Research Article

Multiple Sclerosis Lesion Detection Using Constrained GMM and Curve Evolution

Oren Freifeld,[1] Hayit Greenspan,[1] and Jacob Goldberger[2]

[1] *Department of Biomedical Engineering, Tel-Aviv University, Tel Aviv 69978, Israel*
[2] *School of Engineering, Bar-Ilan University, Ramat-Gan 52900, Israel*

Correspondence should be addressed to Oren Freifeld, orenf@eng.tau.ac.il

This paper focuses on the detection and segmentation of Multiple Sclerosis (MS) lesions in magnetic resonance (MRI) brain images. To capture the complex tissue spatial layout, a probabilistic model termed Constrained Gaussian Mixture Model (CGMM) is proposed based on a mixture of multiple spatially oriented Gaussians per tissue. The intensity of a tissue is considered a global parameter and is constrained, by a parameter-tying scheme, to be the same value for the entire set of Gaussians that are related to the same tissue. MS lesions are identified as outlier Gaussian components and are grouped to form a new class in addition to the healthy tissue classes. A probability-based curve evolution technique is used to refine the delineation of lesion boundaries. The proposed CGMM-CE algorithm is used to segment 3D MRI brain images with an arbitrary number of channels. The CGMM-CE algorithm is automated and does not require an atlas for initialization or parameter learning. Experimental results on both standard brain MRI simulation data and real data indicate that the proposed method outperforms previously suggested approaches, especially for highly noisy data.

1. Introduction

Multiple sclerosis (MS) is the most common nontraumatic neurological disease in young adults. It is an inflammatory demyelinating disease that is primarily associated with axonal loss and formation of lesions in the central nervous system [1]. Magnetic Resonance Imaging (MRI) is the leading diagnostic tool in the context of MS. The extensive use of this imaging technique has significantly improved our understanding of the disease [2]. MRI is currently being used for MS diagnosis, assessment of disease progression, and evaluation of the efficiency of drug therapy [3]. The most common quantitative parameter is the *lesion burden* (or *load*) of the disease expressed in terms of the number and volume of the brain lesions. The MRI measured lesion burden is correlated with clinical findings [4–8]. However, it should be stated that the exact relationship between MRI measured lesions and pathological findings is not entirely clear [9–11]. MRI is also capable of providing measures and visualization of subclinical activity. Last but not least, it is sensitive to temporal variations that are typical to this disease [12].

A correct segmentation of healthy tissues is crucial for successful detection and segmentation of MS lesions [13]. Moreover, one of the characteristics of the disease is a decrease in the volume of healthy tissues, especially gray matter (GM). Having said that, it is clear why, in spite of the fact that the present work focuses on MS lesion detection and segmentation, the topic of healthy tissue segmentation should be addressed as well.

MRI brain image segmentation is a challenging task, even with healthy patients. The challenge stems from several reasons. First, MRI images present various noise artifacts such as intratissue noise, partial-volume effects and thermal noise [14]. These artifacts may cause unrealistic results, where tissue regions as well as lesion regions may appear granular, fragmented, or violating anatomical constraints. Second, the geometry of the brain tissues is complex and cannot be modeled by simple geometric primitives. Third, there is a significant variability in tissue geometry between patients. Finally, intensity distributions of different tissues can present a notable overlap. Segmentation of abnormal structures like MS lesions, is a difficult task. Lesions differ

from each other in size, location and shape. Also, as each MRI modality reflects different physical properties, the exact lesion boundaries may vary between different MRI channels. To complicate matters, the number of voxels that are associated with MS lesions is significantly smaller than the number of voxels that are associated with healthy tissues. Thus, simple clustering algorithms (e.g., K-means), fail to discern MS-lesion voxels from the rest of the image.

With notable exceptions (e.g., [15, 16]), segmentation of MRI brain images of MS patients is usually addressed by probabilistic voxel-based intensity modeling. Some of the salient aspects of these works are described below.

1.1. Parametric and Nonparametric Modeling. Several authors suggested to model the intensity of the brain tissues using a mixture of Gaussian probability density functions [17–22]. In this model, heron termed *Gaussian Mixture Model* (GMM), the intensity of a voxel is assumed to be drawn from the specific normal distribution associated with its parenting tissue. The GMM belongs to the category of parametric classification methods, since it is fully defined by its parameter set. The present work also belongs to this category. Other researchers have suggested using nonparametric methods such as the K-Nearest Neighbor (KNN) classification rule, or Parzen windows [23, 24]. These methods avoid making any assumption on the probability density function that is underlying the data. Instead, they try to locally estimate the probability density function, based on some labeled prototypes. These prototypes are given manually by a human expert, so such methods are not completely automated. A combination of parametric and nonparametric methods has been suggested recently [13].

1.2. MS Lesions: Explicit Model or Outliers? While several authors (e.g., [13, 23]) model the MS lesions as a distinguished class in addition to the healthy tissue classes (CSF, GM, WM), another approach is to model the lesions as outliers. For example, Van Leemput et al. [19] treat voxels that are not well explained (by either of the healthy tissue models) as candidates to be classified as lesions. Among these candidates, the separation into lesion and nonlesion voxels is made according to contextual information and a set of rules. It is also possible to treat lesions as outliers only in an intermediate stage, and then to build an explicit model for them [22].

1.3. Incorporating Spatial Information. A segmentation method that relies solely on the voxel intensity is unlikely to produce sufficient results. As such, many efforts have been made to incorporate spatial information within the segmentation process. One approach is the use of a statistical atlas [17–19, 25]. The atlas provides an a-priori probability regarding expected tissue location. A similar approach is the template-driven segmentation [20, 23, 24] that employs a deformable digital anatomical atlas to extract white matter masks using nonlinear registration. The main purpose is to reduce the number of misclassifications, by assuming that lesions are usually placed within the WM tissue. Atlas/template based approaches make a powerful tool

for image segmentation; however, they suffer from several disadvantages. They require registration of the atlas/template to the brain target. Image registration is a computationally intensive procedure, and not always applicable. Also, registration errors usually result in segmentation errors. Registration is especially problematic when the patient presents significant anatomical variability with respect to the atlas/template, which is often the case in MS patients. Another conventional method to improve segmentation smoothness and immunity to noise is to use a Hidden Markov Random Field (HMRF), thus modeling neighboring voxel interactions [13, 18, 19]. 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 [15] regional properties are used in a multifeature and multiscale approach combining segmentation and classification (supervised) for MS lesion analysis. In [26] a 4D space (T_1 intensity and spatial features) multiGaussian model is used to model each tissue. A large number of Gaussians is used per brain tissue 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 [26] to provide accurate segmentation of T_1 simulated and real MRI brain images into the three healthy tissues, in particular in noisy and low-contrast MRI images, without the need for coregistration of the input image and an atlas. In the current work we propose an extension to the CGMM model that can handle multisequence MRI data with a focus on MS lesion detection and segmentation. We also suggest a new and efficient way for approximating the CGMM computations. Another focus of this work is utilizing the CGMM as a good initialization for a level set method that provides a refined lesion boundary. We show that direct implementation of curve evolution methods to the complicated MRI image provides poor brain segmentation and no detection of MS lesions. The lesion detection obtained using the CGMM makes it possible to apply active contour techniques for lesion detection. A preliminary version of this paper appeared in [27].

The paper is organized as follows. Section 2 presents the CGMM framework. Lesion detection based on the CGMM and lesion boundary refinement via a level-set method, hereon termed CGMM-CE, are described in Sections 3 and 4, respectively. Algorithm validation on both simulated and real brain volumes is shown in Section 5. Discussion and conclusions are presented in Section 6.

2. The CGMM Segmentation Framework

The complex spatial layout of an MRI brain image poses a challenge for conventional intensity based GMM modeling schemes. To accommodate the spatial complexity, we model

an image as if its voxels were drawn independently from a mixture of *many* Gaussians:

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

such that x is the 3D position information included in the feature vector, $I(x)$ is the intensity vector associated with the voxel located at position x, n is the number of components in the mixture model, μ_i and Σ_i are the mean and the covariance of the ith Gaussian component f_i, and α_i is the ith mixture coefficient. Considering the location of the voxel as an additional feature enables us to incorporate spatial information into the probabilistic model. Each Gaussian in the GMM provides a probabilistic model for a specific small area in the MRI image.

