

Simultaneous Whole-Brain Segmentation and White Matter Lesion Detection Using Contrast-Adaptive Probabilistic Models

Oula Puonti¹(✉) and Koen Van Leemput^{1,2}

¹ Department of Applied Mathematics and Computer Science,
Technical University of Denmark, Kongens Lyngby, Denmark
oupu@dtu.dk

² Martinos Center for Biomedical Imaging, MGH,
Harvard Medical School, Boston, USA

Abstract. In this paper we propose a new generative model for simultaneous brain parcellation and white matter lesion segmentation from multi-contrast magnetic resonance images. The method combines an existing whole-brain segmentation technique with a novel spatial lesion model based on a convolutional restricted Boltzmann machine. Unlike current state-of-the-art lesion detection techniques based on discriminative modeling, the proposed method is not tuned to one specific scanner or imaging protocol, and simultaneously segments dozens of neuroanatomical structures. Experiments on a public benchmark dataset in multiple sclerosis indicate that the method’s lesion segmentation accuracy compares well to that of the current state-of-the-art in the field, while additionally providing robust whole-brain segmentations.

1 Introduction

Conditions that affect the integrity of the white matter, including small vessel disease and multiple sclerosis, form a significant health concern. Lesions in the white matter are frequently associated with memory impairment, headaches, depression, muscle weakness, and many other conditions. Because magnetic resonance (MR) imaging can visualize lesion formation with much greater sensitivity than clinical observation, the ability to reliably and efficiently detect white matter lesions from MR scans is of great value to diagnose disease, track progression, and evaluate treatment. Quantifying the independent contribution of white matter lesions to clinical disability is also important for enhancing our understanding of disease mechanisms, and for facilitating efficient testing in clinical trials.

Because of considerable intra- and inter-rater variabilities in manual annotations, and because of the sheer amount of imaging data acquired in clinical trials, there is a strong need for computational tools that can analyze brain images with white matter lesions in a fully automated fashion. Although many partial solutions have been proposed (e.g., [1]), a generally applicable tool that works robustly across disease states and imaging centers remains an open problem [2]. Many of the best performing methods for lesion segmentation currently

use extended spatial neighborhoods to provide rich contextual information, using a *discriminative* approach in which the specific intensity characteristics of training images are explicitly used to encode the relationship between image appearance and segmentation labels (e.g., [3–5]). However, because of the dependency of MR intensity contrast on the scanner platform and pulse sequence, and because there exists no standardized clinical MR protocol to study white matter damage, such discriminative methods do not generalize well to cases where the target and training data come from different scanners or centers. Furthermore, these methods do not provide segmentations of the non-lesioned parts of the brain into various cortical and subcortical structures, although regional atrophy patterns convey vital clinical information in diseases such as multiple sclerosis [6].

In this paper, we propose a novel method for jointly segmenting white matter lesions and a large number of cortical and subcortical structures from multi-contrast MR data. The method combines a previously validated method for whole-brain segmentation of healthy brain scans [7] with a novel spatial model for lesion shape and occurrence that is conditioned on surrounding neuroanatomy. In particular we propose to use a restricted Boltzmann machine (RBM) [8] to provide much richer spatial models than the low-order Markov random fields (MRFs) that have traditionally been used in the field for spatial regularization of lesion segmentations [9]. By using a generative rather than a discriminative formulation, the method is able to completely separate models of anatomy (which are learned from manual segmentations of training data) from intensity models (which are estimated on the fly for each individual scan being segmented). Because the *intensities* of training data are never used, the model can be applied to images with new contrast properties without needing new training data.

We test our approach on publicly available data from the MICCAI 2008 MS lesion segmentation challenge [10], demonstrating the feasibility of the method. Compared to related work for simultaneous whole-brain and lesion segmentation [11], the proposed method segments considerably more structures, and learns spatial lesion models automatically from training data rather than relying on a set of hand-crafted rules to remove false positive detections.

2 Modeling Framework

We build upon a previously published generative modeling approach [7], in which a forward probabilistic image model is “inverted” to obtain automated segmentations. In the following we first briefly summarize the existing whole-brain segmentation method we build upon; then introduce the proposed RBM lesion model; describe how we integrate it within the model for whole-brain segmentation; and specify how we use the resulting model to obtain automated segmentations.

