White Matter MS-Lesion Segmentation Using a Geometric Brain Model

Maddalena Strumia, Frank R. Schmidt, Constantin Anastasopoulos,
Cristina Granziera, Gunnar Krueger, Thomas Brox

***Abstract*—Brain magnetic resonance imaging (MRI) in patients with Multiple Sclerosis (MS) shows regions of signal abnormalities, named plaques or lesions. The spatial lesion distribution plays a major role for MS diagnosis. In this paper we present a 3D MS-lesion segmentation method based on an adaptive geometric brain model. We model the topological properties of the lesions and brain tissues in order to constrain the lesion segmentation to the white matter. As a result, the method is independent of an MRI atlas. We tested our method on the MICCAI MS grand challenge proposed in 2008 and achieved competitive results. In addition, we used an in-house dataset of 15 MS patients, for which we achieved best results in most distances in comparison to atlas based methods. Besides classical segmentation distances, we motivate and formulate a new distance to evaluate the quality of the lesion segmentation, while being robust with respect to minor inconsistencies at the boundary level of the ground truth annotation.**

***Index Terms*—Multi label segmentation, Graph-Cuts, topological constraints, connectivity.**

I. Introduction

Multiple Sclerosis (MS) is an inflammatory demyelinating disease of the central nervous system and its hallmarks are demyelination plaques, inflammation, axonal damage, and edema [1]. MS imaging pathology, which is routinely used for disease diagnosis, shows typical abnormalities in the MR signal of the White Matter (WM), known as lesions. Accordingly, lesion identification is a crucial process in quantifying the burden of the disease. In clinical routine the lesions are manually processed. This process is time consuming, leading to subjective variability, and to a lack of reproducibility.

Recently, a variety of methods have been proposed to segment WM lesions for MS patients (cf. [2], [3] for an overview).

M. Strumia is with the Department of Radiology Medical Physics, University Medical Center Freiburg 79106, Germany and with the German Cancer Research Center (DKFZ) and the German Cancer Consortium (DKTK), Heidelberg 69121, Germany (email: maddalena.strumia@uniklinik-freiburg.de).

F. R. Schmidt is with the Computer Vision Group, TU München, Department of Computer Science, Garching 85748, Germany and with the BIOSS Centre of Biological Signalling Studies University Freiburg, Germany.

C. Anastasopoulos is with the Department of Neuropediatrics and of Neuroradiology, University Medical Center Freiburg 79106, Germany.

C. Granziera is with the Neuro-immunology Unit and Laboratoire de Recherché en Neuroimagérie (LREN), Neurology Division, Department of Clinical Neurosciences, Centre Hospitalier Universitaire Vaudois and University of Lausanne, Switzerland.

G. Krueger is with the Advanced Clinical Imaging Technology group, Siemens Healthcare IM BM PI, Lausanne, Switzerland.

T. Brox is with the BIOSS Centre of Biological Signalling Studies University Freiburg, Germany and with the Department of Computer Science, University of Freiburg, Germany.

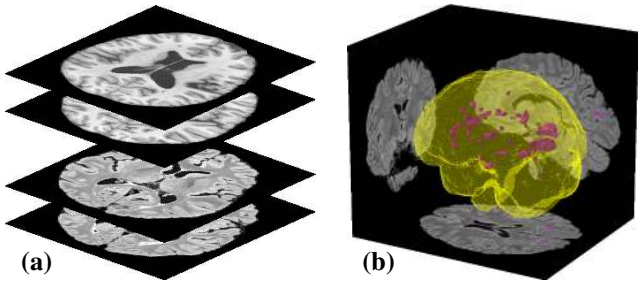

Fig. 1. **Lesion segmentation.** Starting from a 3D image stack (a) containing $c > 0$ image contrasts (here the $c = 2$ is for FLAIR and T_1-w images), we consider a geometric brain model to find the WM lesions of an MS patient (b).

In general most of the methods use an atlas precomputation, a local classifier or a combination of the two paradigms.

Many methods rely strongly on an atlas driven pre-segmentation that is then post-processed [4], [5], [6]. Even though an atlas works reasonably well, the exact boundary between WM and Grey Matter (GM) is difficult to detect due to brain inter-variability. As a result, the detected WM can contain areas of GM that might be segmented as lesion. In addition, since the exact location of lesions cannot be pre-learned, it is hard to apply a purely atlas-driven approach to lesion segmentation.

Local classifiers can be seen as a binary decision process which discriminates between lesions and non-lesion voxels. These classifiers do not consider the geometric structure of the brain and assume either a good pre-selection of the WM [7], [8] or that lesions can be easily detected as outliers with respect to a dictionary based brain model [9], [10].

In this work, we propose an automated segmentation method that uses a geometric model of the brain instead of an atlas (cf. Fig. 1). Instead of learning the location of different tissue types, we propose topological priors such as the connectivity of GM. In addition, we model the relative locations of each tissue type, e.g. the WM lesions are surrounded by WM and the WM is surrounded by GM (cf. l.h.s. of Fig. 3). As a result, we restore the benefits of a geometric model without the artifacts that an atlas registration would entail.

To benefit from these advantages, we have to minimize an energy that is an instance of an NP-hard problem class. In this paper we combine different approximation methods [12], [13] in order to find a meaningful local minimum of this energy. To make our method applicable to a wide range of images obtained by different MR systems and protocols, we represent our parameters with respect to more general hyper-parameters.

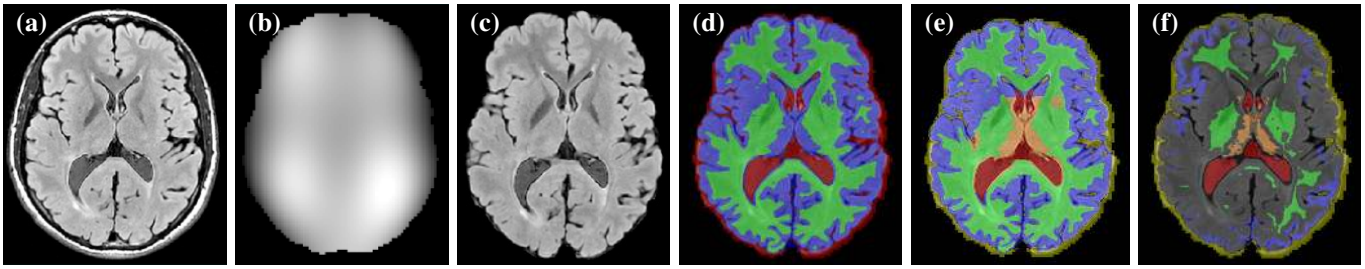

Fig. 2. **Pre-processing.** Based on the original images, such as the FLAIR image (a), a bias field (b) is computed [11] in order to correct (c) the input images. The bias field computation also provides a coarse segmentation into GM, WM and CSF (d). These segments are refined to obtain segments for SGM and V as well (e). Strongly eroding these segments $\hat{S}_\ell^{(0)}$, provides the segments $\hat{S}_\ell$ that we treat as hard constraints (f) in our framework (cf. Section III-A).

These hyper-parameters are estimated in a hyper-parameter learning approach, for which we can find the global optimum.

To test the proposed method and compare it against state-of-the-art methods, we evaluate the lesion segmentation results with standard distances. We also introduce an additional distance that is more robust to minor changes in the lesion annotation.

A. Related Work

Atlas based methods usually start with a registration of the patient data to an atlas in order to determine the locations of the major brain tissue such as GM, WM, and the Cerebro-Spinal Fluid (CSF). Due to the high inter-brain variability this step may result in an erroneous tissue classification.

Constraining the lesion segmentation into or close to the WM is a common procedure when dealing with atlas based approaches [14], [15]. Both papers perform an initial atlas based segmentation of GM, WM, CSF. Kamber et al. [14] showed that confining the lesions into plausible white matter locations, calculated from a set of healthy volunteers, largely improves the results. As a drawback, lesions which are close to the cortex or the subcortical GM are missed. Val Leemput et al. [15] relaxed this inclusion by imposing that the lesions should be located in the vicinity of the WM.

Lesion-TOADS (Topology-preserving Anatomical Segmentation) [5] is an atlas based method that uses the information from a topological and an intensity atlas created from healthy volunteers. As a result, lesions are ignored in a first segmentation step. In a second step, a two class separation within the WM mask is performed to separate lesions from WM. The quality of the resulting lesion segmentation depends on the first segmentation step. In contrast to this approach, we model the location of WM lesions while segmenting the rest of the brain. In addition, we are able to incorporate topological information without the use of an atlas.

The method proposed in [4] is less sensitive to atlas-driven misclassification. As a first step, it extracts the WM and GM tissues from an initial atlas driven segmentation. In a second step, outliers within the WM mask are extracted as long as they differ from the intensity model of GM. A morphological post-precessing results in a lesion segmentation. An extension, the Lesion Segmentation Tool (LST), was proposed by Schmidt et al. [6]. It uses the outliers as initialization of a region growing approach for lesion segmentation. Both approaches are very sensitive to local information.

