

Segmentation and Feature Extraction Techniques, with Applications to MRI Head Studies

Edward A. Ashton, Michel J. Berg, Kevin J. Parker, Jeffrey Weisberg, Chang Wen Chen, Leena Ketonen

To obtain a three-dimensional reconstruction of the hippocampus from a volumetric MRI head study, it is necessary to separate that structure not only from the surrounding white matter, but also from contiguous areas of gray matter—the amygdala and cerebral cortex. At present it is necessary for a physician to manually segment the hippocampus on each slice of the volume to obtain such a reconstruction. This process is time consuming, and is subject to inter- and intra-operator variation as well as large discontinuities between slices. We propose a novel technique, making use of a combination of gray scale and edge-detection algorithms and some *a priori* knowledge, by which a computer may make an unsupervised identification of a given structure through a series of contiguous images. This technique is applicable even if the structure includes so-called false contours or missing contours. Applications include three-dimensional reconstruction of difficult-to-segment regions of the brain, and volumetric measurements of structures from series of two-dimensional images.

Key words: magnetic resonance imaging; three-dimensional; features extraction.

INTRODUCTION

Magnetic resonance imaging is able to produce high-resolution two dimensional images of the human brain. A series of such images may be used to reconstruct a three dimensional volume of the brain or one of its internal structures. A necessary first step to such a reconstruction, however, is a successful segmentation of the desired structure from the surrounding tissue. Several studies (1–3) have shown that it is possible to separate the various major classes of brain tissue—white matter, gray matter, and cerebrospinal fluid (CSF). However, many of the brain's structures are contiguous to other structures of the same tissue type. The hippocampus, for example, is composed of gray matter. It is partially surrounded by white matter, but is contiguous over much of its area to other gray matter structures.

When a physician makes an identification of the hippocampus on an MRI scan, two types of knowledge are used. The differentiation between white and gray matter provides some percentage of the boundary. General

knowledge of the size, location, and shape of the structure then enables the formation of a boundary across regions where no clear distinction of tissue types may be made. It is the goal of our algorithm to allow a computer access to these same two types of information. The first type, differentiation of tissue types, is a well-explored topic. Our algorithm makes use of a modification of the adaptive clustering algorithm of Pappas (4), for which we have derived an improved image model, to provide a smooth, connected tissue segmentation. It then makes use of the active contour model of Kass *et al.* (snakes) (5) to provide the second type of information—size, location, and shape of the desired feature. The combination of these two techniques is important, because in the past the Kass algorithm has generally been limited to a semi-automated technique, requiring manual manipulation on each segmented image, as in (6). Prior gray scale segmentation eliminates irrelevant image features, allowing the snake algorithm to operate efficiently without supervision.

The primary application of this algorithm is in the quick and accurate *in vivo* volume measurement of the hippocampus and amygdala. Jack *et al.* and others (7, 8) have shown that such a measurement may be an important aid in the diagnosis of intractable temporal lobe epilepsy. Jack *et al.* (9) and Kesslak *et al.* (10) have also shown that this data may predict dementia of the Alzheimer type. Current techniques for obtaining these volumes include manual tracing, thresholding, and random marking. These processes require extensive human interaction, and are time-consuming and subject to variation.

In this paper, our algorithm is tested for inter- and intra-operator variability. The results of our algorithm's identification of a phantom structure from a series of simulated MRI scans are then tested against known correct results. We also compare our algorithm's performance in identifying the hippocampus on an actual series of MRI brain scans against a physician's manual identification.

In the following sections we will address the problems of grayscale and contour segmentation. We will then present our experimental procedure and results.

GRAY SCALE SEGMENTATION

Bayesian Formulation

Gray scale segmentation is the process of separating a given image into regions of similar statistical behavior. Among the simplest and most obvious ways of accomplishing this is the well-known *k*-means algorithm, which is an optimal (in terms of minimum mean squared error) pixel-by-pixel scalar quantization of the image into *k* levels. However, *k*-means type algorithms impose no

MRM 33:670–677 (1995)

From the Department of Electrical Engineering, the Department of Neurology (M. J. B.), the Department of Radiology (L. K.), University of Rochester, Rochester, New York.

Address correspondence to: Edward Ashton, Ph.D., Department of Electrical Engineering, University of Rochester, Rochester, NY 14627.

Received September 2, 1994; revised January 17, 1995; accepted January 17, 1995.

