MEDIMA 150
DISK / 24/6/04
ARTICLE IN PRESS
No. of pages: 9
DTD 4.3.1 / SPS-N

Available online at www.sciencedirect.com
SCIENCE DIRECT®

ELSEVIER

Medical Image Analysis xxx (2004) xxx–xxx

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

A brain tumor segmentation framework based on outlier detection ☆

Marcel Prastawa [a,*], Elizabeth Bullitt [c], Sean Ho [a], Guido Gerig [a,b]

[a] *Department of Computer Science, University of North Carolina, CB #3175, Sitterson Hall, Chapel Hill, NC 27599, USA*
[b] *Department of Psychiatry, University of North Carolina, Chapel Hill, NC 27599, USA*
[c] *Department of Surgery, University of North Carolina, Chapel Hill, NC 27599, USA*

Abstract

This paper describes a framework for automatic brain tumor segmentation from MR images. The detection of edema is done simultaneously with tumor segmentation, as the knowledge of the extent of edema is important for diagnosis, planning, and treatment. Whereas many other tumor segmentation methods rely on the intensity enhancement produced by the gadolinium contrast agent in the T1-weighted image, the method proposed here does not require contrast enhanced image channels. The only required input for the segmentation procedure is the T2 MR image channel, but it can make use of any additional non-enhanced image channels for improved tissue segmentation. The segmentation framework is composed of three stages. First, we detect abnormal regions using a registered brain atlas as a model for healthy brains. We then make use of the robust estimates of the location and dispersion of the normal brain tissue intensity clusters to determine the intensity properties of the different tissue types. In the second stage, we determine from the T2 image intensities whether edema appears together with tumor in the abnormal regions. Finally, we apply geometric and spatial constraints to the detected tumor and edema regions. The segmentation procedure has been applied to three real datasets, representing different tumor shapes, locations, sizes, image intensities, and enhancement.

Keywords: Automatic brain segmentation; Brain tumor segmentation; Level-set evolution; Outlier detection; Robust estimation

1. Introduction

Automatic brain tumor segmentation from MR images is a difficult task that involves various disciplines covering pathology, MRI physics, radiologist's perception, and image analysis based on intensity and shape. There are many issues and challenges associated with brain tumor segmentation. Brain tumors may be of any size, may have a variety of shapes, may appear at any location, and may appear in different image intensities. Some tumors also deform other structures and appear together with edema that changes intensity properties of the nearby region. For many human experts, manual segmentation is a difficult and time consuming task, which makes an automated brain tumor segmentation method desirable. There are many possible applications of an automated method, it can be used for surgical planning, treatment planning, and vascular analysis. It has been shown that blood vessels in the brain exhibit certain characteristics within pathological regions (Bullitt et al., 2003). An objective and reproducible segmentation procedure coupled with vascular analysis would allow us to study the relation between pathologies and blood vessels and may function as a new diagnostic measure.

The challenges associated with automatic brain tumor segmentation have given rise to many different approaches. Automated segmentation methods based on artificial intelligence techniques were proposed in (Clark et al., 1998; Fletcher-Heath et al., 2001). The two methods do not rely on intensity enhancements provided by the use of contrast agents. A particular limitation of the two methods is that the input images are restricted to the T1, T2, and PD MR image channels. Additionally, the methods require a training phase prior to segmenting a set of images. Other methods are based on statistical pattern recognition techniques, for example the method proposed by Kaus et al. (1999). This method combines the information from a registered atlas tem-

☆ Supported by NIH-NIBIB R01 EB000219 and NIH-NCI R01 HL69808.
* Corresponding author. Tel.: +919-962-1836.
E-mail address: prastawa@cs.unc.edu (M. Prastawa).

1361-8415/$ - see front matter © 2004 Published by Elsevier B.V.
doi:10.1016/j.media.2004.06.007

plate and user input to supervise training of a the classifier, demonstrating the strength of combining voxel-intensity with geometric brain atlas information. This method was validated against meningiomas and low-grade gliomas. Gering et al. (2002) proposed a method that detects deviations from normal brains using a multi-layer Markov random field framework. The information layers include voxel intensities, structural coherence, spatial locations, and user input. Cuadra et al. (2002) presented high-dimensional warping to study deformation of brain tissue due to tumor growth. This technique relies on a prior definition of the tumor boundary whereas the method we propose in this paper focuses on automatically finding tumor regions.

Previous work on automatic brain tumor segmentation generally uses the enhancement provided by the gadolinium contrast agent in the T1 channel or constrained to blobby shaped tumors with uniform intensity. Even though the intensity enhancement can aid the segmentation process, it is not always necessary to obtain good results. In fact, the use of a contrast agent can be problematic. Typically, tumors are only partially enhanced and some tumors are not enhanced at all. Blood vessels also generally appear enhanced by the contrast agent. These inconsistencies create an ambiguity in the image interpretation, which makes the T1-enhanced image channel a less than ideal feature for tumor segmentation.