The Model of Population and Subjects (MOPS) method [16] extends the classical atlas based approach using local intensity models instead of global ones. This method performs a tissue segmentation of GM, WM, CSF coupling global and local Gaussian Mixture Models (GMM). The global GMMs are computed from the patient data. The local models are extracted from a dataset of healthy volunteers, estimating a Gaussian Mixture Model (GMM) per voxel. The lesions are initially identified as the voxels with a low probability of belonging to any brain tissue, and subsequently, a graph-cut segmentation is performed.

It is also common to train lesion classifiers. An atlas based brain initial segmentation is often combined with a pixel-wise classification. In the presence of small lesions these classifiers may mix lesion and WM information. Warfield et al. [17] proposed an iterative process that combines image registration with k-nearest-neighbor classification. The atlas, against which the patient data is registered, is constructed from a single volunteer by manually annotating relevant anatomic structures. The labels, together with additional features extracted from the patient data, are classified as GM, WM, CSF or lesion. Additionally, the localization of GM and WM is improved by employing a region growing technique. This allows for a high penalty for misclassified GM and WM voxels, which ensue in a more precise lesion segmentation. Zijdenbos et al. [7] created a pipeline based on an Artificial Neural Network (ANN) classifier to discriminate lesion and non lesion voxels, using multiple images and three tissue priors GM, WM, and CSF calculated from an atlas of healthy volunteers. Geremia et al. [8] proposed a method in which a lesion classifier is trained in a random forest framework. While the classifier takes a neighborhood into account, the final decision is made for every pixel independently. Cabezas et al. [18] combine these classifiers with others in a Gentleboost framework.

Patch-based methods have also been developed for brain lesion segmentation. The size of the patch plays an important role in the representation of lesion and healthy tissue. In addition a threshold value needs to be estimated to distinguish the lesion voxels from false positives. This may produce a noisy lesion segmentation, since there is no neighborhood information considered. Weiss et al. [9] sparsely encode the patient image using $3 \times 3 \times 3$-patches as a dictionary and

	CSF	GM	WM	SGM	L	V
CSF	0	$\lambda_{\mathrm{CSF,GM}}$	∞	∞	∞	∞
GM	$\lambda_{\mathrm{CSF,GM}}$	0	$\lambda_{\mathrm{GM,WM}}$	∞	∞	∞
WM	∞	$\lambda_{\mathrm{GM,WM}}$	0	$\lambda_{\mathrm{WM,SGM}}$	$\lambda_{\mathrm{WM,L}}$	∞
SGM	∞	∞	$\lambda_{\mathrm{WM,SGM}}$	0	∞	$\lambda_{\mathrm{SGM,V}}$
L	∞	∞	$\lambda_{\mathrm{WM,L}}$	∞	0	∞
V	∞	∞	∞	$\lambda_{\mathrm{SGM,V}}$	∞	0

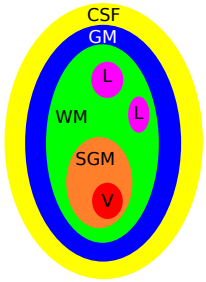

Fig. 3. **Geometric model.** The relationship between different tissue layers ℓ_1 and ℓ_2 is modeled by weighting their interface with $\lambda_{\ell_1,\ell_2} \geq 0$. An infinite cost excludes two labels from having a common boundary. The λ-values are based on the geometric model that is depicted on the right hand side (cf. Section III-B).

the reconstruction error gives the lesion segmentation. Guizard et al. [10] propose a non-local mean approach introducing a rotational invariant distance measure (RMNMS) to account for the diversity of MS lesions.

Recently, also non atlas based methods have been proposed for lesion segmentation. Gao et al. [19] propose a regularized segmentation which encodes each tissue type with a constant intensity value while taking a global bias correction field into account. A fuzzy segmentation is computed by formulating a global energy functional that is convex in each variable. A local minimum is computed by an iterative minimization approach. A graph-cut based segmentation was proposed in Garcia-Lorenzo et al. [20], where the authors use a binary segmentation to discriminate the lesions from the healthy tissues. Post-processing steps are also included to remove lesions appearing in external CSF and close to the brain borders. While these methods are free from any atlas misclassification, there is no geometric information encoded. We believe that encoding geometric information helps to obtain a more robust segmentation.

II. Pre-Processing

In the following, we treat a medical image as a mapping $I\colon \Omega \to \mathbb{R}^c$ that assigns to each voxel $x \in \Omega$ of the image domain Ω exactly $c \in \mathbb{N}$ image contrasts[1] $I(x) \in \mathbb{R}^c$. The goal of a brain segmentation is to find a labeling $f\colon \Omega \to \mathcal{L}$ that assigns to each voxel x a label $f(x) \in \mathcal{L}$ of the label space $\mathcal{L} := \{\mathrm{CSF}, \mathrm{GM}, \mathrm{WM}, \mathrm{SGM}, \mathrm{V}, \mathrm{L}\}$. These labels correspond to sulcal CSF (CSF), Grey Matter (GM), White Matter (WM), Subcortical Grey Matter nuclei (SGM), Ventricles (V) and Lesions (L), respectively. In addition, we denote $S_\ell = \{x \in \Omega | f(x) = \ell\}$ as the segment of the label ℓ. Note that in our setting, the label V anatomically includes the two lateral ventricles.

In a first step, we align the c image contrasts with a rigid registration, which is estimated using the default parameters of the ITK libraries [21]. Subsequently we remove the skull in the images via the BET toolkit [22] and perform a bias-field correction by extending [11] to $c > 1$ image contrasts.

In a second step, we compute a partial segmentation that we will use as a hard constraint for our method. The bias correction [11] already provides us with a combination of the six labels in $\mathcal{L}$, *i.e.*, $\mathrm{csf} = \{\mathrm{CSF}, \mathrm{V}\}$, $\mathrm{gm} = \{\mathrm{GM}, \mathrm{SGM}\}$, WM as shown in Fig. 2(d). In order to separate V from CSF, we perform a binary graph cut segmentation [23] on the csf mask. Similarly, we also perform the separation of SGM from GM on the gm mask via graph cut. This provides for an initial labeling $f_0\colon \Omega \to \mathcal{L} - \{\mathrm{L}\}$ and an initial segmentation $\left(\hat{S}_\ell^{(0)}\right)_{\ell \in \mathcal{L}}$, where $\hat{S}_\ell^{(0)} := \{x \in \Omega | f_0(x) = \ell\}$ (cf. Fig. 2(e)). Since the method [11] performs a three class segmentation, the lesions are inside $\hat{S}_{\mathrm{WM}}^{(0)}$, $\hat{S}_{\mathrm{GM}}^{(0)}$, and $\hat{S}_{\mathrm{SGM}}^{(0)}$. To eliminate these misclassified voxels we remove independently from these segments the voxels with an intensity higher than the median value in the FLAIR image. Afterwards, we erode each segment with a sphere of radius 2mm, resulting in the partial segmentation $\hat{S}_\ell$. Due to the strong erosion, it covers less than 15% of the whole image domain (cf. Fig. 2(f)). While the $\hat{S}_\ell$ provides a strong partial segmentation, it will guide the segmentation process.

The following section describes the energy that we want to minimize to obtain a full brain segmentation that provides us with a reliable lesion segmentation. In the optimization process, described in Section IV, we rely on a few voxels that guide the minimization. These voxels are represented by the partial segmentation $\hat{S}_\ell$ and we treat them as hard constraints (cf. Section III-A).

III. Geometric Brain Model

Given an input image $I\colon \Omega \to \mathbb{R}^c$, we are looking for the most likely brain labeling $f\colon \Omega \to \mathcal{L}$. Using Bayes' theorem, it is common to write the posterior probability $P(f|I)$ as

$$P(f|I) = \frac{P(I|f) \cdot P(f)}{P(I)} \tag{1}$$

The maximization of (1) is equivalent to the minimization of

$$-\log P(f|I) = -\log(P(I|f)) - \log(P(f)) + \text{const.}, \tag{2}$$

where $P(I|f)$ describes the likelihood that the image I originates from the labeling f, and $P(I)$ denotes the probability that the observed image is possible. We treat this probability as a constant, since our brain segmentation is not affected by $P(I)$. To estimate this likelihood, a standard way is to compute the probabilistic intensity models for each tissue type (cf. Section III-A). Hence, we refer to $E_{\mathrm{App}}(f) := -\log(P(I|f))$ as the *appearance energy*.

[1] Here, $c = 2$ may reflect a setup of FLAIR and T_1-w image contrasts.