2.1 Existing Whole-Brain Segmentation Method

Let $\mathbf{D} = (\mathbf{d}_1, \dots, \mathbf{d}_I)$ denote a matrix collecting the (log-transformed) intensities in a multi-contrast brain MR scan with I voxels, where the vector

$\mathbf{d}_i = (d_i^1, \dots, d_i^N)^T$ contains the intensities in voxel i for each of the available N contrasts. Furthermore, let $\mathbf{l} = (l_1, \dots, l_I)^T$ be the corresponding segmentation, where $l_i \in \{1, \dots, K\}$ denotes the one of K possible segmentation labels assigned to voxel i . A generative model then consists of a prior segmentation probability distribution $p(\mathbf{l})$ that encodes prior knowledge about human neuroanatomy, and a segmentation-conditional probability distribution $p(\mathbf{D}|\mathbf{l})$ that measures how probable the observed MR intensities are for different segmentations. In [7] the segmentation prior is parametrized by a sparse tetrahedral mesh with node positions $\boldsymbol{\theta}_l$. Assuming conditional independence of the labels between voxels given $\boldsymbol{\theta}_l$, the prior is given by:

$$p(\mathbf{l}) = \int_{\boldsymbol{\theta}_l} p(\mathbf{l}|\boldsymbol{\theta}_l)p(\boldsymbol{\theta}_l)d\boldsymbol{\theta}_l$$

$$\text{where } p(\mathbf{l}|\boldsymbol{\theta}_l) = \prod_{i=1}^I p(l_i|\boldsymbol{\theta}_l),$$

and $p(\boldsymbol{\theta}_l)$ is a topology-preserving deformation prior. The prior model is learned from manual annotations in 39 subjects as described in [7].

For the segmentation-conditional distribution $p(\mathbf{D}|\mathbf{l})$, a Gaussian mixture model (GMM) is associated with each neuroanatomical label to model the relationship between segmentation labels and image intensities. The smoothly varying intensity inhomogeneities (“bias fields”) that typically corrupt MR scans are modeled as a linear combination of spatially smooth basis functions that are added to the local voxel intensities. Letting $\boldsymbol{\theta}_d$ denote all bias field and GMM parameters with prior $p(\boldsymbol{\theta}_d) \propto 1$, the resulting segmentation-conditional distribution is given by:

$$p(\mathbf{D}|\mathbf{l}) = \int_{\boldsymbol{\theta}_d} p(\mathbf{D}|\mathbf{l}, \boldsymbol{\theta}_d)p(\boldsymbol{\theta}_d)d\boldsymbol{\theta}_d,$$

$$\text{where } p(\mathbf{D}|\mathbf{l}, \boldsymbol{\theta}_d) = \prod_{i=1}^I p(\mathbf{d}_i|l_i, \boldsymbol{\theta}_d)$$

$$\text{and } p(\mathbf{d}|l, \boldsymbol{\theta}_d) = \sum_{g=1}^{G_l} w_{lg} \mathcal{N}(\mathbf{d} - \mathbf{C}^T \boldsymbol{\phi}^i | \boldsymbol{\mu}_{lg}, \boldsymbol{\Sigma}_{lg}).$$

Here $\mathcal{N}(\cdot)$ denotes a normal distribution; G_l is the number of Gaussian distributions associated with label l ; and $\boldsymbol{\mu}_{lg}$, $\boldsymbol{\Sigma}_{lg}$, and w_{lg} are the mean, covariance, and weight of component g in the corresponding mixture model. Furthermore, $\boldsymbol{\phi}^i$ evaluates the bias field basis functions at the i^{th} voxel, and $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_N)$ where \mathbf{c}_n denotes the parameters of the bias field model for the n^{th} MR contrast.