Edema surrounding tumors and infiltrating mostly white matter was most often not considered as important for tumor segmentation. We showed previously (Moon et al., 2002; Prastawa et al., 2003) that edema can be segmented using a prior for edema intensity and restriction to the white matter region. The extraction of the edema region is essential for diagnosis, therapy planning, and surgery. It is also essential for attempts that model brain deformation due to tumor growth. The swelling produced by infiltrating edema usually has distinctly different tissue property characteristics than tumor. Our new scheme presented here is based on the detection of "changes from normal" and will thus systematically include segmentation of edema. Differential identification of the two abnormal regions tumor and edema is clinically highly relevant. Even though the primary therapeutic focus will be on the tumor region, the edema region may require secondary analysis and treatment.

Our method combines the model of the normal tissues and the geometric and spatial model of tumor and edema. It relies on the information provided in the T2 image channel for identifying edema, and it can make use of additional image channels to aid the segmentation. For our datasets, we use only the T1 and T2 image channels. Tumor and edema are treated as intensity abnormalities or outliers. After identifying the abnormalities, an unsupervised clustering technique is applied to the intensity features before utilizing geometric and spatial constraints. We will demonstrate that this method can segment tumors with or without intensity enhancements and automatically detects the presence of edema, thus overcoming limitations of our previous method (Moon et al., 2002; Prastawa et al., 2003). Our approach offers a means of approaching lesions of multiple types and image intensities, and, with a single method, lesions that enhance and do not, and that may or may not be surrounded by edema.

2. Method

The automated segmentation method that we have developed is composed of three major stages, as shown in Fig. 1. First, it detects abnormal regions, where the intensity characteristics deviate from the expectation. In the second stage, it determines whether these regions are composed of both tumor and edema. Finally, once the estimates for tumor and edema intensity parameters are obtained, the spatial and geometric properties are used for determining proper sample locations. The details of each stage are discussed in the following subsections.

2.1. Detection of abnormality

Before identifying tumor and edema, it is necessary to first detect regions that have properties that deviate from the expected properties of a normal, healthy brain. In our segmentation method, this involves finding the intensity parameters for healthy classes and the abnormal class. The initial parameters for the healthy brain classes are obtained by sampling specific regions based on the probabilistic brain atlas shown in Fig. 2 (Evans et al., 1993).

The atlas is aligned with the subject image data by registering the atlas template image with the subject image. The registration is performed using affine transformation with the mutual information image match

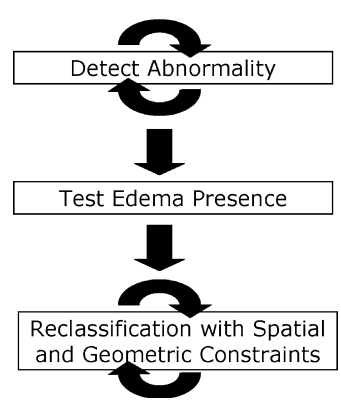

Fig. 1. The three major stages of the segmentation method.

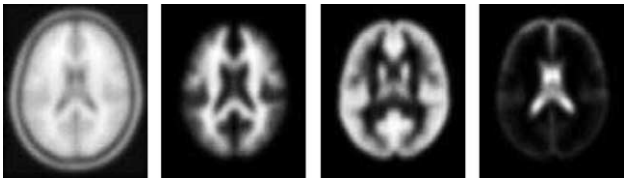

Fig. 2. The digital brain atlas provided by the International Consortium for Brain Mapping (ICBM). From left to right: the T1 template image and probability values of white matter, gray matter, and csf.

measure (Maes et al., 1997). After alignment, the samples for each healthy class (white matter, gray matter, and cerebrospinal fluid (csf)) are obtained by randomly selecting the voxels with high atlas probability values. For our data, the set of training samples is constrained to be the voxels with probabilities higher than a threshold $\tau = 0.85$.

The training data for the healthy classes generally contain unwanted samples due to contamination with samples from other tissue types, particularly tumor and edema. The pathological regions are not accounted for in the brain atlas and they therefore occupy regions that are marked as healthy. The contaminants are data outliers, and they are removed so that the training samples for the healthy classes are not contaminated. The samples are known to be contaminated if their characteristics differ from prior knowledge. The intensities for healthy classes are known to be well clustered and can be approximated using Gaussians (Fig. 3).

Handling data outliers is a crucial step for atlas based image segmentation. Cocosco et al. (2003) developed a segmentation method for healthy brains that builds the Minimum Spanning Tree from the training samples and iteratively breaks the edges to remove false positives (pruning). They showed that pruning the training samples results in significant improvement of the segmentation quality. We use a robust estimate of the mean and covariance of the training data to determine the outliers to be removed.

The robust estimator that we use is the Minimum Covariance Determinant (MCD) estimator. It is defined to be the mean and covariance of an ellipsoid covering at least half of the data with the lowest determinant of covariance. The method is highly robust, with a high breakdown point. The breakdown point is the fraction of the data that must be moved to infinity so that the estimate also moves to infinity. The MCD estimate has a breakdown point of 0.5, more than half of the data needs to be contaminated to make the estimate be unreasonable.