The spatial shape of the tissues is highly nonconvex. However, since we use a mixture of many components, each Gaussian component models a small local convex region. The intravariability of the intensity features within a tissue is significantly less than the intervariability among different tissues. It is therefore sufficient to model the intensity variability within a tissue by a single Gaussian (in the intensity features). 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. We also assume that locally, the intensity and spatial features are uncorrelated. The above assumptions impose the following structure on the mean and covariance of the Gaussian components:

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

where x is the spatial feature vector, I is the intensity feature vector, $\pi(i)$ is the tissue that is linked to the ith Gaussian component, μ_i^x and Σ_i^x are the spatial mean and covariance parameters of the ith Gaussian component, and μ_j^I and Σ_j^I are the intensity mean and covariance parameters shared by all the Gaussian components that are linked to the jth tissue. The Gaussian component f_i has, therefore, the following form:

$$f_i(x, I(x) \mid \mu_i, \Sigma_i) = \mathcal{N}\left(x; \mu_i^x, \Sigma_i^x\right) \times \mathcal{N}\left(I(x); \mu_{\pi(i)}^I, \Sigma_{\pi(i)}^I\right). \qquad (3)$$

The main advantage of the CGMM framework is the ability to combine, in a tractable way, a local description of the spatial layout of a tissue with a global description of the tissue's intensity. The multiGaussian spatial model makes our approach much more robust to noise than intensity-based methods. Note that no prior atlas information is used in the modeling process. An illustration of the CGMM model applied to the three tissues (CSF, GM, and WM) is shown in Figure 1(a). In case of MS lesion segmentation we consider the lesion matter in the CGMM modeling as a fourth tissue in addition to the three healthy tissues (Figure 1(b)).
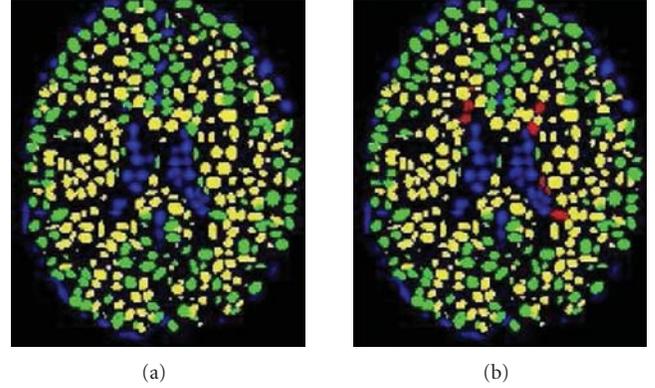

(a)   (b)

Figure 1: Illustration of the GMM representation. (a) Segmenting the brain image into three main tissues. (b) Adding a fourth class for MS lesions; Color Legend: Blue-CSF, Green-GM; Yellow-WM; Red-MS.

2.1. Parameter Estimation via the EM Algorithm. Given an MRI image with T voxels: $\{(x_t, I(x_t)) \mid t = 1, \dots, T\}$, the likelihood of a given CGMM parameter set is

$$\prod_{t=1}^{T} \sum_{i=1}^{n} \alpha_i f_i(x_t, I(x_t) \mid \mu_i, \Sigma_i). \qquad (4)$$

Initially, the model parameters are, of course, unknown. To find the maximum-likelihood parameters we utilize the classical Expectation-Maximization (EM) algorithm [28]. The EM algorithm handles the parameter estimation task by iterating the E and M steps. In the E-step, it treats the parameter set as given, and estimates the probabilities of each sample to be drawn from each Gaussian. Informally, this step can be seen as soft segmentation (but here, the number of segments is n, the number of Gaussians, and not K, the number of the tissues). The M-step performs parameter reestimation. Thus, the equations take the following form.

E-Step:

$$w_{it} = p(i \mid x_t, I(x_t))$$
$$= \frac{\alpha_i f_i(x_t, I(x_t) \mid \mu_i, \Sigma_i)}{\sum_{j=1}^{n} \alpha_j f_j\left(x_t, I(x_t) \mid \mu_j, \Sigma_j\right)} \quad i = 1, \dots, n, \ t = 1, \dots, T$$
$$(5)$$

w_{it} can be viewed as the posterior probability that voxel t was created using Gaussian i.

M-Step. Let

$$n_i = \sum_{t=1}^{T} w_{it} \quad i = 1, \dots, n,$$
$$k_j = \sum_{i \mid \pi(i)=j} n_i \quad j = 1, \dots, k \qquad (6)$$

with n_i and k_j interpreted as the expected number of voxels associated with the ith Gaussian and kth tissue, respectively,

and T is the number of voxels in the MRI image. Parameter reestimation is then defined as follows.

Spatial parameters:

$$\mu_i^x = \frac{1}{n_i} \sum_{t=1}^{T} w_{it} x_t,$$

$$\Sigma_i^x = \frac{1}{n_i} \sum_{t=1}^{T} w_{it} (x_t - \mu_i^x)(x_t - \mu_i^x)^\top, \qquad (7)$$

$$\alpha_i = \frac{n_i}{T}.$$

Intensity parameters:

$$\mu_j^I = \frac{1}{k_j} \sum_{i|\pi(i)=j} \sum_{t=1}^{T} w_{it} I(x_t),$$

$$\Sigma_j^I = \frac{1}{k_j} \sum_{i|\pi(i)=j} \sum_{t=1}^{T} w_{it} \left(I(x_t) - \mu_j^I \right) \left(I(x_t) - \mu_j^I \right)^\top. \qquad (8)$$

Note that while the spatial parameters are estimated separately for each Gaussian, the estimated intensity parameters are the same for all the Gaussians that belong to the same tissue. The grouping function π that links between the Gaussian components and the tissues is learned in the initialization step and is not altered by the EM iterations. Therefore, the affiliation of a Gaussian component to a tissue remains unchanged. However, since the learning is performed simultaneously on all the tissues, voxels can move between tissues during the iterations.

Once the optimal set of parameters is obtained using the EM algorithm, we can compute soft tissue segmentation maps. Moving from soft segmentation to hard segmentation is straight forward, using the Maximum-A-Posteriori (MAP) criterion. Tissue segmentation is achieved by the affiliation of each voxel to the tissue that maximizes the a posteriori probability:

$$\text{tissue-label}(\text{voxel}_t) = \arg \max_j p(\text{tissue} = j \mid x_t, I(x_t))$$

$$= \arg \max_j \sum_{i|\pi(i)=j} \alpha_i f_i(x_t, I(x_t) \mid \mu_i, \Sigma_i). \qquad (9)$$

Figure 2 illustrates the segmentation induced from the CGMM model and shows how the EM iterations improve the segmentation quality. Since a lot of the algorithmic effort is spent on finding a good initialization (see Section 2.2), the EM needs only few iterations to converge.

2.2. Model Initialization. In Section 2.1, we saw how the EM algorithm is applied to the MRI data to find the maximum-likelihood parameter-set for the CGMM. However, the EM is notoriously known to get stuck in local maxima. Hence, appropriate initialization is required. In the first step the K-means clustering is done based only on the intensity features (T_1, T_2, and PD), which gives a good initial segmentation

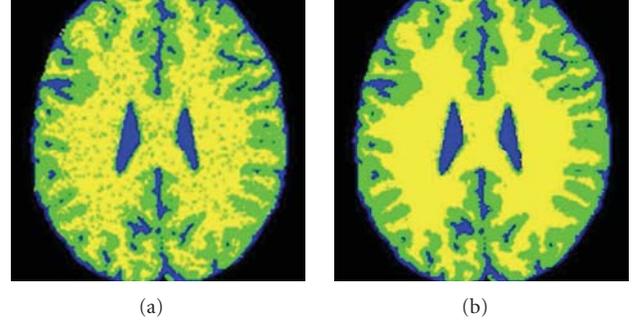(a)   (b)

FIGURE 2: Illustration of the influence of the EM iterations. As the parameter estimation improves with the iterations, the segmentation induced from the model becomes more and more robust to noise. Blue-CSF; Green-GM; Yellow-WM. (a) Initial segmentation. (b) Segmentation after convergence (five EM iterations).

into the three major tissue classes. We utilize T_1 to give tissue labels to the three groups.