0740-3194/95 \$3.00

Copyright © 1995 by Williams & Wilkins

All rights of reproduction in any form reserved.

spatial constraints, and therefore are easily fooled by additive noise. In other words, these algorithms fail to take into account any *a priori* knowledge of the connected nature of the segmented image.

The most common solution to this problem is to incorporate spatial information by modelling the image as a Markov random field (MRF) or Gibbs random field (GRF). The GRF is used more frequently, primarily because the MRF is defined in terms of local properties, making the derivation of a global joint distribution difficult (11). Such a derivation is critical for our application because, given a segmentation x and a data set y , we wish to maximize the *a posteriori* conditional probability $P(x|y)$. This is given by Bayes' theorem as

$$P(x|y) \propto P(y|x)P(x) \quad [1]$$

In this equation, $P(x)$ is the *a priori* probability of the segmentation, which is derived entirely from the random field model. $P(y|x)$ is a measure of how well the data matches the current segmentation. These two terms, then, provide conflicting constraints on the segmentation. $P(y|x)$ seeks the most probable solution based on conformity to Gaussian statistics, while $P(x)$ seeks to conform the global distribution to the prediction of the random field. The relative weighting of these terms will determine the characteristics of the final segmentation.

The Random Field Model

Assume that a random field x is defined over an $N \times N$ lattice (square array) of sites denoted by S , with the individual sites (pixels) denoted by $s_{1,1}, s_{1,2}, s_{1,3}, \dots, s_{N,N}$. For segmentation purposes, each site on x may take on any integer value from 1 to k , where k is the number of classes (tissue types) into which the image is being segmented. Let w be a realization of x . In order to define the random field, it is necessary first to define a neighborhood system G_s on the lattice S . The neighborhood system determines which pixels on the lattice will directly affect the classification of site s . In this work, a first order, or four-point, neighborhood system is used, so that only pixels which share a side with site s will be considered its neighbors.

A *clique* is defined as a set of sites on S such that all points are mutual neighbors. In a four-point neighborhood system, there are four two-point cliques, given by $(s_{n,m}, s_{n-1,m}), (s_{n,m}, s_{n+1,m}), (s_{n,m}, s_{n,m-1})$ and $(s_{n,m}, s_{n,m+1})$. There is also a single one-point clique, $(s_{n,m})$. However, as this adds no information, it is disregarded.

The random field x is an MRF with respect to G if

$$P(x = w) > 0 \text{ for all } w \quad [2]$$

and

$$P(X_s = x_s | X_r = x_r, r \neq s) \\ = P(X_s = x_s | X_r = x_r, r \in G_s) \quad [3]$$

i.e., all configurations must have a non-zero probability, and the conditional probability of a particular site must depend only on its immediate neighbors as defined by G . This model has been used in a number of applications,

including (12, 13). However, it has associated with it certain difficulties which make the Gibbs distribution more attractive. Primary among these is that the MRF is defined only in terms of local characteristics, making the calculation of a global joint distribution very difficult.

In contrast, the GRF provides a global model for an image, specifying a probability mass function of the form:

$$P(X = x) = e^{-U(x)} / Z \quad [4]$$

where $U(x)$ is the Gibbs potential, or Gibbs energy, defined by

$$U(x) = \sum_c V_c(x) \quad [5]$$

$V_c(x)$ is the clique potential, given by

$$V_c(x_i, x_j) = \begin{cases} -\beta & \text{if } x_i = x_j \\ +\beta & \text{otherwise} \end{cases} \quad [6]$$

and Z is a normalizing constant, called the partition function, given by

$$Z = \sum_w e^{-U(x=w)}. \quad [7]$$

Since Z is the sum of the numerator exponents over all possible vectors x , it is extremely difficult to compute. This intractability prevents realization of a GRF using direct calculation, thus requiring the use of stochastic relaxation methods.

Using the Gibbs formulation, it is possible to calculate the global probability of a given image configuration based on the clique potentials, which depend solely on the local properties of the pixels.

Finally, it should be noted that a unique GRF exists for every MRF as long as the Gibbs field is defined in terms of cliques on a neighborhood system (14). In other words, it is possible to uniquely specify an MRF by specifying Gibbs potentials. This means that nothing is lost through making use of the more mathematically tractable Gibbs model.

