Partial Volume Segmentation of Cerebral MRI scans with Mixture Model Clustering

Aljaž Noe[1] and James C. Gee[2]

[1] Faculty of Electrical Engineering, University of Ljubljana
Tržaška 25, SI-1000 Ljubljana, Slovenija
aljaz.noe@uni-lj.si
[2] Department of Radiology, University of Pennsylvania
1 Silverstein, 3400 Spruce Street, Philadelphia, PA 19104, USA
gee@rad.upenn.edu

Abstract. A mixture model clustering algorithm is presented for robust MRI brain image segmentation in the presence of partial volume averaging. The method uses additional classes to represent partial volume voxels of mixed tissue type in the data with their probability distributions modeled accordingly. The image model also allows for tissue-dependent variance values and voxel neighborhood information is taken into account in the clustering formulation. The final result is the estimated fractional amount of each tissue type present within a voxel in addition to the label assigned to the voxel. A parallel implementation of the method is evaluated using both synthetic and real MRI data.

1 Introduction

A fundamental operation in many applications of medical image analysis remains image segmentation, the object of which is to associate with each image voxel a particular class based on its attributes, neighborhood information, or geometric characteristics of objects belonging to the class. This classification is then used by or to constrain higher-level image analysis and processing algorithms, thus robust and accurate image segmentation is a key element of many medical imaging applications.

In this work we consider the problem of segmenting magnetic resonance (MR) images, which is made difficult by the existence of partial volume (PV) averaging and intensity shading artifacts due to limited spatial resolution of the scanner and RF field inhomogeneity, respectively. To improve the quantitative precision of our segmentation, we focus on the former factor and develop a method for determining the fractional content of each tissue class for so-called partial volume voxels of mixed tissue type. Of specific interest in the current work are the primary tissue constituents of the brain: gray (GM) and white matter (WM) as well as cerebrospinal fluid (CSF).

To our knowledge, two general approaches have been applied to address the problem of partial volume (PV) segmentation. A *mixel model* [1, 2] assumes that every voxel in an image is a PV voxel, consisting of a mixture of pure tissue

classes. The object of segmentation in this case is to determine the relative fraction of each tissue class present within every image voxel. Because of the number of parameters that must be estimated at each voxel, either multi-channel data and/or additional constraints are required to obtain the segmentation solution.

A second approach [3, 4] to dealing with partial volume voxels has been to marginalize over the variables describing the fractional portions of each *pure tissue class*. This produces an additional, new set of *partial volume classes*, with which each image voxel may be associated. In this way, partial volume voxels may be separately identified using existing "binary" segmentation algorithms. However, an additional estimation step is necessary to obtain the fractional amount of the pure tissues in each voxel. In the current work, this method is used to adapt the maximum likelihood mixture model clustering algorithm [5–7] for segmentation of partial volume voxels in MR images of the brain.

2 Image Model

We generalize the image model proposed in [3, 4] to account for tissue-dependent intensity variations. Experiments on MRI data show that differences in intensity variation across tissue type are not insignificant: the intensity values for CSF voxels always having the largest amount of variability, followed by GM and WM.

Let $\mathbf{I}_i = (I_{i,1}, I_{i,2}, \ldots I_{i,M})^T$ be the M-channel observation of the i-th voxel in an input image. Voxels of pure tissue class are described by a particular intensity distribution associated with the image appearance of that tissue type. Partial volume voxels, on the other hand, are represented as a linear combination of the intensity distributions associated with the K possible tissue types that can be found in those voxels:

$$\mathbf{I}_i = \sum_{k=1}^{K} t_{i,k} \mathbf{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k); \quad \sum_{k=1}^{K} t_{i,k} = 1\,, \tag{1}$$

where the voxel intensity I for pure tissue class k is represented as an M-element column vector of random variables, which are distributed according to the multivariate Gaussian distribution $\mathbf{N}$ with $\boldsymbol{\mu}_k = (\mu_{k,1}, \mu_{k,2}, \ldots, \mu_{k,M})^T$ the vector of mean intensity values (M channels) for pure tissue class k, and $\boldsymbol{\Sigma}_k$ is the associated M by M covariance matrix for the M-channel observation. Term $t_{i,k}$ represents the fraction of pure tissue class k that is present at the i-th voxel. Note that the mean intensity values ($\boldsymbol{\mu}_k$) and variances ($\boldsymbol{\Sigma}_k$) do not change with spatial location i; that is, we assume that any intensity shading artifacts in the MRI data are first removed in a preprocessing step.

