Fully Automatic Segmentation of the Brain in MRI

M. Stella Atkins, Blair T. Mackiewich *, Ken Whittall **

* School of Computing Science, Simon Fraser University

Burnaby, British Columbia, Canada V5A 1S6

** Department of Radiology, Vancouver Hospital, UBC site

This research has been partially funded by the Natural Sciences and Engineering Research Council of Canada.

Abstract

A robust fully automatic method for segmenting the brain from head MR images has been developed, which works even in the presence of RF inhomogeneities. It has been successful in segmenting the brain in every slice from head images acquired from several different MRI scanners, using different resolution images and different echo sequences.

The method uses an integrated approach which employs image processing techniques based on anisotropic filters and "snakes" contouring techniques, and a-priori knowledge, which is used to remove the eyes, which are tricky to remove based on image intensity alone. It is a multi-stage process, involving first removal of the background noise leaving a head mask, then finding a rough outline of the brain, then refinement of the rough brain outline to a final mask.

The paper describes the main features of the method, and gives results for some brain studies.

Keywords: Magnetic Resonance Imaging, Intracranial Boundary Detection, Nonlinear Anisotropic Diffusion, Active Contour Models

1 Introduction

Magnetic Resonance Imaging (MRI) provides detailed 3D images of living tissues, and is used for both brain and body human studies. Data obtained from MR images is used for detecting tissue deformities such as cancers and injuries, and more recently, functional MRI is used for measuring function such as blood flow. For example, we have been using MRI brain images of Multiple Sclerosis (MS) patients in studies to evaluate the effect of drugs [13]. The MRI scans of MS patients show bright spots in the proton density (PD) and T2 images, which correspond to MS lesions. The volume of these lesions changes over time, and it has been shown that the size of the lesions corresponds with the progress of the disease [23],[9]. In order to perform good quantitative studies, the regions of interest must be well defined. In traditional methods, a radiologist manually outlines the region of interest using a mouse or cursor. More recently, computer-assisted methods have been used [6], [32]; some of these methods are surveyed recently in [14].

We describe here our new fully automatic method for segmenting the brain from the head in MR images. The key to any automatic method is that it must be robust, so that it produces reliable results on every image acquired from any particular MR scanner using any echo sequence. Our method is so robust, that it successfully was able to segment the brain in every slice of head images from 3 different MRI scanners, using different resolution images and different echo sequences. Our method works in the presence of radio frequency (RF) inhomogeneity and it addresses the partial volume effect in a consistent reasonable manner.

The remainder of the paper describes the method we use, and gives some results. We then compare our method with other published methods, and provide a summary with an outline of future work, including extending the method to segment lungs in MR chest images.

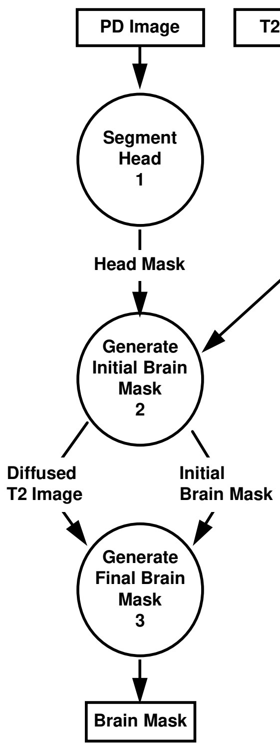

Figure 1: A simplified data flow diagram representing automatic intracranial boundary detection.

2 Method

2.1 Overview

We use a 3-stage method to segment images as shown in Fig. 1. First we remove the background noise, then we generate an initial mask for the region(s) of interest, then we refine the mask for the final segmentation. Each stage (bubble) in the diagram has been implemented using the WiT visual programming environment [2], which aids prototype development and enables experimentation [3].

The first stage, *Segment Head*, uses intensity histogram analysis to remove background noise and provide a head mask defining the head.

The second stage, *Generate Initial Brain Mask*, produces a mask that approximately identifies the intracranial boundary. It uses nonlinear anisotropic diffusion of the T2-weighted data set and spatial information provided by the head mask to identify regions corresponding to the brain. The nonlinear anisotropic diffusion effectively counters RF inhomogeneity by attenuating non-brain regions.

With the initial brain mask as a seed, the third step, *Generate Final Brain Mask*, locates the intracranial boundary using an active contour model algorithm. The active contour model algorithm consistently tracks the edge of the brain, even in the presence of partial volume effects.

Stages 1 and 3 require no a-priori information about the organ to be segmented. Knowledge about the organ of interest is used only in the second stage, in which the initial brain mask is generated. We use two expert pieces of information. One allows us to eliminate non-brain tissues such as the eyes, using morphology, by exploiting the simple fact that the brain centroid must be near the centroid of the slice. The other item of knowledge is that the brain tissues appear in MR images with relatively high intensity and when the T2 (or PD) image is filtered with a special anisotropic diffusion filter, the non-brain tissues can be attenuated, and hence the brain tissues can be segmented using a simple threshold. The methodology is described in more detail below, and in full detail in [20]. Although seemingly complex, the method has proven so robust that it works even in the presence of RF inhomogeneity, which would not be the case if a simple threshold technique were used to identify brain tissue from non-brain tissue. Similar techniques have been applied to chest MRI to segment the lungs.

2.2 Segment Head

The head mask is generated using the method suggested by Brummer at al. to determine the "best" threshold level for removing background noise in PD-weighted MR images [5]. By assuming the MR scanner produces normally distributed white noise, Brummer et al. show that background noise in the reconstructed PD-weighted MR volumes has a Rayleigh distribution:

$$p_{noise}(f) = \frac{f}{\sigma^2} \exp\left(-\frac{f^2}{2\sigma^2}\right) \tag{1}$$

where f is the noise intensity and σ is the standard deviation of the white noise. Fig. 2 shows that this distribution is easily visible in the low intensity range of the uncorrected MRI histogram.