Dubes and Jain have shown that it is possible to achieve dramatic image reconstruction results using multi-parameter GRF modeling and algorithms such as the Gibbs sampler (11) if all parameters of the Gibbs distribution and degrading noise are known. However, there is currently no algorithmic method for deriving these parameters from an arbitrary data set. We therefore assume a simple *pairwise interaction model*. This model stipulates only that there should be a dependence in classification between neighboring pixels, and requires only a weighting constant, β , which determines the relative importance of two terms in the Bayesian formulation. Unfortunately, β must be optimized experimentally for a class of images. In this work, $\beta = 1.0$ was used for all segmentations.

The Image Model

Returning to Eq. [1], we now need a functional form for $P(y|x)$. This will be dependent upon our model for the image data. Chang *et al.* (15) have shown that a satisfactory segmentation of a volumetric MRI head study may

be achieved by modeling each tissue type in the volume as a slowly varying intensity μ_x . The observed data in a given region is then given by the original intensity corrupted by spatially invariant white Gaussian noise. Our experiments indicate, however, that the assumption of spatially invariant noise is not accurate. It is true that the data acquisition system (imaging and postprocessing) will degrade the image in a way that may be accurately modeled by additive Gaussian noise with a given standard deviation, σ , which will be particular to that imaging system. Allowing this to be the sole model for image noise, however, ignores the fact that the tissue we wish to have classified as white matter, for instance, is not perfectly uniform with regard to proton density and magnetic susceptibility. Small local inhomogeneities will produce irregularity in the image, which will be *tissue dependent*. Additionally, nonuniformity of coil sensitivity will add a *spatially dependent* signal variation. We model this degradation as additive Gaussian noise with local standard deviation σ_i , where i is the tissue class and s is the lattice site. The overall noise in the system, therefore, is Gaussian, with variance given by

$$\sigma_{x_s}^2 = \sigma_i^2 + \sigma^2. \quad [8]$$

Under this assumption, the conditional probability may be written as

$$P(y|x) \propto \prod_s \left(\frac{1}{\sigma_{x_s}} \right) \exp - \sum_s \left[\frac{1}{2\sigma_{x_s}^2} (y_s - \mu_{x_s})^2 \right] \quad [9]$$

where σ_{x_s} is the overall noise variance for class x_s .

Substituting into Eq. [1], and summing the exponentials, we obtain

$$P(x|y) \propto \prod_s \left(\frac{1}{\sigma_{x_s}} \right) \exp - \left[\sum_s \left[\frac{1}{2\sigma_{x_s}^2} (y_s - \mu_{x_s})^2 \right] + \sum_c V_c(x) \right] \quad [10]$$

We wish to maximize this function, or, equivalently, to maximize its natural logarithm, given by

$$\ln(P(x|y)) \propto - \sum_s \left[\ln(\sigma_{x_s}) + \frac{1}{2\sigma_{x_s}^2} (y_s - \mu_{x_s})^2 \right] - \sum_c V_c(x). \quad [11]$$

This is achieved in an iterative fashion, using Besag's iterated conditional modes (ICM) (16). We consider the entire volume to be a single three-dimensional GRF with a lower sampling frequency in the z -direction (between frames) than in the x and y directions (within a single frame), and employ a shrinking parameter estimation window, as outlined in (4, 15).

The application of this improved image model provides significantly superior results in two areas. First, in prior formulations, σ was an operator-determined parameter. Because it is difficult to know the noise statistics of an image *a priori*, this parameter needed to be experimentally determined for each image in order to ensure optimality. Second, and more importantly, allowing σ to vary with tissue type and region produces significantly more accurate segmentation, particularly in the case of

images which have been degraded by field inhomogeneities. This point is demonstrated in Fig. 1 and 2.

CONTOUR SEGMENTATION

Snakes

Once tissue segmentation has been accomplished, our algorithm makes use of the active contour model of Kass *et al.* (snakes) (5) to provide the second required type of information: the general size, shape and location of the desired structure. The snake algorithm models the segmentation problem as a physical system. Our implementation begins with an operator-provided initial rough estimate of the location and shape of the hippocampus, i.e., a rough outline of the target on a single slice of the volume. This closed loop of pixels may be thought of as a string of ball bearings, each of which is connected to its nearest neighbors by a spring with spring constant c_1 . The image may then be thought of as an inverted surface, with areas of high intensity corresponding to valleys, and areas of low intensity corresponding to peaks. If we then edge-map the image using a Marr-Hildreth operator, the result will be a largely flat surface, with deep valleys corresponding to edges. The string of ball bearings is then laid down on the surface, and the bearings "roll" into nearby valleys (edges), while stretching smoothly across the boundary points where edges are not available. When the contour reaches a stable point of minimum energy, the target should be outlined by a smooth, closed contour.