Given the intensity K-means clustering of the voxels into the three tissues, we want to model each tissue with many small locally convex Gaussians. it is important that small Gaussians will be allowed, as lesions are often very small, and we would like the Gaussians layout to capture them. We have found empirically that randomly selecting 1/20 of the voxels as initial Gaussian centers yields a resolution that is high enough for detecting the MS lesions. To ensure that small isolated areas are explicitly represented by local Gaussians, we first apply a three-dimensional connected component process to each tissue. If a connected component contained less than three voxels, it was ignored and disregarded as noise. For each connected component (of voxels all from the same tissue) a subset of voxels is randomly chosen and a new Gaussian is formed at each of the chosen voxels. To compute initial covariance values for each Gaussian we assign each voxel to the nearest center with the same tissue label. The Gaussian covariance (both intensity and spatial) is computed based on all the voxels that are assigned to the Gaussian center. As a result of the initialization process each Gaussian is linked to one of the tissues and each voxel is affiliated with a selected Gaussian.

2.3. Efficient EM Based on the KDT Transform. The strength of the proposed model relies on the highly detailed spatial representation obtained by using many small Gaussians such that each Gaussian represents a small coherent and convex region. Of course, the number of Gaussians that are required to cover the entire image increases as the size of the Gaussians decreases. The main implementation problem is the large number of Gaussians, that leads to high computational complexity of the EM algorithm. The computational complexity involved with the E-step is $O(nT)$ where T is the number of voxels and n is the number of Gaussian components. A typical number of voxels in 3D MRI brain images is on the order of 10^6. The number of Gaussians in the CGMM framework may be as high as 10^3 or even 10^4 when segmentation of MS lesions is also required. This fact hinders the E-step computations significantly.

It is clear that voxels that are very far away from the spatial mean of a particular Gaussian i, have a small chance to originate from this Gaussian. Due to numerical limitations, the probability density of each Gaussian vanishes whereas the actual distance from the spatial mean is still finite. Thus, for practical purposes, there is no need to compute in the E-step the weights w_{it} for every Gaussian i and every voxel t. Instead, for every voxel we maintain a small list of its nearest Gaussians. Our experiments showed that even a small number of Gaussian neighbors (say 10) gives rise to very accurate approximations of the model and does not cause any degradation in the segmentation performance. The computational complexity is $O(10T)$ and therefore using this procedure, the number of computations in each E-step is reduced by a few orders of magnitude.

We note in passing that this approach is superior to maintaining a list of voxels for every Gaussian. The reason is that in the latter, we may find that there are voxels that are explained by no Gaussian at all. Admittedly, maintaining a list of nearest Gaussians for every voxel may result in Gaussians that have, if any, very little voxels. However, this should not compromise the results, since such Gaussians will simply shrink and die off during the EM process.

The main problem is efficiently finding the list of nearest Gaussians for each voxel. In this section we suggest an efficient way to find this list. The method is based on the K-Distance Transform (KDT) [29], (and can be applicable to other GMM applications defined on a grid in which the number of Gaussians is large). KDT, introduced by Warfield [29] and later improved by Cuisenaire [30], is an extension to the Distance Transform (DT). Originally, the KDT was used for fast implementation of the kNN classification rule. The KDT input is a discrete Euclidean space containing prototypes and nonprototypes. Each prototype is given a unique identifier. The output is a set of identifier maps. The first map is a first Nearest-Neighbor (NN) map. In effect, each element in this map contains the identifier of the prototype that is the nearest to this element. The second map is a second NN map and so on. The brute force approach for KDT computation is very simple and extremely inefficient: for every space element, compute its distances from all the given prototypes, sort them by the distances, and select only the nearest neighbors. Clearly, this approach is not applicable to large datasets due to its high complexity. Warfield proposed an efficient way for fast KDT computation, that is based on local propagations, rather than direct computations. The main point is that local propagations, implemented using several raster scans, require a number of computations that is only linear in the number of elements in the discrete space [29].

In our work, the input prototypes can be taken as the spatial means of the Gaussians and applying the KDT algorithm, we obtain for every voxel a list of its nearest Gaussians. We can use the Mahalanobis distance as the metric of choice. Note that the Mahalanobis distance depends not only on the locations of the prototype (in effect, the spatial mean of the Gaussian) and the voxel, but also on the covariance matrix of the Gaussian. Experiments indicate that while EM iterations modify the spatial parameters to some extent, they are unlikely to cause a Gaussian to wander far away from its original vicinity. Thus, it is enough to apply the KDT for computing the list of nearest Gaussians only once. There is in fact a chance that the list of the nearest Gaussians will change slightly, but the Gaussians that will be replaced are likely to be far from the voxel so that their contribution is not important. Here too, our experiments indicate that any discrepancy between the approximated model and the true model is intangible.

Note that in the present work, the importance of the KDT is primarily in approximating the exact computation of the probabilistic model. However, it also gives us contextual information. Not only is it possible to know which are the K nearest Gaussians for every voxel, but also their tissue affiliation. Hence, the KDT immediately gives additional useful information: for every voxel it is possible to know, which tissue its nearest Gaussians relate to. Such knowledge can be beneficial to determine if a voxel is located within a WM area and is used in the following section to obtain a fast implementation of rules for detecting MS lesions.

3. Lesion Detection Utilizing the CGMM

A major focus of the current work is to extend the CGMM framework to MS lesion segmentation. Section 2.2 described that one of the very first initialization steps was intensity-based K-means clustering. As mentioned earlier, simple clustering methods for MS lesion detection are usually inapplicable, due to the small number of lesion voxels. We present a novel approach to MS lesion detection, that exploits the Gaussians layout of the CGMM. First, the model is initialized with three tissue classes, and its parameters are learned. Due to the fact that MS lesion voxels are significantly outnumbered by the healthy tissue voxels, the clusters still succeed in representing the three healthy tissues (CSF, GM, and WM). The lesion voxels are of course misclassified at this point as one or more of the healthy tissues. Moreover, at this stage there are *misclassified Gaussians*, that is, Gaussians labeled as healthy tissues that are supposed to be labeled as MS. The purpose of the current stage is to identify these Gaussians, and change their labels accordingly. A major distinction of the present work from previous works is that *MS lesion detection is performed on the Gaussian level rather than on the voxel level.*

For each Gaussian a decision is made, based on its features, whether it should in fact be labeled as MS, or remain labeled as one of the healthy tissues. Both supervised and unsupervised approaches can be used to deal with this task. For example, a rule based system can be used to distinguish the MS lesion Gaussians from the normal tissue Gaussians. A Gaussian for which all these conditions hold, is labeled as MS lesion. An example of such a rule set for Gaussians labeled as GM is the following:

(1) T_2 mean-intensity of the Gaussian $> T_2$ mean-intensity of the GM tissue $+ \epsilon_{GM,T_2}$,

(2) T_2 mean-intensity of the Gaussian $> PD$ mean-intensity of the GM tissue $+ \epsilon_{GM,PD}$,

(3) a large Mahalanobis distance between the mean-intensity of the Gaussian and the mean-intensity of the CSF/GM/WM tissue,

(4) the majority of Gaussian's K nearest Gaussians are labeled as WM,

where ϵ_{GM,T_2} and $\epsilon_{GM,PD}$ are thresholds that can be tuned and optimized. The first two rules reflect the general appearance of the MS lesions. The rules that rely on Mahalanobis distance imply that only Gaussians that are not well explained by the healthy tissue models are suspected as lesion Gaussians. The last rule incorporates contextual information by reflecting our expectation to find lesions within the WM tissue. These rules are similar to rules used by Van Leemput et al. [19]. However, note that at this point decisions are made at the Gaussian level rather than the voxel level.

Clearly, a rule-based system is only one option for MS Gaussian detection. For example, a more supervised approach can be used: given examples marked by a human expert, standard classifiers can be applied to the Gaussian features. While we do not claim that the rule-based system suggested here is optimal (although it is intuitive and easy to implement), we do suggest that the detection on the Gaussian level provides greater strength in comparison to voxel-based detection.

