the top right image shows the edges of the brain in the manual segmentation overlaid on the input image (also in white). The bottom left image shows our segmentation overlaid with the edges of the brain from the manual segmentation, and the bottom right shows the manual segmentation overlaid with the edges of the brain in our segmentation. The purpose of the top two images of the figure is to give an idea of how well the boundaries of our segmentation and the manual segmentation agree with the original greyscale images, and the purpose of the bottom two images is to illustrate how well these two sets of edges agree with each other.

4.3.2. Qualitative results

We are now generating routine segmentations for surgical planning at the Surgical Planning Lab of Brigham and

Women's Hospital. It takes us about 20 minutes to process a single case, which makes the method attractive to clinicians who could otherwise spend 4–6 h creating the segmentations with less automatic tools. Typically clinicians overlay these segmentations with manually segmented pathological areas of the brain to create the final visualizations. Figure 10 shows two views of the segmented white matter surface for two different normal controls. Figure 11 shows a view of a patient with subdural hematoma as well as a tumor.

4.4. Generalization to other anatomical structures

Through in this paper we have developed the details for the segmentation of a particular anatomical structure (the brain) there are other structures that share the characteristics of the brain that have been exploited in this process, and hence, this

Table 5. Measuring false negatives.

Case no	Mean d	% at $d = 0$	% at $d = 1$	% at $d = 2$	% at $d > 2$
1	1.20	10.37	72.57	12.39	4.67
2	1.34	21.41	69.02	4.46	5.11
3	1.51	7.44	66.06	15.97	10.54
4	1.69	10.97	72.20	11.81	5.03
5	1.08	12.49	78.51	6.56	2.43
6	1.46	6.46	69.14	16.03	8.37
7	1.44	6.67	65.49	18.50	9.34
8	1.32	8.34	67.57	19.73	4.36
9	2.09	5.33	48.32	28.47	17.88
10	1.10	14.91	69.71	11.58	3.80
11	1.34	8.76	63.61	22.36	5.27
12	1.25	14.82	71.71	9.71	3.76
13	1.17	14.21	70.97	10.74	4.07
14	1.17	13.08	75.40	8.00	3.52
15	1.00	17.91	72.09	7.48	2.53
16	1.48	9.00	55.81	26.81	8.39
17	2.26	8.87	56.99	24.13	10.02
18	1.43	11.28	56.98	25.07	6.68
19	1.42	15.12	68.81	10.04	6.03

The second column shows the mean distance d (in pixels) between a pixel that falls on the boundary of the brain in our segmentation and the nearest brain-boundary pixel in the manual segmentation. The rest of the columns show the percentage of brain-boundary pixels in our segmentation at different distances from the nearest brain-boundary pixels in the manual segmentation.

method may be extensible to those structures also. The kidney and the heart are two examples of such structures. In order to extend our process to other structures, one could apply the following methodology:

- Identify the various relevant intensity distributions in the data using EM segmentation.
- Incorporate topological information using customized morphological operators or other techniques.
- Incorporate spatial information using deformable contours or other customized spatial priors.

The key word in the above method is ‘customize’, since most of the effort in applying this segmenter to any other application will be hidden in doing exactly this. In some sense, this customization step corresponds to the construction of a model for each anatomical structure, which cannot be bypassed. However, explicit shape modeling promises a more elegant approach to creating general purpose segmentation methods, and is an attractive direction to pursue.

4.5. Conclusions

We have presented a method for segmentation of brain tissue in MR images. This method combines three techniques into an effective tool for segmentation. The EM segmentation step




Figure 10. Top, top view of a reconstructed white matter surface; bottom, left view of the same reconstructed white matter surface.

provides reliable intensity-based classification of the data into various tissue classes. Further processing is needed because brain structures are not defined by unique intensities in MRI. The morphology step leverages the tissue labels generated by the EM segmentation to binarize the image into hypotheses for brain and non-brain regions. It is important to note here that a less powerful method for intensity-based classification could




Figure 11. Side view of a reconstructed brain surface of a patient overlaid with manually segmented pathology—subdural hematoma (red) and tumor (green).

lead to gross errors in using binary morphology. In the cases that the morphology step does not isolate the brain, snakes (or related deformable models) are used as a fast manual editing tool. We have created an implementation that has been used to segment well over 100 brain scans, with results that are comparable to segmentations performed manually by experts.