The subtraction of the best fit Rayleigh curve, $r(f)$, from the volume histogram, $h(f)$, produces a bimodal distribution:

$$g(f) = h(f) - r(f) \tag{2}$$

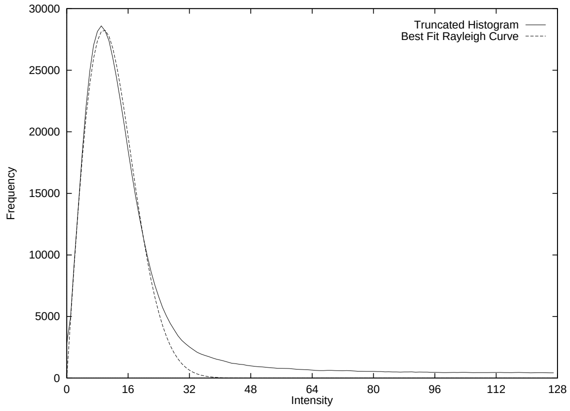

Figure 2: A truncated histogram of a PD-weighted MR volume. The background noise at the low end of the histogram is characterized by a Rayleigh distribution.

(a) (b) (c)

Figure 3: A head mask produced using automatic thresholding and morphology. (a) Original image. (b) Initial head mask. (c) Final head mask after morphology.

A minimum error threshold, τ, can be determined from $g(f)$ by minimizing an error term, ε_τ:

$$\varepsilon_\tau = \sum_{f=0}^{\tau-1} g(f) + \sum_{f=\tau}^{\infty} r(f) \tag{3}$$

Fig. 3 shows the results of automatically thresholding an MR volume using this method. The binary image in Fig. 3(b) produced by thresholding the volume at τ, contains speckle outside the head region and that there are misclassified regions within the head. This "noise" is easily removed using standard morphological operations.

2.3 Generate Initial Brain Mask

Generate Initial Brain Mask is a three step process. It first uses 2D nonlinear anisotropic diffusion to attenuate non-brain regions in each MRI slice. Next, it automatically thresholds the diffused MR volume to produce a binary mask. Finally, misclassified, non-brain regions are removed from the binary mask using morphology and spatial information provided by the brain mask.

2.3.1 Nonlinear Anisotropic Diffusion

Nonlinear anisotropic diffusion filters are iterative, "tunable" filters introduced by Perona and Malik [24]. Gerig et al. used such filters to enhance MR images [10]. Sapiro and Tannenbaum used a similar technique to perform edge preserving smoothing of MR images [25]. Others have shown that diffusion filters can be used to enhance and detect object edges within images [21, 24, 1].

Perona and Malik formulate the anisotropic diffusion filter as diffusion process that encourages intraregion smoothing while inhibiting interregion smoothing:

$$\frac{\partial}{\partial t} I(\overline{x}, t) = \nabla \bullet (c(\overline{x}, t) \nabla I(\overline{x}, t)) \tag{4}$$

In our case, $I(\overline{x}, t)$ is the MR image. $\overline{x}$ refers to the image axes (i.e. x, y) and t refers to the iteration step. $c(\overline{x}, t)$ is called the *diffusion function* and is a monotonically decreasing function of the image gradient magnitude. It allows for locally adaptive diffusion strengths: edges are selectively smoothed or enhanced based on the evaluation of the diffusion function. We use the following diffusion function in our algorithm [24]:

$$c(\overline{x}, t) = \exp\left(- \left(\frac{|\nabla I(\overline{x}, t)|}{\sqrt{2}\,\mathcal{K}} \right)^2 \right) \tag{5}$$

$\mathcal{K}$ is referred to as the *diffusion constant* or the *flow constant*, and the behavior of the filter depends on $\mathcal{K}$. The greatest flow is produced when the image gradient magnitude is close to the value of $\mathcal{K}$. Therefore, by choosing $\mathcal{K}$ to correspond to gradient magnitudes produced by noise, the diffusion process can be used to reduce noise in images. Assuming an image contains no discontinuities, object edges can be enhanced by choosing a value of $\mathcal{K}$ slightly less than the gradient magnitude of those edges. Similarly, an edge can be smoothed by choosing a value of $\mathcal{K}$ slightly larger than the gradient magnitude of that edge.

The discrete diffusion process consists of updating each pixel in the image by an amount equal to the flow contributed by its four nearest neighbors. [1] We use 2D nonlinear anisotropic diffusion to attenuate non-brain regions in each MRI slice[2]. As illustrated in Fig. 4, this attenuation enables the brain to be segmented using a simple threshold. Without such attenuation, a threshold-based segmentation would be difficult, if not impossible.

2.3.2 Automatic Threshold

Once each T2-weighted MRI slice has been diffused to attenuate non-brain voxels, we segment the brain using a single automatic threshold. The threshold level is determined by fitting a gaussian curve to the histogram of the diffused volume, then choosing a level at 2 standard deviations below the mean [5]. Fig. 5 shows the voxel intensity histogram of a diffused volume with the best fit gaussian curve and threshold level overlaid. Fig. 6 shows a slice of the binary mask produced by the threshold.

2.3.3 Mask Refinement

The binary mask produced by automatic thresholding contains misclassified regions, such as the eyes (see Fig. 6(b)). These regions are removed from each slice using morphology and spatial information provided by the head mask. First, holes are filled within each region of the mask. Next, binary erosion is performed to separate weakly connected regions. After erosion, regions whose centroids fall outside a bounding box defined by the head mask are eliminated. This bounding box is illustrated in Fig. 7.

Finally, binary dilation is performed on the remaining regions to return them to their original size. The results of these steps are shown in Fig. 8.

Again, similar techniques were used to obtain an initial lung segmentation from the masked chest, with just different morphology parameters to ensure that the centroid of each lung was appropriately located within the chest. Details are given in the later section on Results.

[1] Eight nearest neighbors can be used if the flow contribution of the diagonal neighbors is scaled according to their relative distance from the pixel of interest [10]. Anisotropic data can be handled similarly.

[2] We also attempted to attenuate non-brain tissues using 3D nonlinear anisotropic diffusion over the entire MRI volume. 3D diffusion, however, increased the partial volume effect, blurring the edges of the brain [20].