With this model segmentation proceeds by estimating $\hat{\mathbf{l}} = \arg \max_{\mathbf{l}} p(\mathbf{l}|\mathbf{D})$, using the approximation $p(\mathbf{l}|\mathbf{D}) \simeq p(\mathbf{l}|\mathbf{D}, \hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l)$ where $\{\hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l\}$ are the parameter values that maximize $p(\boldsymbol{\theta}_d, \boldsymbol{\theta}_l|\mathbf{D})$. These values are estimated using coordinate ascent, where the atlas deformation parameters $\boldsymbol{\theta}_l$ are optimized with a conjugate gradient (CG) algorithm, and the remaining parameters $\boldsymbol{\theta}_d$ with a generalized

expectation-maximization (GEM) algorithm [7]. The optimization is done iteratively in an alternating fashion, keeping the deformation parameters fixed while optimizing the intensity model parameters and vice versa until convergence. The GMM parameters are initialized based on the structure probabilities given by the segmentation prior after affine registration to the target scan. We emphasize that the intensity model parameters are learned *given* the target scan and thus automatically adapt to its intensity properties. In [7] the intensity-adaptiveness was demonstrated on several datasets acquired with different sequences, scanners and field strengths.

2.2 Spatial Lesion Prior Using a Convolutional RBM

In order to model the spatial configuration of white matter lesions, we employ a restricted Boltzmann machine (RBM) [8], a specific type of MRF in which long-range voxel interactions are encoded through local connections to hidden units, which effectively function as feature detectors. Letting $\mathbf{z} = (z_1, \dots, z_I)^T$ denote a binary lesion map, where $z_i \in \{0, 1\}$ indicates whether the voxel is part of a lesion, a RBM prior on \mathbf{z} is defined by

$$p(\mathbf{z}) = \sum_{\mathbf{h}} p(\mathbf{z}, \mathbf{h}), \quad \text{with} \\ p(\mathbf{z}, \mathbf{h}) \propto \exp \left[-E_{\text{RBM}}(\mathbf{z}, \mathbf{h}) \right],$$

where $\mathbf{h} = (h_1, \dots, h_J)^T$, $h_j \in \{0, 1\}$ denotes a vector of J binary hidden units, and the RBM “energy” is defined as:

$$E_{\text{RBM}}(\mathbf{z}, \mathbf{h}) = -\mathbf{b}^T \mathbf{z} - \mathbf{c}^T \mathbf{h} - \mathbf{h}^T \mathbf{W} \mathbf{z}.$$

The parameters of this model include the vectors \mathbf{b} and \mathbf{c} (which bias individual visible and hidden units to take on certain values), as well as the weight matrix \mathbf{W} (which models the interaction between the hidden and visible units). The attractiveness of this specific MRF model arises from the presence of the hidden units, which increase the expressive power of the model, as well as the property that the values of \mathbf{z} are independent of one another given \mathbf{h} and vice versa, which greatly facilitates inference computations. Specifically, for each hidden unit h_j and lesion z_i the conditional distributions are written as [12]:

$$p(h_j = 1 | \mathbf{z}) = \sigma \left(c_j + (\mathbf{W} \mathbf{z})_j \right) \\ p(z_i = 1 | \mathbf{h}) = \sigma \left(b_i + (\mathbf{h}^T \mathbf{W})_i \right),$$

where $\sigma(x) = (1 + \exp(-x))^{-1}$.

In order to scale this framework to model full-sized images, we use a convolutional approach that imposes a repeated, sparse spatial structure on the parameters [12]. For the sake of clarity of presentation, in the following we describe the case for one-dimensional images, although the technique generalizes readily into three dimensions. In the convolutional RBM a set of P filters $\{\mathbf{f}^p\}_{p=1}^P$, $\mathbf{f}^p = (f_1^p, \dots, f_Q^p)^T$ is defined, each of size $Q \ll I$. The parameter matrix \mathbf{W} is then restricted to be of the form

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}^1 \\ \vdots \\ \mathbf{W}^P \end{pmatrix}, \quad \text{where} \quad \mathbf{W}^p = \begin{pmatrix} f_1^p & \dots & f_Q^p & 0 & \dots & 0 \\ 0 & f_1^p & \dots & f_Q^p & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & f_1^p & \dots & f_Q^p \end{pmatrix},$$

so that each filter detects the same specific feature in different parts of the image, and inference can be done efficiently using convolution. Similarly, in the parameter vector \mathbf{c} each filter output shares the same bias across the image [12]. In our implementation we do not put such a restriction on the visible biases \mathbf{b} , as this allows modeling spatially varying prior probabilities of lesion occurrence.

