LESION DETECTION IN NOISY MR BRAIN IMAGES USING CONSTRAINED GMM AND ACTIVE CONTOURS

Oren Freifeld, Hayit Greenspan

Bio-medical Engineering, Tel-Aviv University

Jacob Goldberger

School of Engineering, Bar-Ilan University

ABSTRACT

This paper focuses on the detection and segmentation of multiple sclerosis (MS) lesions in magnetic resonance images. The proposed method performs healthy tissue segmentation using a probabilistic model for normal brain images. MS lesions are simultaneously identified as outlier Gaussian components. The probablistic model, termed constrained-GMM, is based on a mixture of many spatially-oriented Gaussians per tissue. The intensity of a tissue is considered a global parameter and is constrained to be the same value for a set of related Gaussians per tissue. An active contour algorithm is used to delineate lesion boundaries. Experimental results on both standard brain MR simulation data and real data, indicate that our method outperforms previously suggested approaches especially for highly noisy data.

Index Terms— MRI, MS-lesions, statistical modeling

1. INTRODUCTION

Multiple sclerosis (MS) is the most common non traumatic neurological disease in young adults. MRI is currently being used for diagnosis of MS, assessment of disease progression, and evaluation of the efficiency of drug therapy. The most common quantitative parameter is the burden (or load) of the disease expressed in terms of the number and volume of the brain lesions. Traditionally, the multiple sclerosis lesion (MSL) segmentation task is treated via statistical voxel-based intensity modeling (e.g. [6, 7, 8]). In [6], the intensity of normal brain tissues is modeled by mixture of Gaussian probability distribution functions (GMMs). The model consists of a single Gaussian per tissue type and lesions are treated as outlier voxels. Voxels are labeled as MSLs if they are brighter than the average gray matter (GM) voxel, and if most of their 3-D neighbors belong to MSLs or to the white matter (WM) tissue. Noise present in the MRI input may cause unrealistic results, where tissue regions as well as lesion regions may appear granular, fragmented, or violating anatomical constraints. In most works, extensive pre-processing is thus applied including noise removal and intensity normalization. Atlas registration is ofetn used (e.g. [6, 8]) requiring a computationally intensive co-registration procedure. Another conventional method to improve segmentation smoothness and immunity to noise is to use a Hidden Markov Random Field (HMRF)[6], thus modeling neighboring voxel interactions. Smoother structures are obtained in the presence of moderate noise as long as the HMRF parameters controlling the strength of the spatial interactions are properly selected. HMRF based algorithms are computationally intractable unless some approximation is used which still requires computationally intensive algorithms.

In very recent works new approaches are starting to emerge. In [1] regional properties are used in a multi-feature and multiscale approach combining segmentation and classification (supervised) for MSL analysis. In [4] a 4-D space (T1 intensity and spatial features) multi-Gaussian model is used to model each tissue. A large number of Gaussians is used per brain tissue in order to capture the complex spatial layout. The intensity of a tissue is considered a global feature and is incorporated into the model by parameter tying of all related Gaussians. The model, termed constrained GMM (CGMM), was shown in [4] to provide accurate segmentation of T1-weighted simulated and real MR brain images into the three healthy tissues, in particular in noisy and low-contrast MR images, without the need for co-registration of the input image and an atlas. In the current work we propose an extension to the CGMM model: a generalization of the model to handle multi-sequence MRI data with a focus on MSL detection and segmentation.

2. THE CGMM SEGMENTATION FRAMEWORK

The complex spatial layout of an MRI brain image provides a challenge for conventional GMM modeling schemes. In order to accommodate the spatial complexity, we model an image as a mixture of *many* Gaussians:

$$f(I(x), x|\Theta) = \sum_{i=1}^{n} \alpha_i f_i(I(x), x; \mu_i, \Sigma_i) \qquad (1)$$

such that $I(x)$ is the intensity vector associated with the x-th voxel, x is the 3-D position information included in the feature vector, n is the number of components in the mixture model, μ_i and Σ_i are the mean and the covariance of the i-th Gaussian component f_i, and α_i is the i-th mixture coefficient.