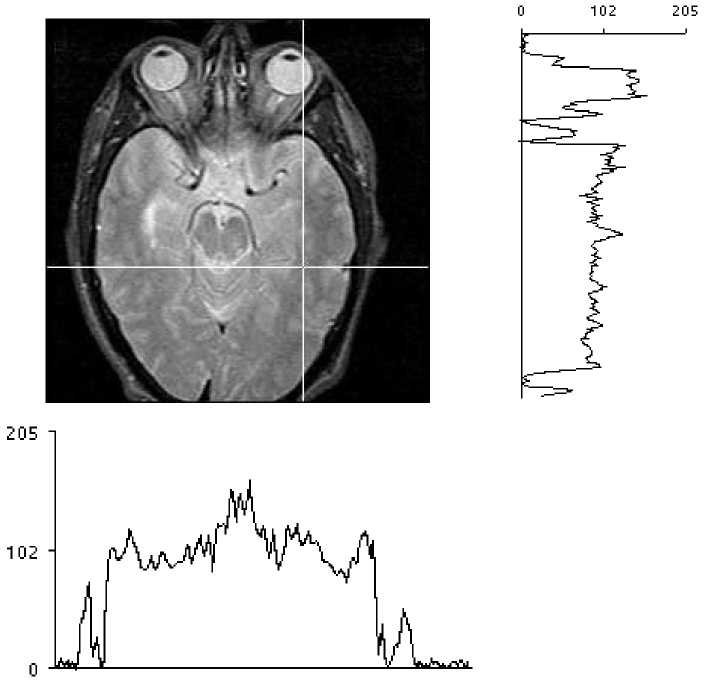

(a)

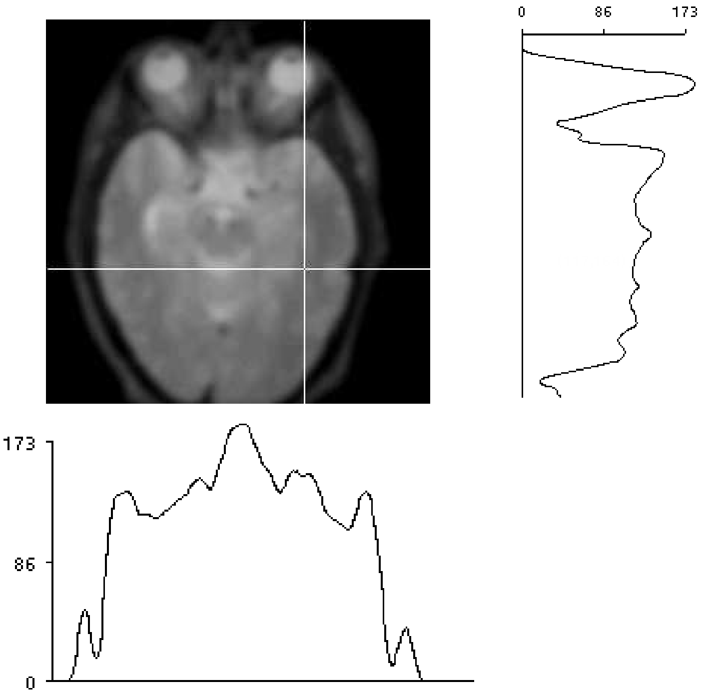

(b)

Figure 4: Intracranial boundary detection using 2D nonlinear anisotropic diffusion filtering. (a) Original image. (b) 2D diffused image. Diffusion attenuates non-brain tissues, enabling a simple threshold to segment the brain.

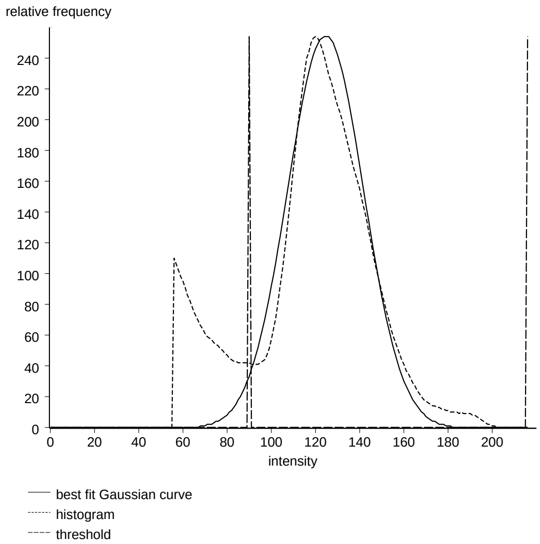

Figure 5: A histogram of the diffused T2-weighted MR scan with the best fit gaussian curve and threshold levels overlaid.

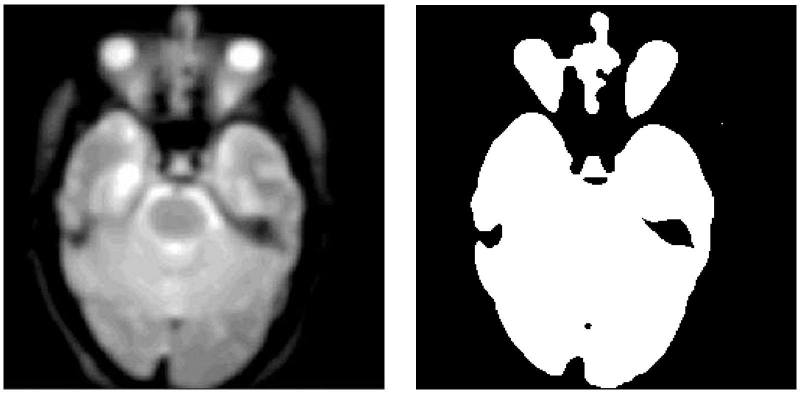

(a) (b)

Figure 6: A binary mask produced by automatic thresholding. (a) The diffused image slice. (b) The corresponding binary mask.

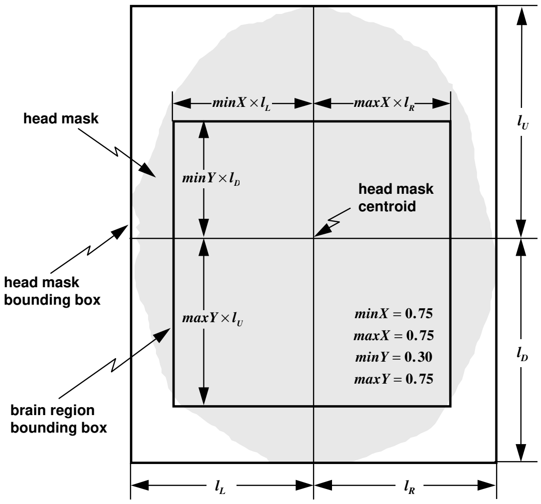

Figure 7: Spatial information from the head mask is used to eliminate regions that are unlikely to correspond to brain tissue. Features whose centroids fall outside the *brain region bounding box* are discarded. The dimensions of the *brain region bounding box* were chosen experimentally to produce good results for all data sets.