Following the MS lesion Gaussian-detection stage, all the relabeled Gaussians now form a fourth class named MS lesion (MSL), with its own global intensity parameters. The EM is now applied to CGMM with four tissue types and segmentation of the tissues and the lesions is obtained as explained in Section 2. An illustration of the relabeling of the Gaussians can be seen in Figure 1(b).

4. The CGMM-CE Framework for Lesion Boundary Refinement

The segmentation induced from the CGMM yields good results in terms of lesion detection. However, there is usually an overestimation in the size of the lesions. This is most likely due to the fact that the area in the proximity of a lesion is somewhat different from the other tissues. In addition, the segmentation step of the CGMM model is done voxel-wise and does not take into account the smoothness of the lesion boundaries. Hence, a postprocessing 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 [31] based on the Chan-Vese model [32]. In [31] a curve evolution equation for object-background image segmentation is derived 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 voxels around it are the background. The optimal lesion segmentation is defined by minimizing the energy:

$$F(\mu_1, \Sigma_1, \mu_2, \Sigma_2, \partial\Omega) = \sum_{i=1}^{2} \int_{\Omega_i} e_i(I(x)) + \text{length}(\partial\Omega), \quad (10)$$

where Ω_1 are the lesion region voxels, Ω_2 are the nonlesion voxels, $\partial\Omega$ is the curve separating the two regions, μ_i, Σ_i are the Gaussians parameters of Ω_i and $e_i(I) = -\log f(I \mid \mu_i, \Sigma_i)$. In level-set formulation, the functional from (10) takes a new form:

$$E(\Theta, \phi) = \sum_{i=1}^{2} \int_{\Omega} \left[e_i(I(x)) \chi_i(\phi(x)) \right] dx + \text{length}(\partial\Omega), \quad (11)$$

where ϕ is the level-set function, Θ is the parameter set $\{\mu_1, \Sigma_1, \mu_2, \Sigma_2\}$, and χ_1, χ_2 are the characteristic functions of regions 1 and 2, respectively: $\chi_1(\phi) = H_\epsilon(\phi)$, $\chi_2(\phi) = 1 - H_\epsilon(\phi)$ such that $H_\epsilon(\cdot)$ is a regularized Heaviside function. The length term is $\int_{\Omega} |\nabla H_\epsilon(\phi(x))| dx$. Note that the level-set formulation enables performing integration on the entire image domain, Ω. The energy (local) minimum can be obtained utilizing the level-set theory via a Gradient Descent (GD) method. From the GD equation, $\phi_t = -\delta F/\delta \phi$, we obtain the following curve evolution equation:

$$\phi_t(x) = \left(\nu \cdot \text{div}\left(\frac{\nabla\phi}{|\nabla\phi|} \right) + e_2(I(x)) - e_1(I(x)) \right) \delta_\epsilon(\phi(x)). \quad (12)$$

Since brain images are complex, and since lesions are relatively small, it is impractical to use the standard initial contour (small circles) [31, 32] to find the three tissues and the lesions, simultaneously. The method is highly sensitive to the initial conditions for complicated four-phase images (such as brain images) as a gradient-descent optimization finds only local minima of the functional. Figure 3(a) shows a segmentation obtained from a four-phase curve-evolution using a standard circle-grid initialization [31]. It can be seen that (even with extensive manual tuning of the curve-evolution parameters) due to the complexity of the image and the added noise, we obtain poor brain segmentation results and no detection of MSL.

Using segmentation results of the three tissues and lesions obtained by CGMM just for initialization of a four-phase CE is not sufficient. Even with the improved CGMM initialization, the four-phase CE fails, as the image is too complex and the lesions are relatively small. Figure 3(c) shows an example of MSL segmentation results based on a CGMM segmentation followed by five iterations (that were needed until convergence) of four-phase curve evolution (the initial CGMM segmentation is shown in Figure 3(b)). As expected, and in agreement with Rousson and Deriche [31], the lesions gradually disappeared and the segmentation failed due to the complexity of the image.

In our approach we use a two-phase CE: lesions (object) versus nonlesion (background). We use the CGMM not just a starting point but also for background and object modeling. We model the lesion in the curve-evolution process using an intensity based Gaussian that is extracted from the lesion component of the CGMM model and is kept fixed during curve evolution iterations. We model the background using an intensity based GMM that is extracted from the CGMM model. Since there are three tissues that belong to Ω_2, we define $e_2(I(x))$ to be $\min(e_{CSF}, e_{GM}, e_{WM})$.
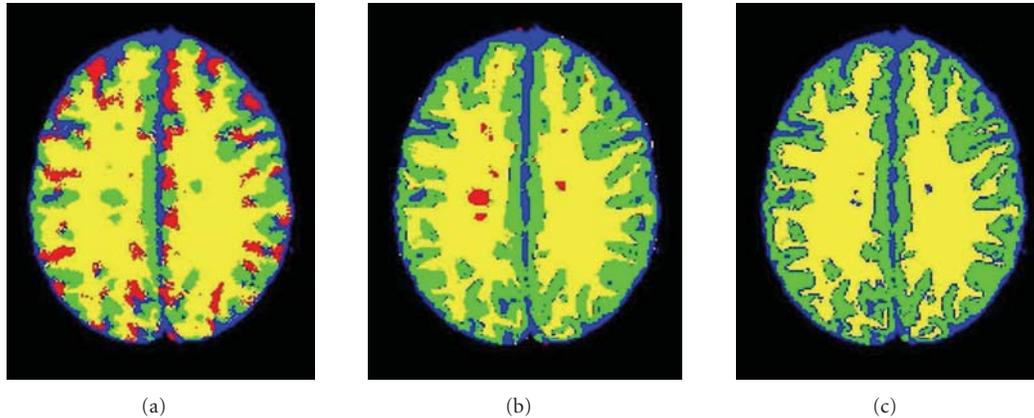
(a)   (b)   (c)

FIGURE 3: Illustration of MSL segmentation (BrainWeb data 9% noise, slice 105) using a four-phase curve-evolution (CE). Blue: CSF; Green: GM; Yellow: WM; Red: MSL. (a) Segmentation based a standard small-circles initialization followed by 30 iterations of four-phase curve evolution. (b) CGMM initial segmentation. (c) CGMM initial segmentation followed by five iterations of a four-phase curve-evolution.

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. In [31] the energy is minimized with respect to both the Gaussian parameters and the boundary position. In our case, we avoid the parameter reestimation step. Instead, we use the global intensity parameters learned via the CGMM model. This enables us to continue enforcing the global constraints of the CGMM in the active contour step. 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, whereas parameter estimation during the curve evolution is likely to cause the system to learn the parameters of the false lesion. To summarize, three modifications are suggested in the current work, in applying the Rousson-Deriche framework to the MSL delineation task.

(i) The CGMM lesion segmentation is used as the initial contour.

(ii) The background is modeled by a GMM (CSF, GM, WM) instead of a single Gaussian.

(iii) Parameter estimation during the curve evolution process is avoided; CGMM intensity parameters are used instead.

The lesion detection obtained using CGMM, therefore, makes it possible to apply active contour techniques for lesion delineation.

We dub the proposed method, which is based on CGMM followed by a curve evolution refinement, CGMM-CE. The steps of the CGMM-CE algorithm that were presented in Sections 2, 3, and 4 are summarized below.

(i) Apply intensity based K-means to obtain an initial clustering of the voxels into three tissues.

(ii) Form the initial spatial Gaussians based on clustering of the connected components of each tissue. Set the Gaussian labels according to the tissue they belonged to.

(iii) Utilize fast KDT to compute the nearest Gaussians for each voxel.

(iv) Apply the EM algorithm to find the probabilistic model for the three tissues.

(v) Find the lesion Gaussian using Gaussian outlier detection.

(vi) Apply the EM algorithm to find the probabilistic model for the four "tissues" (three healthy tissues + lesions).

(vii) Find lesion segmentation based on the probabilistic model.

(viii) Utilize curve evolution techniques to refine the lesion boundaries.

5. Experimental Results

Lesion segmentation validation is a challenging task. First, with notable exceptions there are hardly any data sets that are publicly available for the research community. Second, there is no clear convention for the best measures to quantify the results. Third, the nature of the lesions is almost never binary. It is a challenge to decide where the lesion ends and the healthy tissue starts, a fact that introduces variability in human-expert markings. Manual segmentation of large data sets is a labor and time-consuming task. This makes it difficult to produce ground truth segmentation for large data sets. In the current work, we demonstrate results on a the Brainweb online database [33] which is considered to be a standard benchmark upon which MRI brain image segmentation algorithms are tested and compared. An initial result on a real brain, is shown as well.