Once the target has been positively identified in a single slice, that information may be passed to the adjacent slices, where it serves as the initialization for the new snake. For this particular application of this process there are significant difficulties in the passing of contours from slice to slice, owing largely to the complex shape of the hippocampus. Further discussion of these difficulties and their solution is given below.

This intuitive explanation is, of course, somewhat oversimplified. More rigorously, the overall energy at each pixel in the snake contour is given by

$$E = e_s + e_i + c_1 e_{3d} \quad [12]$$

In this equation, e_s is the value of site s on the edge map. The internal spline energy, e_i , is a measure of first and




Fig. 1. Fixed σ adaptive clustering applied to a degraded MRI brain scan. Note overdevelopment of gray matter on the right and underdevelopment on the left, causing partial loss of both hippocampi.




Fig. 2. Modified adaptive clustering applied to a degraded MRI brain scan. Note symmetrical development of the cortex and of both hippocampi.

second order discontinuities in the contour at a given point, given by

$$e_1 = (c_2|v_x(s)|^2 + c_3|v_{xx}(s)|^2)/2 \quad [13]$$

Here the first term, $v_x(s)$ limits first order discontinuity (point separation), i.e., the distance from the point under examination to its next nearest neighbors. The second term, $v_{xx}(s)$ limits second order discontinuity (bending). This term is more difficult to measure. In this work it is approximated by $1/\rho$, where ρ is the local (over 5 points) radius of curvature.

The final term in Eq. [12] e_{3d} , is a measure of how far a pixel has moved from its position in the previous frame. This term is not included in the original snake formulation, and is intended to provide continuity and smoothness between frames. This is most important for accurate 3-d reconstruction, as will be shown in our experimental results.

Like the gray scale segmentation algorithm, the snake algorithm requires parameters which cannot be determined theoretically. In this case these parameters are the weighting constants c_1, c_2 , and c_3 . While this might appear to be a serious drawback, our experiments (as well as the many applications of snakes found in the segmentation literature) indicate that these parameters can be optimized a single time for application to a broad class of images. In this work, values of $c_1 = 0.1$, $c_2 = 1.0$, and $c_3 = 0.05$ were used.

The Kass algorithm has been used previously in three-dimensional segmentation problems (6, 17, 18). It has not, however, been used in conjunction with gray scale segmentation. The specifics of this problem make this a significant innovation, because the low contrast between white and gray matter and relatively high levels of noise found in MRI images, particularly those with a small slice thickness, make the number of local minima in the original images so great that the snake algorithm becomes virtually useless with gradient-ascent type energy minimization. The problem is compounded by partial volume effects, which frequently cause tissues to appear to fade into one another with no discernible boundary and render standard noise reduction techniques ineffective. This situation is illustrated in Fig. 3.

One possible limitation of this two step approach is that errors in the initial tissue segmentation could produce errors in the resultant contours. These errors might




Fig. 3. (a) Physician's manual hippocampal identification. (b) Identification by snake algorithm on unsegmented MRI data (initial contour in white, final contour in black.) (c) Identification by snake algorithm on segmented MRI data (initial contour in white, final contour in black.) Note that the contour in (b) has moved very little, due to a lack of discernible edges.

then propagate through the volume. However, this will only occur under two conditions—first, if a small error is present in the same position on several successive slices, or second, if a very large error pulls the snake far from its correct position. The first case is the more likely of the two. Fortunately, small displacements are unlikely to affect the positioning of the snake on successive correctly segmented frames. The second case could be more damaging, in that a very large placement error could propagate throughout the volume. However, such a large segmentation error is extremely unlikely, and has not been observed in any of our experiments on either simulated or real data.

It is possible that under certain circumstances this algorithm could be effective on unsegmented images in conjunction with a broad feature-space search algorithm such as simulated annealing (19). However, the huge computational requirements of such algorithms far outweigh the benefits of avoiding the segmentation step.

The snake algorithm has also generally been used either for applications where the desired feature changes little between slices, as in (17), or in conjunction with interactive registration, as in (6). Extraction of the hippocampus require displacement compensation. We have derived a simple and efficient system for accomplishing this without human interaction.