We automatically learn appropriate values for the parameters $\{\mathbf{W}, \mathbf{b}, \mathbf{c}\}$ from manually annotated training data, i.e., binary lesion maps for a number of different subjects. For this purpose, we use the persistent contrastive divergence (PCD) learning algorithm, which performs stochastic gradient ascent on the log-likelihood of the training data using approximate gradients computed with Markov chain Monte Carlo (MCMC) sampling [13].

2.3 Joint Model

We incorporate the RBM lesion model into the whole-brain segmentation framework by assuming that a lesion can only occur in a voxel when its underlying neuroanatomical label is white matter ($l = \text{wm}$), effectively changing its status from healthy white matter ($z = 0$) into white matter lesion ($z = 1$). Towards this end, we define a joint segmentation prior on both \mathbf{l} and \mathbf{z} :

$$\begin{aligned} p(\mathbf{l}, \mathbf{z}) &= \int_{\boldsymbol{\theta}_l} p(\mathbf{l}, \mathbf{z} | \boldsymbol{\theta}_l) p(\boldsymbol{\theta}_l) d\boldsymbol{\theta}_l, \quad \text{where} \\ p(\mathbf{l}, \mathbf{z} | \boldsymbol{\theta}_l) &= \sum_{\mathbf{h}} p(\mathbf{l}, \mathbf{z}, \mathbf{h} | \boldsymbol{\theta}_l) \quad \text{and} \\ p(\mathbf{l}, \mathbf{z}, \mathbf{h} | \boldsymbol{\theta}_l) &\propto \exp \left[-E_{\text{RBM}}(\mathbf{z}, \mathbf{h}) + \sum_{i=1}^I \log p(l_i | \boldsymbol{\theta}_l) - \sum_{i=1}^I \phi(l_i, z_i) \right], \end{aligned}$$

where in abuse of notation $p(l_i | \boldsymbol{\theta}_l)$ refers to the deformable atlas of the whole-brain segmentation model, and $\phi(l, z)$ evaluates to zero when $l = \text{wm}$ or $z = 0$, and infinity otherwise. The role of $\phi(l, z)$ is to restrict lesions to appear only inside white matter – without it the model would devolve into simply

$p(\mathbf{l}, \mathbf{z}) = p(\mathbf{l})p(\mathbf{z})$. In similar vein, we define an intensity model which is conditional on both \mathbf{l} and \mathbf{z} :

$$p(\mathbf{D}|\mathbf{l}, \mathbf{z}) = \int_{\boldsymbol{\theta}_d} p(\mathbf{D}|\mathbf{l}, \mathbf{z}, \boldsymbol{\theta}_d) p(\boldsymbol{\theta}_d) d\boldsymbol{\theta}_d,$$

$$\text{where } p(\mathbf{D}|\mathbf{l}, \mathbf{z}, \boldsymbol{\theta}_d) = \prod_{i=1}^I p(\mathbf{d}_i|l_i, z_i, \boldsymbol{\theta}_d)$$

$$\text{and } p(\mathbf{d}|l, z, \boldsymbol{\theta}_d) = \sum_{g=1}^{G_l} w_{lg} \mathcal{N}(\mathbf{d} - \mathbf{C}^T \boldsymbol{\phi}^i | \boldsymbol{\mu}_{lg}, \gamma^z \boldsymbol{\Sigma}_{lg}).$$

This model preserves the original segmentation-conditional GMMs for voxels without lesions ($z = 0$), but widens the variances of the Gaussian components by a user-specified factor $\gamma > 1$ otherwise. Such wide distributions aim to capture the fact that lesions often do not have a clearly defined intensity profile in MR, e.g., ranging from iso-intense to white matter to intensities similar to CSF in T1-weighted contrasts.