A fast algorithm for computing the MCD estimate is described in (Rousseeuw and Van Driessen, 1999). The algorithm first creates several initial subsets, where the elements are chosen randomly. From each subset, the algorithm determines different initial estimates of the robust mean and covariance. The estimates are then refined by performing a number of C-step operations on each initial selections. A single C-step operation consists of the following steps:

(1) Given a subset of the data, compute the mean and covariance of the elements in the subset.
(2) Compute Mahalanobis distances of the data elements in the whole set.
(3) Sort points based on distances, smallest to largest.
(4) Select a new subset where the distances are minimized (e.g., first half of the sorted data points).

An illustration of a single C-step iteration is shown in Fig. 4. A C-step operation will result in a subset selection that yields a determinant of covariance less or equal to the one obtained from the previous subset. The iterative applications of C-steps yield final estimates with the smallest determinant of covariance. From all the final estimates computed with different initial selections, the mean and covariance estimate with the smallest determinant of covariance is chosen as the robust estimate. Given the robust mean and covariance, samples that are further than three standard deviations are considered as outliers (Fig. 5). The inliers of the healthy brain tissue class samples are used as training samples for estimating the corresponding density functions.

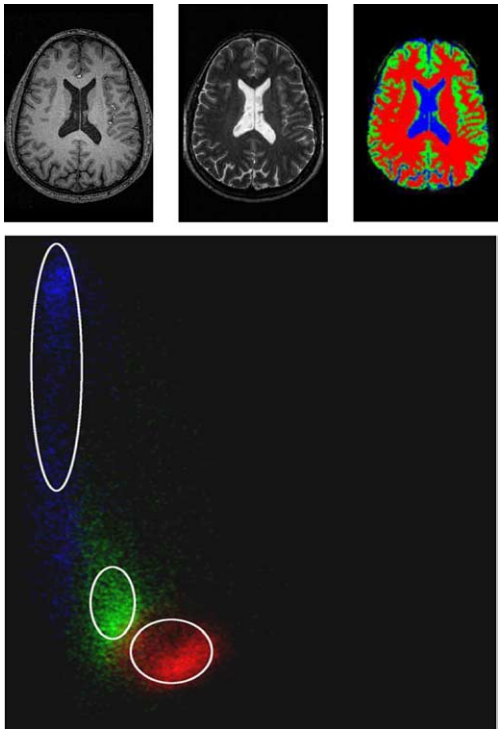

Fig. 3. Example healthy dataset. Top, from left to right: T1 image, T2 image, and segmentation labels (from brightest to darkest: white matter, gray matter, and csf). Bottom: the intensity histogram for the three classes, the horizontal axis represents T1 intensities and the vertical axis represents T2 intensities. The intensity features for each class is tightly clustered and can be approximated with a Gaussian. (This figure is available in color, see the online version.)

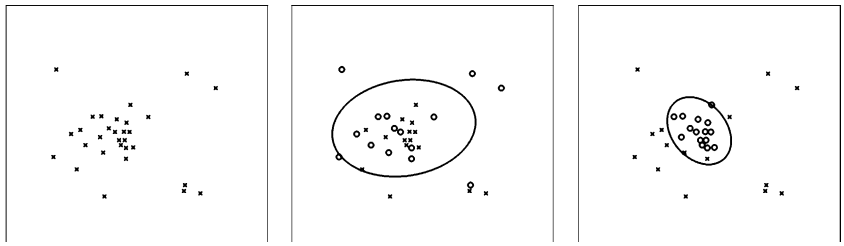

Fig. 4. An illustration of a single C-step iteration, a key component of the MCD robust estimation algorithm. Left: original 2-D data. Center: random selection of a subset of the data (marked with circles). Right: the selection after a C-step iteration, where the closest points to the previous mean and covariance estimate is selected. The ellipsoidal curves in the center and right plots show the locations one standard deviation away from the mean and covariance estimate, which are computed from the selected points.

The specific aim at this stage is to compute the density estimates and posterior probabilities for the class labels $\Gamma = \{$white matter, gray matter, csf, abnormal, non-brain$\}$. A parametric density function is not ideal for the case of tumor segmentation. Tumors do not always appear with uniform intensities, particularly in the case where some tissues inside the tumor are necrotic tissues. We therefore make no assumption regarding the intensity distributions and use a non-parametric model for the probability density functions. The density functions are approximated using kernel expansion or Parzen windowing (Duda et al., 2001). Given the vector of intensities $\vec{I}(x)$ at location x, the probability density function on intensity for the class label Γ_j is

$$p(\vec{I}(x)|\Gamma_j) = \frac{1}{N}\sum_{i=1}^{N} K_\lambda(\vec{I}(x) - \vec{T}_i),$$

where K_λ is the multivariate Gaussian kernel with standard deviation λ and $\vec{T}_i$ is a class training sample. The kernel bandwidth λ chosen for our dataset is 4% of the intensity range for each channel.