Motion Estimation

The thickness of an MRI slice (3 mm in the series we examined) in comparison to the pixel size (0.7 mm) and the complex shape of the hippocampus make simple passing of contours from one slice to the next, as in (17),

problematic. In order to obtain maximum structural information we use coronal sections, with the result that the major axis of the hippocampus is not perpendicular to the image planes. Because of this, in some cases the location of the hippocampus changes by as much as half the object size from one slice to the next. The large errors associated with misaligned initial contours are illustrated in Fig. 4e and 4f. We therefore rely on a block motion estimator to obtain a motion vector between slices, under the assumption of global motion of a 64×64 pixel block around the hippocampus (20).

The motion vector is obtained by first finding the maximum point on a cross-correlation of this block from slice n with the corresponding block from slice $n - 1$. Under the assumption of global motion, the motion is then given by the vector from the origin of the correlation space to the location of its maximum.

Once the motion vector has been obtained, the location of each pixel of the previous contour is adjusted accordingly. This adjustment provides dramatically more accurate results, as illustrated in Fig. 5.

RESULTS AND DISCUSSION

In the first set of experiments, the sensitivity of the snake algorithm to the initial contour was studied. There are two cases for which the initial contour sensitivity is relevant. The first is the response to the operator-entered contour, which will determine the inter- and intra-operator variability. Because one of our prime motivations is to limit such variability, we expect this sensitivity to be




Fig. 4. Results of contour segmentation on a single slice of segmented MRI data with various initial contours. Initial contours are shown in white, Final contours are shown in black.




Fig. 5. (a) Segmented hippocampal region of slice 7 of 13 from MRI series, with final contour shown in black. (b) Hippocampal region of slice 8. Initial contour, passed without motion estimation from slice 7, shown in white. Final contour shown in black. (c) Hippocampal region of slice 8. Initial contour, passed with motion estimation from slice 7, shown in white. Final contour shown in black.

very small. The second case involves the response to contours passed from a previous slice to the current one. This will not affect variability. It will, however, have an impact on the overall accuracy of the system.

In order to investigate the first case, we provided the snake algorithm with multiple initial contours for the segmentation of a single slice. The first contour provided a reasonably accurate estimate of object size and shape. Successive initial contours degenerated to a level of accuracy well below that which a reasonable operator could be expected to provide. A sampling of the results, with initial contours shown in white and final results in black, is given in Fig. 4. Note that available edges are successfully located in the resulting segmentation until Fig. 4e, where the initial contour clearly excludes a significant portion of the hippocampus, and Fig. 4f, where the initial contour clearly includes the amygdala. Discounting such gross operator errors, this experiment indicates considerable robustness in the algorithm. Insensitivity to initial input is one of the primary advantages of this technique, and is a result of the prior segmentation which eliminates spurious edges which might draw the snake away from the correct result, thus making manual manipulation of the contours, as in (6), unnecessary.

Case two was investigated, and a statistical measure of operator sensitivity was obtained, by providing our algorithm with multiple input contours for the segmentation of an MRI brain sequence which will be described in more detail below. This resulted in a coefficient of variance of 3.38% in the calculation of the area of the initial slice. This declined to 2.85% on the succeeding slice,

and the coefficient of variance for the overall calculation of the hippocampal volume was 2.70%.

Our second set of experiments involved a 28-slice simulated MRI phantom. This phantom mimics the statistics of the hippocampal region, with a background of gray level 190, an object of gray level 155, and a statistically identical region which must be bridged by the active contour algorithm. As in the actual hippocampus, the major axis of the phantom is skew to the image planes. The center slice of this phantom is given in Fig. 6. Increasing levels of noise were added to the images in the phantom, and our algorithm's calculation of the volume of the selected structure at each noise level was compared to the known correct results. An identical initial contour was used to begin processing at the center slice for each noise level. As an additional test of robustness, the series with $\sigma = 12.0$ was re-run with an initial contour at the first slice. Results were identical to the series begun at the center slice. Results for a sample slice, with the initial contour shown, are given in Fig. 7. Note that although the shown image is not grayscale segmented, all calculations were performed on segmented images, for reasons cited above. A plot of calculated object area versus true area at each slice and for each noise level is given in Fig. 8. Percentage errors on the simulations, where the correct results were precisely known, were 0.5% for $\sigma = 5.0$, and 2.7% for the series which most closely matched the actual statistics of the MRI scans used in later experiments ($\sigma = 12.0$). Note that there is no significant error until $\sigma = 17.5$, which is beyond the upper range of the noise found in our real MRI data. Object area was used as a comparison metric because a primary application of this algorithm is the accurate calculation of hippocampal volumes. Additionally, variation in volume calculation is the primary metric of accuracy cited in the literature in the evaluation of manual techniques which this algorithm is intended to replace. This evaluation technique is clearly limited, in that a




Fig. 6. One frame of our MRI phantom.




