

Multi Channel MRI Segmentation With Graph Cuts Using Spectral Gradient And Multidimensional Gaussian Mixture Model

Jérémy Lecoer^{*†‡}

Jean-Christophe Ferré^{§*†‡}

D. Louis Collins[¶]

Sean Morrissey^{||*†‡}

Christian Barillot^{*†‡}

ABSTRACT

A new segmentation framework is presented taking advantage of multimodal image signature of the different brain tissues (healthy and/or pathological). This is achieved by merging three different modalities of gray-level MRI sequences into a single RGB-like MRI, hence creating a unique 3-dimensional signature for each tissue by utilising the complementary information of each MRI sequence.

Using the scale-space spectral gradient operator, we can obtain a spatial gradient robust to intensity inhomogeneity. Even though it is based on psycho-visual color theory, it can be very efficiently applied to the RGB colored images. More over, it is not influenced by the channel assignment of each MRI.

Its optimisation by the graph cuts paradigm provides a powerful and accurate tool to segment either healthy or pathological tissues in a short time (average time about ninety seconds for a brain-tissues classification).

As it is a semi-automatic method, we run experiments to quantify the amount of seeds needed to perform a correct segmentation (dice similarity score above 0.85). Depending on the different sets of MRI sequences used, this amount of seeds (expressed as a relative number in pourcentage of the number of voxels of the ground truth) is between 6 to 16%.

We tested this algorithm on brainweb for validation purpose (healthy tissue classification and MS lesions segmentation) and also on clinical data for tumours and MS lesions detection and tissues classification.

1. INTRODUCTION

Taking advantage of the various protocols that acquire multiple modality images is a current issue (typically T1, T2, PD, DTI or Flair sequences in MR neuroimaging). The data are becoming more and more multi-channel data and their unique and complementary information should be merged together before segmentation to get rid of the inconsistencies one can encounter when segmenting each modality separately. Today, reliable registration methods, using different resolution and time, are available, nevertheless, a simple, robust, fast and reliable segmentation approach still does not exist for such kind of problem.

Multichannel segmentation usually relies on clustering or classification. In the current work, we propose a new and original scale-space approach for segmenting organs and tissues from multidimensional images. We propose a technique that can perform a joint segmentation of three MRI volumes at a time. The aim of this

* INRIA, VisAGeS Unit/Project, IRISA, Rennes, France

† University of Rennes I, CNRS IRISA, Rennes, France

‡ INSERM, U746 Unit/Project, IRISA Rennes, France

§ Department of Neuroradiology, Pontchaillou University Hospital, Rennes, France

¶ Montreal Neurological Institute, McGill University, Montreal, Canada

|| Departement of Neurology, University Hospital Pontchaillou, Rennes, France

technique is to be able to quantify local and/or global variations that are useful indicators of diseases states and evolution.

As the intensity distribution of the interesting tissues follows a Gaussian law in each modality, by merging three volumes into a single “color” MRI - each volume becoming a color channel - the color distribution, thus created, follows also a multidimensional Gaussian law. Each tissue being characterized by a 3-dimensional signature, discriminating each tissue from one another is easier. The main idea, presented here, to segment a volume is to use a scale-space color invariant edge detector - *i.e.* the spectral gradient - in a graph cut framework.

Among all the various energy minimization techniques for segmentation, Greig et al.¹ proposed a method based on partitioning a graph by a minimum cut / maximum flow algorithm inherited from Ford.² Then, Boykov et al.^{3,4} enhanced this approach by improving its computational efficiency, their method is now referred as Graph Cut.

In the following sections, we will first present the spectral gradient and how it can be embedded in graph cuts optimization. Then, we will show validation and results on synthetic and clinical data such as Multiple Sclerosis lesions and brain tumors. Finally, we will discuss on the contributions and future improvements to be made.

2. METHODS

The framework we’ve designed is as follow : from three grey-level MRI sequences, we build a color MRI by assigning each red, green or blue channel to a sequence. Then we compute the spectral gradient and use it in a graph cut framework which requires seeds as input. In the end of this framework, we obtain the segmented structures (*e.g.* brain, MS lesions, tumors). Figure 1 summarizes this framework.