2.4 Generate Final Brain Mask

Given the initial brain mask as a seed, we use an active contour model algorithm, extended from the "Snakes" algorithm introduced by Kass et al. [16], to locate the boundary between the brain and the intracranial cavity. The algorithm deforms the contours defined by the perimeter of the initial brain mask to lock onto the edge of the brain. Each active contour is defined as an ordered collection of n points in the image plane:

$$\begin{aligned} V &= \{v_1, \cdots, v_n\} \qquad (6) \\ v_i &= (x_i, y_i),\ i = \{1, \cdots, n\} \end{aligned}$$

The points in the contour iteratively approach the intracranial boundary through the solution of an energy minimization problem. For each point, v_i, an energy matrix, $E(v_i)$, is computed:

$$E(v_i) = \alpha E_{cont}(v_i) + \beta E_{bal}(v_i) + \gamma E_{int}(v_i) + \kappa E_{grad}(v_i) \qquad (7)$$

where $E_{cont}(v_i)$ is a "*continuity*" energy function that influences the contour to take the shape of a circle, $E_{bal}(v_i)$ is an adaptive "*balloon*" force that pushes the contour outward until it reaches

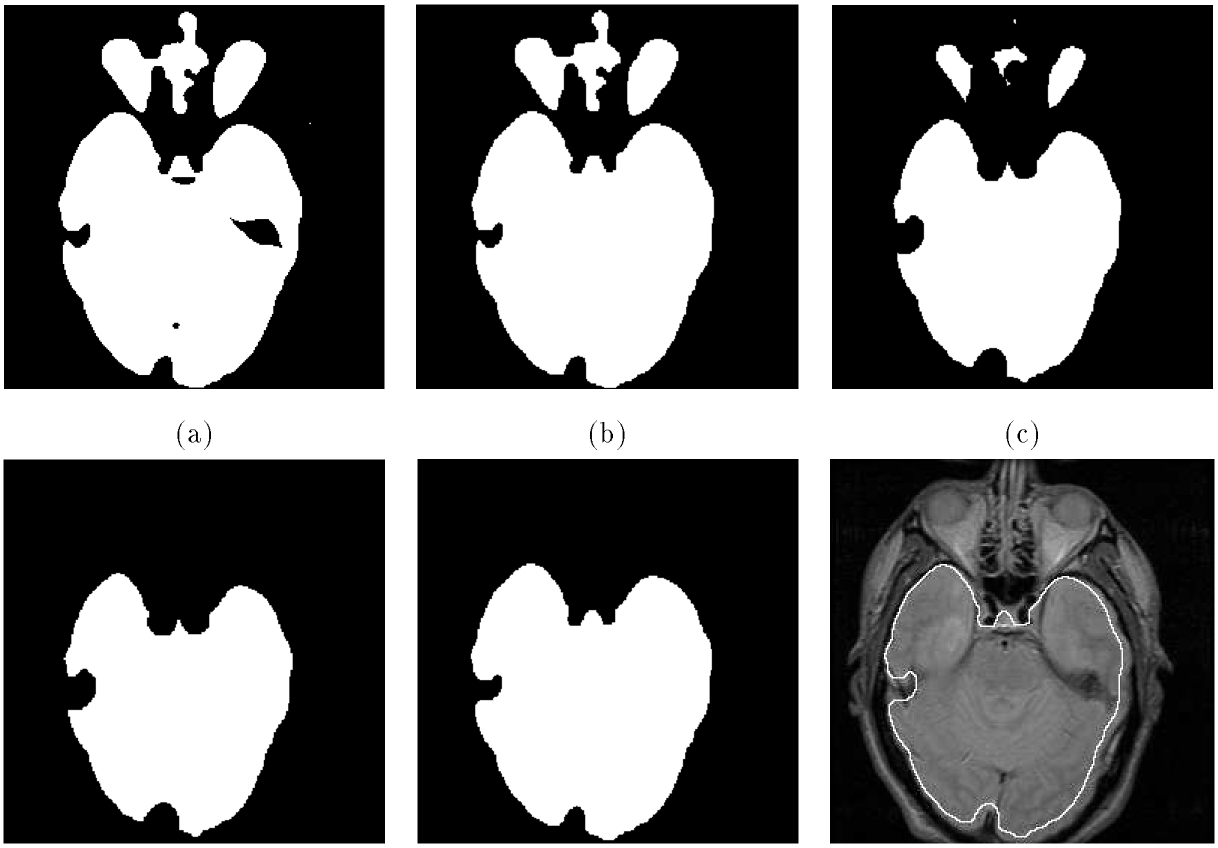

(a) (b) (c) (d) (e) (f)

Figure 8: Elimination of misclassified regions from the initial brain mask. The mask, (a) produced by thresholding, (b) after regions have been filled, (c) after binary erosion, (d) after elimination of non-brain regions, and (e) after binary dilation. (f) shows the initial brain mask overlaid on the PD-weighed slice.

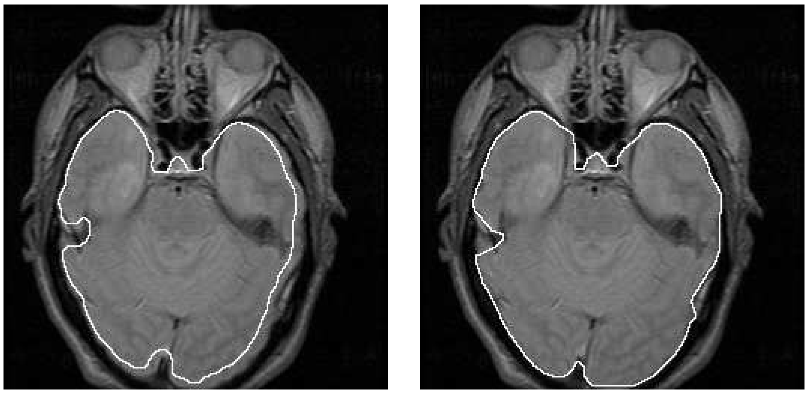

(a) (b)

Figure 9: Refinement of the intracranial contour. (a), the contour defined by the perimeter of the initial brain mask. (b), the intracranial contour detected using the active contour model algorithm.

a strong gradient [7], $E_{int}(v_i)$ is an "*intensity*" energy function, computed from the PD-weighted MRI volume, that influences the contour to move towards low intensity regions, and $E_{grad}(v_i)$ is a "*gradient*" energy function, computed from the diffused T2-weighted MRI volume, that influences the contour to move towards regions of strong gradient. α, β, γ, and κ are scalar constants providing relative weightings of the energy terms. Each v_i is then iteratively moved to the point of minimum energy in its neighborhood, which corresponds to smallest element in $E(v_i)$.