2.4 Inference

Segmentation with the proposed model can be accomplished by first estimating the parameters $\{\hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l\}$ that maximize $p(\boldsymbol{\theta}_d, \boldsymbol{\theta}_l|\mathbf{D})$, and subsequently analyzing $p(\mathbf{l}, \mathbf{z}|\mathbf{D}, \hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l)$, as in the whole-brain segmentation method described in Sect. 2.1 [7]. However, optimization of the model parameters is now complicated by the fact that the RBM model introduces non-local dependencies between the voxels through the weighted connections between the lesions and the hidden units. To side-step this difficulty, during the parameter estimation phase – in which we have no interest in accurately segmenting the white matter lesions – we temporarily replace the RBM energy $E_{\text{RBM}}(\mathbf{z}, \mathbf{h})$ with a simple energy of the form:

$$E_{\text{tmp}}(\mathbf{z}, \mathbf{l}) = - \sum_{i=1}^I [l_i = \text{wm}] (z_i \log(w) + (1 - z_i) \log(1 - w)),$$

where $0 \leq w \leq 1$ is a user-specified parameter which essentially defines a uniform spatial prior probability for lesions to occur *within* white matter. This effectively removes the hidden units from the model, and reduces the form of $p(\boldsymbol{\theta}_d, \boldsymbol{\theta}_l|\mathbf{D})$ to the one of the original segmentation method, so that the same optimization strategy can be used. Compared to the original method, the only difference is that each Gaussian distribution $\mathcal{N}(\cdot | \boldsymbol{\mu}_{lg}, \boldsymbol{\Sigma}_{lg})$ associated with the white matter label $l = \text{wm}$ is replaced with a mixture of the form:

$$(1 - w) \mathcal{N}(\cdot | \boldsymbol{\mu}_{lg}, \boldsymbol{\Sigma}_{lg}) + w \mathcal{N}(\cdot | \boldsymbol{\mu}_{lg}, \gamma \boldsymbol{\Sigma}_{lg}), \quad (1)$$

yielding a distribution with the same mean but heavier tails, making parameter estimation more robust to intensity outliers such as white matter lesions.

The adaptation in the GEM algorithm to enforce the parameter sharing between the two mixture components in Eq. (1) is straightforward.

Once the optimal parameter estimates are found, we replace the temporary energy with the original RBM energy and infer the corresponding whole-brain and lesion segmentation by MCMC sampling from $p(\mathbf{l}, \mathbf{z} | \mathbf{D}, \hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l)$, exploiting the specific structure of the RBM model. In particular, we generate S triplets $\{\mathbf{l}_s, \mathbf{z}_s, \mathbf{h}_s\}_{s=1}^S$ by sampling from the distribution $p(\mathbf{l}, \mathbf{z}, \mathbf{h} | \mathbf{D}, \hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l)$ using block-Gibbs sampling. This is straightforward to implement because each of the conditional distributions factorizes over the voxels (for \mathbf{l} and \mathbf{z}) or the hidden units (for \mathbf{h}). The sampling is performed in two alternating steps: first, we sample the values for the hidden units given the lesions:

$$\mathbf{h}_s \sim \prod_{j=1}^J p(h_j = 1 | \mathbf{z}_{s-1}).$$

Then, given the sampled hidden unit values \mathbf{h}_s , we jointly sample the labels \mathbf{l} and \mathbf{z} from:

$$\{\mathbf{l}_s, \mathbf{z}_s\} \sim \prod_{i=1}^I p(l_i, z_i | \mathbf{d}_i, \mathbf{h}_s, \hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l)$$

where

$$p(l_i, z_i | \mathbf{d}_i, \mathbf{h}, \hat{\boldsymbol{\theta}}_d, \hat{\boldsymbol{\theta}}_l) \propto \begin{cases} p(\mathbf{d}_i | l_i, z_i, \hat{\boldsymbol{\theta}}_d) p(l_i | \hat{\boldsymbol{\theta}}_l) p(z_i | \mathbf{h}), & \text{if } l_i = \text{wm or } z_i = 0 \\ 0, & \text{otherwise.} \end{cases}$$

The initial lesion segmentation, i.e., \mathbf{z}_0 , is obtained as the maximum-a-posteriori estimate using the temporary energy E_{tmp} .

Once we have acquired S triplets, the samples of the hidden units $\{\mathbf{h}_s\}$ are discarded as they are of no interest to us. The “hard” segmentations of \mathbf{l} and \mathbf{z} are obtained by voxel-wise majority voting across $\{\mathbf{l}_s\}$ and $\{\mathbf{z}_s\}$.