2.1 Image model simplification

To determine the fractional amount of specified pure tissue classes within every image voxel, we must solve for $N \times (K-1)$ unknowns $t_{i,k}$ from N vector equations

(1), one for each voxel and there are N voxels in the image. Since each vector $\mathbf{I}_i$ has M components, we have $N \times M$ equations. Assuming that the tissue class parameters ($\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$) are known, a solution can be found if $M \geq (K-1)$. In practice, we are interested in the three classes: CSF, GM and WM. Multi-echo images of high resolution are generally not available and even these would be partially correlated and noisy, so that the problem remains ill posed.

Additional constraints are therefore necessary and same as in [3, 4], we make the assumption that each partial volume voxel is a mixture of only two tissue types, which introduces negligible error in practical applications that use high-resolution T1 data. Formally we define a number of sets $\mathcal{G}_k$ each containing indices of pure classes that are present in the k-th PV class:

$$\mathcal{G}_k = \{k_1, k_2\}; \qquad k = 1 \ldots K_{PV}, \quad k_1, k_2 \in \{1 \ldots K\}\,, \tag{2}$$

where K_{PV} is the number of PV classes in an image while K is the total number of pure tissue classes. For voxels of pure tissue class k and PV voxels consisting of pure classes k_1 and k_2, respectively, (1) reduces to:

$$\mathbf{I}_i = \mathbf{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \tag{3}$$

and

$$\mathbf{I}_i = t_{i,k_1}\mathbf{N}(\boldsymbol{\mu}_{k_1}, \boldsymbol{\Sigma}_{k_1}) + t_{i,k_2}\mathbf{N}(\boldsymbol{\mu}_{k_2}, \boldsymbol{\Sigma}_{k_2})\,; \quad t_{i,k_1} + t_{i,k_2} = 1\,. \tag{4}$$

3 Mixture Model Clustering

To determine the parameters ($\boldsymbol{\mu}_k$, $\boldsymbol{\Sigma}_k$) for the pure tissue classes, an extended version of the maximum likelihood mixture model algorithm [5–7] was developed. First, appropriate probability density functions are described for the pure tissue and PV classes. Second, we introduce a Gibbs model as the weighting function to favor spatially extended classifications [4]. Finally, $P(k|\mathbf{I}_i)$ is determined and used to estimate the parameters $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$.

3.1 Probability density functions

The intensities of voxels belonging to pure tissue class k conform to a multivariate normal distribution: $\mathbf{I}_i = \mathbf{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$. The corresponding probability density function for observing intensity $\mathbf{I}_i$ given tissue class k is therefore given by:

$$P(\mathbf{I}_i|k) = \frac{\exp\Big(-\frac{1}{2}(\mathbf{I}_i - \boldsymbol{\mu}_k)^T {\boldsymbol{\Sigma}_k}^{-1}(\mathbf{I}_i - \boldsymbol{\mu}_k)\Big)}{\sqrt{(2\pi)^M\,|\boldsymbol{\Sigma}_k|}}\,; \quad k = 1 \ldots K\,. \tag{5}$$

The probability density function for PV voxels containing a mixture of pure tissue classes k_1 and k_2 is derived from the model (4) and involves a linear

combination of two Gaussian distributions:

$$P_{PV}(\mathbf{I}_i|k_1,k_2,t) = \frac{\exp\Big(-\frac{1}{2}\big(\mathbf{I}_i-\hat{\boldsymbol{\mu}}_k(t)\big)^T\hat{\boldsymbol{\Sigma}}_k(t)^{-1}\big(\mathbf{I}_i-\hat{\boldsymbol{\mu}}_k(t)\big)\Big)}{\sqrt{(2\pi)^M\left|\hat{\boldsymbol{\Sigma}}_k(t)\right|}}\,; \qquad (6)$$

$$\hat{\boldsymbol{\mu}}_k(t) = t\boldsymbol{\mu}_{k_1} + (1-t)\boldsymbol{\mu}_{k_2}\,,$$
$$\hat{\boldsymbol{\Sigma}}_k(t) = t^2\boldsymbol{\Sigma}_{k_1} + (1-t)^2\boldsymbol{\Sigma}_{k_2}\,.$$