Fig. 7. Our contour segmentation routine applied to one slice of our MRI phantom. Initial contour is shown in white.




Fig. 8. True versus calculated areas for noise standard deviation ranging from 5.0 to 20.0. Noise standard deviations in actual MRI data ranged from 6.0 to 15.0.

completely erroneous segmentation could, through chance, result in an accurate volume calculation. However, visual inspection of the segmentations in our experiments confirms the apparent accuracy of the algorithm.

Our next experiment was conducted on a 13-slice section of a 60-slice coronal SPGR (spoiled grass) MRI head sequence with 35/5/30/2 (repetition time/echo time/flip angle/excitations). Slice thickness was 3.0 mm. Imaging was performed using a General Electric Signa 1.5 Tesla superconducting system. Note that, with regards to motion estimation and contour segmentation, this is a

"worse case," in that studies intended for hippocampal evaluation are frequently performed with a slice thickness of 1.5 mm. Clearly this would limit changes in hippocampal shape and displacement between slices. Manual identifications of the hippocampus were made on each of the slices by a physician. As a comparison, one manually identified cross-section was then presented as an initial contour to our algorithm, which calculated a three-dimensional reconstruction of the hippocampus. A plot of manually versus automatically calculated areas is given in Fig. 9. Calculation time for this series, running on a Sun SPARC 10 workstation, was approximately 45 min. Three representative slices taken from the series, including original images with manual identifications and segmentations with automatic identifications, are given in Fig. 10. This figure shows the close correlation between the manual identifications and the automatic identifications. It should be noted that some variation is accounted for by the manual oversimplification of the shape of the hippocampus.

Finally, three-dimensional reconstructions of the left and right hippocampus as identified by our algorithm, and of the left hippocampus as identified manually, were produced using the Advanced Visual Systems (AVS) volume rendering system. These reconstructions are given in Figs. 11 and 12.

Jack *et al.* (21) have reported a coefficient of variance between operators of 1.9% for a combined thresholding/tracing technique. While this is somewhat smaller than our results given above, it should be noted that we deliberately provided our algorithm with what we considered to be the most widely variant initial contours which



Fig. 9. Manual versus automatic calculated hippocampal areas.




Fig. 10. (a) Slice 2 of 13 from MRI series, with manual hippocampal identification. (b) Segmentation of (a), with automatic hippocampal identification. (c) and (d); slice 6 of 13. (e) and (f); slice 10 of 13.



Fig. 11. Reconstruction of automatically extracted left and right hippocampi.

could reasonably be expected. Because our algorithm is deterministic given an initial contour, and because different initial contours tend to converge as they proceed through the volume (as demonstrated by the decline in variance from the first slice to the final result) a reasonably well trained and consistent operator should be able to achieve very consistent results.

CONCLUSIONS

In trials on both simulations and actual data, our algorithm appears to provide accurate and reproducible results, with reasonable computational complexity. Reproducibility in this context indicates that the algorithm is insensitive to the precise shape and placement of the initial contour, as demonstrated in the previous section.



Fig. 12. Reconstruction of manually extracted left hippocampus.

The primary weakness of our algorithm is demonstrated in Fig. 10f. It is clear in this figure that the segmentation algorithm has grouped a portion of the amygdala with the hippocampus. This is a result of the series beginning with an initial contour at slice 6, which does not contain any portion of the amygdala. The algorithm, therefore, is unaware that a boundary between the two structures at that point should exist. Despite this problem, however, the reconstructions given in Figs. 11 and 12 would tend to suggest that, in some instances, the automatic identification may be more accurate. In particular, the sharp discontinuities in the reconstruction of the manual identifications given in Fig. 12 should be noted. These are the result of the physician working with two-dimensional data only, whereas our algorithm is able to consider the three-dimensional volume as a whole.