The posterior probability is computed using the class prior probability from the atlas $Pr(\Gamma_j, x)$ at location x

$$P(\Gamma_j|\vec{I}(x)) = \frac{p(\vec{I}(x)|\Gamma_j)Pr(\Gamma_j, x)}{p(\vec{I}(x))}.$$

The spatial priors for white matter, gray matter, csf, and non-brain classes are the corresponding atlas probabilities. For the abnormal class, we use a fraction of the sum of white matter and gray matter atlas probabilities since tumor and edema usually appear in these regions and not in the csf regions.

An issue with MR images is the presence of the image inhomogeneity or the bias field. We deal with this by interleaving the segmentation process with bias correction, following the spirit of (Wells et al., 1996). The entire process of detecting the abnormal regions is shown in Fig. 6, a loop that is composed of the following five stages:

(1) Threshold the posterior probabilities and sample the high confidence regions. At the first pass, the atlas probabilities are used in place of the posterior probabilities.
(2) Remove the samples for normal tissues that exceed a distance threshold based on the MCD estimate.
(3) Estimate the non-parametric density for each class labels using kernel expansions. The initial density for the abnormal class is set to be uniform, which makes this class act as a rejection class. The brain voxels with intensity features that are different from those of healthy classes or not located in the expected spatial coordinates will be assigned to this class.
(4) Compute the posterior probabilities.

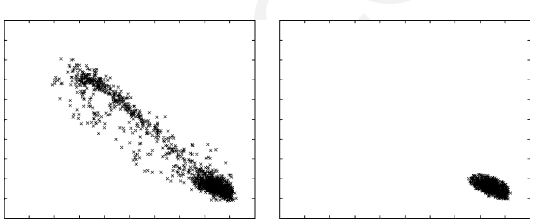

Fig. 5. The white matter training data for a subject with tumor and edema, the horizontal axis represents the T1 intensities and the vertical axis represents the T2 intensities. Left: original samples obtained by atlas-guided sampling which is contaminated with samples from other distributions. Right: remaining samples after trimming using the robust MCD estimate.

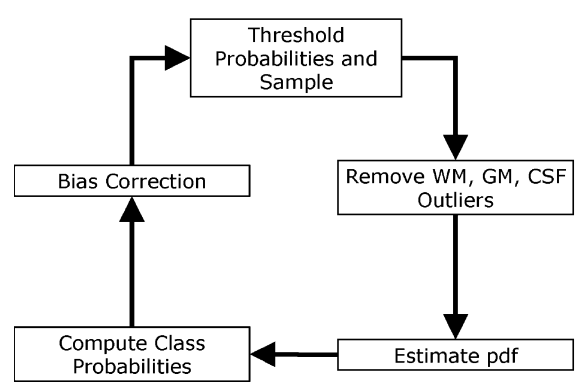

Fig. 6. The process of detecting abnormal regions, the first stage of the method.

(5) Estimate bias field from white matter and gray matter probabilities. Apply correction using the estimated bias field.

The first major segmentation stage detects the abnormal regions by executing the loop for several iterations, obtaining the intensity descriptions for each class. The abnormal class density at different iterations for the Tumor020 data is shown in Fig. 7.

The bias correction method is based on the one developed by Van Leemput et al. (1999). The method uses the posterior probabilities to estimate the homogeneous image. It then computes the bias field estimate, as the log-difference between the homogeneous images and the real subject images. The bias field is modeled as a polynomial, and the coefficients of the polynomial is determined through least squares fitting. The method assumes that the class intensity distributions are approximately Gaussians. We therefore use only the white matter and gray matter probabilities for bias correction, as they generally can be approximated by Gaussians without significant errors.

2.2. Detection of edema

The densities and posterior probabilities computed for the abnormal class in the previous stage give us a rough estimate of how likely it is that some voxels are part of tumor or edema. We assume that the detected abnormal voxels are composed mostly of tumor and possibly edema. Edema is not always present when tumor is present, therefore it is necessary to specifically test the presence of edema. This is done by first obtaining the intensity samples for the abnormal region, the posterior probabilities are thresholded and a subset of the region is selected. The samples are clustered and then we determine whether there exist separate clusters for tumor and edema. The density estimate for tumor (and edema, if present) is obtained by performing kernel expansion on the samples.

Tumor and edema are generally separable given the information in the T2 weighted image (Fig. 8). Edema has high fluid content and therefore appears brighter than tumor in this image channel. To separate the densities, we apply unsupervised clustering to the samples obtained by thresholding. The method we have chosen is k-means clustering with $k = 2$ (Duda et al., 2001). Once we obtain the clusters, we can identify the tumor cluster as the cluster with the T2 mean that has the lower value.

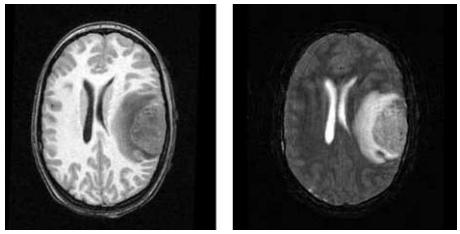

Fig. 8. The T1 image (left) and the T2 image (right) from the Tumor020 data. The tumor and edema on the right part of the brain can be clearly differentiated based on the T2 intensities. As observed in the T2 image, the tumor region (rightmost) is darker than the surrounding edema region, as edema is composed mostly of fluid.