3 Experiments and Results

3.1 Data

We demonstrate the proposed method on the 20 publicly available training cases of the MICCAI 2008 challenge on multiple sclerosis lesion segmentation [10]. This dataset includes 10 subjects scanned at Children’s Hospital Boston (CHB) and another 10 scanned at the University of North Carolina (UNC). For each subject the scan set consists of a T1-weighted, a T2-weighted and a FLAIR scan with isotropic resolution of 0.5mm, along with expert segmentations provided by CHB¹. As a pre-processing step the data was downsampled by a factor of two to a resolution of 1mm isotropic as is customary for this dataset [3, 5, 14]. No further pre- or post-processing, such as intensity normalization or bias field correction, was applied.

¹ Manual segmentations from UNC are now also available, but at the time of the challenge this was not the case [10] so we decided to use only the segmentations provided by CHB.

3.2 Implementation

We closely follow the implementation details of the whole-brain segmentation method described in [7]. Because of the small number of manual segmentations available for training the RBM model, we applied two rotations of 10 and -10 degrees around the three main axes, producing 6 extra training scans per subject. We trained different RBM models with either $P = 20$ or $P = 40$ filters, with sizes of $(Q \times Q \times Q)$, where Q was either 5, 7 or 9. Each model was trained with 5600 gradient steps of size 0.1 in the PCD algorithm [13]. Based on pilot experiments, we found that using two mixture components for white matter worked well (i.e., $G_{\text{wm}} = 2$), provided that one of the Gaussians is constrained to be a near-uniform distribution that can collect model outliers other than white matter lesions (in practice we use a Gaussian with a fixed scalar covariance matrix $10^6 \mathbf{I}$ and weight 0.05). Finally, as the main characteristic of white matter lesions is that they appear hyper-intense compared to normal white matter in FLAIR contrast [2], we decided to only allow voxels to be assigned to lesion in the Gibbs sampling process if their intensity is higher than the estimated white matter mean in FLAIR.

We implemented the algorithm in Matlab, except for the mesh deformation part, which was written in C++, and the RBM convolutions, which were performed on a GPU. In our experiments, estimation of the parameters $\{\hat{\theta}_d, \hat{\theta}_l\}$ was performed on a cluster where each node has two quad-core Xeon 5472 3.0 GHz CPUs and 32 GB of RAM. Only one core was used in the experiments, taking roughly 1.7 h per subject. Gibbs sampling was done on a machine with a GeForce GTX Titan 6 GB GPU. We generated $S = 150$ samples, collected after an initial burn-in of 50 sampling steps, taking approximately 10 min per subject. Thus the full segmentation time for a single target scan is roughly two hours.

3.3 Evaluation Set-Up

In order to compare our results against previous methods on the same data, we use the true positive rate $\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}$ and the positive predictive value $\text{PPV} = \frac{\text{TP}}{\text{TP} + \text{FP}}$ as performance metrics. Here TP, FP and FN count the true positive, false positive and false negative voxels compared to the expert segmentation. Because our method contains four user-specified parameters γ , w , Q and P , which can have a large influence on the obtained results, and because the RBM requires training data to learn its parameters, we perform our evaluation in a cross-validation setting. In particular, we split the available data randomly into five distinct sets, each having 16 training and 4 test subjects. For segmenting each set of 4 test subjects, the remaining 16 are used to train the RBM and to find the best combination (γ, w, Q, P) , defined as the combination maximizing the product of the mean TPR and PPV over the 16 subjects. Using the product as a measure of fitness promotes parameter combinations that provide both sensitive and specific lesion segmentations.




Fig. 1. Example segmentations from two subjects: CHB04 (first row) and CHB08 (second row). From left to right: T1-weighted scan, T2-weighted scan, FLAIR scan, manual segmentation overlaid on the FLAIR scan, and the full segmentation obtained using the proposed method. Lesions are denoted in red (Color figure online).

3.4 Results

Figure 1 shows two examples of the joint whole-brain and lesion segmentations obtained using the proposed method, along with the manual segmentations. Although our method can segment 41 different neuroanatomical structures in total [7], the MICCAI challenge data only includes manual segmentations of lesions, so validation of the automatic segmentations of these structures could not be performed. However, visual inspection of the 20 cases did not reveal any significant failures in the whole-brain segmentation component of the method.