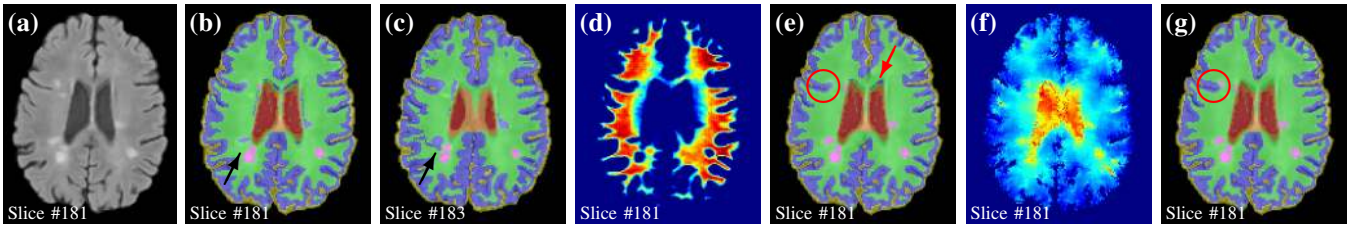

Fig. 4. **Connectivity.** Segmenting (a) without the connectivity constraint may lead to a disconnected GM (arrow in (b)). Enforcing connectivity as in [24] creates a thin artifact connection path (c). [25] uses an appearance based distance map (d) to enforce connectivity, but fails in our multi-labeling setting (arrow in (e)). Using the proposed gradient based distance map (f) instead results in a connected GM (g) and an improved lesion segmentation (circle in (e) and (g)).

$P(f)$ models the prior information of our geometric model, which is the main contribution of this work. Here we model two independent geometric properties, *i.e.*, $P(f) = P_{\text{Loc}}(f) \cdot P_{\text{Geo}}(f)$. Firstly, only certain transitions between tissue types are allowed, resulting in $P_{\text{Loc}}(f)$ (cf. Section III-B). Secondly, we want to enforce some global geometric properties of the brain, resulting in $P_{\text{Geo}}(f)$ (cf. Section III-C). Defining $E_{\text{Loc}}(f) := -\log(P_{\text{Loc}}(f))$ as the *local geometric energy* and $E_{\text{Geo}}(f) := -\log(P_{\text{Geo}}(f))$ as the *global geometric energy*, we seek to minimize the following energy

$$E(f) := E_{\text{App}}(f) + E_{\text{Loc}}(f) + E_{\text{Geo}}(f). \tag{3}$$

In the remainder of this section, we will describe these different energies in detail. The minimization of (3) will be addressed in Section IV.

A. Appearance Energy

In brain lesion segmentation a standard way to represent tissue intensities is to model each tissue type with a Gaussian Mixture Model (GMM) [16], [15]. In this work we follow the same idea and we denote the GMM probability of a tissue $\ell \in \mathcal{L}$ as p_ℓ.

The voxels in the precomputed partial segmentation $\hat{S}_\ell$ (cf. Section II) are hard constraints and will not change their label, so the probability that a voxel x is labeled ℓ is

$$P_\ell(x) = \begin{cases} 1 & \text{if } x \in \hat{S}_\ell \\ 0 & \text{if } \exists \ell' \neq \ell : x \in \hat{S}_{\ell'} \\ p_\ell(I(x)) & \text{otherwise.} \end{cases} \tag{4}$$

For the appearance model P_ℓ, we assume that each voxel x that belongs to the tissue type ℓ is an independent, identically distributed random variable. Thus the *appearance energy* can be written as

$$\begin{aligned} E_{\text{App}}(f) &= -\log\left(\prod_{x\in\Omega} P_{f(x)}(I(x))\right) \\ &= \sum_{x\in\Omega} -\log\left(P_{f(x)}(I(x))\right) \end{aligned} \tag{5}$$

where we assumed $-\log(0) = +\infty$. Due to the used probability P_ℓ in (4), every segmentation f will be penalized by an infinite cost if any voxel $x \in \hat{S}_\ell$ in the partial segment $\hat{S}_\ell$ will be assigned a label $f(x) \neq \ell$. This implies that $\hat{S}_\ell \subset S_\ell$ for all $\ell \in \mathcal{L} - L$, *i.e.*, the segmentation S_ℓ must always contain $\hat{S}_\ell$. Therefore the inclusion $\hat{S}_\ell \subset S_\ell$ becomes a global constraint that is always enforced during the optimization process.

B. Local Geometric Model

To avoid overly noisy labelings it is common to penalize neighboring voxels if their labels disagree. This penalty is represented as a pair-wise term between adjacent voxels. In this work we use these terms to additionally impose our local topological constraints. To this end, we use a 6-neighborhood $\mathcal{N}(x)$ for each voxel and define the geodesic surface area as

$$E_{\text{Loc}}(f) = \sum_{x\in\Omega} \sum_{y\in\mathcal{N}(x)} \lambda_{f(x),f(y)} \cdot g(x,y). \tag{6}$$

where $g : \Omega \times \Omega \to \mathbb{R}^+$ is a gradient based function which provides a high penalty between voxels of similar intensities [23], [26], [27]:

$$g(x,y) = \frac{\exp\left(-\frac{1}{2}\left(\frac{|I(x)-I(y)|}{\sigma}\right)^{1.5}\right)}{\text{dist}(x,y)} + \epsilon, \tag{7}$$

where ϵ is a positive scalar value which guarantees that $g(x,y)$ always exceeds the threshold ϵ. σ guarantees that our method does not depend on the intensity range of the image I and the distance $\text{dist}(\cdot,\cdot)$ between two voxels takes the anisotropic resolution of I into account. The symmetric weighting factor $\lambda_{f(x),f(y)}$ handles different label transitions individually. To favor a smooth segmentation we set $\lambda_{\ell_1,\ell_2} = 0$ if $\ell_1 = \ell_2$. In addition, we only want to allow those segmentations that satisfy the following proposed topological conditions (cf. r.h.s. of Fig. 3):

- **CSF** Since it is only adjacent to GM, we set $\lambda_{\text{CSF},\ell} = \infty$ for $\ell \in \{\text{WM}, \text{SGM}, \text{L}, \text{V}\}$.
- **GM** Since it is surrounded by CSF and it is adjacent to the WM, we set $\lambda_{\text{GM},\ell} = \infty$ for $\ell \in \{\text{SGM}, \text{L}, \text{V}\}$.
- **WM** The WM surrounds the lesions and is adjacent to the GM as well as the SGM. Hence, we set $\lambda_{\text{WM},\ell} = \infty$ only for $\ell \in \{\text{CSF}, \text{V}\}$.
- **L** Since they are only adjacent to WM, we set $\lambda_{\text{L},\ell} = \infty$ for $\ell \in \{\text{CSF}, \text{GM}, \text{SGM}, \text{V}\}$.
- **SGM** Since it is adjacent to the ventricle and the WM, we set $\lambda_{\text{SGM},\ell} = \infty$ for $\ell \in \{\text{CSF}, \text{GM}, \text{L}\}$.
- **V** We model the ventricle as being surrounded by the subcortical GM. Therefore, we set $\lambda_{\text{V},\ell} = \infty$ for $\ell \in \{\text{CSF}, \text{GM}, \text{WM}, \text{L}\}$.

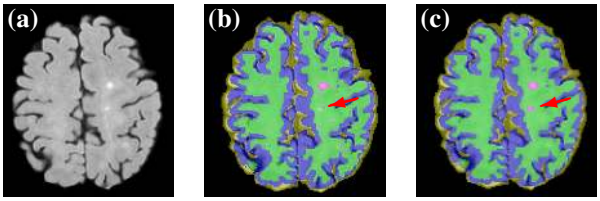

Fig. 5. **Minimal lesion size.** Segmenting (a) without the constraint (9) results in a noisy segmentation (b). After enforcing the constraint, small lesion segments either disappear or grow to satisfy the constraint (arrow in (b,c)).

The different transitions are illustrated in the sketch of Fig. 3, while the values for λ_{ℓ_1,ℓ_2} are summarized in the table of Fig. 3. Note that the six labels allow only for five different transitions, which makes the use of α/β swap in Section IV very appealing.

C. Global Geometric Model

The probability distributions p_ℓ for $\ell \in \{\mathrm{GM}, \mathrm{L}, \mathrm{SGM}\}$ are difficult to discriminate. Therefore, we introduce additional global constraints to the segments S_ℓ of these labels. Anatomically, the GM is continuously connected and we assume in our model that the SGM is also connected (cf. Fig. 3). Therefore, we can enforce a connectivity prior on S_{GM} and S_{SGM}, resulting in the energy E_{CC}. In order to be robust with respect to noise in the WM, we also like to enforce that each lesion contains at least a ball of radius 1 voxel, resulting in the energy $E_{\min}$. These two energies are defined as:

$$E_{\mathrm{CC}}(f) = \begin{cases} \infty & \text{if } S_{\mathrm{GM}} \text{ or } S_{\mathrm{SGM}} \text{ are not connected} \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

and

$$E_{\min}(f) = \begin{cases} 0 & \forall A \subset S_{\mathrm{L}} \text{ is connected component} \\ & \exists x \in A : \forall y \notin A : \|x - y\| > 1 \\ \infty & \text{otherwise} \end{cases} \tag{9}$$

Combining these two energies leads to:

$$E_{\mathrm{Geo}}(f) = E_{\mathrm{CC}}(f) + E_{\min}(f).$$

An example of the minimal size prior is shown in Fig. 5. The lesions with a radius smaller than 1 voxel can disappear or increase in size (cf. Fig. 5(c)).

