Evaluating a segmentation algorithm is meaningful for a particular context only. A way to estimate the results of a method is to consider a complete system realizing a specific task.

N. Ayache, 1996

Chapter 2

Automatic Segmentation of the Brain from MRI–T1 Data

Abstract

A method for automatic segmentation of the cerebral cortex from T1–weighted MR image data has been developed. The starting point is a supervised segmentation technique which has proven highly effective and accurate for visualization purposes. A technique is proposed to automate the required user interaction, *i.e.*, defining a seed point and a threshold range. The thresholds are detected in a region growing process and are defined by the linkages of the brain to other tissues. The method is first evaluated on three simulated datasets by comparing the automated segmentation with the original distributions. The second evaluation is on a total of 30 patient datasets, by comparing the automated segmentations with supervised segmentations carried out by a neuro-anatomy expert for visualization purposes. The comparison between two binary segmentations is performed both quantitatively and qualitatively. The automated segmentations are found to be both accurate and reproducible. The proposed method can, consequently, be used as a default segmentation for visualization purposes in routine clinical procedures.

2.1 Introduction

Segmentation is the grouping of similar voxels into coherent volumetric structures which may be required for quantitative volumetric analysis, for morphological analysis, or for visualization purposes. An important area is the segmentation of the cerebral cortex, *e.g.*, for quantitative analysis of different brain structures (Kohn et al. 1991, Collins et al. 1992, Kikinis et al. 1992), labeling of cortical structures (Sandor and Leahy 1995), localization of electrodes on the surface of the brain (Van den Elsen 1993), establishing abnormal gyration (Shenton et al. 1992), and providing an anatomical framework for functional studies of the cortex (see Chapter 7).

The most widely applied segmentation technique is manual segmentation, which has several disadvantages: *i)* it generally requires a high level of expertise, *ii)* it is time and labor consuming, and *iii)* it is subjective and therefore not reproducible. Studies investigating inter and intra-patient variations in cerebral function or anatomy have repeatedly shown these shortcomings, which explains the demand for automated techniques (Gerig et al. 1992, Kikinis et al. 1992, Shenton et al. 1992) . We focus on a supervised method originally proposed by Höhne and Hanson (1992) that is based on region growing and morphological operations. The aim of this chapter is to automate the remaining user interaction in this established method, thereby obtaining a fully automated segmentation of the brain and evaluate the segmentation results for visualization and volume estimation purposes. We will first give an overview of existing region growing based techniques in the context of MRI brain segmentation followed by a more specific introduction to our segmentation approach.

2.1.1 Brain segmentation from MR images

In this chapter we deal with segmentation of MR brain images. In general, simple segmentation techniques fail dramatically because of two disturbing factors: *i)* The MR acquisition suffers from image gradients and/or RF coil inhomogeneities, and *ii)* different anatomical structures are often linked to each other owing to partial volume effects, noise, imaging artefacts, or by connecting tissue (the optic nerves, blood vessels, etc). The former factor can be compensated for to some extent by reducing the non-uniformity of the intensities. To disconnect unwanted anatomical links, the use of morphological operations (erosion/dilation) is rather effective (Höhne and Hanson 1992).

MRI segmentation techniques can be roughly divided into: clustering methods, region based methods, and edge based methods. For a review on segmentation of MRI data we refer to Clarke et al. (1995) and Niessen (1997). Here we focus on the region based methods with region growing as main contributor.

In region growing based techniques a segment is formed by selecting a seed pixel and continuously adding neighboring pixels that meet certain—generally simple—requirements, *e.g.*, a threshold. Region growing has the advantage over thresholding that the resulting objects will be spatially connected.

In a classical paper, Zucker (1976) gave an overview of several merge-criteria for region growing that were proposed in the early literature on this technique. A more recent comparison of different approaches was given by Jiang et al. (1993). Several techniques have been reported for the use of region growing in medical images. Seeded region growing (Adams and Bischof 1994, Justice et al. 1997) makes use of statistical analysis on specified voxels to steer the region growing. Chang and Li (1994) used a region growing process where the thresholds are dynamically and automatically computed using feature histogram analysis.

In general, region growing methods suffer from connections to other tissues. Combination of region growing with approaches from mathematical morphology (erosion/ dilation) has proven rather effective (Höhne and Hanson 1992). Others have used a region growing method combined with a connectivity threshold (Barillot et al. 1991) or interactive, semi-automatic leak removal (Sekiguchi et al. 1994). A slightly different approach to region growing was reported by Cline et al. (1987) who performed a surface extraction from a user specified seed point on the surface. Weak connecting bridges of a few pixels in width were eliminated.

Although several of these approaches can be considered powerful tools for segmentation, none of them is fully automatic and interaction is required to perform a segmentation task. This usually requires a medical expert to actively control the segmentation process and interactively correct the result(s). One of the most attractive approaches is the interactive segmentation technique of Höhne and Hanson (1992)*. It is based on region growing and morphological operations (especially erosion and geodesic dilation, *i.e.*, dilation within a mask volume), and requires a user defined seed point and a threshold range.