5.1. BrainWeb Simulated MRI Data. The CGMM-CE algorithm was tested on bias-free simulated MRI data taken from the BrainWeb online database [33]. The following coregistered modalities were used: T_1, T_2 and *PD*. Voxel size was $1 \, \text{mm}^3$. Experiments were done on 61 slices (axial slices

<center>(a)   (b)   (c)</center>
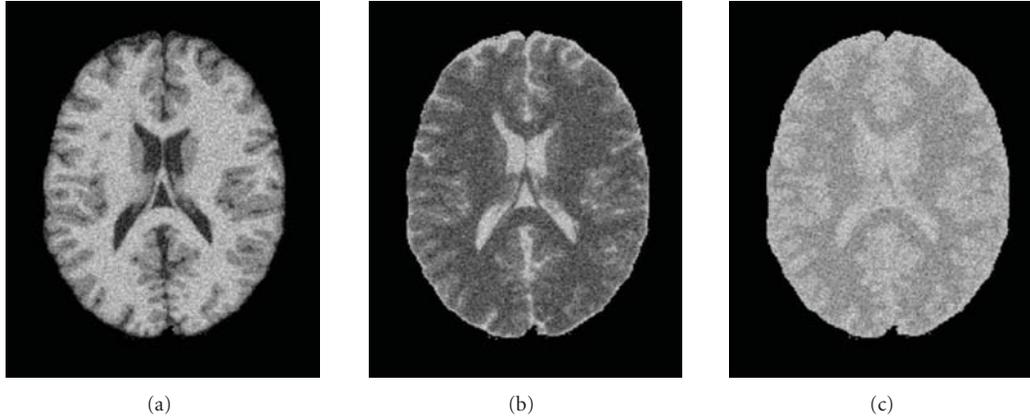
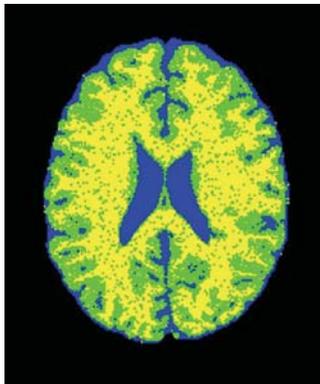FIGURE 4: BrainWeb data with 9% noise level (a) T_1, (b) T_2, (c) Pd.

FIGURE 5: Intensity K-means segmentation. $K = 3$. BrainWeb data, slice 95 with 9% noise. Blue: CSF; Green: GM; Yellow: WM.

TABLE 1: Dice index for lesion segmentation on BrainWeb data for different noise levels.

Algorithm	Noise			
	3%	5%	7%	9%
KVL	0.80	0.73	0.61	0.47
CGMM+CE	0.79	0.79	0.78	0.76

60–120) that contain 93% of the lesion burden. Different levels of noise were added: 3%, 5%, 7% and 9%. According to BrainWeb's web-site, the noise in the simulated images has Rayleigh statistics in the background and Rician statistics in the signal regions. The "percent noise" number (e.g., 5% noise) represents the percent ratio of the standard deviation of the additive white Gaussian noise versus the signal for a reference tissue. In the brain images under observation the noise reference tissues were WM for T_1, and CSF for T_2 and PD.

Figure 4 shows slice 90 with different levels of noise. Visual inspection suggests that images with 9% noise present a challenge even for manual segmentation. We are aware that the noise level in real data is almost always less than 9%. We provide the 9% segmentation result to demonstrate the robustness of the CGMM-CE method.

As preprocessing steps, the brain region was extracted and the intensity distribution of each channel was normalized to have zero mean and unit variance. Then, a K-means intensity-based clustering was performed with $K = 3$, to achieve an initial crude segmentation (see Figure 5). The global intensity parameters of each tissue were initialized as the sample mean and sample covariance of the extracted tissue segment.

Following the CGMM parameter initialization and KDT computation (with 10 neighbors), a subset of the Gaussians was relabeled as MSL. For every Gaussian, a set of decision rules was applied. If all the conditions described in this set held, and in addition the majority of Gaussian's 10 nearest Gaussians were labeled as WM, then the Gaussian was relabeled as MSL. The selected subset formed a fourth class. Optimal MSL parameters were found using the EM algorithm. In all the experiments, it took less than seven iterations to reach convergence. Finally, the curve evolution step was applied to refine the lesion segmentation. The velocity, v, in (12), was chosen as a constant: $v(x, y) = $ const. The constant value was set to be 5, 3, 1, 1 for noise levels 3, 5, 7, 9, respectively (higher velocity for low noise levels).

Performance is evaluated via a comparison with Van Leemput's state-of-the-art algorithm (hereon termed KVL) [19] implemented by the EMS software package-http (http://www.medicalimagecomputing.com/EMS). In the KVL implementation, the statistical brain atlas of the SPM99 (http://www.fil.ion.ucl.ac.uk/spm/) was normalized to the target brain volume images. Note that KVL is the only one who actually made his code/toolbox available, so the comparison is fair. Usually people report on lesion detection on MRI images with 3% noise level. We have not found reported results on 9% (or in fact, even 7%) other than KVL [19].

Figure 6 shows segmentation results on a single slice. Results are shown for both KVL, CGMM, and CGMM-CE methods. Figures 6(a)–6(c) shows a low noise case of 3%, Figures 6(d)–6(f) shows a high noise case of 9%, and Figures 6(g) and 6(h) show the ground truth MS lesion segmentation. Visual inspection indicates that for a low level of noise (3%), both algorithms present results that

(a) KVL   (b) CGMM   (c) CGMM-CE on T_1 (d) KVL

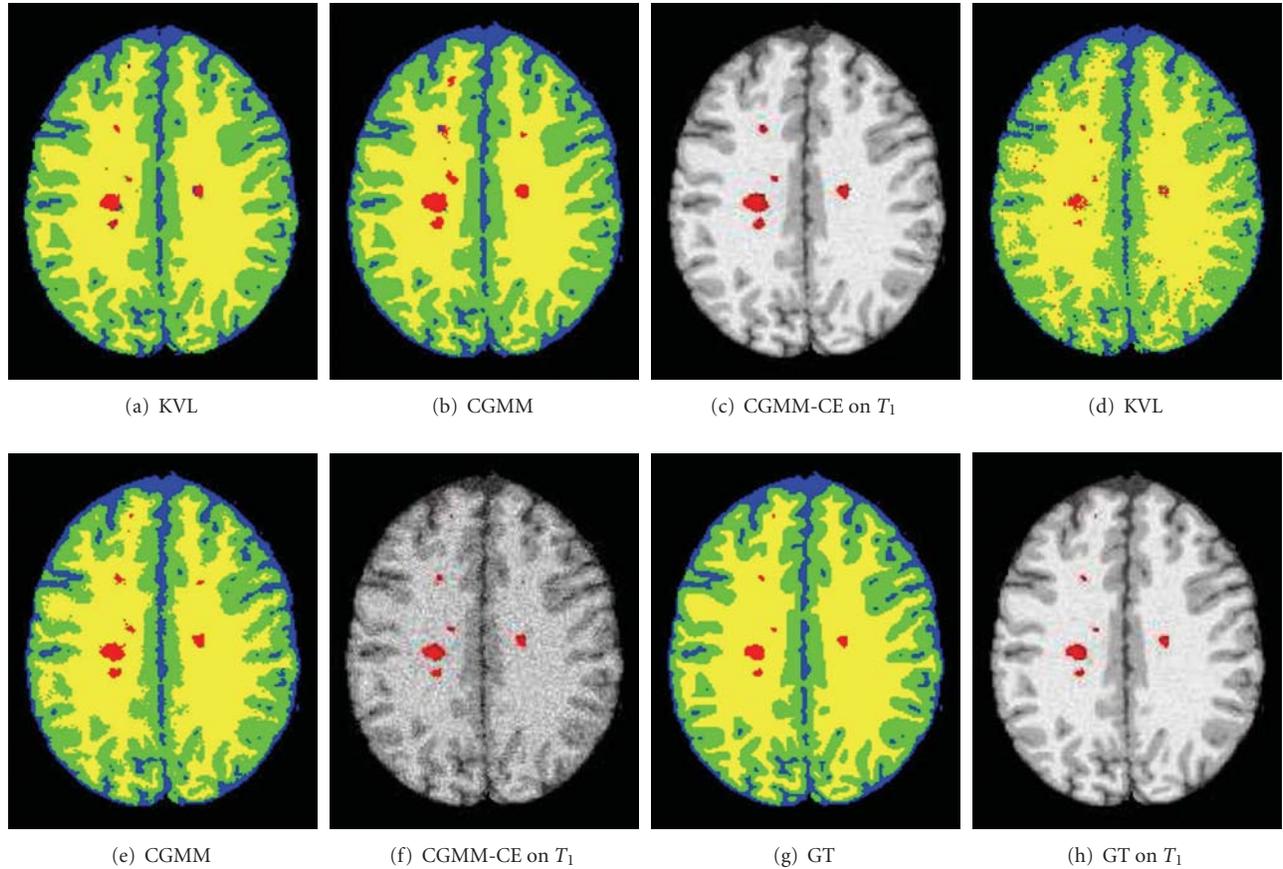(e) CGMM   (f) CGMM-CE on T_1 (g) GT   (h) GT on T_1