The combination of energy functions described above enables the active contour model algorithm to detect the intracranial boundary in all image slices using the same relative energy weightings. The diffusion process ensures that the initial brain mask falls completely inside the brain. Therefore, the adaptive balloon force aides in cases where the initial brain mask is poor. The intensity energy term helps the active contour algorithm produce consistent results where partial volume effects are particularly severe. Finally, computing the gradient energy term from the diffused T2-weighted volume greatly stabilizes the active contour algorithm because the gradient derivatives are small in the diffused volume [20].

Fig. 9 illustrates the result of applying our active contour model algorithm to the MR slice shown in Fig. 8. α, β, γ, and κ were experimentally chosen to produce good results.

3 Results

3.1 Brain Studies

Fig. 10 shows the intracranial contour detected automatically by our algorithm overlaid on selected slices of a PD-weighted MRI data set. The PD and T2 data sets were acquired axially on a GE 1.5 Tesla MRI scanner, with repetition time TR = 2000 ms, and echo times of 35 ms and 70 ms respectively. The data consists of 22 slices with 256 * 256 8-bit pixels per slice. The pixel size is 0.781 mm^2, and the slice thickness is 5 mm.

The intracranial boundary shown in Fig. 10 is accurate in all slices, except for an apparent slight overshoot at the top in slices 4,5 and 6. Careful examination shows that the partial volume effect in these slices is dealt with in a consistent way, and the apparent overshoot is not a problem for subsequent analysis of the brain data.

Our algorithm produced comparable results for many other data sets from two other scanners [20]. In all cases, our algorithm successfully detected the intracranial boundary without user interaction and without parameter modification.

3.2 Previous Work

3.2.1 Automatic Thresholding

Suzuki and Toriwaki use iterative thresholding to distinguish brain tissues from others in axial MR slices [29]. Starting at set values, thresholds for the head and the brain are then iteratively adjusted based on the geometry of resulting masks (ie. the head mask includes the brain mask). This method is ineffective in the presence of RF inhomogeneity and in slices where the brain is not one homogeneous region closely surrounded by the skull.

Li et al. use knowledge-based thresholding in multimodal MRI data to classify voxels into multiple intensity categories [18]. In each axial slice, they compute the centroid of voxels categorized as brain. Next, four points defining a quadrangle are found at the edge of the brain by tracing left, right, up, and down from the centroid to a transition in tissue categories. All voxels outside the quadrangle that are not categorized as brain tissue are then masked to define the intracra-