To determine the validity of the clustering, we use the overlap measure called the Davies–Bouldin index (Davies and Bouldin, 1979). This measure is the ratio of the average within cluster distances and the between cluster distance. Given m candidate tumor samples τ_i with the mean value μ_{tumor}, and n candidate edema samples ϵ_i with the mean value μ_{edema}, the overlap measure is

$$\frac{1}{2}\left(\frac{\frac{1}{m}\sum_{i=1}^{m}\|\tau_i-\mu_{\text{tumor}}\|+\frac{1}{n}\sum_{i=1}^{n}\|\epsilon_i-\mu_{\text{edema}}\|}{\|\mu_{\text{tumor}}-\mu_{\text{edema}}\|}\right).$$

The T2 channel contains most of the information needed for differentiating tumor and edema. Therefore, we have chosen to measure the overlap for only the T2 data of each cluster. If the amount of overlap is larger than a specified threshold, then the tumor density is set to be the density for the abnormal class and the edema density is set to zero.

2.3. Reclassification with spatial and geometric constraints

Once this stage is reached, tumor and edema are already segmented based on atlas priors and intensity characteristics. However, the geometric and spatial properties were not considered and this generally leads to having at least a few false positives. Since there is no model for the intensity distributions of tumor and edema, it is necessary to use geometric and spatial heuristics to prune the samples that are used for estimating the densities. Here, we make use of the prior knowledge that tumor is mostly blobby. For edema, we use the constraint that each edema region is connected to a nearby tumor region. Some edema voxels can be located far

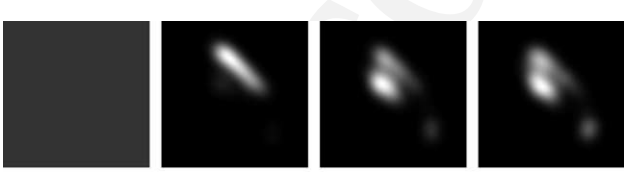

Fig. 7. Snapshots of the estimated probability density function of the abnormal class for the Tumor020 data. Each image shows the result of different iterations of the loop shown in the previous figure. The density is initialized so that all intensities are equally likely. The horizontal axis represents the T1 intensities and the vertical axis represents the T2 intensities. The two high density regions visible at the final iteration are the tumor and edema densities, which have a significant separation along the dimension of the T2 intensities.

away from tumor regions, but they must be connected to a tumor region spatially.

Tumor structures generally appear as blobby lumps, this shape constraint is enforced through region competition snakes (Sebastian et al., 2003; Tek and Kimia, 1995; Tek and Kimia, 1997; Zhu et al., 1995). The tumor posterior probabilities is used as the input for the snake, which is represented as the zero level set of the implicit function ϕ. The level set evolution is governed by the following equation (Ho et al., 2002):

$$\frac{\partial \phi}{\partial t} = \alpha(P(\text{tumor}|\vec{I}(x)) - P(\overline{\text{tumor}}|\vec{\text{I}}(\text{x})))|\nabla\phi| + \beta\nabla\cdot\left(\frac{\nabla\phi}{|\nabla\phi|}\right)|\nabla\phi|.$$

The propagation term is represented by α. It is modulated by the difference of the posterior probabilities for the tumor class and the non-tumor class ($P(\text{tumor}|\vec{I}(x))$ and $P(\overline{\text{tumor}}|\vec{\text{I}}(\text{x}))$), so that the direction of the propagation is determined by the sign of the difference. The probability that a voxel is part of brain and not part of tumor is represented by $P(\overline{\text{tumor}}|\vec{I}(x))$, more explicitly

$$P(\overline{\text{tumor}}|\vec{I}(x)) = P(\text{wm}|\vec{I}(x)) + P(\text{gm}|\vec{I}(x)) + P(\text{csf}|\vec{I}(x)) + P(\text{edema}|\vec{I}(x))$$

The snake shrinks when the boundary encloses part of the regions not part of tumor and expands when the boundary is inside the tumor region. We apply a smoothing on the snake contour using mean curvature flow, and the strength of this smoothing is controlled by the β term. The initial level set function is obtained by performing a distance transform on the segmented tumor objects.

Edema, if present, is always contiguous with the tumor. With this prior knowledge, we therefore assume that edema is located near tumor structures. Each segmented edema object must have a voxel that is no further than some small distance from tumor regions. This test can be done efficiently by using the connected component algorithm and mathematical morphology. We first generate a binary image representing the segmented edema region. Then, we use this image as an input for the connected component algorithm to determine the individual edema objects. Each object is then dilated with a small structuring element, and then compared against the segmented tumor regions The objects that share at least a voxel with a tumor region is considered valid. Edema samples from these regions are kept, while other edema samples are discarded.