Future work in this area will include experimentation with larger scanned-phantom data sets and more extensive human studies, as well as extending this technique to other imaging modalities, including CT and ultrasound.

REFERENCES

1. H. E. Cline, W. E. Lorensen, R. Kikinis, F. Jolesz, Three-dimensional segmentation of MR images of the head using probability and connectivity, *J. Comput. Assist. Tomogr.* **14**, 1037-1045 (1990).
2. M. Bomans, K-L. Hoehne, U. Tiede, M. Riemer, 3-D segmentation of MR images of the head for 3-D display, *IEEE Trans. Med. Imaging*, **9**, 177-183 (1990).
3. E. Ashton, M. Berg, K. Parker, C. W. Chen, J. Weisberg, L. Ketonen, Segmentation and features extraction techniques, with applications to biomedical images, in "International Conference on Image Processing, Austin, Texas, 1994.
4. T. N. Pappas, An adaptive clustering algorithm for image segmentation, *IEEE Trans. Signal Processing*, **SP-40**, 901-914 (1992).
5. M. Kass, A. Witkin, D. Terzopoulos, Snakes: active contour models, *Int. J. Comput. Vision* **1**, 321-331 (1988).
6. I. Carlbom, D. Terzopoulos, K. Harris, Computer assisted registration, segmentation, and 3d reconstruction from images of neuronal tissue sections, *IEEE Trans. Med. Imaging*, **13**, 351-362 (1994).
7. C. R. Jack, F. W. Sharbrough, C. K. Twomey, G. D. Cascino, K. Hirschorn, W. Marsh, A. Zinsmeister, B. Scheithauer, Temporal lobe seizures: lateralization with MR volume measurements of the hippocampal formation, *Neuroradiology* **175**, 423-429 (1990).
8. F. Cendes, F. Andermann, P. Gloor, A. Evans, M. Gotman, C. Watson, D. Melanson, A. Olivier, T. Peters, J. Cendes, G. Leroux, MRI volumetric measurement of amygdala and hippocampus in temporal lobe epilepsy, *Neurology* **43**, 719-725 (1993).
9. C. R. Jack, R. Petersen, P. O'Brien, E. Tangalos, MR-based hippocampal volumetry in the diagnosis of Alzheimer's disease, *Neurology* **42**, 183-188 (1992).
10. J. Kesslak, O. Nalcioglu, C. Cotman, Quantification of magnetic resonance scans for hippocampal and parahippocampal atrophy in Alzheimer's disease, *Neurology* **41**, 51-54 (1991).
11. R. C. Dubes, A. K. Jain, Random field models in image analysis, *J. Appl. Stat.* **16**, 131-164 (1989).
12. C. Bouman, B. Liu, Multiple resolution segmentation of textured images, *IEEE Trans. Pattern Anal. Mach. Intel.* **13**, 99-113 (1991).
13. E. J. M. Rignot, R. Chellappa, Algorithms for segmentation of complex-amplitude SAR data, *NASA Tech Briefs*, **17**, 28-30 (1993).
14. J. Besag, Spatial interaction and the statistical analysis of lattice systems, *J. R. Stat. Soc.* **36**, 192-326 (1974).
15. M. M. Chang, M. I. Sezan, A. M. Tekalp, Bayesian segmentation of MR images using 3-D Gibbsian priors, in SPIE Conference on Image and Video Processing, San Jose, CA, (1993).
16. J. Besag, On the statistical analysis of dirty pictures, *J. R. Stat. Soc.* **48**, 259-302 (1986).
17. I. Herlin, N. Ayache, Features extraction and analysis methods for sequences of ultrasound images, *Image Vision Comput.* **10**, 673-682 (1992).
18. L. Cohen, I. Cohen, Deformable models for 3d medical images using finite elements and balloons, in "Proc., Conference on Computer Vision and Pattern Recognition '92", pp. 592-598.
19. S. Geman, D. Geman, Stochastic relaxation, gibbs distribution and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intel.* **6**, 721-741 (1984).
20. R. Manduchi, G. A. Mian, Accuracy analysis for correlation-based image registration algorithms, in "International Symposium on Circuits and Systems '93," pp. 834-837.
21. C. R. Jack, M. D. Bentley, C. K. Twomey, A. Zinsmeister, MR imaging based volume measurements of the hippocampal formation and anterior temporal lobe: validation studies, *Neuroradiology* **176**, 205-209 (1990).