As in [3, 4], we then marginalize (6) over t to obtain the probability density function for PV classes:

$$P(\mathbf{I}_i|k) = \int_0^1 P_{PV}(\mathbf{I}_i|k_1,k_2,t)dt\,; \qquad (7)$$
$$k = K+1\ldots K+K_{PV}\,, \quad k_1,k_2\in\mathcal{G}_{k-K}\,, \quad k_1\neq k_2\,.$$

To generalize the notation, we have numbered the PV classes from $K+1$ to $K+K_{PV}$, so that $P(\mathbf{I}_i|k)$ expresses the probability density for both pure tissue and PV classes. The integral in (7) does not have a closed form solution and must therefore be evaluated by numerical integration.

Shape of PV probability density function depends largely on parameters of the gaussians that define pure classes, which are present in PV class. This is particularly noticeable, when the difference in mean values of two pure classes is of the same order as standard deviation. Sample probability density functions are shown in Fig. 1.

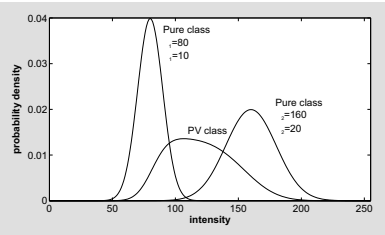

Fig. 1. Sample probability density functions for pure and partial volume voxels

3.2 Weighting functions

In [5, 6] the probability density function for class k is weighted by the current estimate of the voxel count for that class. This weighting is used to update the probabilities in a manner similar to that of a Bayesian prior. Here we introduce an alternative weighting function that favors segmentations which are spatially

extended. Specifically, we use the familiar Potts model that is also applied in [4]:

$$P_i(k) = \frac{1}{Z} \exp\left(-\beta \cdot \sum_{j \in \mathcal{N}_i} \frac{\delta(k, k_j)}{d(i,j)} \right); \quad k = 1 \ldots K + K_{PV} \,, \tag{8}$$

$$k_j = \arg\max_{k'} \left(P(\mathbf{I}_j | k') \right) \,,$$

where

$$\delta(k_1, k_2) = \begin{cases} -2 & \text{if } k_1 = k_2 \\ -1 & \text{if } k_1 \in \mathcal{G}_{k_2 - K} \vee k_2 \in \mathcal{G}_{k_1 - K} \\ -1 & \text{if } \mathcal{G}_{k_1 - K} \cap \mathcal{G}_{k_2 - K} \neq \{\} \\ +1 & \text{otherwise} \end{cases}; \quad k_1, k_2 \in \{1 \ldots K + K_{PV}\}; \tag{9}$$

k_i is the current ML class estimate for the voxel at location i in the image; k is the class for which we are updating the prior; $\mathcal{N}_i$ is the set of D18 neighborhood voxels of voxel i; β is a parameter of the distribution, controlling the amount of influence the weighting function should exert on the likelihood function; and Z is a normalizing constant. Function $d(i,j)$ represents the distance between voxels i and j, which is used to limit the influence of distant neighborhood voxels.

3.3 Parameter Estimation

Given the probability density and weighting functions, the conditional probability $P(k|\mathbf{I}_i)$ is calculated, from which an estimate of the parameters $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$ for each pure tissue class k can then be estimated as follows:

$$P(k|\mathbf{I}_i) = \frac{P_i(k) P(\mathbf{I}_i|k)}{\sum_{k'=1}^{K+K_{PV}} P_i(k') P(\mathbf{I}_i|k')}; \quad k = 1 \ldots K + K_{PV} \tag{10}$$

$$\boldsymbol{\mu}_k = \frac{\sum_{i=1}^{N} P(k|\mathbf{I}_i) \cdot \mathbf{I}_i}{h_k}; \qquad h_k = \sum_{i=1}^{N} P(k|\mathbf{I}_i) \,, \quad k = 1 \ldots K \,. \tag{11}$$

$$\boldsymbol{\Sigma}_k = \frac{\sum_{i=1}^{N} P(k|\mathbf{I}_i) \cdot \mathbf{I_i} \cdot \mathbf{I}_i^T}{h_k} - \boldsymbol{\mu}_k \cdot \boldsymbol{\mu}_k^T$$