The key aspect of our system is that it divides the segmentation task into three simpler steps, each of which exploits a different set of constraints of the problem, and more importantly, the result of each step facilitates the one that follows it. Specifically, the success of binary mathematical morphology hinges on prior accurate binarization of the data into object and background, and by employing EM segmentation prior to the morphological operations, we generate reliable class labels for the data, thereby facilitating the morphology step.

Second, removal of most connecting elements between the brain and non-brain structures in the morphology step makes the process of initializing the true brain contour easier on the user, thereby facilitating the final, deformable contour step of our system.

ACKNOWLEDGEMENTS

This paper describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for this work was provided in part by ARPA under ONR contracts N00014-91-J-4038 and N00014-94-01-0994, and in part by a grant from IBM. We thank the anonymous reviewers for their constructive comments.

REFERENCES

- Ayache, N. (1995) Medical computer vision, virtual reality and robotics. *Image Vision Comput.*, 13, 295–313.
- Brummer, M. E., Mersereau, R. M., Eisner, R. L. and Lewine, R. R. J. (1993) Automatic detection of brain contours in MRI data sets. *IEEE Trans. Med. Imaging*, 12, 153–166.
- Chakraborty, A., Staib, L. and Duncan, J. (1994) An integrated approach to boundary finding in medical images. In *Proc. IEEE Workshop on Biomedical Image Analysis*, Seattle, WA. IEEE.
- Chang, C., Wang, W., Rangarajan, A., Constable, R., Zubal, I. and Gindi, G. (1994) Segmentation of mr brain scans using atlas-guided deformable contours. In *AAAI Spring Symposium: Applications of Computer Vision in Medical Image Processing*, Palo Alto, CA. AAAI.
- Cline, H., Lorensen, W., Kikinis, R. and Jolesz, F. (1990) Three-Dimensional Segmentation of MR Images of the Head Using Probability and Connectivity. *J. Comp. Assis. Tomogr.*, 14, 1037–1045.
- Cohen, L. (1991) On active contour models and balloons. *Computer Vision, Graphics and Image Processing: Image Understanding*, 53, 211–218.
- Cohen, L. and Cohen, I. (1993) Finite element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Trans. PAMI*, 15, 1131–1147.
- Collins, D., Peters, T., Dai, W. and Evans, A. (1992) Model based segmentation of individual brain structures from mri data. In Robb, R. (ed.), *Proc. First Conf. on Visualization in Biomedical Computing*, pp. 10–23. SPIE.
- Cootes, T., Hill, A., Taylor, C. and Haslam, J. (1994) Use of active shape models for locating structures in medical images. *Image Vision Comput.*, 12, 355–366.
- Gerig, G., Martin, J., Kikinis, R., Kubler, O., Shenton, M. and Jolesz, F. (1991) Automating segmentation of dual-echo mr head data. In *XIIIth Int. Conf. on Information Processing in Medical Imaging*, pp. 175–187, Springer-Verlag.
- Gobel, J., Snell, J., Hinckley, K. and Kassell, N. (1994) A real-time system for 3d neurosurgical planning. In Robb, R. (ed.), *Proc. Third Conf. on Visualization in Biomedical Computing*. SPIE.
- Grimson, E. et al. (1994) An automatic registration method for frameless stereotaxy, image guided surgery, and enhanced reality visualization. In *Proc. Computer Society Conf. on Computer Vision and Pattern Recognition*, Seattle, WA. IEEE.
- Grzeszczuk, R. and Levin, D. (1994) Brownian strings: segmenting images with stochastically deformable contours. In Robb, R. (ed.), *Proc. Third Conf. on Visualization in Biomedical Computing*. SPIE.