For the analysis of medical volume data, we have used the Höhne approach incorporated into ANALYZE^{TM} (Robb and Hanson 1996) to segment, *e.g.*, the aorta from CTA (Balm et al. 1997), and the brain from MRI (Chapters 3-7). The method proved to be very practical for reasons of speed, simplicity, and intuitiveness. The technique first extracts a base volume from a dataset using an operator defined threshold (range), thereby encapsulating the required object. A series of erosions is applied to remove undesirable linkages with other structures. Subsequent connection to a user defined seed point and geodesic dilation (to counteract the erosion) results in a segmentation of the required object. A considerable part of the process can be automated, thereby achieving an enormous decrease in time and labor consumption compared to manual segmentation. However, the process requires little, yet vital interaction from an experienced user, *i.e.*, a seed point must be selected and a threshold range must be set. Especially the latter is crucial to the success of the process (Zucker 1976, Chang and Li 1994), which renders the segmentation results subjective and not reproducible.

In this chapter we propose a method called CACTUS (Completely Automatic

---

*We will refer to this technique as the Höhne approach

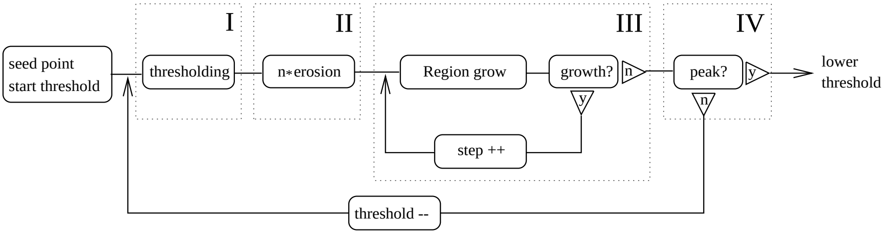

Figure 2.1 *Schematic of the algorithm for detection of the lower threshold T_{low}.*

Computer Technique for Unsupervised Segmentation) to automate the required user interaction of the Höhne approach for segmentation of the brain from MRI. We evaluated the method for volume estimation and visualization purposes. In Section 2.2 the method is described and the phantom and patient data used for evaluation of the method are presented. Validation consisted of the calculation of difference and similarity measures between CACTUS results and reference data. Furthermore, neuroanatomy experts qualitatively evaluated the CACTUS results and their remarks have been used throughout the results and discussion sections.

2.2 Methods and materials

2.2.1 CACTUS

The segmentation method CACTUS first selects a starting threshold value (T_{start}) and a seed point. For segmentation of the brain from MR images this selection is (in most cases) a trivial task.

T_{start} is determined by calculating the second peak in the histogram of grey values of the MRI data (the first and highest peak of the histogram refers to background voxels). A volume of possible seed points is defined by thresholding with T_{start} and applying two erosions with a 6-voxel structuring element. A spherical search from the middle of the dataset is initiated and the first encountered point in the seed volume is marked as the seed point.

The threshold range detection has T_{start} and the seed point as input to first detect the lower threshold (T_{low}) and then the upper threshold (T_{up}) in a similar way. To detect T_{low} (T_{up}) the threshold setting decreases (increases) from T_{start}. The actual algorithm for the detection of T_{low} (see Figure 2.1) and T_{up} consists of four parts executed at each threshold setting: I) thresholding, II) erosion, III) region growing, and IV) peak detection.

In part I the active threshold setting is applied to extract a base volume from the grey value data. This base volume is eroded twice (part II) with a 6-voxel structuring element to break small linkages. A region growing process (part III) is initiated from the seed point (or seed volume, see later) and grows one layer of voxels at each itera-

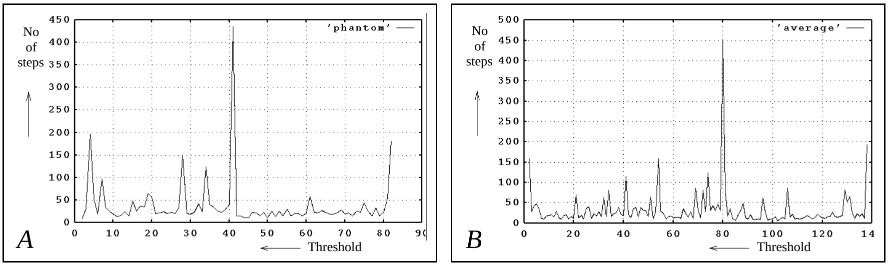

Figure 2.2 *Detection of T_{low}. The number of growing steps with a decreasing threshold value for a phantom (A) and a typical patient (B) dataset are presented. For both cases the peak was detected which resulted in a T_{low} of 42 for the phantom and 81 for the patient data. For results of the corresponding segmentations see Figures 2.4 and 2.5.*

tion of the region growing process with the base volume as a mask (*i.e.*, no growing of the volume outside the eroded base volume is permitted). Each iteration or addition of one layer of voxels will be referred to as one growing step. Eventually the growing stops for that threshold setting yielding the corresponding region growing volume. The total number of growing steps is evaluated in the peak detection part (IV) where it is compared with the growing steps at the preceding threshold settings (a number of five sufficed). When no peak is detected, the threshold setting is decreased (for T_{low}) or increased (for T_{up}) by a fixed number (a step size of 1 is generally recommended) and the region growing process is repeated using the current region growing volume as the seed volume.