IV. Optimization

Since the energy (3) is not submodular [24], we search for a local optimum which will depend on a good initialization. Following a coarse-to-fine approach, we subsample the image $I : \Omega \to \mathbb{R}^c$ and its image domain Ω by a factor of $s^{(1)}_{x_1} \times s^{(1)}_{x_2} \times s^{(1)}_{x_3}$, where $x = (x_1, x_2, x_3)$, to obtain $I_1 : \Omega_1 \to \mathbb{R}^c$. By subsampling I_1 with another factor of $s^{(2)}_{x_1} \times s^{(2)}_{x_2} \times s^{(2)}_{x_3}$ we obtain $I_2 \colon \Omega_2 \to \mathbb{R}^c$. The subsampling factors $s^{(1)}_{x_1}, s^{(1)}_{x_2}, s^{(1)}_{x_3}, s^{(2)}_{x_1}, s^{(2)}_{x_2}, s^{(2)}_{x_3} \geq 1$, and their values depend on the resolution of the data.

In Section IV-A we compute an initialization for the labeling $f|\Omega_2$ (cf. Fig. 6(b)). In Sections IV-B we improve $f|\Omega_2$ (cf. Fig. 6(c)). After finding a local minimum for $f|\Omega_2$ (cf. Fig. 6(d)), we iteratively increase the resolution and find local minima for $f|\Omega_1$ and finally for f itself (cf. Fig. 6 (e-f)). The λ_{ℓ_1,ℓ_2} values of Section III-B are fixed parameters with the exception of $\lambda_{\mathrm{WM,L}}$. In Section IV-C we learn globally optimal hyper-parameters to choose the correct value for $\lambda_{\mathrm{WM,L}}$.

A. Computing the initial labeling

A multi-labeling problem is in general NP-hard, but it becomes polynomial if a total ordering $\leq_{\mathcal{L}}$ of $\mathcal{L}$ is given such that λ_{ℓ_1,ℓ_2} is convex with respect to this ordering. Note that the λ_{ℓ_1,ℓ_2} values, in Equation (6), are constant for each pairs of labels $(\ell_1, \ell_2) \in \mathcal{L}^2$. However, λ can also be interpreted as a function from the label space $\mathcal{L}^2$ to its value: $\lambda : \mathcal{L}^2 \to \mathbb{R}$. Ishikawa studied λ functions that can be written as $\lambda_{\ell_1,\ell_2} = F(|\ell_1 - \ell_2|)$ for a convex function F [13]. The polynomial runtime of [13] was achieved by first allowing an independent binary decision for each label ℓ and thus expanding the label space to the set of its subsets $2^{\mathcal{L}}$. Therefore, the multi-labeling $f \colon \Omega \to \mathcal{L}$ becomes the extended multi-labeling $\hat{f} \colon \Omega \to 2^{\mathcal{L}}$. Secondly, Ishikawa introduced the inclusion constraints

$$\ell \in \hat{f}(x) \Rightarrow \ell' \in \hat{f}(x) \qquad \text{for all } \ell' \leq_{\mathcal{L}} \ell,$$

which ensure that $\hat{f}(x)$ has the form $\{\ell' \in \mathcal{L} | \ell' \leq_{\mathcal{L}} \ell\}$ for some $\ell \in \mathcal{L}$. He showed that the data terms can be chosen in a way that an optimum of the original multi-labeling $f \colon \Omega \to \mathcal{L}$ is derived from the optimal $\hat{f} \colon \Omega \to 2^{\mathcal{L}}$ via $f(x) := \max \hat{f}(x)$.

Delong and Boykov proved in [28] that a simplified problem of what we consider in (3) fits into the Ishikawa formulation, which can be globally optimized. To achieve this, we have to change the label space and we can only consider the energy $E_0(f) := E_{\mathrm{App}}(f) + E_{\mathrm{Loc}}(f)$. For the altered label space, we can either choose a smaller label space or an extended label space. We can now consider the maximal lower bound $\check{\mathcal{L}} \subset \mathcal{L}$ and the minimal upper bound $\hat{\mathcal{L}} \supset \mathcal{L}$ that fit into the Ishikawa construction. Since the label set $\check{\mathcal{L}} = \mathcal{L} - \{\mathrm{L}\}$ ignores the lesion information, we rather use

$$\hat{\mathcal{L}} := \{\mathrm{CSF}, \mathrm{GM}, \mathrm{WM}, \mathrm{L}, \{\mathrm{L}, \mathrm{SGM}\}, \mathrm{SGM}, \mathrm{V}, \mathrm{F}_2\}. \tag{10}$$

This produces two labels that we did not intend to model. As a result, the lesions and the SGM are not separated, leading to the *false label* $\mathrm{F}_1 := \{\mathrm{L}, \mathrm{SGM}\}$. The *false label* F_2 is an artifact from the ventricle model and is embedded inside V (cf. Fig.6(b)).

We use the global optimum $\hat{f} \colon \Omega_2 \to \hat{\mathcal{L}}$ of $E_0(\cdot)$ to find a feasible labeling $f \colon \Omega_2 \to \mathcal{L}$ with respect to $E(\cdot)$, *i.e.*, a labeling that satisfies $E(f) < \infty$. Since F_2 uses the same appearance model as V, we replace F_2 with V. Since F_1 combines L and SGM, we have to separate these two labels. In order to disallow transition between SGM and L, some of the WM has to appear between SGM and L. To solve this, we perform α/β-swapping [12] in the region identified by $S_{\mathrm{L}} \cup S_{\mathrm{SGM}} \cup S_{\mathrm{F}_1}$ with respect to the label space $\{\mathrm{WM}, \mathrm{SGM}, \mathrm{L}\}$. This leads to a mapping $f \colon \Omega_2 \to \mathcal{L}$ with $E_0(f) < \infty$, but it may still violate some global geometrical

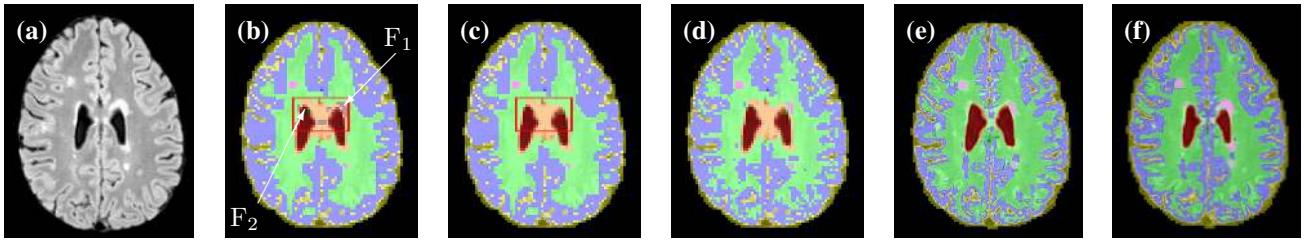

Fig. 6. **Method overview.** Given the input images including the FLAIR image (a), minimizing $E_0(f)$ produces the initial labeling (b), which contains the *false labels* F_1 (grey) and F_2 (dark green). After removing the false labels (c), the method improves the computed minima in a coarse-to-fine approach (d-f).

constraints, resulting in $E(f) = \infty$. In this case, the minimal lesion size constraint cannot be violated due to our subsampling that results in a lesion diameter that is larger than 1 voxel in the original resolution. Consequently, either S_{GM} or S_{SGM} are not connected. Without loss of generality we assume that S_{GM} is not connected.

A connectivity prior is in general NP-hard. Vicente et al. [24] proposed a connectivity constraint for binary image segmentation and for user selected points in the image. These conditions limit the automated process and the multi-labeling segmentation. In a recent approach, Stühmer et al. [25] proposed a connectivity prior introducing a distance measure based on the image intensities. This is a feasible approach if the connected object has a distinguishable intensity distribution from the other labels/objects. Since the GM, SGM, and the L have similar intensity distributions, we use a distance function based on the gradient values (cf. Fig. 4(f)) which is able to correctly connect the GM tissue and recover some L labels (cf. Fig. 4(g)). After encoding these extra constraints, we obtain an initial segmentation $f\colon \Omega_2 \to \mathcal{L}$ that satisfies $E(f) < \infty$ (cf. Fig. 6 (c)).

B. Local Energy Minimization

In Section IV-A we showed how to obtain a feasible initialization $f\colon \Omega_2 \to \mathcal{L}$, *i.e.*, a labeling that satisfies $E(f) < \infty$. In this section we improve our labeling in order to find a local minimum. Therefore, we perform all possible α/β-swaps on the whole label space $\mathcal{L}$. Note that even though we have to deal with six labels, only five of the 15 possible α/β-swaps are necessary (cf. Fig. 3).

Every α/β-swaps does not increase the energy, and it halts as soon as the energy does not improve between two iterations. Upon convergence, we improve the GMMs p_ℓ of Section III-A in an EM-like fashion. Since we initialize the GMMs with respect to the hard constraint $\hat{S}_\ell$ (cf. Section II), we want to avoid large jumps in the intensity model during GMM optimization. The GMM optimization usually starts with an initial distribution p_ℓ and improves it via gradient descent to the distribution q_ℓ. The idea that we follow here is to use only "half of the gradient update". This still reduces the overall energy without drastically changing the intensity models.