- Haralick, R., Sternberg, S. and Zhuang, X. (1987) Image analysis using mathematical morphology. *IEEE Trans. PAMI*, 9, 532–550.
- Hohne, K. H. and Hanson, W. (1992) Interactive 3d segmentation of mri and ct volumes using morphological operations. *J. Comp. Assis. Tomogr.*, 16, 285–294.
- Huppi, P. S., Tsuji, M., Kapur, T., Barnes, P., Jakab, M., Zientara, G., Kikinis, R. and Jolesz, F. (1995a). Quantification of changes in postnatal brain development in preterm infants using adaptive segmentation of mri data. In *Proc. Third Annual Meeting of the Society of Magnetic Resonance*, Vol. I, abstract 117.
- Huppi, P. S., Tsuji, M., Kapur, T., Barnes, P., Zientara, G., Kikinis, R., Jolesz, F. and Volpe, J. (1995b). 3d-mri, a new measure of brain development in preterm infants. *Pediatr. Res.*, 37, abstract.
- IBM (1993) *IBM POWER Visualization System: Programmer's Technical Reference*, 1st edn. Number SC38-0487-00.
- Kikinis, R. et al. (1992) Quantitative Analysis of Brain and Cerebrospinal Fluid Spaces with MR Imaging. *J. Magn. Res. Imag.*, 2, 619–629.
- Mellor, J. P. (1995) *Real Time Camera Calibration*. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Morocz, I. A., Gudbjartsson, H., Kapur, T., Zientara, G. P., Smith, S., Muza, S., Lyones, T. and Jolesz, F. A. (1995) Quantitation of diffuse brain edema in acute mountain sickness using 3d mri. In *Proc. Third Annual Meeting of the Society of Magnetic Resonance*, Vol. I, abstract 1241.
- Pelizzari, C., Chen, G., Spelbring, D., Weichselbaum, R. and Chen, C. (1989) Accurate three dimensional registration of ct, pet and/or mr images of the brain. *J. Comp. Assis. Tomogr.*, 13, 20–26.
- Sandor, S. and Leahy, R. (1994) A 3d morphological algorithm for automated labelling of the cortex in magnetic resonance brain images. In *AAAI Spring Symposium Applications of Computer Vision in Medical Image Processing*, Palo Alto, CA. AAAI.
- Serra, J. (1982) *Image Analysis and Mathematical Morphology*. Academic Press, London.
- Shenton, M., et al. (1992) Abnormalities of the Left Temporal Lobe and Thought Disorder in Schizophrenia. *N. Engl. J. Med.*, 327, 604–612.
- Shi, P., Amini, A., Robinson, G., Sinusas, A., Constable, C. and Duncan, J. (1994) Shape-based 4d left ventricular myocardial function analysis. In *Proc. IEEE Workshop on Biomedical Image Analysis*, Seattle, WA. IEEE.
- Staib, L. and Duncan, J. (1992) Boundary finding with parametrically deformable contour models. *IEEE Trans. PAMI*, 14, 1061–1075.
- Szeliski, R. and Lavalley, S. (1993) Matching 3D Anatomical Surfaces with Non-Rigid Deformations Using Octree-Splines. In Vemuri, B. (ed.), *Proc. SPIE, 2031, Geometric Methods in Computer Vision*.
- Wells, W., Kikinis, R., Grimson, E. and Jolesz, F. (1994) Statistical intensity correction and segmentation of magnetic resonance image data. In Robb, R. (ed.), *Proc. Third Conf. on Visualization in Biomedical Computing*. SPIE.
- Wells, W., Kikinis, R., Grimson, E. and Jolesz, F. (1995) Adaptive segmentation of mri data. In Ayache, N. (ed.), *Proceedings of the First Conference on Computer Vision, Virtual Reality and Robotics in Medicine*, pp. 59–69. Springer.
- Wells, W., Kikinis, R., Grimson, E. and Jolesz, F. (1996) Adaptive segmentation of mri data. *IEEE Trans. Med. Imaging*, to appear.
- Westbook, C. and Kaut, C. (1993) *MRI in Practice*. Blackwell Scientific, Oxford.
- Witkin, A., Kass, M. and Terzopoulos, D. (1988) Snakes: Active contour models. *Int. J. Comp. Vision*, 1 321–331.