This research was funded in part by a grant from the Israeli Ministry of Science, for strategic research directions.

In the following paragraph, to simplify the presentation, we will consider the lesion matter as a fourth tissue in addition to the white matter (WM), gray matter (GM) and cerebrospinal fluid (CSF). The spatial shape of the tissues is highly non-convex. However, since we use a mixture of many components, each Gaussian component models a small local convex region. The intra variability of the intensity feature within a tissue is significantly less than the inter-variability among different tissues. It is therefore sufficient to model the intensity variability within a tissue by a single Gaussian (in the intensity feature). To incorporate this insight into the model, each Gaussian is linked to a single tissue and all the Gaussians related to the same tissue share the same intensity parameters. The above assumptions impose the following structure on the mean and variance of the Gaussian components:

$$\mu_i = \begin{pmatrix} \mu_i^x \\ \\ \mu_{\pi(i)}^I \end{pmatrix} \quad , \quad \Sigma_i = \begin{pmatrix} \Sigma_i^x & 0 \\ \\ 0 & \Sigma_{\pi(i)}^I \end{pmatrix} \qquad (2)$$

where x is the spatial feature vector, I is the intensity feature vector, $\pi(i) \in 1,..,4$ is the tissue that is linked to the i-th Gaussian component and μ_j^I and Σ_j^I , $j = 1,..,4$ are the mean and variance parameters of all the Gaussian components that are linked to the j-th tissue.

3. LESION DETECTION UTILIZING THE CONSTRAINED GMM

The EM algorithm is applied on the MR data to find the maximum-likelihood parameter-set for the CGMM. The EM however, is notoriously known to get stuck on local maxima; hence it requires appropriate initialization. Utilizing the CGMM model for lesion detection, the problem is even more severe since the lesions are small and isolated. We present a novel semi-supervised top-down initialization method for the CGMM that utilizes our knowledge about the number of tissues of interest, as well as our knowledge about the general appearance of MSL in different modalities of MRI. This method is an improved variation on the initialization method from [4]. The initialization is divided into three steps.

In the first step, K-means clustering is done based on the intensity features ($T1$, $T2$ and PD), which gives a rough segmentation into the three major tissue classes (WM, GM and CSF). We assume the WM voxels to be of highest $T1$ intensity, GM voxels to be of second highest $T1$ intensity, and the CSF voxels to be of lowest $T1$ Intensity. This rough but important initial clustering determines an initial group of voxels that belong to each tissue. We use this segmentation to calculate the approximate (spatial) position of the Gaussians in each tissue.

The second step is a top-down procedure. We start by iteratively splitting regions (per tissue) until we obtain convex regions that are suitable for Gaussian modeling (see [4] to see how the convexity is obtained). In [4] each split was done according to K-means clustering based only on spatial features. Here, we do the same for large regions; for regions smaller than a pre-defined threshold, K-means is performed on spatial and intensity features, with higher weight given to the intensity features. We can use the intensity features only when the region is small enough. Otherwise, we will miss the desired spatial structure. This modification helps in separating MSL regions from the regions of the tissue to which they were misclassified. We still don't know at this step which regions, if any, correspond to MSL. Once the regions are determined, the initial estimation for the Gaussian parameters is calculated as described in [4]. An illustration of this step can be seen in Figure 1a.

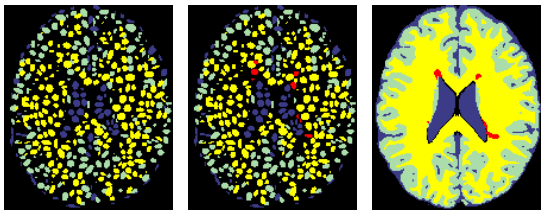

(a) (b) (c)

Fig. 1. Brainweb data: CGMM illustration on slice 96 with 9% noise. (a) A 2-D (2σ) spatial projection of Gaussian components before MSL detection (b) after MSL detection (c) Ground Truth. Color legend: Gray - GM, Yellow-WM, Blue - CSF, Red -MSL.