In the case of purely Gaussian models $p_\ell = G(\mu_p, \Sigma_p)$ and $q_\ell = G(\mu_q, \Sigma_q)$, we can compute an average with respect to the Bhattacharyya distance and obtain the Gaussian $G(\mu, \Sigma)$ with

$$\mu = \frac{\mu_p + \mu_q}{2} \qquad \Sigma = \Sigma_p U \Lambda^{\frac{1}{2}} U^\top,$$

where $U\Lambda U^\top$ is the diagonalization of $\Sigma_p^{-1}\Sigma_q$. In the case of Gaussian mixture models $p_\ell = G(\mu_p, \Sigma_p, \alpha_p)$ and $q_\ell = G(\mu_q, \Sigma_q, \alpha_q)$ with multiple Gaussians that are mixed with respect to the mixture coefficients α, we first match the models of p and q via the linear assignment approach [29] using the pairwise Bhattacharyya distance as a data term. Afterwards, we compute the average of the matched models as above and set $\alpha = \frac{\alpha_p + \alpha_q}{2}$.

The local optimization with respect to α/β-swaps and GMM re-estimation is iterated until convergence. Note that in the GMM re-estimation the hard constraints $\hat{S}_\ell$ are always used since $\hat{S}_\ell \subset S_\ell$ is always enforced (cf. Section III-A). Upon convergence, we extend the resolution from Ω_2 to Ω_1 and re-iterate until a locally optimal labeling $f|\Omega_1$ is found. In a final step, we optimize in the original resolution and obtain a local optimal $f\colon \Omega \to \mathcal{L}$. The complete method is summarized in Algorithm 1, and the process is shown in Fig. 6. For higher resolutions, the minimal lesion size constraint (9) may be violated. In this case, we enforce the center of mass c_A of each connected component $A \subset S_{\mathrm{L}}$ such that $f(c_A) = \mathrm{L}$ implies $f(x) = \mathrm{L}$ for all $\|x - c_A\| \leq 1$ voxel, which can be easily integrated into the graph cut optimization [28]. After introducing these constraints, small lesions will either be automatically removed or extended to include a ball with a diameter of 3 voxels.

Algorithm 1

1: Find initial labeling according to Section IV-A
2: **repeat**
3: Local optimization w.r.t. α/β-swaps and GMM re-estimation
4: **until** convergence
5: Increase resolution if possible and **goto Line 2**

C. Hyper-parameter Learning

The GMM parameters *i.e.*, means and covariances, are optimized for every tissue with an Expectation Maximization algorithm [30]. The remaining parameters are the five different λ terms in Section III-B. Since we are mainly interested in lesion segmentation, we empirically fix every λ value except for $\lambda_{\mathrm{WM,L}}$. In order to make this term adaptive to the contrast in the image, we model it as being affinely dependent

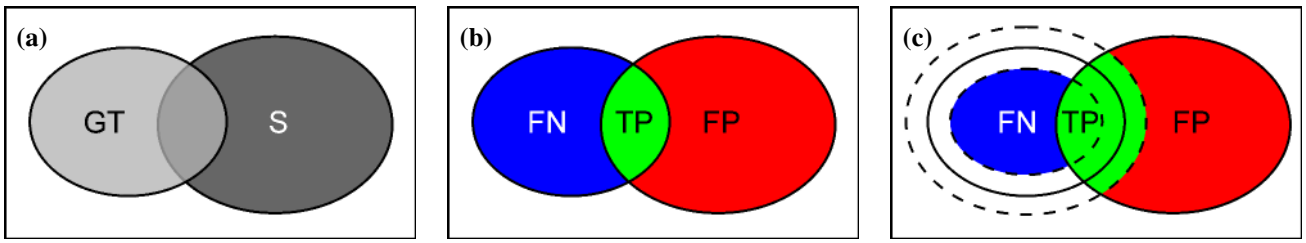

Fig. 7. **Distance Dice.** Given ground truth GT and segmentation S (a), the Dice coefficient D combines FN, FP and TP (b). Introducing a tolerance zone of distance θ (dashed lines) improves these quantities. As a result, we obtain that D_θ is more tolerant with respect to small changes along the boundary (c).

on the maximum magnitude $\mathrm{mg}(I) := \max_{x\in\Omega} \|\nabla I(x)\|$ of the image gradient $\nabla I(x) = (\frac{\partial I}{\partial x_1}, \frac{\partial I}{\partial x_2}, \frac{\partial I}{\partial x_3})$, *i.e.*, $\lambda_{\mathrm{WM,L}} = m \cdot \mathrm{mg}(I) + b$ where m and b are the affine hyper-parameters that we would like to choose optimally.

We can find the optimal values for m and b by using the parametric max flow framework [31]. It provides many finite intervals for $\lambda_{\mathrm{WM,L}}$ for each of the patients of the training dataset. For each $\lambda_{\mathrm{WM,L}}$ a segmentation is computed, and the corresponding energy value is stored. These intervals tessellate the (m, b)-domain into a finite number of regions. For each region the energy value is computed as the sum of the corresponding energies in the training set. The region with the globally minimal energy provides us with the optimal values for the hyperparameter m and b.

V. Experiments

In Section V-A we describe the commonly used measures to evaluate the quality of a segmentation method. In addition, we introduce a new measure called *Distance Dice* D_θ. The goal of this new measure is to be more robust with respect to perturbations within a small distance θ along the lesion segmentation boundaries.

We initially evaluated our MS lesion segmentation method on the MS grand challenge clinical dataset [32] (cf. Section V-B). The challenge was proposed in 2008 and the provided images have a different signal to noise ratio and contrast to noise ratio compared to images created with more modern hardware and acquisition protocols. Therefore, we additionally tested the methods on an in-house database (cf. Section V-C). Each method that we compared against performed better on the in-house dataset than on the MICCAI dataset (compare Tables I and II).

For all the experiments we used T_1-w and FLAIR images of each patient. We empirically set $\lambda_{\mathrm{CSF,GM}} = 1$, $\lambda_{\mathrm{GM,WM}} = 2$, $\lambda_{\mathrm{WM,SGM}} = 10$ and $\lambda_{\mathrm{SGM,V}} = 1$. The parameters of Equation (7) are empirically set to $\epsilon = 10^{-2}$, and $\sigma = \left(-\frac{1}{2}\frac{\mathrm{mg}(I)^{1.5}}{\log(\epsilon)}\right)^{\frac{1}{1.5}}$. The subsampling factors for the MICCAI database have been set to $s^{(1)}_{x_1} = s^{(1)}_{x_2} = s^{(1)}_{x_3} = s^{(2)}_{x_1} = s^{(2)}_{x_2} = s^{(2)}_{x_3} = 2$. For the in-house database we set $s^{(1)}_{x_3} = 1$ and $s^{(1)}_{x_1} = s^{(1)}_{x_2} = s^{(2)}_{x_1} = s^{(2)}_{x_2} = s^{(2)}_{x_3} = 2$. Note that the there is no downsampling in the through-plane x_3 axis for I_1 since the in-house data has a lower resolution in the through-plane direction (Section V-C).

A. Measures

Validation measures often used in segmentation are: Volume Difference (VD), average symmetric Surface Distance (SD), True Positive Rate (TPR) and False Positive Rate (FPR). VD and SD are defined on the voxel-level, while TPR and FPR are defined on the lesion-level [32]. VD is defined as $\mathrm{VD} = \frac{|V_S - V_{GT}|}{V_{GT}}$ where V_S and V_{GT} denote the volume of the Segmentation (S) and of the Ground Truth (GT) annotation, respectively. The optimal value is zero. SD is defined as

$$SD = \frac{\sum_{x\in\partial\mathrm{GT}} d(x, \partial\mathrm{S}) + \sum_{y\in\partial\mathrm{S}} d(\partial\mathrm{GT}, y)}{|\partial\mathrm{S}| + |\partial\mathrm{GT}|}$$

where ∂S and ∂GT represent the surface voxels of S and GT respectively, and d represents the Euclidean distance. The optimal value for SD is zero. TPR is defined as $TPR = \frac{|S\cap GT|}{|GT|}$, where S and GT are the segmentation and manual annotation respectively; the best value is 100%. The $FPR = \frac{|S| - |S\cap GT|}{|S|}$ and the best value is zero.

Additionally, commonly used measures are the Precision $P = \frac{TP}{TP+FP}$ where TP=true positive, FP=false positive and FN=false negative, the Recall $R = \frac{TP}{TP+FN}$ and the Dice coefficient $D = \frac{2\cdot TP}{FP+FN+2\cdot TP}$. These measures are defined on the voxel level and the best value is 100%. Note that the true positive rate on the voxel level (TPRv) is equivalent to the recall R, and the FPRv is equivalent to $100 - P$ where P is the precision. Since the GT is often inaccurate at the boundary of a lesion, we introduce a distance-based Dice coefficient D_θ where θ defines a tolerance zone on the border (internal and external) of the GT. The voxels inside the tolerance zone are neither considered as FN nor as FP, resulting in $D_\theta \geq D$ (cf. Fig. 7). It is similar to the *slack borders* measure that was introduced in [33] to evaluate lung segmentations. This measure ignores the information around the boundary for TP as well as for FN and FP. Therefore, it may be smaller than the Dice coefficient. The distance Dice D_θ on the other hand always provided a higher or equal measure than the classical Dice coefficient.