To illustrate the peak detection process we will focus on the detection of T_{low}. Figure 2.2 shows the typical results of the region growing process for a phantom dataset and a patient dataset (see Section 2.2.2) starting from T_{start}. Normally, the process terminates once a peak is detected, but for illustration purposes peak detection was turned off and the process continues until the threshold setting has reached the value one. For both the phantom and the patient study the peaks (*resp.* at values 41 and 80) are easily detectable.

A sharp increase in the number of growing steps signals the inclusion of a new segment through a linkage between two structures; the previous threshold setting is selected as T_{low}. This is the critical decision for the automated threshold detection. The peak detection is performed twice to detect T_{low} and T_{up}. Crucial in the process is the use of erosions with each threshold before growth is permitted. This effectively removes small linkages caused by anatomy and the partial volume effect and thus ensures a dramatic increase in the number of iterations growth per threshold when a new segment is incorporated.

We have to note that T_{up} is detected by going upward from T_{start} *and* using the T_{low}. Furthermore, the parameters of CACTUS for the region growing, erosion and dilation processes were set to two erosions, three (is number of erosions plus one) geodesic dilations and a 6-voxel structuring element. This was based on information from the ANALYZE^{TM} reference manual and our experiences with segmentation procedures.

The final segmentation is performed by applying the seed point and the threshold range [T_{low} - T_{up}] for segmentation of the brain using the Höhne approach.

2.2.2 Phantom and patient data

Simulations of MRI data from a phantom (Cocosco et al. 1997) are available at the McConnell Brain Imaging Centre at the Montreal Neurological Institute[†]. Each voxel of the phantom can contain multiple tissue types, thus providing a fuzzy (*i.e.*, partial volume voxel) model. To simplify the comparison between segmentation results we binarized the Montreal phantom by setting all voxels with a probability for brain (white and grey matter) higher than or equal to 50 % to one and other voxels to zero.

The Montreal phantom dataset offers the possibility to simulate differences in pulse sequence, slice thickness, noise (options: 0, 3, 5, 7, and 9% of the standard deviation of the mean intensity of white matter) and non-uniformity (RF) (options: 0, 20, and 40% resulting in a (varying) RF field of: 1, 0.9-1.1, and 0.8-1.2). Three MRI–T1 datasets with a slice thickness of 1 mm were acquired using; *i)* 3% noise and an RF of 0% (phantom code: 3n00RF), *ii)* 3% noise and an RF of 20% (3n20RF), and *iii)* 9% noise and an RF of 40% (9n40RF). Note: 3% noise and an RF of 20% are denoted typical, while 9% noise and an RF of 40% are extreme (Cocosco et al. 1997).

The patient data are T1-weighted 3D gradient-echo MR scans (voxel size $1 \times 1 \times 1.2$ mm) of 30 patient. These patients are cases from the Department of Child Psychiatry at the University Hospital Utrecht diagnosed with the Gilles de la Tourette Syndrome, Attention-Deficit Hyperactivity Disorder, autistic behavior, and/or Obsessive Compulsive Disorder. Upon initial screening, no gross abnormalities were detected in the MRI data.

2.2.3 Evaluation methodology

The evaluation of the phantom studies can be done in an objective fashion, since the segmentation obtained with the above method can be compared with the binarized distribution. This reference set is called the *binarized phantom* segmentation.

With patient data an objective reference is generally not available. However, we had already segmented the patient datasets for visualization purposes in another project (see Chapter 6) and we decided to use these segmentation results as reference sets for comparison. The reference segmentations had been performed using the supervised method of Höhne as explained in Section 2.1 and all work was done by an experienced operator who had required on average 15 minutes to perform an adequate

---

[†] For more details see web site http://www.bic.mni.mcgill.ca/brainweb

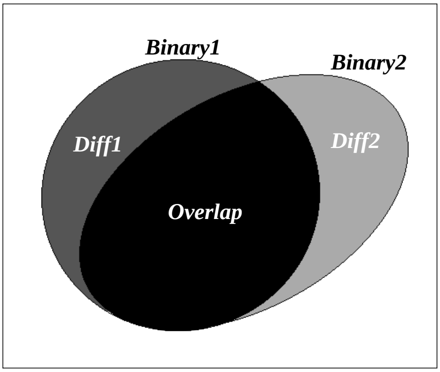

$$AC = \frac{Diff1}{Overlap} \quad (1)$$

$$MC = \frac{Diff2}{Overlap} \quad (2)$$

$$S = \frac{2 \text{ x } Overlap}{Binary1 + Binary2} \quad (3)$$

Figure 2.3 *Difference and similarity measures to compare two binary segmentation sets, Binary1 and Binary2.*

segmentation of the brain from MRI data. These reference sets are called *supervised patient* segmentation.

In order to quantitatively assess the correspondence between two binary segmentation sets, we applied two difference measures, Additional CACTUS (*AC*) and Missing CACTUS (*MC*), and a Similarity measure *S* (Zijdenbos et al. 1994) (see Figure 2.3).