The third step detects the Gaussians that should be associated with MSL, and creates a fourth class from them. Intensity parameters are computed for each Gaussian and for each tissue (by averaging the parameters of the Gaussians that are linked to the tissue). A set of rules is then used to distinguish the MSL blobs from the normal tissue blobs. A Gaussian (blob) for which all these conditions hold, is labeled as MSL. An example of such a set of rules: T1 mean-intensity per blob $>$ T1 mean-intensity of the GM tissue, T2 mean-intensity per blob $>$ T2 mean-intensity of the GM tissue, PD mean-intensity per blob $>$ PD mean-intensity of the GM tissue -ϵ_1 and the Euclidean distance between the mean-intensity vectors of the blob and the GM tissue $> \epsilon_2$.

All the re-labeled blobs now form a fourth class named MSL. Rules similar to the first three rules above, were used in other works. For example, in [6], the T2 and PD grey values of outlier voxels were compared to the GM mean (along with additional rules). However, note that at this point we take decisions at the blob level rather than the pixel level. This makes the lesion detection much more robust to noise.

Following initialization, CGMM parameters are learned via the EM algorithm. Once the EM process is finalized, decisions are made at a voxel level, and a hard segmentation is induced from the model parameters. The CGMM-based lesion detection process is illustrated in Figure 1.

4. LESION BOUNDARY DELINEATION BASED ON ACTIVE CONTOURS

The segmentation induced from the CGMM yields good results in terms of detection of lesions. However, there is usually an overestimation in the size of the lesions. This is likely the result of the fact that the area in the proximity of a lesion is somewhat different from the other tissues. In addition, the segmentation step, after the lesion model was learned, is done pixel-wise without considering the smoothness of the lesion boundaries. Hence a post-processing step for refining the lesion boundary is needed. We use a modified version of the variational framework for segmentation of vector valued images, suggested by Rousson and Deriche [5]. They derived a curve evolution equation for object-background image segmentation that is based on statistical properties of the two regions. The intensity in each region is modeled using a Gaussian distribution. In our case, the lesion is the object and pixels around it are the background. The optimal lesion segmentation is obtained by minimizing the energy:

$$F(\mu_1, \Sigma_1, \mu_2, \Sigma_2, \partial\Omega) = \sum_{i=1}^{2} \int_{\Omega_i} e_i(I(x)) + \text{length}(\partial\Omega) \quad (3)$$

where Ω_1 are the lesion region pixels, Ω_2 are the non-lesion pixels, $\partial\Omega$ is the curve separating the two regions, μ_i, Σ_i are the Gaussians parameters of Ω_i and $e_i(I) = -\log f(I|\mu_i, \Sigma_i)$. The energy (local) minimum can be obtained utilizing the level-set theory via a gradient descent (GD) method. From the GD equation, $\phi_t = -\frac{\delta F}{\delta \phi}$, we obtain the following curve evolution equation:

$$\phi_t(x) = \left(v \cdot div(\frac{\nabla \phi}{|\nabla \phi}|) + e_2(I(x)) - e_1(I(x)) \right) \delta_\epsilon(\phi(x))$$

In [5] energy is minimized with respect to both the Gaussian parameters and the boundary position. In our case we avoid the parameter re-estimation step. Instead, we use the global intensity parameters learned via the CGMM model. This enables us to continue enforcing in the active contour step the global constraints of the CGMM. For example, if one of the slices has an initial contour that encircles only false positives, then the constraint will usually cause it to die off, while parameter estimation during the curve evolution is likely to cause the system to learn the parameters of the false lesion. Since there are three tissues that belong to Ω_2 , we define $e_2(I(x))$ to be $\min(e_{CSF}, e_{GM}, e_{WM})$. This definition implies that the last two terms in the curve evolution equation are the log-likelihood ratio between MSL and the most likely normal-brain tissue.