FIGURE 6: Segmentation results on BrainWeb data, slice 105. Blue: CSF; Green: GM; Yellow: WM; Red: MSL. (a)–(c) 3% noise, (d)–(f) 9% noise, (g) and (h) ground truth.

are close to the ground truth. For strong noise scenarios (9%), the CGMM produces a better and smoother healthy tissue segmentation. Also note that in terms of detection of whole MS lesions (e.g., number of false positives), the CGMM presents more successful results. Following the curve evolution step, the CGMM-CE framework presents further improved results.

The need for boundary refinement following the CGMM segmentation, and the result of applying the curve evolution process are shown as a zoom-in in Figure 7. Two slice examples are shown including for each is the ground-truth (left), the CGMM segmentation result (center), and the refined CGMM-CE segmentation (right). In the bottom left of Figure 7(e) we can see that two distinct lesions were joined together in the CGMM segmentation. Figure 7(f) shows that the refined contour succeeded in separating these two lesions. Comparing these two images further, we can also see that the overestimation of several of the lesions has significantly decreased. Similar results have been observed in other such slices examined. Note that in Figure 7(c) the refined lesion split to two very close adjacent lesions, this is not a general feature of the CE, as the fact that the CGMM (see Figure 7(b)) produced a lesions whose two parts are almost disconnected. The reason for this is that only three main healthy tissues are modeled, but in fact the data presents about 9 or 10 healthy tissues, most of them are

typically ignored. This particular lesion is on top such a tissue—the Glilal matter and therefore few pixels are missed, which resulted in a final insignificant splitting.

A quantitative comparison between the algorithms was made using the Dice overlap metric which measures the overlap between the automatic segmentation and the ground truth for each tissue j. This performance index is often used in assessment of medical image segmentation algorithms, and is given by

$$\text{Dice Index} = \frac{2V_{ae}^{j}}{V_{a}^{j} + V_{e}^{j}}, \qquad (13)$$

where V_{as}^{j} is the number of voxels that are assigned to tissue j by both the ground truth and the automated algorithm, V_{a}^{j} and V_{e}^{j} denote the number of voxels assigned to tissue j by the algorithm and the ground truth, respectively. Dice values vary between zero (no agreement with the ground truth segmentation) and one (perfect agreement with ground truth segmentation).

Table 1 shows MS lesion Dice results for KVL and the CGMM-CE methods. In the 3% case, KVL is slightly better. For higher values of noise, the performance improvement of the CGMM-CE framework over the KVL is notable. It should be noted that for lesion segmentation, Dice values above 0.70 are considered to be clinically stratifying [16]. Moreover, a
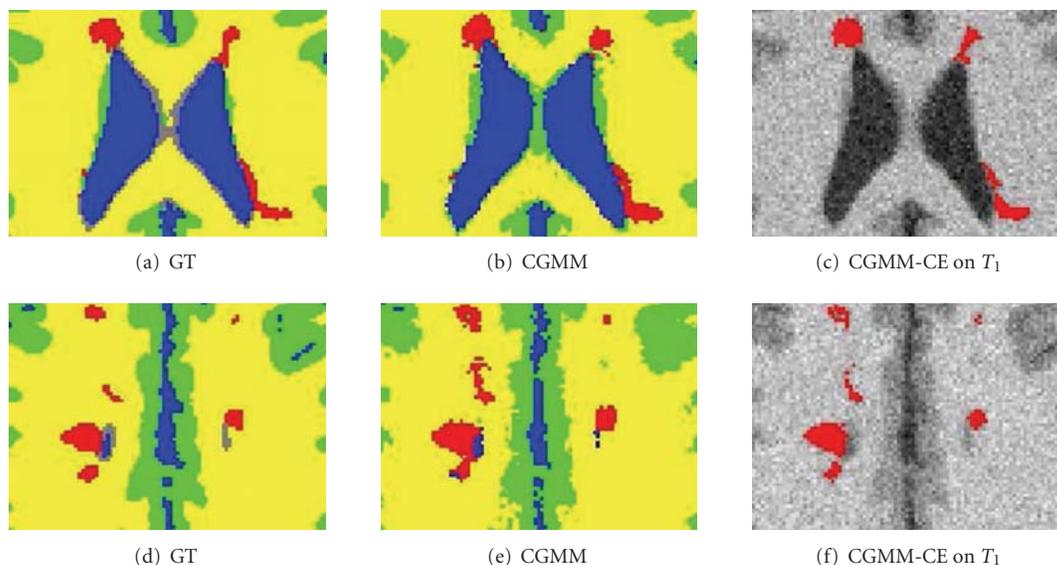
(a) GT

(b) CGMM

(c) CGMM-CE on T_1

(d) GT

(e) CGMM

(f) CGMM-CE on T_1

Figure 7: Zoom in of segmentation results. Blue: CSF; Green: GM; Yellow: WM; Red: MSL. Upper row: slice 97. Lower row: slice 104.

(a)

(b)

(c)
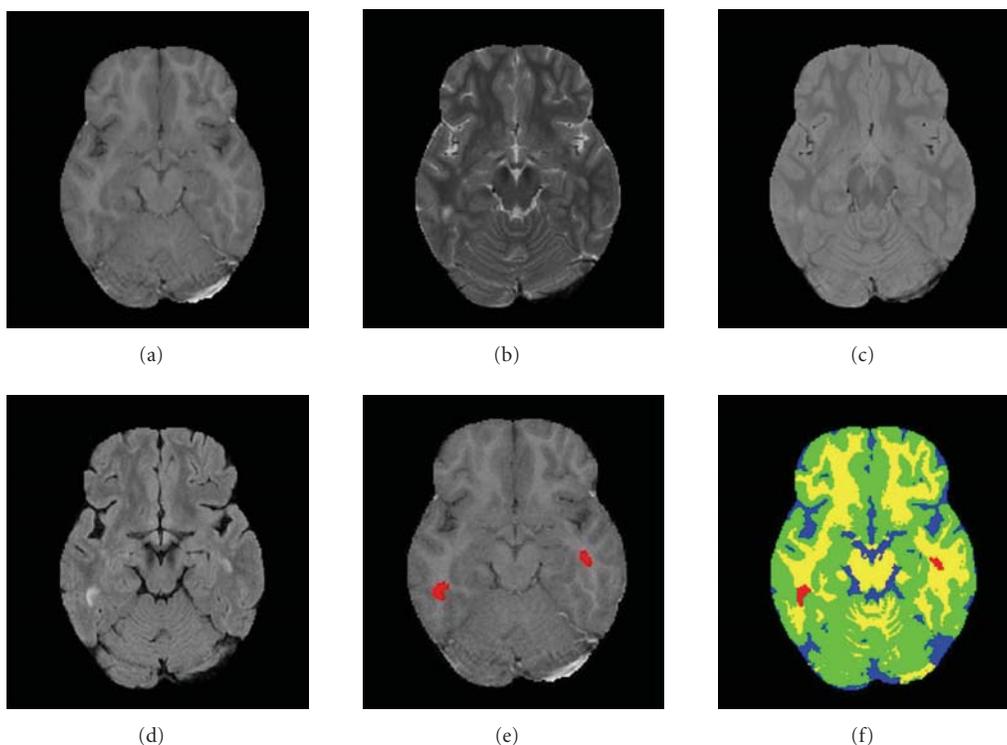

(d)

(e)

(f)

Figure 8: Real brain images, example 1. (a) T_1 (b) T_2 (c) PD (d) FF (e) Manual segmentation of MS lesions (red) overlayed on T_1 (f) CGMM-CE segmentation. Blue: CSF; Green: GM; Yellow: WM; Red: MSL.

comparison between two different segmentations given by two different human experts, is likely to yield Dice index lower than 0.70 [18].