B. MICCAI Dataset

The MS lesion challenge proposed in 2008 [32] consists of 20 patient data in the training set and 22 in the test set. The patients were scanned at the University of North Carolina (UNC) and at the Children's Hospital of Boston (CHB). While the acquisition protocol consists of T_1-w, T_2-w and FLAIR images, we only used the T_1-w and FLAIR

TABLE I
COMPARATIVE RESULTS ON THE MICCAI MS GRAND CHALLENGE TEST SET [32].

Method	Hospital	VD		SD		TPR		FPR		Score
		1 rat.	2 rat.	1 rat.	2 rat.	1 rat.	2 rat.	1 rat.	2 rat.	
Lesion-TOADS [5]	UNC	45.6	128.2	8.3	12.7	44.1	61.5	68.1	72.8	79.96
	CHB	75.3	57.6	6.5	5.3	52.2	52.2	79.2	69.1	
LST [6]	UNC	63.3	110.5	9.5	13.5	**48.5**	59.7	73.6	73.4	80.00
	CHB	50.4	69.3	5.9	4.4	50.0	**55.9**	78.8	65.8	
Geremia et al. [8]	UNC	**42.1**	56.7	8.0	10.7	43.5	**66.8**	74.7	82.8	82.07
	CHB	47.6	49.7	**4.2**	**2.0**	**56.1**	54.0	78.0	64.3	
MOPS [16]	UNC	29.4	**54.1**	9.9	11.3	36.1	62.3	46.2	54.7	84.46
	CHB	43.3	53.0	5.2	6.4	45.8	45.0	42.7	38.9	
RMNMS [10]	UNC	51.2	56.6	**5.7**	**7.3**	44.9	63.1	39.1	53.2	**86.11**
	CHB	**43.2**	**47.9**	5.4	4.4	48.4	46.0	46.3	34.7	
Proposed	UNC	77.7	215.1	11.5	13.1	30.8	49.7	**37.8**	**46.7**	83.93
	CHB	45.0	55.7	4.5	4.8	41.6	39.1	**32.8**	**21.2**	

images for the proposed method. All the images have an isotropic resolution of $0.5 \times 0.5 \times 0.5$ mm^3. The training set consists of 10 CHB and 10 UNC patients. For the CHB patient one manual annotation is provided while there are two annotations for the UNC patients. The test set is evaluated online with two manual annotations. We used the training set to estimate the hyper-parameters for the $\lambda_{\mathrm{L,WM}}$ and the test set to evaluate the performances of our method. We evaluate this data with respect to VD, SD, TPR and FPR. The website of the challenge also provides a score based on average values of the four measures. We compared our method with Lesion-TOADS [5], LST [6], the method proposed by Geremia et al. [8], the MOPS method [16], and the RMNMS method [10]. The results are shown in Table I, the bold values represent the best value for each metric and each dataset. All the methods have comparable results, but our method provides the best performance with respect to the FPR, *i.e.*, provides the most conservative segmentation with respect to all other state of the art methods. In Fig. 8 we show visual results of the proposed method for two patients from the test set. CHB15 is presented in the first three columns and UNC03 in the last three.

C. In-House Database

The in-house database consists of data from 20 MS patients (mean age 37.2; age range 28-55 years). The database consists of images from two different datasets: the first one contains 5 patient data, while the second one has 15 patient data. The lesion load of the first dataset is between 1.01-9.8 cm^3, while for the second dataset is between 0.9-24 cm^3. We used the first dataset for the hyper-parameter learning phase and the second one for the testing phase. The studies have been approved by the local ethic committee and all the patients provided informed consent for the analysis of the imaging data. For both datasets the brain images have been acquired with a 3T Siemens Trio MR System. The acquisition protocol for both datasets consists of a 3D MPRAGE T_1-w image and a 3D FLAIR image. The data from the first dataset has a spacing of $0.5 \times 0.5 \times 1.2$ mm^3 per voxel, the MPRAGE acquisition parameters are TE/TR/TI=2.15/1390/800 ms and the FLAIR acquisition parameters are TE/TR/TI=388/5000/1800 ms. The data from the second dataset has a spacing of $0.49 \times 0.49 \times 1$ mm^3 per voxel, the MPRAGE acquisition parameters are TE/TR/TI=2.98/1300/800 ms and the FLAIR acquisition parameters are TE/TR/TI=394/5000/1800 ms. The manual annotation was performed by an expert neuro-radiologist on the FLAIR images.

We performed the testing on the in-house dataset of 15 patients and compared it against Lesion-TOADS [5] and LST [6]. For both methods we fine-tuned the parameters to get the best results (cf. Fig. 9). In Table II we show the results for each of the presented measures, the bold values represent the best value for each metric. We observe that the proposed method performs best in terms of VD, TPR, R, D and $D_{0.5\mathrm{mm}}$. The LST method performs best for SD and P measures, while Lesion-TOADS has the best performances for the FPR. In general the D_θ scores improve with respect to D for all methods.

TABLE II
COMPARATIVE RESULTS ON THE IN-HOUSE DATASET OF 15 PATIENTS.

Method	VD	SD	TPR	FPR	P	R	D	$D_{0.5\mathrm{mm}}$
L-TOADS [5]	70.5	5.8	45.2	**48.3**	36.4	37.4	30.5	45.1
LST [6]	60.6	**2.6**	61.5	56.3	**73.2**	36.3	47.3	70.2
Proposed	**36.5**	2.8	**70.4**	56.2	62.4	**50.5**	**52.4**	**73.2**

D. Robustness of the Proposed Method

We tested the robustness of the method against increasing levels of noise. We artificially added Rician noise to both FLAIR and the T_1-w images on three representative patients with low, moderate and high lesion loads, respectively. The percentage of the added noise is based on the brightest tissue, *i.e.*, GM for the FLAIR and WM for the T_1-w. The tested noise levels are 1%, 3%, 5%, 7%, and 9%, and we compared it against Lesion-Toads and LST. For both methods, as well as for the proposed method, we used the optimal parameters estimated for these patient data without noise. In Fig. 11 the noise level versus the Dice coefficient D and $D_{0.5}$ are plotted for LST, Lesion-TOADS and the proposed method. The performances of the proposed method are minimally affected for noise levels $1\%-5\%$, while for 9% of the noise level D has a decrease of 20% in the patient data with low and moderate lesion loads. Note that both the D and $D_{0.5}$ values for 0% of noise are the same as for 1% for all the methods. We

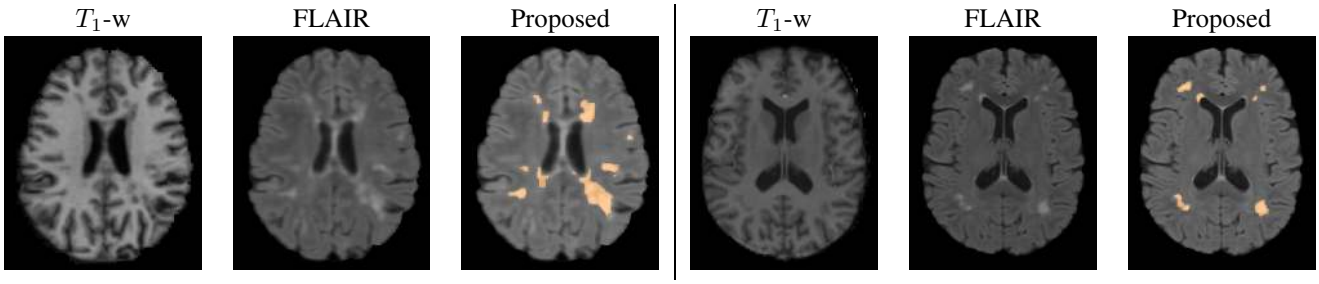

Fig. 8. **MICCAI dataset.** Using the T_1-w and FLAIR images, we obtain for the proposed method a lesion segmentation (shown in orange as overlay of the FLAIR image). The first three columns show the patient CHB15, and the others the patient UNC03 of the MICCAI MS challenge test dataset.