Since brain images are complex, and since lesions are relatively small, it is impractical to use the standard initial contour (small circles) as in [5, 2] and others. Hence, we use the segmentation of the lesions from the CGMM model as an initialization for the curve. Note that the method proposed in [5] can not be directly applied on the intensity information to find the three tissues and the lesions, simultaneously. The method is highly sensitive to the initial condition for complicated N-phases images (such as brain images) as the GD method can find only local minima of the functional [5]. The lesion detection obtained using the CGMM, therefore, makes it possible to apply active contour techniques for lesion delineation.

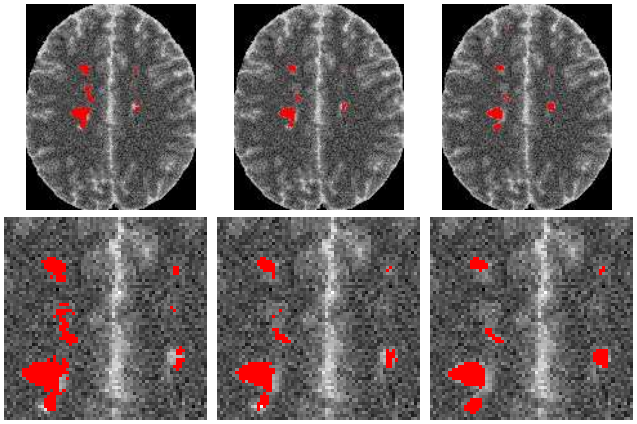

Fig. 2. Brainweb data: lesion segmentation results on slice 104 with 9% noise. Lesions are shown overlaid on T2. Second row: lesion area magnification.

5. EXPERIMENTAL RESULTS

The algorithm was tested on bias free simulated MRI data [1] [3] with the following co-registered modalities: T1, T2 and PD. Voxel size was $1mm^3$. The experiments were done on 61 slices (slices 60-120) which contain 93% of the lesion burden in these data sets. The level of noise was 3,5,7 and 9 (percent ratio of the std. of the white Gaussian noise vs the signal for a reference tissue). Brain region was extracted in a pre-processing step. Figure 2 shows slice no. 104 from the simulated data with noise level of 9%. In the initialization step region splitting based solely on spatial features was done for regions with at least 280 voxels and region splitting based also on intensity (with weight 4 times higher than the spatial features) was done on regions with a number of voxels between 70 and 280. MSL blobs detection was done according to the rules described in Section 3, with PD threshold, ϵ_1, empirically set to 0.2 and ϵ_2 set to 0.4 (similar results were obtained for values varying from 0.3 to 0.5). EM iterations for the CGMM were performed, following which we applied the active contour (AC) step to refine the lesion segmentation. The velocity, v, in the curve evolution equation, was chosen as a constant:5,3,1,1, per noise level (set higher for low noise levels).