The final segmentation is obtained by reclassifying the image using the iterative steps similar to the one described in Section 2.1, with some modifications (Fig. 9). The outlier removal stage is removed and there are additional steps where these geometric and spatial constraints are enforced. The entire loop is performed several times, after going through one loop the tumor and edema probabilities at the voxel locations that do not pass the tests to zero. This way, the segmentation for these locations are determined based on the next best candidate class. The tumor shape constraint is disabled at the last fitting stage. This is done to obtain the proper boundary for the tumor structures, which may not be entirely smooth. For instance, gliomas typically have a general blobby shape and ragged boundaries.

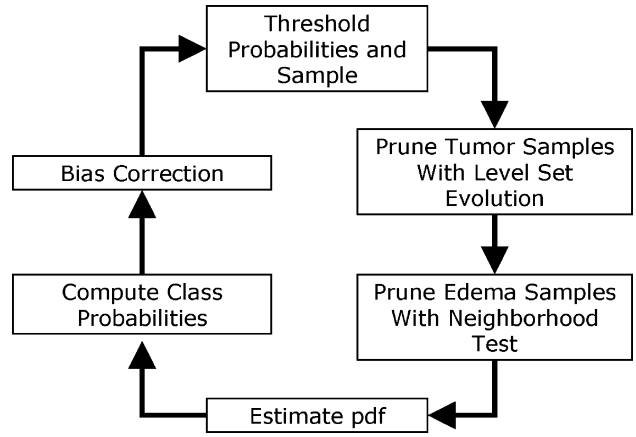

Fig. 9. The third stage of the method where the image is reclassified using tumor geometric properties and edema spatial relation.

3. Results

We have applied the method to three real datasets, representing different tumor shapes, locations, sizes, image intensity, and enhancement, as shown in Fig. 10. Tumor020 has a partially enhancing tumor that causes a large deformation of the normal structures. Tumor025 contains a large, partially enhancing tumor inside the brain stem. Tumor033 contains a low grade tumor which is not highlighted in the T1-enhanced channel.

Two sets of segmentations are done manually by one human rater at different times. The volumes of the manually segmented structures are shown in Table 1. The first set of manual segmentations is considered to be the gold standard for validating the automatic segmentation method. We used the VALMET segmentation validation tool (Gerig et al., 2001) to generate five validation metrics. The volume overlap metric is the normalized voxel intersection count for the pair of segmentations A and B: $(A \cap B)/(A \cup B)$, otherwise known as Jaccard's similarity coefficient (Jaccard, 1912). The other metrics are the maximum Hausdorff surface distance and the average surface distances (inside/negative, outside/positive, and absolute).

The intra-rater variability is shown in Table 2. The surface distance metrics indicate that the manual seg-

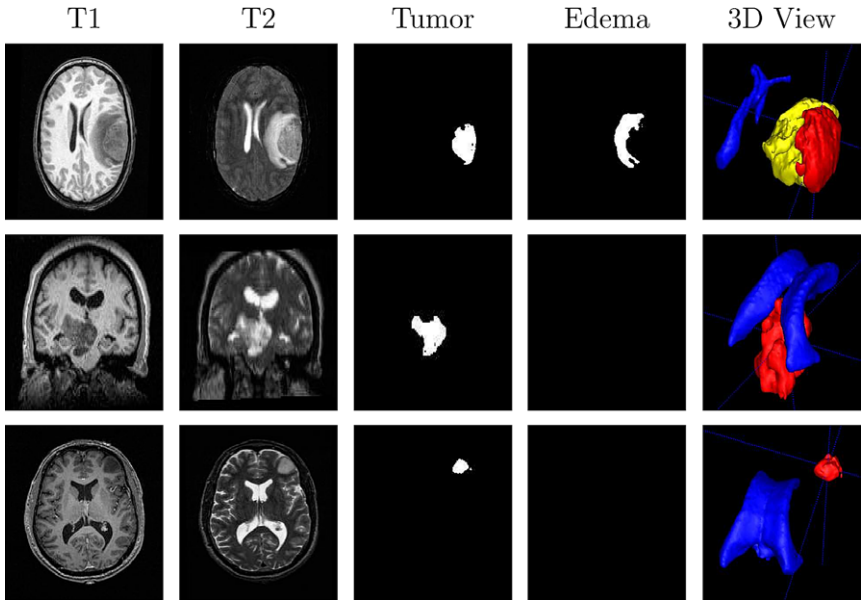

Fig. 10. The datasets and the generated segmentation results. The last column shows the 3D views of the segmented structures: medium gray represents tumor, bright gray represents edema, and dark gray represents ventricles. From top to bottom: Tumor020, Tumor025, Tumor033. These results illustrate that our method does differential segmentation for tumor and edema, which works also in cases where no edema is present. (This figure is available in color, see the online version.)