AC equals the fraction of segmented voxels in the CACTUS data that are not segmented in the corresponding reference set (binarized phantom segmentation or supervised patient segmentation), divided by the overlap, *i.e.*, voxels segmented in both sets. *MC* is the equivalent for segmented voxels in the reference set that are missing in the CACTUS results. *S* expresses the fraction of overlap divided by the total number of segmented voxels in both files; the factor of 2 ensures that a perfect segmentation has a similarity value of 100%.

The qualitative visual inspection of the visualization results was performed by three neuro-anatomy experts. Typical examples of 2D and 3D visualizations that were used, are shown in Section 2.3.

2.3 Results

Intermediate results of the CACTUS method for a phantom and a patient dataset are shown in Figures 2.4 and 2.5 using representative slices. Frames *A* correspond to the original grey value slices and Frames *B* are the result of a segmentation using the peak value. Comparison with the results obtained using a threshold value set to one

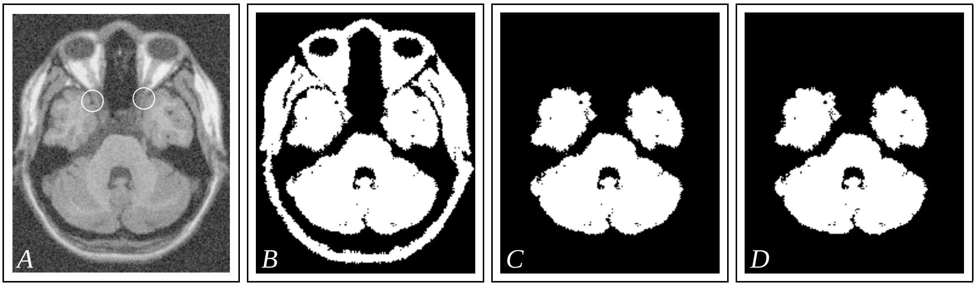

Figure 2.4 *Intermediate results of the CACTUS method for the phantom set with extreme noise and RF inhomogeneity. A representative slice of the original grey value data is shown in Frame (A). Frame (B) is the result when applying a* T_{low} *of 41 (see Figure 2.2(A)) for segmentation of these grey value data. Frame (C) presents the segmentation result with a* T_{low} *of 42 showing the exclusion of the undesired segment(s). Frame (D) is the final result of the CACTUS method, i.e., a segmentation result using* T_{low} *(42), no* T_{up} *was detected. The circles in Frame (A) indicate the relatively low grey value connections that define* T_{low}.

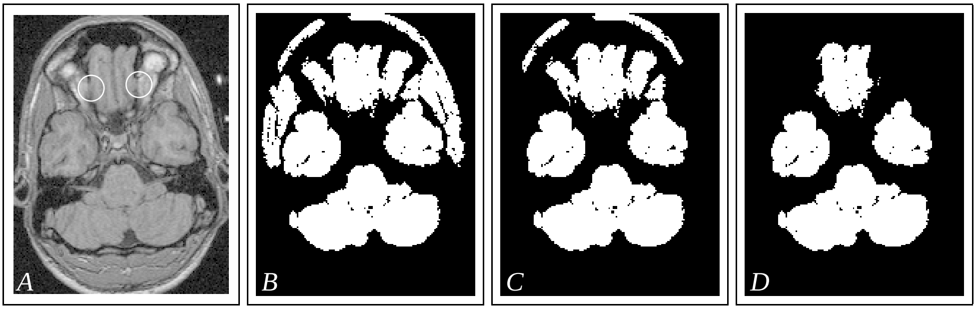

Figure 2.5 *Intermediate results of the CACTUS method for a patient set. See Figure 2.4 for an explanation of the frames.* T_{low} *for Frame (B) is 80, and for Frame (C) 81 (see Figure 2.2). Frame (D) is the final result using both* T_{low} *(81) and* T_{up} *(124). Here, differences between Frames (C) and (D) show the importance of the upper threshold to break the linkages owing to relatively high grey values. The circles in Frame (A) indicate the relatively high grey value connections that define* T_{up}.

higher (Frames *C*) excludes an undesired segment illustrating the importance of the detected peak. Owing to high grey value linkages some segments may still be connected to the brain (see Figure 2.5*C*), but these segments are excluded using T_{up} as can be seen in the final result in Figure 2.5*D*.

The presented slices clearly indicate the connections between the brain and other tissues. These connections are detected in the region growing process and determine the threshold range. T_{low} and T_{up} are defined by a relatively low, *resp.* high grey value connection (see indicated circles in Figures 2.4 and 2.5).