These parameter estimates then yield new probability density functions and the process is repeated until the voxel count in each pure tissue class does not change from one iteration to the next. This approximation to maximizing the likelihood [5, 6] can also be seen as a special case of an Expectation Maximization (EM) algorithm, suited for Gaussian distributions. In steady state at least a local likelihood maximum is guarantied.

3.4 Initialization

Based on extensive experimentation on real and simulated MR images, we have found that the clustering algorithm can be made robust to initialization values by

specifying a sufficiently large class variance. Therefore, without additional prior information about the mean and variance values of the classes, the parameters are initialized as follows:

$$\mu_{k,l} = \frac{k}{K+1}\left(\max_i\left(I_{i,l}\right) - \min_i\left(I_{i,l}\right)\right) + \min_i\left(I_{i,l}\right) \tag{12}$$
$$\sigma_{k,l,l} = \left(\max_i\left(I_{i,l}\right) - \min_i\left(I_{i,l}\right)\right)^2 \;; \quad k = 1\ldots K,\; l = 1\ldots M$$

Initial mean intensity values are equally distributed between the minimum and maximum intensity values found in the image. Diagonal elements of the covariance matrix are all set to the image intensity range, whereas off-diagonal elements are set to zero.

4 Partial Volume Tissue Classification

The clustering algorithm determines $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$ by iterating over the estimation of $P(k|\mathbf{I}_i)$, until convergence is achieved. Once the intensity distribution and all class parameters are known for each tissue type, the fractional portion t_{i,k_1} for a PV voxel at location i consisting of tissues k_1 and k_2 can then be obtained from (4) via maximum likelihood estimation (MLE):

$$t_{i,k_1} = \frac{(\boldsymbol{\mu}_{k_1} - \boldsymbol{\mu}_{k_2})^T(\mathbf{I}_i - \boldsymbol{\mu}_{k_2})}{(\boldsymbol{\mu}_{k_1} - \boldsymbol{\mu}_{k_2})^T(\boldsymbol{\mu}_{k_1} - \boldsymbol{\mu}_{k_2})}\,. \tag{13}$$

To allow us to produce a segmentation without having to specify a threshold for distinguishing between partial volume and pure tissue voxels, we need to modify (13) to include the information about pure tissue classes. We can write:

$$t^*_{i,k} = P\left(k\,|\mathbf{I}_i\right) + \sum_{k'} P\left(k' + K\,|\mathbf{I}_i\right)\frac{(\boldsymbol{\mu}_k - \boldsymbol{\mu}_{k_2})^T(\mathbf{I}_i - \boldsymbol{\mu}_{k_2})}{(\boldsymbol{\mu}_k - \boldsymbol{\mu}_{k_2})^T(\boldsymbol{\mu}_k - \boldsymbol{\mu}_{k_2})}\,; \tag{14}$$
$$k_2 \in \mathcal{G}_{k'} \wedge k_2 \neq k$$

where summation index k' runs over all PV classes that contain pure class k (for which $k \in \mathcal{G}_{k'}$ is true). We must also normalize the portions of pure classes so that they sum to unity over all classes k:

$$t_{i,k} = \frac{t^*_{i,k}}{\sum_{k'=1}^{K} t^*_{i,k'}}\;; \quad k = 1\ldots K \tag{15}$$

5 Implementation

Two preprocessing steps must be performed prior to clustering and segmentation. First, we extract the brain parenchyma from the MR image of the head using the Brain Extraction Tool – details of the method can be found in [8]. Intensity

shading artifacts in the extracted image are then removed with the MNI-N3 method [9].