Table 1
Volumes of the segmented structures, from the first set of manual segmentation results

Dataset	Tissue type	Volume (mm^3)
Tumor020	Tumor	35578.6
Tumor020	Edema	64860.6
Tumor025	Tumor	24742.4
Tumor033	Tumor	3661.5

mentations are reliable. The overlap metrics are also high, with the exception of the Tumor033 segmentation. This is likely due to the small size of the tumor. The quantitative validation of the automatic segmentation method is shown in Table 3. The level of agreement based on surface distances is similar for all tumors. However, the varying overlap values demonstrate that the overlap metric is sensitive to the size and complexity of the segmented objects (Fig. 10). The level of agreement with the manual result for edema is lower than tumor. This is mainly due to the ambiguity in determining the edema boundary, especially the tumor–edema boundary. For each case, the time required for the automatic segmentation method is about 1 h 30 min on a 2 GHZ Intel Xeon machine. The automatic segmentation process is done with little user intervention. The user only needs to specify several parameters before the segmentation begins. These parameters include the atlas probability threshold, the level set evolution settings, kernel width for Parzen windowing, and the distance threshold for outlier detection.

4. Discussion

4.1. Application areas

The automatic segmentation method proposed in this paper can process a wide variety of tumors since it does not rely on contrast enhancement. It segments the whole brain, including healthy tissue types, and automatically identifies edema. Defining the edema region can be useful for surgical planning, definition of radiation therapy fields, and, since the edema region indicates the volume over which the tumor exerts obvious chemical

Table 2
Validation metrics comparing the two sets of manual segmentation results done by the same human rater, demonstrating the intra-rater variability of the manual segmentations

Dataset	Tissue type	Overlap (%)	Hausdorff (mm)	Inside (mm)	Outside (mm)	Absolute (mm)
Tumor020	Tumor	89.0	3.98	0.32	1.17	0.54
Tumor020	Edema	75.5	13.1	0.48	1.4	0.75
Tumor025	Tumor	81.2	4.1	0.21	1.31	0.73
Tumor033	Tumor	59.4	5.22	0.42	2.06	1.51

Table 3
Validation metrics of the automatic tumor segmentation results against the first set of manual segmentation results

Dataset	Tissue type	Overlap (%)	Hausdorff (mm)	Inside (mm)	Outside (mm)	Absolute (mm)
Tumor020	Tumor	80.0	16.79	1.28	2.16	1.64
Tumor020	Edema	68.2	12.80	0.63	2.43	1.75
Tumor025	Tumor	79.2	17.85	1.01	3.70	1.44
Tumor033	Tumor	70.6	8.60	0.25	2.47	1.85

effects, identification of areas of interest to multiple investigators interested in tumor growth and treatment. Knowing the edema region can also be useful for surgical planning and radiation therapy. Often, edema regions need to be treated to reduce the risk of recurrence.

4.2. Future work

The segmentation method presented in this paper detects abnormal regions in the brain based on the atlas and image intensities. Other properties can also be used for this process. This can include geometric properties such as curvature or brain asymmetry (Joshi et al., 2003). Although the contrast enhanced image channel leads to ambiguous information, there are cases where it leads to more straightforward identification of brain tumors, assuming that enhanced blood vessels and noise can be properly identified. Robust estimation schemes other than the MCD may be necessary for these extensions.

A potential issue that is not handled by the proposed method is large deformation of brain structures. When there is large deformation, the brain atlas used may lead to incorrect sampling. In this case, the atlas based samples would be severely contaminated and the MCD algorithm may not yield correct results. The spatial priors here would also limit the segmentation quality, as the segmentation output cannot differ greatly from the atlas. The current method can still handle some level of deformation due to the use of outlier detection, but it would be helpful to explicitly account for these deformations using deformable registration.

An issue that goes together with the issue of knowing the deformation induced by tumor is the problem of determining the possible shapes of brain tumors. The shape model for tumor enforced using region competition snake constrains the segmented tumor to have rather smooth shapes. The notion of spatial coherence for brain tumors need to be properly enforced in order to segment wider varieties of brain tumors. This is a difficult issue because tumors can appear in many different sizes and shapes.

5. Conclusion

This paper presents a new approach for automatic segmentation of tumors and adjoining edema from non-enhancing multichannel MRI (T2 weighted channel explicitly required). Most methods so far have been applicable only to enhancing, homogeneous tumors. Furthermore, they require user-guidance in training a supervised classifier or to obtain a rough outline of the region of interest. Here, we show that robust estimation and outlier detection can be a promising new concept for detecting abnormalities in the brain.

The presented technique automatically identifies the presence of edema. Our collaborating clinicians confirm that this is a highly relevant feature, as the edema region often may require secondary analysis and treatment after the primary focus to the tumor region. The technique uses a concept that detects *difference from normal* and uses non-parametric estimates for distributions rather than traditional mixture Gaussian models. The technique also makes use of other features besides intensity: the shape of brain tumor and location of edema. In the future, we will improve this framework so that it can segment wider varieties of brain tumors with and without edema.

References

Bullitt, E., Gerig, G., Pizer, S.M., Aylward, S.R., 2003. Measuring tortuosity of the intracerebral vasculature from MRAimages. IEEE Transactions on Medical Imaging 22, 1163–1171.

Clark, M.C., Hall, L.O., Goldgof, D.B., Velthuizen, R., Murtagh, F.R., Silbiger, M.S., 1998. Automatic tumor-segmentation using knowledge-based techniques. IEEE Transactions on Medical Imaging 117, 187–201.