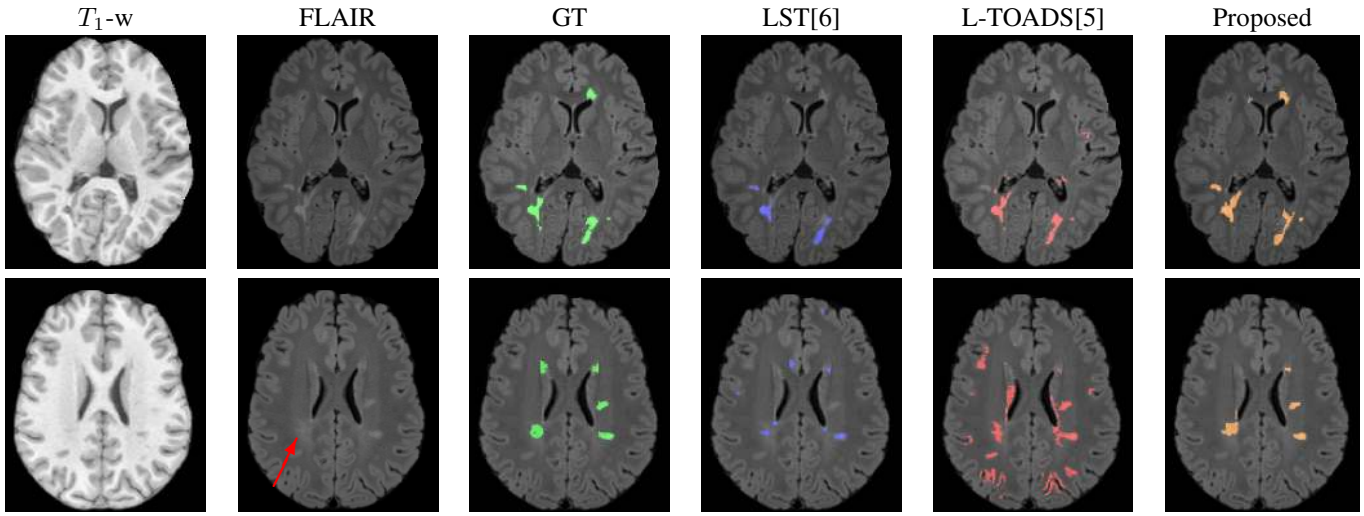

Fig. 9. **In-house dataset.** Using T_1-w and FLAIR images, we obtain different segmentations shown as overlay of the FLAIR image. The ground truth GT (green) is shown beside the results obtained with LST (blue), Lesion-TOADS (red) and the proposed method (orange). Each row presents a different patient and the red arrow, in the second row, points to a dirty-appearing WM region.

observe that with high lesion load all the methods seem not to be affected by the different noise levels. In Fig. 10 visual results with 1% and 9% noise levels for the moderate lesion load patient data are shown.

VI. Summary and Discussion

In this article we have presented a new method for 3D lesion segmentation without using an MRI-atlas. We have proposed a geometric model of the brain that is less sensitive to spatial variabilities of the brain's anatomy. Therefore, the effect of model-overfitting is less pronounced at the interface of neighboring brain regions. We have proposed to minimize an energy which models the appearance of the tissues and the topological structure of the brain like the ordering of the different tissue layers and the connectivity of the grey matter.

While we have used Gaussian mixture models to model the tissues appearance, it is straightforward to extend our method to more descriptive appearance models involving SVMs, Random Forests or convolutional neural networks. These models would change Equation (5) and take more information such as local textures into account.

One limitation of the proposed method is that the geometric model is anatomically not entirely correct, since the ventricles may also be directly adjacent to the WM and the lesions, which we did not model. However, there are signals from the ventricle that are not correctly suppressed. These voxels have a high intensity value in the FLAIR image and they can therefore be represented by the SGM-model. In practice, our model may create a thin artifact layer of subcortical GM around the ventricles. While the resulting labeling may not be anatomically correct, it slightly affects the lesion segmentation. Fig. 8 as well as the first row of Fig. 9 illustrate that the lesions around the ventricles can be segmented with the proposed method. Note that the method is expected to have similar results with lesions in the corpus callosum.

The proposed method is biased towards lesions appearing in the WM, however lesions which are appearing in the subcortical GM can still be segmented. In our database we did not find any lesion in the main deep GM structures (basial ganglia, thalamii). However, in order to satisfy the geometric constraints the method is supposed to create a "rim" of WM around these lesions. Additionally the proposed method may not entirely segment the dirty-appearing WM, which in FLAIR images typically appears as a smooth hyper-intense region in the periventricular WM, Fig. 9 (second row). Since the boundaries of these regions are hard to detect, they may be partially segmented as healthy WM. This limitation might be overcome by lowering the value of $\lambda_{\mathrm{L,WM}}$. Since the training

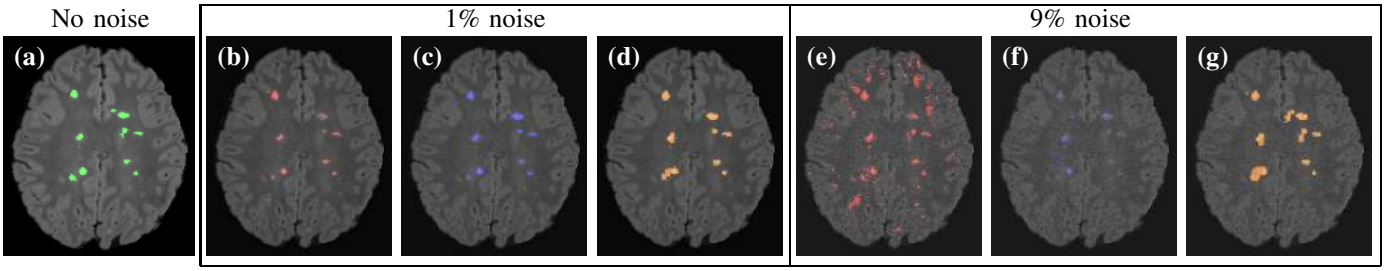

Fig. 10. **Segmentation with noise.** In (a) the original FLAIR image with GT. We show results with 1% of additive Rician noise (b)−(d), and with 9% (e)−(g). We compared with respect to Lesion-Toads (red, (b), (e)), LST (blue, (c), (f)), and the proposed method (orange, (d), (g)).

set did not contain extended regions of dirty-appearing-WM, the learning process could not take this effect into account.

The proposed connectivity constraint uses a distance function based on the intensities' gradient. It is more general and efficient compared to the formulations proposed in the literature [24], [25]. The connectivity prior is always imposed as a hard constraint, but our formulation is designed for multi-labeling segmentation while the one proposed in [25] is meant for binary segmentation. We have shown that [25] fails to enforce a connectivity prior in our multi-labeling setting, whereas our method properly connects the GM (cf. Fig. 4). Imposing this constraint, we were able to recover some lesions that might have been misclassified as GM otherwise.

The second global constraint that we have added is the minimum lesion size which is fixed to a diameter of 3 voxels for the high resolution step. As shown in Fig. 5, we were able to remove noisy signals in the WM, while extending small lesions to the pre-determined size.

The overall energy (3) is minimized following a coarse-to-fine approach. Based on [13], we have globally optimized our energy on the low resolution, while ignoring the minimal size constraint (9) for lesions. Delong and Boykov [28] showed that combining an inclusion constraint with an exclusion constraint results in an energy that is no longer submodular. Hence, we could not include the exclusion constraint between L and SGM in our framework. Instead, two *false labels*, F_1 and F_2, were modeled (cf. Section IV-A). With the α/β-swap algorithm [12] we have been able to repair the correct labels and to find a feasible solution of our energy. This feasible initialization is locally minimized with further α/β-swaps (cf. Section IV-B).

Our method is based on few parameters, the λ_{l_1,l_2} values for each neighboring tissue. These values are empirically estimated and kept constant for all our experiments with the exception of $\lambda_{\mathrm{WM,L}}$. In order to make this term adaptive to the contrast in the image, we have modeled it as being affinely dependent on the maximum magnitude of the image gradient (cf. Section IV-C). In the hyper-parameter learning step, for which we can guarantee global optimality, we have estimated the affine parameters m and b in order to obtain the best result with respect to the training set.

The proposed method uses FLAIR and T_1-w images. The T_1-w image provides more information about the contrast between different tissue types, while the FLAIR image is more sensitive towards lesion intensities. Additional image contrasts can be easily added to our framework. It will only affect the computation of the GMMs, which can still be computed efficiently for multiple image contrasts. The size of the used graphs on the other hand is virtually not affected by the amount of contrasts that are available. The run time of the proposed method is ca. 3 hours for the in-house database and ca. 7 hours for the MICCAI data.

The lesion load of the in-house database (0.9 cm^3 – 24 cm^3) showed that the proposed method is robust to mild, moderate and severe lesion loads. Additionally, the robustness of the method against noise has been tested for three representative patients with low, moderate and high lesion loads, respectively. In moderate and low lesion loads the performances of the proposed method are slightly affected with low noise levels, while with 9% of noise the Dice coefficient D has a drop of 20%. We observed that most of the lesions are still detected, but the number of FP increases showing an over-segmentation, Fig. 10(g). A possible explanation of this behavior is that the λ values are estimated on data without artificially added noise, therefore they are not optimal for the noisy images.