A parallel version of the clustering algorithm was implemented by subdividing the image into a number of segments, which are then processed in separate threads, one for each processor available. All threads are synchronized at 3 time points: before and after the calculation of the weighting values and before the estimation of the new class parameters. The algorithm is outlined below:

1. Initialization
   - Select number of pure tissue classes K.
   - Define PV classes, represented as sets $\mathcal{G}_k$, by specifying the corresponding combinations of pure tissue classes.
   - Set initial estimates of class parameters ($\boldsymbol{\mu}_k$,$\boldsymbol{\Sigma}_k$) using (12).
2. Calculate the probability densities for all classes using (5) and (7) in multiple threads. Wait until all threads complete their processing before proceeding.
3. Calculate the weighting values in multiple threads using (8). Wait until all threads complete their processing before proceeding.
4. Calculate the updated probabilities using (10) for each class k and the new estimates for the class parameters using (11). Wait until all threads complete their processing before proceeding.
5. If the current voxel count for each pure tissue class is different from that found in the previous iteration, or the maximum allowed number of iterations have not been processed, return to step 2. We terminate the loop when the change in $\sum_{k=1}^{K} h_k$ between iterations is less than 1 or number of iterations is 50. h_k is defined in (11).
6. Segment the image by determining the fractional amount of each tissue type with every image voxel using (14) and (15).

The algorithm was implemented in standard ANSI C code. Moreover, our multithreading support conforms to both the POSIX-PTHREADS standard and Microsoft Windows multithreading API, so the code will compile on both UNIX and Windows workstations.

6 Experimental Results

The segmentation algorithm was evaluated using both synthetic and real data. In each of the reported experiments, β was set to 0.3 and algorithm convergence usually occurred after 10-20 iterations.

6.1 Synthetic Image

We constructed a square, 100 by 100, image and subdivided the image into 3 regions, each separated by a vertical boundary. The left and right most regions were considered pure "tissues" and their image values were drawn from normal distributions with the following mean and variance values, respectively: $\mu_1 = 70$, $\Sigma_1 = 10$ and $\mu_2 = 150$, $\Sigma_2 = 20$. The middle strip of the image, 30 pixels

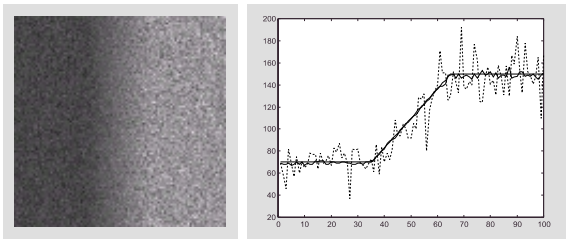

Fig. 2. Synthetic data. (Left) Image to be segmented. (Right) Plotted in dashed line is the horizontal intensity profile obtained along line 50 of the image; solid thin line is the mean horizontal intensity profile averaged over all lines in the image; and in solid thick line is the ideal horizontal profile (without noise added).

wide, contained partial volume pixels, which modeled a smooth linear transition between the two pure classes. The synthetic image is shown in Fig. 2.

The following are the estimated mean and variances for the tissue classes: $\mu_1 = 70.35$, $\Sigma_1 = 10.05$; $\mu_2 = 148.34$, $\Sigma_2 = 19.22$. The segmentation results or t values for the first "tissue" class are shown in Fig. 3. The figure also shows the squared error between the ideal and estimated t values for the class – the total error was $E_1 = 26.65$, where

$$E_{i,k} = \left(t_{i,k} - t_{i,k}^{ideal}\right)^2 , \quad E_k = \sum_{i=1}^{N} E_{i,k} \, . \tag{16}$$

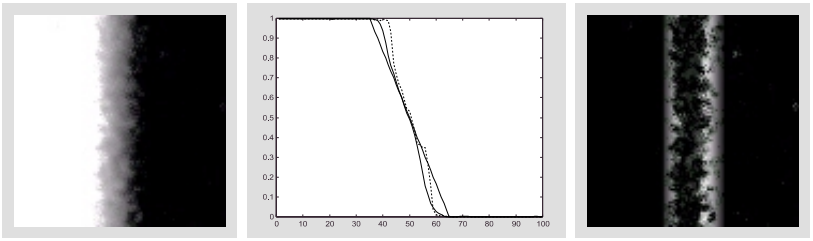

Fig. 3. Segmentation results for the synthetic data. (Left) Fractional values t for the first class at each voxel plotted as an 8-bit gray-scale image with intensity = 0 corresponding to $t = 0.0$ and intensity = 255 to $t = 1.0$. (Middle) Plotted in dashed line is the horizontal intensity profile obtained along line 50 of the segmentation; solid thin line is the mean horizontal intensity profile averaged over all lines in the segmentation image; and thick line is the ideal horizontal profile. (Right) Image of pointwise squared error between estimated and ideal t values for the first class