Phantom	*AC*	*MC*	*S*
3n00RF	6.0 %	0.3 %	96.9 %
3n20RF	1.0	2.8	98.1
9n40RF	10.3	0.4	94.8

Table 2.1 *Difference and similarity results for the three phantom datasets with respect to the binarized phantom segmentation (see Figure 2.3).*

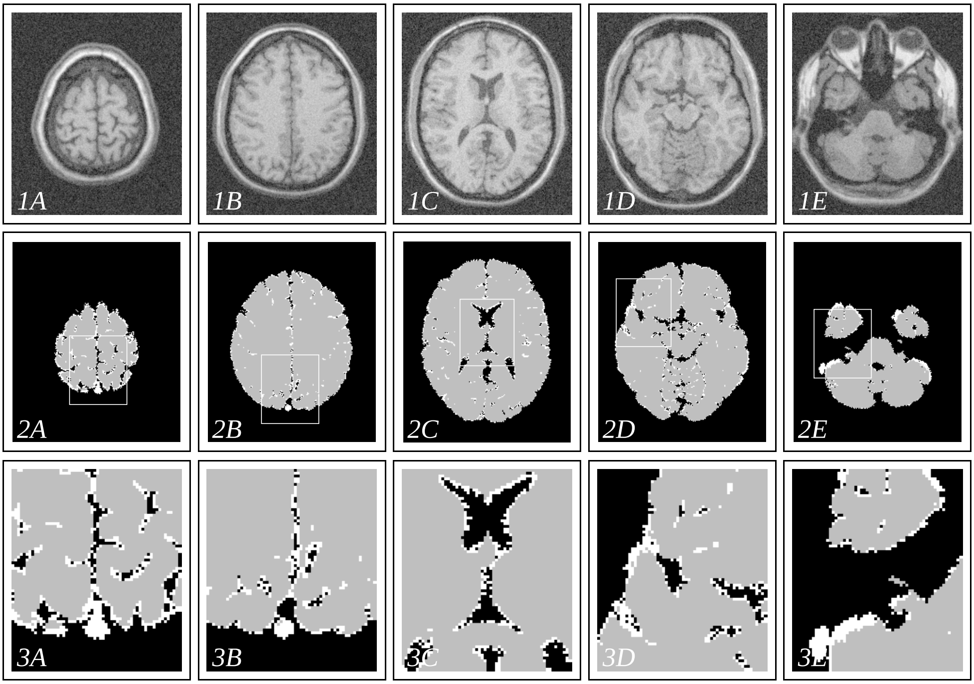

Figure 2.6 *Results of the CACTUS method for a phantom set compared with the binary phantom segmentation. Top row: Representative slices of the original grey value data of the third phantom data set (highest level of noise and RF inhomogeneity). The second row presents in grey the overlap between the CACTUS result and the binary phantom segmentation. The extra voxels segmented by the CACTUS method are shown in white. The third row shows enlargements of the indicated details in the second row.*

2.3.1 Phantom data

Table 2.1 shows the difference and similarity measures of CACTUS for the three phantom sets compared with the binarized phantom segmentation. Similarities are high and we decided to present the resulting images and renderings of the phantom set with the highest level of noise and RF inhomogeneities (see also Figure 2.4 for intermediate results). Presentation of the results for the phantom set with the more

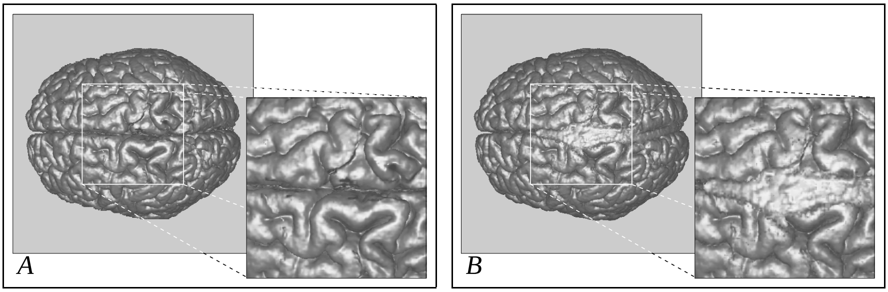

Figure 2.7 *Comparison of volume renderings for the high noise, high RF inhomogeneity phantom dataset with zoomed details of the longitudinal cerebral fissure. On the left (A) the binarized phantom segmentation is used for the rendering of the surface of the brain. On the right (B) the CACTUS results are shown.*

realistic settings (3% noise and 20% RF inhomogeneity) was regarded pointless because the differences between the CACTUS segmentation and the binarized phantom segmentation are very small and hardly noticeable in both the 2D and 3D visualizations.

Despite the uncommon high values of noise and RF non-uniformity, Figures 2.6 and 2.7 show that the similarity between the binarized phantom segmentation and the CACTUS results is high. However, two problems can be noted. The first is the inclusion of sinuses and meninges, especially in the longitudinal cerebral fissure. The second problem is the ragged border caused by the high noise content. The 2D images and 3D renderings show that these problems have a limited impact on the overall visualization quality.

2.3.2 Patient data

In Table 2.2 difference and similarity measures are presented for the 30 patient datasets. The initial similarity measures of patient sets XIV and XXIII were not consistent with the rest of the patient sets. Investigation revealed that no peak was detected in case XIV, whereas in case XXIII the first peak was missed. In Table 2.2 the results of these studies after (manual) reduction of the peak detection level (see Section 2.4) are shown.

The similarity measures of the other 28 sets are high (average of 98.0 % with a standard deviation of 1.3). We selected three datasets for presentation of the results, *i.e.*, the sets with the lowest (XII), average (XXI, see also Figure 2.5), and highest (XXIX) similarity.