The computation time (using Pentium (R) 4 CPU 3.40 GHz, 1.49 GB of RAM) for each significant part of the algorithm is initialization: 3 minutes, KDT: 2 minutes, MS Gaussian detection: 1 minute, and EM iteration: 1 minute. The curve evolution computation time was insignificant.

This is partially due to the fact that the CGMM provided a starting point that is very close to the local minima of the functional and thus the number of curve evolution iterations was usually small. Also, since we used the CGMM intensity parameters, there was no need to reestimate these parameters for every iteration. Thus, the entire CGMM-CE process takes less than a quarter of an hour for BrainWeb data, slices 60–120.
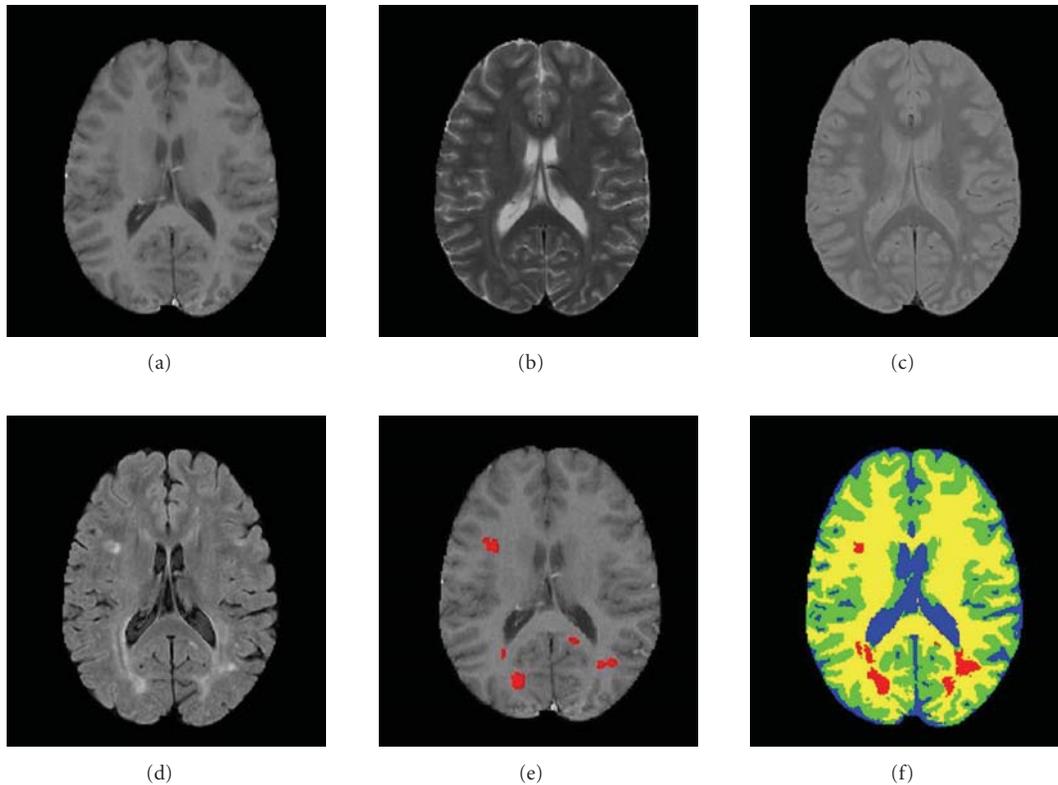
FIGURE 9: Real brain images, example 2. (a) T_1, (b) T_2, (c) PD, (d) FF, (e) Manual segmentation of MS lesions (red) overlayed on T_1, (f) CGMM-CE segmentation. Blue: CSF; Green: GM; Yellow: WM; Red: MSL.

5.2. Results on Real Data. In what is still a work in progress, the algorithm was also tested on real MRI data. These experiments were performed in a joint effort with the MS Center at Sheba Medical Center, Israel. Here, ground truth segmentation of the healthy tissues is not available. However, we have manual MS lesion segmentation, provided by a human expert. In these experiments, the Fast Flair (FF) modality, was added to T_1, T_2, and PD which were previously used. Note that the BrainWeb data [33] does not offer this modality. The FF provides a good contrast between CSF and the other healthy tissues (GM and WM). To some extent, it also provides a good contrast between MS lesions and the three healthy tissues. The main purpose of FF in MS is to suppress false positive from CSF signals. Preprocessing steps include conventional procedures such as extraction of the brain region (also know as skull removal); coregistration of the images from the different channels; and bias filed correction. Once the preprocessing is finalized, we extract from each voxel both intensity and spatial features. Unlike the isotropic voxel used in the BrainWeb data experiment, here the voxel size was $1\,mm \times 1\,mm \times 3\,mm$. The set of decision rules for Gaussian detection was similar to the one used in the BrainWeb dataset, with the addition of rules that take the FF into account (e.g., lesions are brighter than all the other tissues). Results on real data are shown in Figures 8 and 9. In both examples it is possible to see a smooth and visually plausible segmentation of the healthy tissues. With regard to the lesion detection and delineation, the main lesions are in fact detected (except for one in Figure 9).

6. Discussion

In [26], the CGMM was shown to be a successful framework for healthy tissue segmentation of T_1 MRI brain images. In the current work, the CGMM-CE framework was presented, as a methodology to detect and delineate MS lesions within multichannel input data. A novel model approximation scheme has been suggested to accelerate computing time by several orders of magnitude. Finally, a curve evolution technique was applied to refine the segmentation of the MS lesions. CGMM was shown here to provide a successful parametric framework for healthy tissue and MS lesion segmentation in abnormal brains, without requiring prior registration to a brain atlas. Unlike most MS lesion segmentation algorithms, here detection of MS lesions is done on the Gaussian-level rather than the voxel-level. Besides the fact that it ensures smoother segmentation and an inherent robustness to noise, Gaussian-level reasoning incorporates spatial information in a simple and intuitive way (unlike MRF that is often used in the literature).

One of the advantages of the suggested rule-based system for MS Gaussian detection, in addition to its simplicity, is that it enables the user (the radiologist) to tune the parameters in a clear and understandable way. Also, users can easily add or remove rules to match the rules that they himself apply for manual lesion segmentation. For example, they can decide to exclude all Gaussians that are close to the centerline of the brain. While similar rules were used in the literature, here the rules are applied to Gaussians (that

correspond to small regions) rather than to voxels. Of course, the rule-based system for MS Gaussian detection is only one option for classifying Gaussians. For example, it is possible to use the markings given by a human expert to train classifiers such as a support vector machine.

The Gaussian KDT enables a fast implementation of the CGMM framework and provides additional contextual information. The KDT can be useful for other GMM applications in which the number of Gaussians is large. The curve evolution technique suggested by Rousson and Deriche [31] was utilized to refine the CGMM lesion segmentation. Since brain images are complex, this technique cannot be directly applied to the MRI data. The CGMM segmentation provides an excellent initial contour, so the Gradient-Descent limitations are reduced. The use of the optimal MSL intensity parameters from the CGMM, eliminates the need to reestimate them in each curve evolution, thus speeding the process. In experiments with simulated data, both visual and quantitative comparisons with the KVL algorithm have demonstrated better results, especially as the noise increases. Typical results of MSL segmentation reported in the literature refer to 3% noise only (e.g., the sate-of-the-art algorithm suggested by Akselrod-Ballin et al. [15]). The proposed algorithm performs well with increased noise up to 9%. Experiments with real MRI data yielded promising results. We are currently extending the experiments to additional real brain datasets, in a joint effort with the MS center at Sheba hospital. Also planned is the validation across multiexpert ground-truth data. Some additional work needs to be done to address the very small size lesions, and to achieve more accurate delineation. Multiple-expert markings need to be compared as well, in order to estimate the observer variability on this data set, which can then be compared to the automated algorithm performance.