[1] BrainWeb: http://www.bic.mni.mcgill.ca/brainweb

Method	3%	5%	7%	9%
KVL	80	73	61	47
CGMM	57	61	65	65
CGMM+AC	77	77	75	73

Table 1. Dice overlap metric results (%) on BrainWeb data

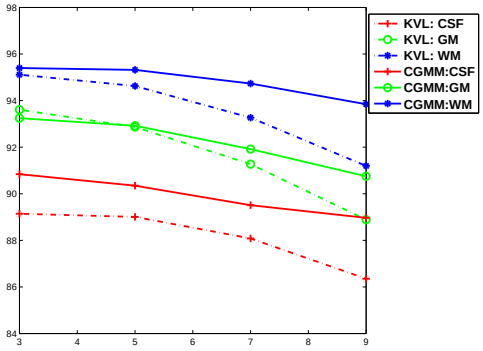

Fig. 3. MS patient normal tissue segmentation. Dice(%) vs noise level

Figure 2 shows lesion segmentation results using the CGMM framework. Performance is evaluated via a comparison with K. Van-Leemput's algorithm (hereon termed KVL) [6] implemented by the EMS software package [2]. In the KVL implementation the statistical brain atlas of the SPM99 [3] was normalized to the target brain volume images. The comparison between the algorithms is made using the Dice metric which measures the overlap between the automatic segmentation and the ground truth for each tissue k, as $\frac{2V_{as}^k}{V_a^k+V_e^k}$, where V_{as}^k is the number of voxels that are assigned to tissue k by both the ground truth and the automated algorithm, V_a^k and V_e^k denote the number of voxels assigned to tissue k by the algorithm and the ground truth, respectively. Table 1 shows Dice metric results for KVL, CGMM and CGMM followed by AC algorithms. The CGMM performance improvement over the KVL in the presence of noise is notable. Adding on the AC step shows a substantial increase in performance. The performance of the CGMM for segmentation of the three healthy tissues, on the registered multi-modal (T1,T2,PD) MRI information, is demonstrated in Figure 3. CGMM shows an improved performance over the KVL. CGMM Results (9% case) are: $94\%, 91\%$ and 89% for WM, GM and CSF tissues, respectively. For KVL : $92\%, 89\%$ and 86%, respectively. An increase is also evident over the T1 channel only [4] (in particular for the CSF): $93\%, 87\%$ and 78%, respectively.

Promising results of the proposed algorithm on real MRI data can be seen in Figure 4. In this experiment, a Fast Flair (FF) sequence was added to T1,T2 and PD, and a similar set of rules that is also based on the FF channel was used. The

[2]http://www.medicalimagecomputing.com/EMS

[3]http://www.fil.ion.ucl.ac.uk/spm/

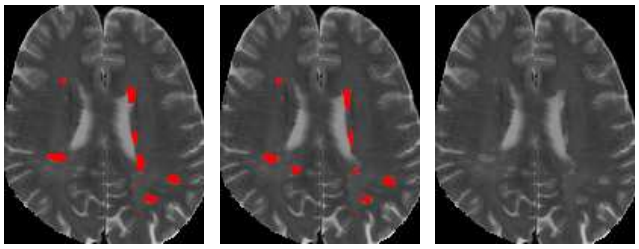

(a) (b) (c)

Fig. 4. Lesion segmentation results on real data: (a) Result using CGMM+AC, (b) Ground truth, (c) T2 channel.

Ground truth is based on a human expert.

To conclude, we presented a fully automated algorithm for tissue classification and MS lesion detection and segmentation in extremely noisy MR brain images. Quantitative comparison with KVL algorithm has demonstrated better results, especially as the noise increases. Experiments on real MRI data yielded good results. The CGMM was shown to provide a parametric framework good enough for healthy tissue and MSL segmentation in abnormal brains, without requiring a priori registration to a brain atlas. We are currently extending the experiments to additional real brain datasets, in a joint effort with the MS unit at Sheba hospital. Also planned is the validation across multi-expert ground-truth data.

6. REFERENCES

[1] A. Akselrod-Ballin et al. An integrated segmentation and classification approach applied to multiple sclerosis analysis. In *Computer Vision and Pattern Recognition*, pages 1122–29, 2006.

[2] T. Chan and L. Vesse. An active contour model without edges. In *Theories in Computer Vision, Lecture Notes in Computer Science*, pages 141–151, 1999.

[3] D. Collins et al. Design and construction of a realistic digital brain phantom. *IEEE Trans. Med. Imaging*, 17(3):463–81, 1998.

[4] H. Greenspan, A. Ruf, and J. Goldberger. Constrained guassian mixture model framework for automatic segmentation of MR brain images. *IEEE Trans. Med. Imaging*, 25(9):1233–45, 2006.

[5] M. Rousson and R. Deriche. A variational framework for active and adaptative segmentation of vector valued images. In *IEEE Workshop on Motion and Video Computing (WMVC)*, 2002.

[6] K. Van Leemput et al. Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE Trans. Med. Imaging*, 20(8):677–88, 2001.

[7] X. Wei et al. Quantitative analysis of MRI signal abnormalities of brain white matter with high reproducibility and accuracy. *JMRI*, 15:203–9, 2002.

[8] A. Zijdenbos, R. Forghani, and A. Evans. Automatic pipeline analysis of 3d MRI data for clinical trials: application to MS. *IEEE Trans. Med. Imaging*, 21(10):1280–91, 2002.