Patient number	*AC*	*MC*	*S*
I	6.8 %	0.0 %	96.6 %
II	0.0	0.6	99.6
III	8.4	0.0	95.9
IV	2.4	0.3	98.6
V	3.2	0.0	98.4
VI	8.6	0.0	95.8
VII	4.3	0.0	97.8
VIII	4.8	0.0	97.6
IX	3.5	0.0	97.4
X	1.4	0.5	99.0
XI	5.0	0.0	97.5
XII	11.7	0.0	94.4
XIII	0.0	0.7	99.6
XIV*	5.6	0.9	96.8
XV	8.3	0.0	95.9
XVI	0.0	2.9	98.5
XVII	0.1	3.7	98.1
XVIII	2.9	0.0	98.5
XIX	2.8	0.2	98.4
XX	0.3	2.6	98.5
XXI	4.2	0.0	97.9
XXII	5.4	0.2	97.2
XXIII*	3.0	0.0	98.5
XXIV	0.7	0.0	99.6
XXV	2.3	0.0	98.8
XXVI	3.6	0.0	98.2
XXVII	2.1	0.0	98.9
XXVIII	1.0	0.4	99.3
XXIX	0.0	0.2	99.8
XXX	0.2	1.2	99.3

Table 2.2 *Results for the patient datasets (see also Table 2.1). The two patients indicated with an asterisk can be considered outliers and are explained separately in the text.*

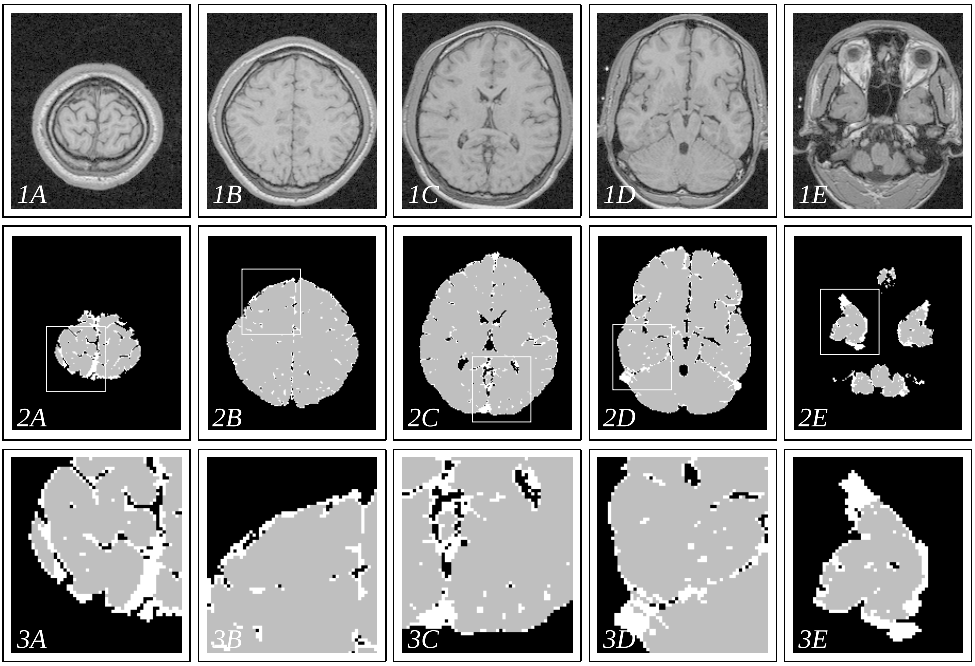

Figure 2.8 *Results of the CACTUS segmentation compared with the supervised segmentation using the Höhne method for patient dataset XII (lowest similarity).*

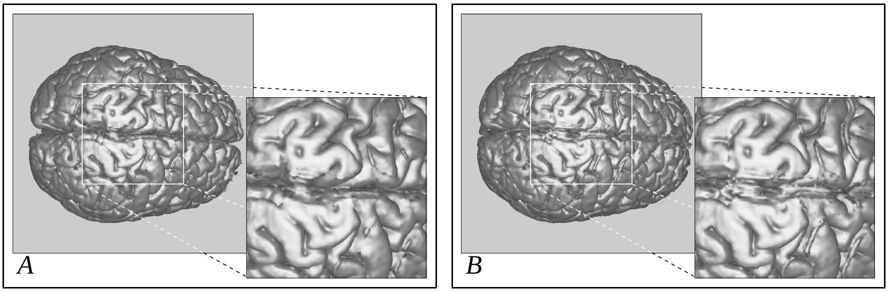

Figure 2.9 *Comparison of volume renderings for patient dataset XII (lowest similarity) with an enlargement of the longitudinal cerebral fissure. On the left (A) the images resulting from the supervised segmentation using the Höhne method, on the right (B) the corresponding image of the CACTUS method.*

The differences in the segmentations between the CACTUS method and the supervised (Höhne) method for case XII (lowest similarity) are shown in Figure 2.8. The differences are most conspicuous in the vessels (especially the sinuses) and meninges on the surface of the brain.