To be more robust with respect to small changes along the borders of the segmentations, we have introduced a new metric D_θ (cf. Section V). We tested and compared the proposed method on the MS grand challenge [32] and on an in-house dataset. The in-house datasets have a different image quality and contrast compared to the challenge data, which may be due to more modern hardware and optimized protocols. This can be observed in the better performance of the proposed method, of LST [6] and Lesion-TOADS [5] for the in-house database. The results on the MS grand challenge (cf. Table I) showed that our method is competitive with state of the art methods and outperforms all methods with respect to the FPR. Therefore, we can conclude that our method provides the most cautious segmentation. The results on the in-house database (cf. Table II) showed that the proposed method has the best scores in most of the metrics and that in general the D_θ scores improved with respect to D for all methods.

As future work we will use the proposed method to assess the lesions progression in longitudinal studies.

REFERENCES

[1] B. Weinshenker, B. Bass, G. Rice, J. Noseworthy, W. Carriere, J. Baskerville, and G. Ebers, "The natural history of multiple sclerosis: a geographically based study I. clinical course and disability," *Brain*, vol. 112, no. 1, pp. 133–146, 1989.

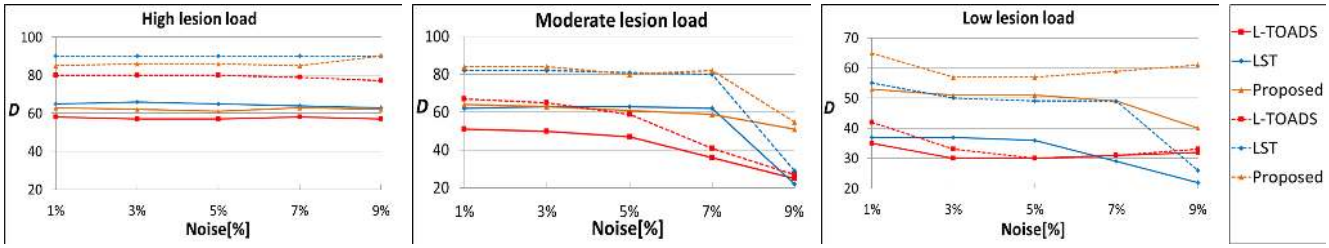

Fig. 11. **Noise robustness.** D (solid lines) and $D_{0.5}$ (dashed lines) values for increasing levels of noise. Lesion-TOADS (red), LST (blue), and the proposed method (orange) for three representative patients with high (left), moderate (center), and low (right) lesion loads.

[2] D. García-Lorenzo, S. Francis, S. Narayanan, D. L. Arnold, and D. L. Collins, "Review of automatic segmentation methods of multiple sclerosis white matter lesions on conventional magnetic resonance imaging," *Medical Image Analysis*, vol. 17, no. 1, pp. 1–18, 2013.

[3] X. Lladó, A. Oliver, M. Cabezas, J. Freixenet, J. C. Vilanova, A. Quiles, L. Valls, L. Ramió-Torrentà, and À. Rovira, "Segmentation of multiple sclerosis lesions in brain MRI: a review of automated approaches," *Information Sciences*, vol. 186, no. 1, pp. 164–185, 2012.

[4] J.-C. Souplet, C. Lebrun, N. Ayache, and G. Malandain, "An automatic segmentation of T2-FLAIR multiple sclerosis lesions," in *MIDAS Journal*, 2008.

[5] N. Shiee, P.-L. Bazin, A. Ozturk, D. S. Reich, P. A. Calabresi, and D. L. Pham, "A topology-preserving approach to the segmentation of brain images with multiple sclerosis lesions," *Neuroimage*, vol. 49, no. 2, pp. 1524–1535, 2010.

[6] P. Schmidt, C. Gaser, M. Arsic, D. Buck, A. Förschler, A. Berthele, M. Hoshi, R. Ilg, V. J. Schmid, C. Zimmer, B. Hemmer, and M. Mühlau, "An automated tool for detection of FLAIR-hyperintense white-matter lesions in MS," *Neuroimage*, vol. 59, no. 4, pp. 3774–3783, 2012.

[7] 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.

[8] E. Geremia, O. Clatz, B. H. Menze, E. Konukoglu, A. Criminisi, and N. Ayache, "Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images," *Neuroimage*, vol. 57, no. 2, pp. 378–390, 2011.

[9] N. Weiss, D. Rueckert, and A. Rao, "Multiple sclerosis lesion segmentation using dictionary learning and sparse coding," *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 735–742, 2013.

[10] N. Guizard, P. Coupé, V. S. Fonov, J. V. Manjón, D. L. Arnold, and D. L. Collins, "Rotation-invariant multi-contrast non-local means for MS lesion segmentation," *NeuroImage: Clinical*, vol. 8, pp. 376–389, 2015.

[11] C. Li, C. Gatenby, L. Wang, and J. C. Gore, "A robust parametric method for bias field estimation and segmentation of MR images," *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 218–223, 2009.

[12] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 23, no. 11, pp. 1222–1239, 2001.

[13] H. Ishikawa, "Exact optimization for Markov random fields with convex priors," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 25, no. 10, pp. 1333–1336, 2003.

[14] M. Kamber, R. Shinghal, D. L. Collins, G. S. Francis, and A. C. Evans, "Model-based 3-D segmentation of multiple sclerosis lesions in magnetic resonance brain images," *IEEE Transactions on Medical Imaging (TMI)*, vol. 14, no. 3, pp. 442–453, 1995.

[15] 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 (TMI)*, vol. 20, no. 8, pp. 677–688, 2001.

[16] X. Tomas-Fernandez and S. K. Warfield, "A model of population and subject (MOPS) intensities with application to multiple sclerosis lesion segmentation," *IEEE Transactions on Medical Imaging (TMI)*, vol. 34, no. 6, pp. 1349–1361, 2015.

[17] 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.

[18] M. Cabezas, A. Oliver, S. Valverde, B. Beltran, J. Freixenet, J. C. Vilanova, L. Ramió-Torrentà, À. Rovira, and X. Lladó, "BOOST: A supervised approach for multiple sclerosis lesion segmentation," *Journal of neuroscience methods*, vol. 237, pp. 108–117, 2014.

[19] J. Gao, C. Li, C. Feng, M. Xie, Y. Yin, and C. Davatzikos, "Non-locally regularized segmentation of multiple sclerosis lesion from multi-channel MRI data," *Magnetic Resonance Imaging*, 2014.

[20] D. García-Lorenzo, J. Lecoeur, D. L. Arnold, D. L. Collins, and C. Barillot, "Multiple sclerosis lesion segmentation using an automatic multimodal graph cuts," *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 584–591, 2009.

[21] L. Ibanez, W. Schroeder, L. Ng, J. Cates *et al.*, *The ITK software guide*. Kitware, 2003, vol. 8.

[22] S. Smith, "Fast robust automated brain extraction," *Human Brain Mapping*, vol. 17, no. 3, pp. 143–155, 2002.

[23] Y. Boykov and M.-P. Jolly, "Interactive graph cuts for optimal boundary & region segmentation of objects in ND images," in *Proc. of IEEE Int. Conf. on Computer Vision (ICCV)*, vol. 1. IEEE, 2001, pp. 105–112.

[24] S. Vicente, V. Kolmogorov, and C. Rother, "Graph cut based image segmentation with connectivity priors," in *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2008.

[25] J. Stühmer, P. Schröder, and D. Cremers, "Tree shape priors with connectivity constraints using convex relaxation on general graphs," *Proc. of IEEE Int. Conf. on Computer Vision (ICCV)*, pp. 2336–2343, 2013.

[26] Y. Boykov and O. Veksler, "Graph cuts in vision and graphics: Theories and applications," *Handbook of mathematical models in computer vision*, pp. 79–96, 2006.

[27] P. Kohli and P. H. Torr, "Efficiently solving dynamic markov random fields using graph cuts," *Proc. of IEEE Int. Conf. on Computer Vision (ICCV)*, vol. 2, pp. 922–929, 2005.

[28] A. Delong and Y. Boykov, "Globally optimal segmentation of multi-region objects," *Proc. of IEEE Int. Conf. on Computer Vision (ICCV)*, pp. 285–292, 2009.

[29] J. Munkres, "Algorithms for the assignment and transportation problems," *SIAM Journal on Imaging Sciences*, vol. 5, no. 1, pp. 32–38, 1957.

[30] R. R. Curtin, J. R. Cline, N. P. Slagle, W. B. March, P. Ram, N. A. Mehta, and A. G. Gray, "MLPACK: A scalable C++ machine learning library," *Journal of Machine Learning Research*, vol. 14, pp. 801–805, 2013.

[31] V. Kolmogorov, Y. Boykov, and C. Rother, "Applications of parametric maxflow in computer vision," *Proc. of IEEE Int. Conf. on Computer Vision (ICCV)*, pp. 1–8, 2007.

[32] M. Styner, J. Lee, B. Chin, M. Chin, O. Commowick, H. Tran, S. Markovic-Plese, V. Jewells, and S. Warfield, "3D segmentation in the clinic: A grand challenge II: MS lesion segmentation," *In MIDAS Journal*, vol. 2008, pp. 1–6, 2008.

[33] E. van Rikxoort and B. van Ginneken, "Automatic segmentation of the lungs and lobes from thoracic CT scans," *Fourth international workshop on pulmonary image analysis*, pp. 261–268, 2011.