We can see that the errors occur only at the boundaries where the region with PV voxels meets the regions containing pure classes. We contribute this error largely to noise because it decreases when we reduce the amount of noise variance for the pure classes. This also explains the smaller amount of error in the segmentation of the left half of the image, where the noise variance for the first pure class was smaller.

6.2 Simulated T1-weighted Brain Volume

A second, more realistic synthetic dataset of an MRI head scan was created using the Brain-Web simulator [10–13]. Each simulation was a $1mm^3$ isotropic MRI volume with dimensions 181x217x181. Three datasets incorporating different amounts of noise were segmented and the mean absolute error between the ideal and estimated t values over all voxels were as follows:

- 9% noise: GM: 0.08458 (σ=0.11885); WM: 0.04399 (σ=0.08759); CSF: 0.04157 (σ=0.09795)
- 3% noise: GM: 0.05435 (σ=0.08597); WM: 0.02923 (σ=0.06414); CSF: 0.02585 (σ=0.06517)
- 0% noise: GM: 0.03874 (σ=0.06301); WM: 0.01936 (σ=0.03755); CSF: 0.02077 (σ=0.05612)

A segmented slice for the synthetic brain volume with 9% noise level is shown in Fig. 4. Although there appears to be minimal partial volume averaging in the results, the segmentation obtained without using PV classes ($K_{PV} = 0$) had errors about 2 times larger and the algorithm took much longer to converge (> 50 iterations).

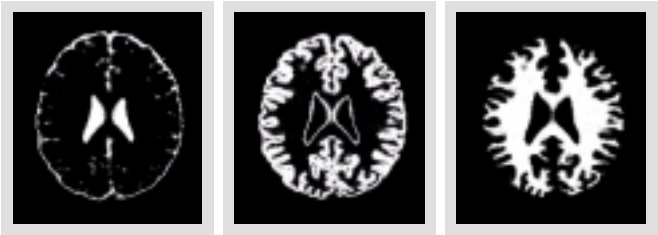

Fig. 4. Partial volume segmentation of simulated brain volume. (Left) Cerebrospinal fluid, where fractional values t at each voxel are plotted as an 8-bit gray-scale image with intensity = 0 corresponding to $t = 0.0$ and intensity = 255 to $t = 1.0$ (Middle) Gray matter. (Right) White matter.

6.3 Manually Segmented Real T1 MR Images of the Brain

Twenty normal brain MRI datasets and their manual segmentations were obtained from the Center for Morphometric Analysis at Massachusetts General Hospital – these IBSR datasets are publicly available on Internet [14]. The volumes were preprocessed to extract brain parenchyma and corrected for intensity

inhomogeneities. However, 7 of the preprocessed volumes still exhibited strong shading artifacts that the MNI-N3 method [9], was unable to remove.

Since the manual segmentations for this set of images do not contain any information about fractional tissue content, we calculated a similarity index for each class by thresholding our partial volume segmentation results. Specifically, we report below the values for the Jaccard similarity $= |S_e \cap S_{ideal}|/|S_e \cup S_{ideal}|$, where S_e and S_{ideal} are the estimated and "true" sets of voxels, respectively, for a given tissue class:

Image	100_23	110_3	111_2	112_2	11_3	12_3*	13_3	15_3*	16_3	17_3
GM	0.8339	0.8216	0.8115	0.7563	0.7982	0.7102	0.8455	0.4029	0.7203	0.7341
WM	0.7528	0.7078	0.7391	0.6799	0.7238	0.5848	0.7777	0.001	0.6408	0.6283

Image	191_3	1_24*	202_3	205_3	2_4*	4_8*	5_8*	6_10*	7_8	8_4
GM	0.8195	0.5693	0.8426	0.8237	0.5334	0.5551	0.3764	0.4333	0.7761	0.7395
WM	0.7408	0.4405	0.7633	0.7687	0.2679	0.2697	0.1309	0.2237	0.6842	0.6657

Table 1. Jaccard similarity measure between estimated and 'true' segmentation of IBSR datasets. With star (*) are marked images, for which the intensity inhomogeneity couldn't be properly corrected.