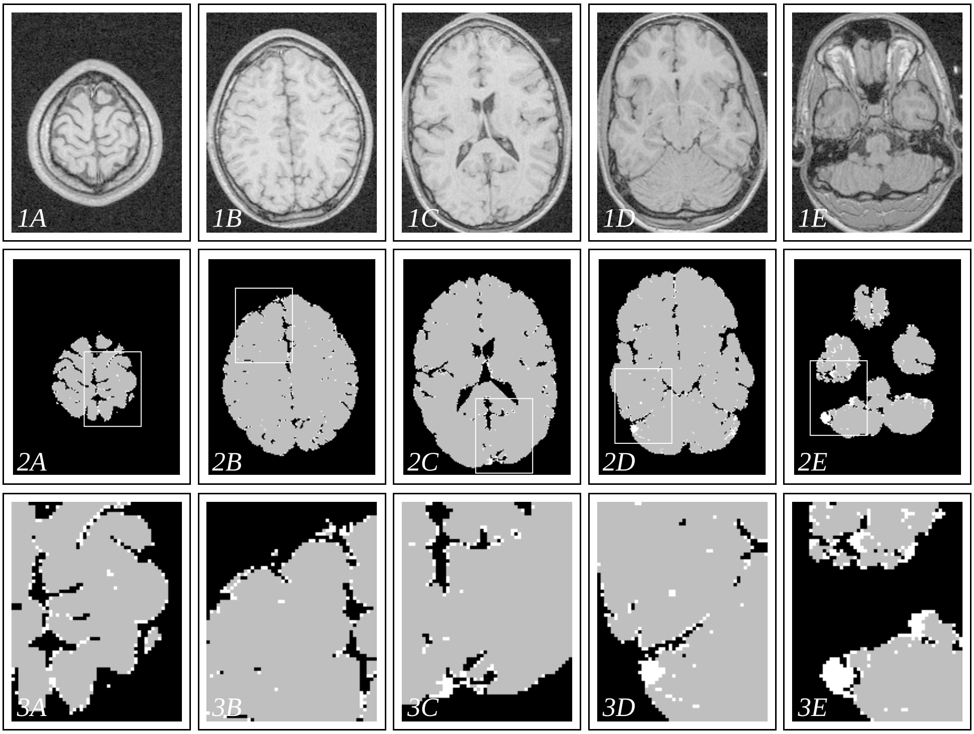

Figure 2.10 *Results of the CACTUS segmentation compared with the supervised segmentation using the Höhne method for patient dataset XXI (average similarity).*

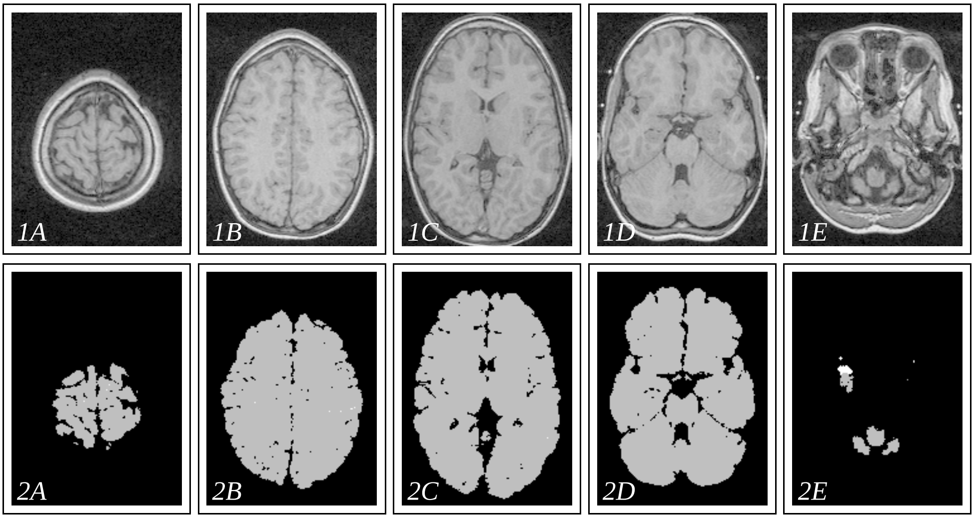

Figure 2.11 *Results of the CACTUS segmentation compared with the supervised segmentation using the Höhne approach for patient dataset XXIX (highest similarity).*

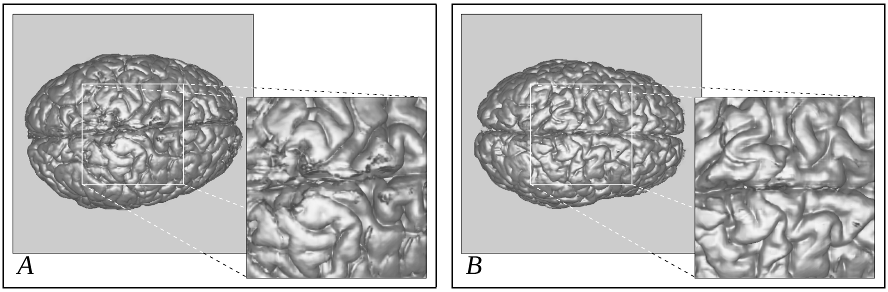

Figure 2.12 *Volume renderings of CACTUS segmentations for two patient sets with an enlargement. Frame (A) shows results for patient dataset XXI (average similarity), and Frame (B) patient dataset XXIX (highest similarity).*

For cases XXI (average similarity, see Figure 2.10) and XXIX (highest similarity, see Figure 2.11), the differences in the segmentations are small. Comparison between the renderings of the CACTUS segmentation results and the supervised segmentation results using the Höhne approach yielded differences that were hardly noticeable. We therefore only show the renderings of the CACTUS segmentation results (see Figure 2.12).