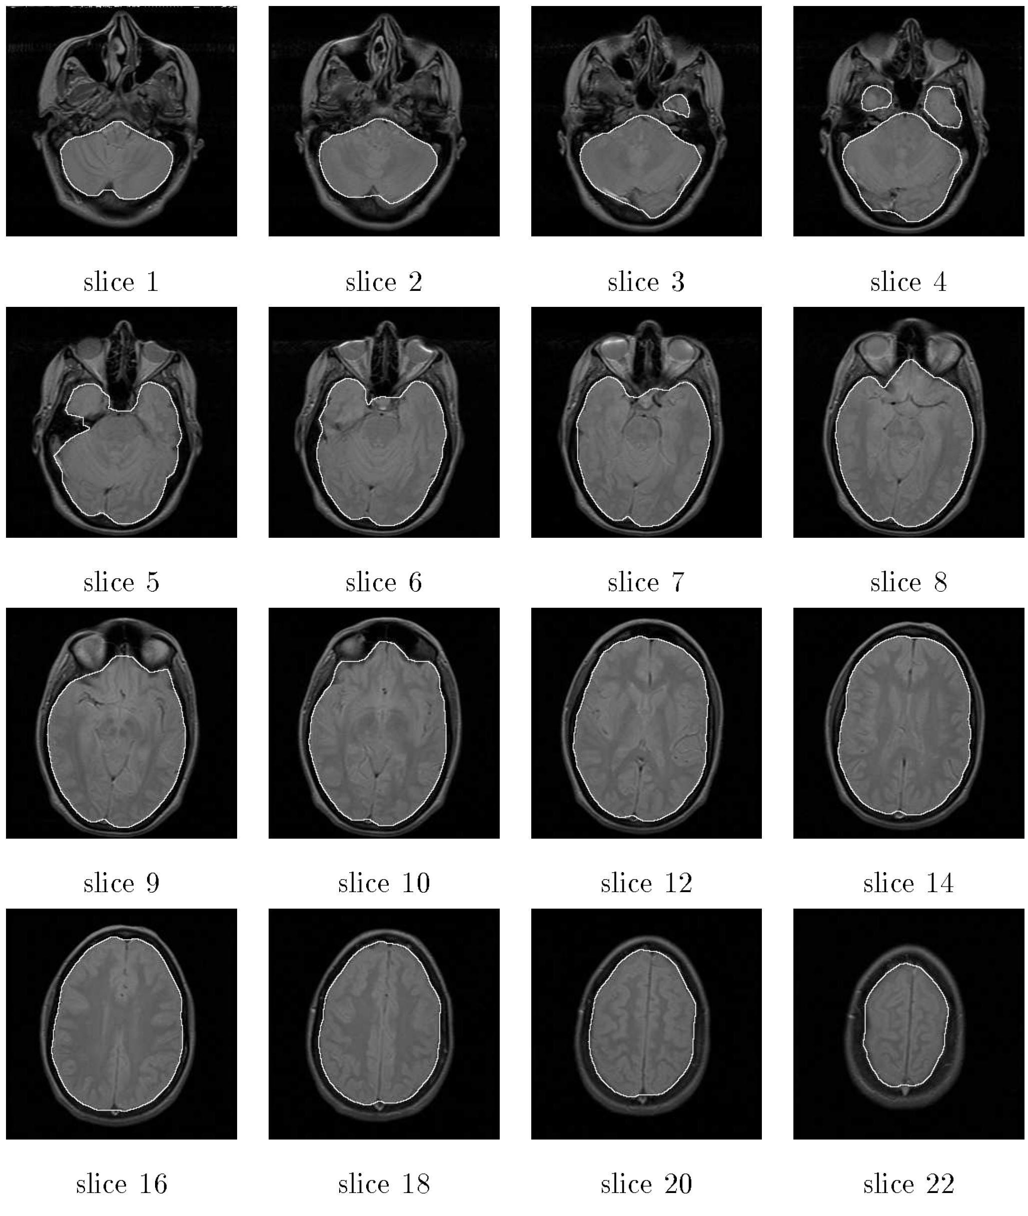

Figure 10: An automatically detected intracranial contour overlaid on selected slices of a PD-weighed MRI scan.

nial contour. Obviously, this method works only in slices where the brain constitutes one fairly homogeneous region.

Brummer et al. use histogram analysis and morphology to generate a 3D brain mask [5]. Using a model of background noise, they first automatically generate a mask of the head and perform intensity correction on the masked volume. Next they create an initial brain mask using an automatic threshold based on a presupposed brain voxel intensity distribution. They then eliminate regions in the brain mask that are too close to the edge of the head. Finally, they use novel morphological operations to clean up the resulting mask. This method misses brain tissue in extreme slices and includes non-brain tissues in others. In some cases it produces errors near the eyes. The method relies on *a priori* intensity correction to deal with RF inhomogeneity, so can't be used retroactively.

3.2.2 Region Growing and Boundary Detection

Pannizzo et al. detect the intracranial boundary in axial MRI slices by tracing a horizontal line outwards from the center of the image [22]. The point, in each direction, at which the voxel intensity under the line drops below a reference threshold is considered to be a point in the intracranial boundary. A running average of voxel intensities under the line is then computed. The intracranial boundary points are relocated to the first voxels with an intensity too far below the average. The running average procedure is repeated for rows in the image. The entire process is then repeated for all columns. The result is a sequence of points defining the intracranial contour. This method can only detect the intracranial in slices where the brain is one homogeneous region and may not cope well with RF inhomogeneity.

Lim and Pfefferbaum [19] and Attwood et al. [4] trace lines radially outward from the approximate center of the head in an approach similar to that proposed by Pannizzo et al. Attwood et al impose a smoothness criterion on the resulting contour to refine the location of the intracranial boundary. Zijdenbos et al. fit a surface to the voxel intensities under the contour produced by [19] in order to produce a local threshold [31, 32]. The radial line tracing algorithm is then repeated using the threshold to choose the voxels at the intracranial boundary. Zijdenbos' method is more robust than the others in the presence of RF inhomogeneity and partial volume effects. Clearly, none of these methods works for slices where the brain is separated into two or more disjoint regions.

3.2.3 Statistical Segmentation

Cline et al. segment the brain from MR images of the head using statistical classification [8]. To segment the brain, samples of brain voxels and non-brain voxels are interactively identified. Bivariate normal distributions, corresponding to the different tissue types in the PD-weighted and T2-weighted MR images, are fitted to the sampled intensities. All the image voxels are then classified according to where their intensities lie in the distributions. Finally, the results are smoothed to remove discontinuities in classified regions. This method requires user interaction, fails in the presence of RF inhomogeneity, and falsely classifies non-brain regions, such as the eyes, as brain.

Lachmann and Barillot isolate brain tissues in MRI slices using several texture analysis methods [17]. They use texture information to create an initial voxel classification. Cluster analysis and Bayesian relaxation is then used to refine the classification. Lachmann and Barillot do not show results for slices containing the eyes or the mouth. Bayesian relaxation-based techniques confuse these features with brain tissue [12, 13, 11].

Stringham et al. use a statistical relaxation method that incorporates gradient magnitude as well as voxel intensity for brain segmentation [28]. This segmentation algorithm is robust in the presence of RF inhomogeneity, but confuses tissues, such as the eyes, with brain tissue. Further, user interaction is required to seed the relaxation process.

Vinkin et al. use probabilistic hyperstacks for 2D image segmentation [30]. A hierarchical stack of decreasing resolution images is produced by progressively low pass filtering the original. Pixels in high resolution images (children) are probabilistically linked to voxels in low resolution images (parents) according to an objective criterion. Features of interest are identified in a low resolution image and propagated to all siblings, weighted by the linked probabilities. Thus, a statistical segmentation is produced. Identifying the features of interest and thresholding the final segmentation must be performed manually in this method.

3.2.4 Active Contour Models

Chakraborty et al. combine statistical segmentation and boundary detection to isolate features in MR images [6]. They first segment the images using a method similar to the ICM algorithm [12]. They then use a parametrically deformable shape model algorithm to find the boundary of interesting features in the segmented image [27]. The shape model algorithm modifies the shape of a pre-defined closed contour to match the shape of a region of interest. This method requires user interaction to seed the segmentation and provide an initial shape contour. Further, the segmentation may fail due to RF inhomogeneity.

Snell et al. use an active surface template to find the intracranial boundary in MRI volumes of the head [26]. The method is based on the active contour model algorithm, "Snakes", first proposed by Kass et al. [16]. Given an initial estimate of an object boundary, "Snakes" approaches the actual boundary by solving an energy minimization problem. In Snell's method, the user identifies points in MR image that correspond to points on a "standard" active surface template of the brain. Based in these points, the template is registered to the image. The "Snakes" algorithm is then used to attract the surface template to the intracranial boundary. Snell's method appears work better than all the other methods discussed herein. However, Snell used high resolution isotropic 3D MRI data to test his algorithm. Such MRI scans are generally not performed clinically. Still, the method requires user interaction and may fail for images that do not contain the entire brain.