References

[1] A. Compston and A. Coles, "Multiple sclerosis," *The Lancet*, vol. 359, no. 9313, pp. 1221–1231, 2002.

[2] D. H. Miller, R. I. Grossman, S. C. Reingold, and H. F. McFarland, "The role of magnetic resonance techniques in understanding and managing multiple sclerosis," *Brain*, vol. 121, no. 1, pp. 3–24, 1998.

[3] D. H. Miller, M. Filippi, F. Fazekas, et al., "Role of magnetic resonance imaging within diagnostic criteria for multiple sclerosis," *Annals of Neurology*, vol. 56, pp. 273–278, 2004.

[4] S. J. Huber, G. W. Paulson, E. C. Shuttleworth, et al., "Magnetic resonance imaging correlates of dementia in multiple sclerosis," *Archives of Neurology*, vol. 44, no. 7, pp. 732–736, 1987.

[5] R. Medaer, E. Nelissen, B. Appel, M. Swerts, J. Geutjens, and H. Callaert, "Magnetic resonance imaging and cognitive functioning in multiple sclerosis," *Journal of Neurology*, vol. 235, no. 2, pp. 86–89, 1987.

[6] D. W. Paty and D. K. Li, "Interferon beta-1b is effective in relapsing-remitting multiple sclerosis," *Neurology*, vol. 43, no. 4, p. 662, 1993.

[7] M. Rovaris, M. Filippi, M. Falautano, et al., "Relation between MR abnormalities and patterns of cognitive impairment in multiple sclerosis," *Neurology*, vol. 50, no. 6, pp. 1601–1608, 1998.

[8] R. A. Sperling, C. R. G. Guttmann, M. J. Hohol, et al., "Regional magnetic resonance imaging lesion burden and cognitive function in multiple sclerosis: a longitudinal study," *Archives of Neurology*, vol. 58, no. 1, pp. 115–121, 2001.

[9] D. Goldberg-Zimring, H. Azhari, S. Miron, and A. Achiron, "3-D surface reconstruction of multiple sclerosis lesions using spherical harmonics," *Magnetic Resonance in Medicine*, vol. 46, no. 4, pp. 756–766, 2001.

[10] D. Goldberg-Zimring, B. Shalmon, K. H. Zou, H. Azhari, D. Nass, and A. Achiron, "Assessment of multiple sclerosis lesions with spherical harmonics: comparison of MR imaging and pathologic findings," *Radiology*, vol. 235, pp. 1036–1044, 2005.

[11] H. S. Kirshner, S. I. Tsai, V. M. Runge, and A. C. Price, "Magnetic resonance imaging and other techniques in the diagnosis of multiple sclerosis," *Archives of Neurology*, vol. 42, no. 9, pp. 859–863, 1985.

[12] A. Shahar and H. Greenspan, "Probabilistic spatial-temporal segmentation of multiple sclerosis lesions," in *Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis*, Lecture Notes in Computer Science, pp. 269–280, 2004.

[13] B. R. Sajja, S. Datta, R. He, et al., "Unified approach for multiple sclerosis lesion segmentation on brain MRI," *Annals of Biomedical Engineering*, vol. 34, no. 1, pp. 142–151, 2006.

[14] D. L. Pham, C. Xu, and J. L. Prince, "A survey of current methods in medical image segmentation," *Annual Review of Biomedical Engineering*, vol. 2, pp. 315–337, 2000.

[15] A. Akselrod-Ballin, M. Galun, M. J. Gomori, et al., "An integrated segmentation and classification approach applied to multiple sclerosis analysis," in *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR '06)*, vol. 1, pp. 1122–1129, 2006.

[16] A. P. Zijdenbos, R. Forghani, and A. C. Evans, "Automatic pipeline analysis of 3-D MRI data for clinical trials: application to multiple sclerosis," *IEEE Transactions on Medical Imaging*, vol. 21, no. 10, pp. 1280–1291, 2002.

[17] G. Dugas-Phocion, M. A. Gonzalez, C. Lebrun, et al., "Hierarchical segmentation of multiple sclerosis lesions in multi-sequence MRI," in *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI '04)*, pp. 157–160, 2004.

[18] R. Harmouche, L. Collins, D. Arnold, S. Francis, and T. Arbel, "Bayesian MS lesion classification modeling regional and local spatial information," in *Proceedings of the International Conference on Pattern Recognition (ICPR '06)*, vol. 3, pp. 984–987, 2006.

[19] K. Van Leemput, F. Maes, D. Vandermeulen, A. Colchester, and P. Suetens, "Automated segmentation of multiple sclerosis lesions by model outlier detection," *IEEE Transactions on Medical Imaging*, vol. 20, no. 8, pp. 677–688, 2001.

[20] X. Wei, S. K. Warfield, K. H. Zou, et al., "Quantitative analysis of MRI signal abnormalities of brain white matter with high reproducibility and accuracy," *Journal of Magnetic Resonance Imaging*, vol. 15, no. 2, pp. 203–209, 2002.

[21] W. M. Wells III, W. Grimson, R. Kikinis, and F. A. Jolesz, "Adaptive segmentation of MRI data," *IEEE Transactions on Medical Imaging*, vol. 15, no. 4, pp. 429–442, 1996.

[22] Y. Wolff, S. Miron, A. Achiron, and H. Greenspan, "Improved CSF classification and lesion detection in MR brain images with multiple sclerosis," in *Medical Imaging 2007: Image Processing*, vol. 6512 of *Proceedings of SPIE*, 2007, 65122P.

[23] S. K. Warfield, M. Kaus, F. A. Jolesz, and R. Kikinis, "Adaptive, template moderated, spatially varying statistical classification," *Medical Image Analysis*, vol. 4, no. 1, pp. 43–55, 2000.

[24] Y. Wu, S. S. Warfield, I. L. Tan, et al., "Automated segmentation of multiple sclerosis lesion subtypes with multichannel MRI," *NeuroImage*, vol. 32, no. 3, pp. 1205–1215, 2006.

[25] M. Kamber, D. L. Collins, R. Shinghal, G. S. Francis, and A.C. Evans, "Model-based 3D segmentation of multiple sclerosis lesions in dual-echo MRI data," in *Visualization in Biomedical Computing*, Proceedings of SPIE, pp. 590–600, 1992.

[26] H. Greenspan, A. Ruf, and J. Goldberger, "Constrained Gaussian mixture model framework for automatic segmentation of MR brain images," *IEEE Transactions on Medical Imaging*, vol. 25, no. 9, pp. 1233–1245, 2006.

[27] O. Freifeld, H. Greenspan, and J. Goldberger, "Lesion detection in noisy MR brain images using constrained GMM and active contours," in *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI '07)*, pp. 596–599, 2007.

[28] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society, Series B*, vol. 39, no. 1, pp. 1–38, 1977.

[29] S. K. Warfield, "Fast k-NN classification for multichannel image data," *Pattern Recognition Letters*, vol. 17, no. 7, pp. 713–721, 1996.

[30] O. Cuisenaire, *Distance transformations: fast algorithms and applications to medical image processing*, Ph.D. thesis, Laboratoire de Telecommunications at Teledetection, Universite Catolique de Louvain, Belgium, 1999.

[31] M. Rousson and R. Deriche, "A variational framework for active and adaptative segmentation of vector valued images," in *Proceedings of the Workshop on Motion and Video Computing*, pp. 56–61, 2002.

[32] T. Chan and L. Vese, "An active contour model without edges," in *Theories in Computer Vision*, Lecture Notes in Computer Science, pp. 141–151, 1999.

[33] C. A. Cocosco, V. Kollokian, and R. K. S. Kwan, "BrainWeb: online interface to a 3D MRI simulated brain database," *NeuroImage*, vol. 5, p. 425, 1997.