2.4 Discussion

Little attention has been given to the determination of T_{start} and the seed point selection, because no significant problems were encountered with these starting parameters. T_{start} was easily extracted from the histogram and the seed point only required a relatively high grey value and a location somewhere in the middle of the brain, which was easy to achieve. Furthermore, we used several white matter voxels as seed points for the region growing process to test the location dependency of the seed point. The location only affected the beginning of the growth graph of the process. The graphs of the different seed points rapidly showed the same pattern leading to identical values for T_{low} and T_{up}.

The applicability of the CACTUS method depends on the sensitivity to the parameters of the morphological operations (number of erosions/ dilations and the structuring element) and the peak detection algorithm. Based on our experiences with the Höhne approach incorporated into ANALYZE^{TM} we used two erosions, three geodesic dilations and a 6-voxel structuring element for all morphological operations. Tests using other parameters for the Höhne approach and CACTUS yielded poorer segmentation results compared with the selected settings.

The peak detection performs a simple comparison of the present number of growing steps with the numbers at the previous threshold settings. An increase is defined as a peak if the present value is 1.5 times the sum of the previous five values. The constants we used were based on observations done with the first four patient sets and proved not very critical, only two cases out of thirty-three (including phantom sets) failed. A failure can be detected easily by retrospective investigation of the growth graph. For the two outliers a satisfactory segmentation was obtained by lowering the peak detection parameter of 1.5 to 1.1 . This adjustment was quickly investigated on the other cases, but did not seem to improve overall results.

The threshold detection depends on the linkages between structures. Small linkages owing to the partial volume effect and small anatomical structures have to be removed by an erosion step (see Figure 2.1), as these interfere with the region growing process. More precisely, the threshold detection depends on linkages that remain present even after two erosions. Acquisition characteristics, *e.g.*, slice thickness, noise and RF, are likely to influence the values of parameters as the form and size of the morphological structuring element or the number of erosions and dilations. The parameters of CACTUS were optimized using the first four patient sets. This explains the higher similarity value for the more realistic phantom dataset when compared to the set with lower RF inhomogeneity. For every acquisition technique and characteristics the parameters of CACTUS can be optimised using several cases. These settings can then be applied for every following case of that acquisition type.

In the studies with patient data, the supervised segmentation method of Höhne is used as a reference, not as a gold standard. When comparing both methods, we have to realize that CACTUS is in fact an automated version of the Höhne technique and both will have some identical problems. First of all, a gradient over the images affects both segmentation methods the same way. Secondly, the application of erosion and subsequent geodesic dilation is a very simple and powerful technique for removal of the (thin) connections between different tissues, but it also affects possible long, thin intra-tissue structures. For instance, the cerebellum manifests some filament-like structures that are easily eroded, but difficult to regain via dilations within a threshold mask. This was most apparent in the top part of the cerebellum (see Figure 2.11 Frames *C*).

In the majority of presented cases we found that CACTUS segmented some additional tissue surrounding the brain (see *AC* compared with *MC* in tables 2.1 and 2.2). This is not problematic for the volume visualizations of the brain surface by virtue of the used grey opacity shader incorporated in our volume rendering package VROOM (Zuiderveld 1995). This shader assigns opacities to the grey value data in a neighborhood of the surface voxel (see also (Levoy 1988)) and ensures that additional tissue surrounding the brain has little to no impact on the final image.

Segmentation of a typical dataset requires about half a hour of CPU time on a processor of an SGI Power Challenge R10K (MIPS R10000, 195 MHz); the speed can be considerably improved by optimizing the implementation and/or using the inherent parallelism of the algorithm.

2.5 Conclusions

We have developed an automatic approach to segment the cerebral cortex from T1-weighted MR images. This fully automated segmentation method CACTUS is based on the simple operations thresholding, region growing, erosion, and geodesic dilation.

The results of CACTUS for patient data have been compared to the corresponding supervised method. The results are quite similar, and in a large number of cases visually almost indistinguishable. Furthermore, we have used phantom data to evaluate the segmentation performance against a gold standard. The results were very accurate, to the extent that CACTUS might prove adequately accurate for quantitiative analysis.

Acknowledgements

We are indebted to our colleagues W.F.C. Baaré, R. Maas, L.C. Meiners, W.J. Niessen, H. Schnack, E.P.A. Voncken, and R. van der Weide. We gratefully acknowledge the research licence of ANALYZE^{TM}, provided by Dr R.A. Robb, Mayo Foundation/ Clinic, Rochester, Minnesota.