In Table 1 we compare our lesion segmentation performance with that of two state-of-the-art lesion segmentation tools: a random forest (RF) classifier [3], which is a discriminative method, and a dictionary-learning approach (DL) [14], which is unsupervised and therefore contrast-adaptive (as is the proposed method). Compared to the winning method [15] of the MICCAI 2008 lesion segmentation challenge, which obtained a mean TPR of 0.21 and a mean PPV of 0.30, all the methods show greatly improved segmentation results. On average the proposed method achieves better results than both the DL and RF approaches, although the improvement over the RF approach is very slight. We note that neither of the two benchmark methods segments other structures than lesions, and that the RF classifier is specifically trained on the contrast properties of this particular data set, and is therefore less generally applicable than the proposed and DL methods. Note that the results of the DL method are not entirely comparable, as the authors used a different set of manual annotations for validating the UNC subjects. This explains the quite large difference in performance of the DL method compared to the two other methods for subjects UNC01 and UNC06.

In very recently published work [5], the authors present a lesion segmentation framework based on deep convolutional encoder networks. This model is somewhat similar to the proposed method in the sense that both use convolutional architectures for learning suitable features for lesion detection automatically. The authors also report results on the MICCAI 2008 dataset, obtaining an average TPR of 0.40 and an average PPV of 0.41 which ties the performance of the proposed method. However, their approach suffers from the same limitations as the RF method, i.e., it is a discriminative method that only segments lesions.

4 Discussion

In this paper we have proposed a method for joint white matter lesion detection and whole-brain segmentation using a novel spatial lesion model. Due to the generative modeling approach, the method is not tied to one specific scanner platform or imaging protocol, and shows good performance when compared to the current state-of-the-art in lesion segmentation. The presented results are significantly limited by the amount of training data, which was very small given the number of parameters and potential expressive power of the RBM model. Future work will involve further experimentation with different RBM training algorithms and sampling strategies, and an extensive performance validation on

Table 1. Quantitative comparison with two state-of-the-art methods.

	DL [14]		RF [3]		Proposed			DL [14]		RF [3]		Proposed	
Patient	TPR	PPV	TPR	PPV	TPR	PPV	Patient	TPR	PPV	TPR	PPV	TPR	PPV
CHB01	0.60	0.58	0.49	0.64	0.75	0.57	UNC01	0.33	0.29	0.02	0.01	0.02	0.01
CHB02	0.27	0.45	0.44	0.63	0.57	0.48	UNC02	0.54	0.51	0.48	0.36	0.75	0.29
CHB03	0.24	0.56	0.22	0.57	0.30	0.69	UNC03	0.64	0.27	0.24	0.35	0.28	0.19
CHB04	0.27	0.66	0.31	0.78	0.59	0.49	UNC04	0.40	0.51	0.54	0.38	0.62	0.40
CHB05	0.29	0.33	0.40	0.52	0.45	0.39	UNC05	0.25	0.10	0.56	0.19	0.50	0.18
CHB06	0.10	0.36	0.32	0.52	0.19	0.50	UNC06	0.13	0.55	0.15	0.08	0.17	0.10
CHB07	0.14	0.48	0.40	0.54	0.34	0.65	UNC07	0.44	0.23	0.76	0.16	0.60	0.26
CHB08	0.21	0.73	0.46	0.65	0.37	0.70	UNC08	0.43	0.13	0.52	0.32	0.27	0.21
CHB09	0.05	0.22	0.23	0.28	0.04	0.55	UNC09	0.69	0.06	0.67	0.36	0.67	0.21
CHB10	0.15	0.12	0.23	0.39	0.19	0.69	UNC10	0.43	0.23	0.53	0.34	0.47	0.48
Mean		DL [14]				RF [3]				Proposed			
		TPR=0.33 PPV=0.37				TPR=0.40 PPV= 0.40				TPR= 0.41 PPV= 0.40			

larger data sets of white matter lesions. We further plan to also evaluate the obtained healthy structure segmentations by quantifying local atrophy patterns in large collections of brain images of MS patients.

Acknowledgements. This research was supported by NIH NCRR (P41-RR14075), NIBIB (R01EB013565), the Lundbeck Foundation (R141-2013-13117), and financial contributions from the Technical University of Denmark.