Cocosco, C.A., Zijdenbos, A.P., Evans, A.C., 2003. A fully automatic and robust brain MRI tissue classification method. Medical Image Analysis 7, 513–527.

Cuadra, M.B., Gomez, J., Hagmann, P., Pollo, C., Villemure, J.-G., Dawant, B.M., Thiran, J.-Ph., 2002. Atlas-based segmentation of pathological brains using a model of tumor growth. In: Dohi, T., Kikinis, R. (Eds.), Medical Image Computing and Computer-Assisted Intervention MICCAI 2002, Springer, pp. 380–387.

Davies, D.L., Bouldin, D.W., 1979. A cluster separation measure. IEEE Transactions on Pattern Analysis and Machine Intelligence 1 (2), 224–227.

Duda, R.O., Hart, P.E., Stork, D.G, 2001. Pattern Classification, second ed. Wiley.

Evans, A.C., Collins, D.L., Mills, S.R., Brown, E.D., Kelly, R.L., Peters, T.M., 1993. 3D statistical neuroanatomical models from 305 MRI volumes. In: Proceedings of the IEEE Nuclear Science Symposium and Medical Imaging Conference, pp. 1813–1817.

Fletcher-Heath, L.M., Hall, L.O., Goldgof, D.B., Murtagh, F.R., 2001. Automatic segmentation of non-enhancing brain tumors in

magnetic resonance images. Artifical Intelligence in Medicine 21, 43–63.

Gerig, G., Jomier, M., Chakos, M., 2001. VALMET: a new validation tool for assessing and improving 3D object segmentation. In: Niessen, W., Viergever, M. (Eds.), Medical Image Computing and Computer-Assisted Intervention MICCAI 2001, vol. 2208. Springer, New York, pp. 516–523.

Gering, D.T., Grimson, W.E.L., Kikinis, R., 2002. Recognizing deviations from normalcy for brain tumor segmentation. In: Dohi, T., Kikinis, R. (Eds.), Medical Image Computing and Computer-Assisted Intervention MICCAI 2002, vol. 2488. Springer, pp. 388–395.

Ho, S., Bullitt, E., Gerig, G., 2002. Level set evolution with region competition: automatic 3-Dsegmentation of brain tumors. In: Katsuri, R., Laurendeau, D., Suen, C. (Eds.), Proceedings of the 16th International Conference on Pattern Recognition. IEEE Computer Society, pp. 532–535.

Jaccard, P., 1912. The distribution of flora in the alpine zone. New Phytologist 11, 37–50.

Joshi, S., Lorenzen, P., Gerig, G., Bullitt, E., 2003. Structural and radiometric asymmetry in brain images. Medical Image Analysis 7, 155–170.

Kaus, M.R., Warfield, S.K., Nabavi, A., Chatzidakis, E., Black, P.M., Jolesz, F.A., Kikinis, R., 1999. Segmentation of meningiomas and low grade gliomas in MRI. In: Taylor, C., Colchester, A. (Eds.), Lecture Notes in Computer Science, MICCAI, vol. 1679. Springer, pp. 1–10.

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multimodality image registration by maximization of mutual information. IEEE Transactions on Medical Imaging 16 (2), 187–198.

Moon, N., Bullitt, E., Van Leemput, K., Gerig, G., 2002. Automatic brain and tumor segmentation. In: Dohi, T., Kikinis, R. (Eds.), Medical Image Computing and Computer-Assisted Intervention MICCAI 2002, vol. 2489. Springer, pp. 372–379.

Prastawa, M., Bullitt, E., Moon, N., Van Leemput, K., Gerig, G., 2003. Automatic brain tumor segmentation by subject specific modification of atlas priors. Academic Radiology 10, 1341–1348.

Rousseeuw, P.J., Driessen, K., 1999. A fast algorithm for the minimum covariance determinant estimator. Technometrics 41 (3), 212–223.

Sebastian, T.B., Tek, H., Crisco, J.J., Kimia, B.B., 2003. Segmentation of carpal bones from ct images using skeletally coupled deformable models. Medical Image Analysis 7 (1), 21–45.

Tek, H., Kimia, B.B., 1995. Image segmentation by reaction–diffusion bubbles. In: ICCV'95, pp. 156–162.

Tek, H., Kimia, B.B., 1997. Volumetric segmentation of medical images by three-dimensional bubbles. Computer Vision and Image Understanding (CVIU) 65 (2), 246–258.

Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999. Automated model-based bias field correction of MR images of the brain. IEEE Transactions on Medical Imaging 18, 885–896.

Wells, W.M., Kikinis, R., Grimson, W.E.L., Jolesz, F., 1996. Adaptive segmentation of MRI data. IEEE Transactions on Medical Imaging 15, 429–442.

Zhu, S., Lee, T.S.,Yuille, A., 1995. Region competition: unifying snakes, region growing, and Bayes/MDL for multi-band image segmentation. In: ICCV'95, pp. 416–423.