4 Conclusion

We have introduced a novel, fully automatic intracranial boundary detection algorithm that has proven effective on clinical MRI data sets acquired from two different scanners. The algorithm consists of three incremental steps. The first step uses histogram analysis to localize the head, providing a region which must completely surround the brain. The second step uses nonlinear anisotropic diffusion and automatic thresholding to create a mask that isolates the brain within the detected head region. Using this mask as a seed, the final step employs an active contour model algorithm to detect the intracranial boundary. This algorithm has proven robust in the presence of RF inhomogeneity and partial volume effects.

Our MRI brain segmentation algorithm is in regular use for studies of Multiple Sclerosis lesions and for studies involving image compression, where the non-brain region is automatically given a higher compression ration than the brain region in the images.

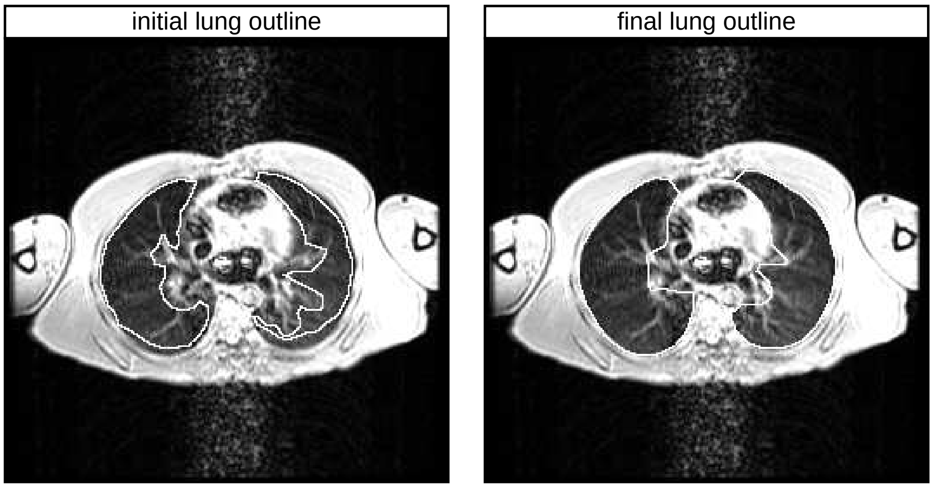

Figure 11: The initial lung boundary and the final automatically detected lung boundary.

4.1 Future Work – Lung Segmentation

The principles generated for automatic brain segmentation will be extended to the problem of lung segmentation for use in studies of lung diseases such as cystic fibrosis and emphysema, where the volume of the lungs is needed. A reliable consistent method for outlining the lungs is required for MR chest images. Early results with MR images are shown in Fig. 11 which shows the results of applying the algorithm to segment the lungs from an MR chest image.

The MR data are multiple spin echo images using a sequence with the repetition time of 2700 ms, and 4 images at 10, 20, 30 and 40 ms. Each image is of just one slice, 10 mm thick. The images are 256 x 256 pixels acquired axially, with 16 bits/pixel. The pixel size is 1.72 x 1.72 mm^2. In some echoes, the heart boundaries are dark, which leads to some misclassification of empty heart chambers as lung. This lung is of a healthy athlete, and does not show abnormalities.

Our initial work on segmenting the lung had to provide a way to remove background noise, which is high in these lung images because of flow artifacts from the heart. We removed the noise by subtracting the echo at 20 ms from the echo at 10ms, and clipping all negative values to zero (these were mostly the noise pixels in the 20 ms echo). This image contains noise from the 10 ms echo, plus reduced-intensity data points. We then subtracted this "noise" image from the original 10 ms echo to obtain an image which contained much less noise in the center. The remaining noise was then thresholded using the identical technique we used for segmenting background noise from the head images. The thresholded image was eroded and dilated as for the head mask, giving us an initial chest mask ready for use in step 2.

The problem of finding an initial lung mask was transformed into the initial brain mask problem, by using the chest mask to remove all the background, and then inverting the lung data, so that the lung tissue appeared bright. The initial lung mask was obtained by fitting a gaussian to the brightest intensity peak of the histogram obtained from anisotropic blurring of the masked, inverted data. We used all four echoes. The lung threshold at 2 standard deviations below the mean was used to provide the initial lung mask. After applying the threshold, we removed non-lung tissue such as the heart, using a morphology technique similar to that used for separating non-brain tissue from the brain; in this case, by specifying that the centroid of each lung is appropriately located within the chest. We found that this initial lung mask was very sensitive to the threshold, but as long as the threshold was high enough so that the initial lung contour was found inside the lung, the final step usually nullified this effect, as can be seen in Fig. 11.

Other research workers identifying lung volumes such as [15] using CT lung data, obtain lung outlines which resemble ours in that the major blood vessels are included within the lung. The problem of inclusion of blood vessels in the calculation of lung volume is a subject of ongoing research. Our method provides a consistent objective measure of lung volume which may be supplemented if necessary by a calculation of blood volume.

Acknowledgment

The authors wish to thank Dr. Don Paty, Dr. David Li, Dr. Alex MacKay, Andrew Riddehough and Keith Cover of the University of British Columbia MS/MRI Group and to Charlie DeCarli from the NIH, for providing MRI data on which the algorithms were tested. We also wish to thank Dr. Alex Mackay for his very helpful comments during the preparation of this manuscript.

References

[1] Luis Alvarez, Pierre-Louis Lions, and Jean-Michel Morel. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Number. Anal.*, 29(3):845–866, June 1992.

[2] T. Arden and J. Poon. *WiT User's Guide*. Logical Vision Ltd., Burnaby, Canada, October 1993. Version 4.1.

[3] M.S. Atkins, T. Zuk, B. Johnston, and T. Arden. Role of visual languages in developing image analysis algorithms. In *Proceedings of the IEEE Conference on Visual Languages*, pages 262–269, St. Louis, MO, October 1994. IEEE.

[4] C. I. Attwood, G. D. Sullivan, and K. D. Baker. Recognizing cortical sulci and gyri in MR images. In P. Mowforth, editor, *Proceedings of the British Machine Vision Conference*, 1991.