References

1. Tomas-Fernandez, X., Warfield, S.: A model of population and subject (MOPS) intensities with application to multiple sclerosis lesion segmentation. *IEEE Trans. Med. Imaging* **34**(6), 1349–1361 (2015)
2. García-Lorenzo, D., Francis, S., Narayanan, S., Arnold, D.L., Collins, D.L.: Review of automatic segmentation methods of multiple sclerosis white matter lesions on conventional magnetic resonance imaging. *Med. Image Anal.* **17**(1), 1–18 (2013)
3. Geremia, E., Menze, B.H., Clatz, O., Konukoglu, E., Criminisi, A., Ayache, N.: Spatial Decision Forests for MS Lesion Segmentation in Multi-Channel MR Images. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) *MICCAI 2010, Part I. LNCS*, vol. 6361, pp. 111–118. Springer, Heidelberg (2010)
4. Karimaghloo, Z., Rivaz, H., Arnold, D.L., Collins, D.L., Arbel, T.: Adaptive voxel, texture and temporal conditional random fields for detection of Gad-Enhancing multiple sclerosis lesions in brain MRI. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *MICCAI 2013, Part III. LNCS*, vol. 8151, pp. 543–550. Springer, Heidelberg (2013)
5. Brosch, T., Yoo, Y., Tang, L.Y.W., Li, D.K.B., Traboulsee, A., Tam, R.: Deep Convolutional Encoder Networks for Multiple Sclerosis Lesion Segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015, Part II. LNCS*, vol. 9351, pp. 3–11. Springer, Heidelberg (2015)

6. Filippi, M., Rocca, M.A., Arnold, D.L., Bakshi, R., Barkhof, F., De Stefano, N., Fazekas, F., Frohman, E., Wolinsky, J.S.: EFNS guidelines on the use of neuroimaging in the management of multiple sclerosis. *Eur. J. Neurol.* **13**(4), 313–325 (2006)
7. Puonti, O., Iglesias, J.E., Van Leemput, K.: Fast, sequence adaptive parcellation of brain MR using parametric models. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *MICCAI 2013, Part I. LNCS*, vol. 8149, pp. 727–734. Springer, Heidelberg (2013)
8. Smolensky, P.: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1, pp. 194–281. MIT Press, Cambridge (1986)
9. Van Leemput, K., Maes, F., Vandermeulen, D., Colchester, A., Suetens, P.: Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE Trans. Med. Imaging* **20**(8), 677–688 (2001)
10. Styner, M., Lee, J., Chin, B., Chin, M., Commowick, O., Tran, H., Markovic-Plese, S., Jewells, V., Warfield, S.: 3D segmentation in the clinic: a grand challenge II: MS lesion segmentation. *MIDAS J.*, 1–5 (2008)
11. Shiee, N., Bazin, P.L., Ozturk, A., Reich, D.S., Calabresi, P.A., Pham, D.L.: A topology-preserving approach to the segmentation of brain images with multiple sclerosis lesions. *NeuroImage* **49**(2), 1524–1535 (2010)
12. Lee, H., Grosse, R., Ranganath, R., Ng, A.Y.: Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In: *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009*, pp. 609–616. ACM, New York (2009)
13. Tieleman, T.: Training restricted Boltzmann machines using approximations to the likelihood gradient. In: *Proceedings of the 25th International Conference on Machine Learning, ICML 2008*, pp. 1064–1071. ACM, New York (2008)
14. Weiss, N., Rueckert, D., Rao, A.: Multiple sclerosis lesion segmentation using dictionary learning and sparse coding. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *MICCAI 2013, Part I. LNCS*, vol. 8149, pp. 735–742. Springer, Heidelberg (2013)
15. Souplet, J., Lebrun, C., Ayache, N., Malandain, G.: An automatic segmentation of T2-FLAIR multiple sclerosis lesions. In: *The MIDAS Journal - MS Lesion Segmentation (MICCAI 2008 Workshop)* (2008)

Brainlesion: Glioma, Multiple Sclerosis, Stroke and
Traumatic Brain Injuries

First International Workshop, Brainles 2015, Held in
Conjunction with MICCAI 2015, Munich, Germany,
October 5, 2015, Revised Selected Papers

Crimi, A.; Menze, B.; Maier, O.; Reyes, M.; Handels, H.
(Eds.)

2016, IX, 298 p. 117 illus., Softcover

ISBN: 978-3-319-30857-9