It is evident in the results above which of the volumes were the seven for which intensity inhomogeneity was still prevalent in the data. Excluding these volumes, the mean Jaccard index was 0.7837 and 0.6985 for GM and WM, respectively, and these are superior to the results reported for other methods in [4, 15].

7 Conclusion

We have presented an algorithm for partial volume segmentation of MR images of the brain. Experimental results are comparable or superior to other published algorithms. Our method is an extension of a probabilistic clustering algorithm [5, 6], to accommodate partial volume voxels and to allow class-dependent model values for the intensity variance. Although the convergence properties of the original technique are generally unknown, we have observed robust performance from our implementation as a function of the estimates used to initialize the class parameters. In the current work, the weighting function was augmented to favor spatially contiguous regions in the segmentation but other possibilities are being examined, including the use of prior anatomic information as in [7]. Another, more important feature that is under implementation is the simultaneous correction of intensity inhomogeneities to not only obviate the need for this preprocessing step but to improve on existing techniques.

References

1. Hwan Soo Choi, David R. Haynor, and Yongmin Kim. Partial volume tissue classification of multichannel magnetic resonance images - a mixel model. In *IEEE Transactions on Medical Imaging*, volume 10, pages 395–407, September 1991.
2. Lucien Nocera and James C. Gee. Robust partial volume tissue classification of cerebral MRI scans. In K. M. Hanson, editor, *SPIE Medical Imaging*, volume 3034, pages 312–322, February 1997.
3. David H. Laidlaw, Kurt W. Flescher, and Alan H. Barr. Partial-volume bayesian classification of material mixtures in MR volume data using voxel histograms. In *IEEE Transactions on Medical Imaging*, volume 17, pages 74–86, February 1998.
4. David W. Shattuck, Stephanie R. Sandor-Leahy, Kirt A. Schaper, David A. Rottenberg, and Richard M. Leahy. Magnetic resonance image tissue classification using a partial voume model. 2000. Submitted.
5. John A. Hartigan. *Clustering algorithms*. John Wiley & Sons, Inc., New York, 1975.
6. Richard O. Duda and Peter E. Hart. *Pattern classification and scene analysis*. John Wiley & Sons, Inc., New York, 1973.
7. J. Ashburner and K. Friston. Multimodal image coregistration and partitioning - a unified framework. In *Neuroimage*, volume 6, pages 209–217, October 1997.
8. Stephen M. Smith. Robust automated brain extraction. In *Sixth Int. Conf. on Functional Mapping of the Human Brain*, page 625, 1998.
9. John G. Sled, Alex P. Zijdenbos, and Alan C. Evans. A nonparametric method for automatic correction of intensity nonuniformity in MRI data. In *IEEE Transactions on Medical Imaging*, volume 17, pages 87–97, February 1998.
10. `http://www.bic.mni.mcgill.ca/brainweb/`.
11. C. A. Cocosco, V. Kollokian, R.K.-S. Kwan, and A.C. Evans. BrainWeb: Online Interface to a 3D MRI Simulated Brain Database. In *NeuroImage*, volume 5, May 1997.
12. R.K.-S. Kwan, A.C. Evans, and G. B. Pike. *An Extensible MRI Simulator for Post-Processing Evaluation*, volume 1131 of *Lecture Notes in Computer Science*, pages 135–140. Springer-Verlag, May 1996.
13. D. Louis Collins, A.P. Zijdenbos, V. Kollokian, J.G. Sled, N.J. Kabani, C.J. Holmes, and A.C. Evans. Design and Construction of a Realistic Digital Brain Phantom. In *IEEE Transactions on Medical Imaging*, volume 17, pages 463–468, June 1998.
14. `http://neuro-www.mgh.harvard.edu/cma/ibsr`.
15. Jagath C. Rajapakse and Frithjof Kruggel. Segmentation of MR images with intensity inhomogeneities. In *Image and Vision Computing*, volume 16, pages 165–180, 1998.