[5] Marijn E. Brummer, Russell M. Mersereau, Robert L. Eisner, and Richard R. J. Lewine. Automatic detection of brain contours in MRI data sets. *IEEE Transactions on Medical Imaging*, 12(2):153–166, June 1993.

[6] Amit Chakraborty, Lawrence H. Staib, and James S. Duncan. An integrated approach to boundary finding in medical images. In *Proceedings of the IEEE Workshop on Biomedical Image Analysis*, pages 13–22, Los Alamos, CA, June 1994. IEEE Computer Society Press.

[7] Vikram Chalana, Wendy Costa, and Yongmin Kim. Integrating region growing and edge detection using regularization. In *Proceddings of the SPIE Conference on Medical Imaging*. SPIE, 1995.

[8] Harvey E. Cline, William E. Lorensen, Ron Kikinis, and Ferenc Jolesz. Three-dimensional segmentation of MR images of the head using probability and connectivity. *Journal of Computer Assisted Tomography*, 14(6):1037–1045, November/December 1990.

[9] D.W. Paty et al. Interferon beta-1b in the treatment of multiple sclerosis: Final outcome of the randomized controlled trial. *Neurology*, 45:1277–1285, July 1995.

[10] Guido Gerig, Olaf Kübler, Ron Kikinis, and Ferenc A. Jolesz. Nonlinear anisotropic filtering of MRI data. *IEEE Transactions on Medical Imaging*, 11(2):221–232, June 1992.

[11] B. Johnston. 3D multispectral stochastic image segmentation. Master's thesis, The University of British Columbia, Computer Science Department, Vancouver, B.C., January 1994.

[12] B. Johnston, M.S. Atkins, and K.S. Booth. Partial volume segmentation in 3D of lesions and tissues in magnetic resonance images. In *Proceedings of the SPIE-Medical Imaging 1994*,

volume 2167, pages 28–39, Bellingham, WA, 1994. SPIE–The International Society for Optical Engineering.

[13] B. Johnston, M.S. Atkins, B. Mackiewich, and M. Anderson. Segmentation of multiple sclerosis lesions in intensity corrected multispectral MRI. *IEEE Transactions on Medical Imaging*, 15(2):154–169, April 1996.

[14] Micheline Kamber, Rajjan Shinghal, D. Louis Collins, Gordon S. Francis, and Alan C. Evans. Model-based 3D segmentation of multiple sclerosis lesions in magnetic resonance brain images. *IEEE Transactions on Medical Imaging*, pages 442–453, 1995.

[15] K. Kanazawa, N. Niki, H. Ohmatsu, and N. Moriyama. Computer assisted disgnosis of lung cancer using helical X-ray CT. *IEEE Workshop on Biomedical Image Analysis*, pages 261–267, June 1994.

[16] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, pages 321–331, 1988.

[17] F. Lachmann and C. Barillot. Brain tissue classification from MRI data by means of texture analysis. In *Proceedings of the SPIE: Medical Imaging VI: Image Processing*, volume 1652, pages 72–83, Chapel Hill, N.J., 1992.

[18] C. Li, D.B. Godlgof, and L.O. Hall. Knowledge-based classification and tissue labeling of MR images of human brain. *IEEE Transactions on Medical Imaging*, 12(4):740–750, December 1993.

[19] K. O. Lim and A. Pfefferbaum. Segmentation of MR brain images into cerebrospinal fluid spaces, white and gray matter. *Journal of Computer Assisted Tomography*, 13(4):588–593, July/August 1989.

[20] Blair Mackiewich. Intracranial boundary detection and radio frequency correction in magnetic resonance images. Master's thesis, Simon Fraser University, Computer Science Department, Burnaby, B.C., August 1995.

[21] N. Nordström. Biased anisotropic diffusion – a unified generalization and diffision aapproach to edge detection. *Image Vision Comput.*, 8(4):318–327, 1990.

[22] F. Pannizzo, M.J.B. Stallmeyer, J. Friedman, R.J. Jennis, J. Zabriskie, C. Pland, R. Zimmerman, J.P. Whalen, and P.T. Cahill. Quantitative MRI studies for assessment of multiple sclerosis. *Magnetic Resonance in Medicine*, 24:90–99, 1992.

[23] D.W. Paty and D.K.B. Li. Interferon beta-1b is effective in relapsing-remitting multiple sclerosis. ii. MRI analysis results of a multicenter, randomized, double-blind, placebo-controlled trial. *Neurology*, 43:662–667, April 1993.

[24] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990.

[25] Guillermo Sapiro and Allen Tannenbaum. Edge preserving geometric smoothing of MRI data. Technical report, University of Minnisota, April 1994. Dept. of Electrical Engineering.

[26] J. W. Snell, M. B. Merickel, J.M. Ortega, J. C. Goble, J. R. Brookeman, and N. F. Kassell. Segmentation of the brain from 3D MRI using a hierarchical active surface template. In *Proceedings of the SPIE Conference on Medical Imaging*. SPIE, 1994.

[27] L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformale models. *IEEE Trans, PAMI*, November 1992.

[28] R.R. Stringham, W.A. Barrett, and D.C. Taylor. Probabilistic segmentation using edge detection and region growing. In *Proceedings of the SPIE: Visualization in Biomedical Computing 1992*, volume 1808, pages 40–51, Chapel Hill, N.J., 1992.

[29] Hidetomo Suzuki and Jun ichiro Toriwaki. Automatic segmentation of head MRI images by knowledge guided thresholding. *Computerized Medical Imaging and Graphics*, 15(4):233–240, July-August 1991.

[30] Koen L. Vincken, André S. E. Koster, and Max A. Viergever. Probabilistic multiscale image segmentation - set-up and first results. In *Proceedings of the SPIE: Medical Imaging VI: Image Processing*, volume 1808, pages 63–77, Chapel Hill, N.J., 1992.

[31] Alex P. Zijdenbos. *MRI Segmentation and the Quantification of White Matter Lesions*. PhD thesis, Vanderbilt University, Electrical Engineering Department, Nashville, Tennessee, December 1994.

[32] Alex P. Zijdenbos, Benoit M. Dawant, Richard A. Margolin, and Andrew C. Palmer. Morphometric analysis of white matter lesions in MR images: Method and validation. *IEEE Transactions on Medical Imaging*, 13(4):716–724, December 1994.