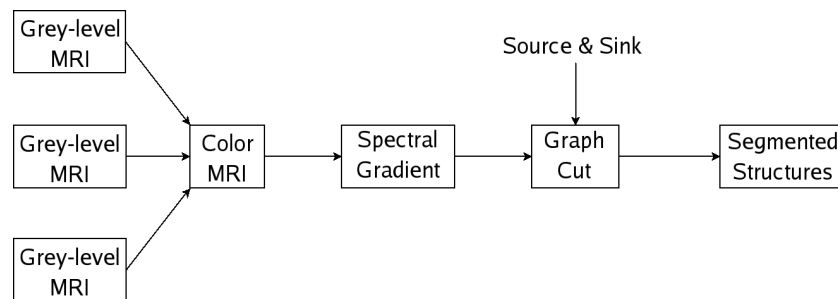


Figure 1. Our framework

As it can be seen on the figure above, the graph cut takes two different inputs : the Spectral Gradient - a boundary term - and the source and sink, a regional term which constitutes the interactive part of the algorithm. In the following sub-sections, we will explain the general framework of the Graph Cut, the mathematical observations used to compute the segmentation and, finally, how we’ve combined them all to achieve our goal.

2.1 Introduction to Graph Cut

According to the scheme by Boykov et al.,^{5,6} the segmentation problematic is described by a directional flow graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ which represents the image. The node set is defined by two particular nodes called terminal nodes - also known as “source” and “sink” - which respectively represents the class “object” and “background”. The other nodes correspond to the 3D volume voxels and the directed weighted edges connecting the nodes somehow encode the similarity between the two considered voxels. This method is semi-automatic as the user need to select two sets of voxels of the image, the object set v_o containing voxels of the object and the background set v_b containin voxels of the background.

Let the set \mathcal{P} contain all the voxels p of the image, the set \mathcal{N} be all the pair $\{p, q\}$ of the neighboring elements of \mathcal{P} and $\mathcal{V} = (V_1, V_2, \dots, V_{|\mathcal{P}|})$ be a binary vector where each V_p can be one of the two labels “object”

or “background”. Therefore, the vector \mathcal{V} defines a segmentation. The energy we want to minimize by the graph cut has the form given by :

$$E(\mathcal{V}) = \alpha \cdot \sum_{p \in \mathcal{P}} R_p(V_p) + \sum_{\substack{\{p,q\} \in \mathcal{N} \\ V_p \neq V_q}} B_{\{p,q\}} \quad (1)$$

The term $R_p(\cdot)$, commonly referred as the regional term, expresses how the voxel p fits into given models of the object and background. The term $B_{\{p,q\}}$, known as the boundary term, reflects the similarity of the voxels p and q . Hence, it is large when p and q are similar and close to zero when they are very different. The coefficient α is used to adjust the importance of the region and boundary terms.

The edge connecting a pair of neighboring voxel is called n -link and its cost can be based on various metrics such as local intensity gradient, Laplacian zero-crossing, gradient direction or other criteria with the restriction that this cost cannot be negative. It's the border term $B_{\{p,q\}}$ of equation 1. The simplest implementation of the weight cost of an edge between neighboring voxels p and q is then $w_{\{p,q\}} = K \cdot \exp(-\frac{|I_p - I_q|}{\sigma^2})$ where I_p and I_q are the intensities at voxel p and q . This weight function forces the segmentation boundaries at places with high intensity gradient.

All the v_o nodes are connected to the source and all background seeds (v_b) are connected to the sink, those two sets of links are called t -link (*terminal* links). Those t -links are the regional term $R_p(V_p)$ of equation 1. The simplest implementation involves infinite cost to all t -links between the seeds and the terminals. In more advanced work, the cost is based on how the intensity of voxel p fits into given intensity models (*e.g* histograms) of the object and background, hence giving a piece of regional information.

The graph is now completely defined and the segmentation contour is drawn by finding the minimum cut of this graph. A cut in a graph is a subset of edge that divides the graph into two parts : the source- and the sink-part, hence separating the object from the background in a binary segmentation ; the cost of a cut being the sum of the cost of its edges. The minimum cut is thus the cut with the minimum cost and can be computed in polynomial time using max-flow algorithm,² push-relabel technic⁷ or the now classical Boykov-Kolmogorov method.⁸

2.2 Image Observation : Maths At Work

As explained above, we need two kinds of information to complete our task : some on the region which we obtain from multivariate mixture model and some on the border which is provided by the spectral gradient. We will present in the following parts those two concepts.

2.2.1 Multivariate Gaussian Mixture Model : The Basics

To exploit the complementary piece of information contained in three registered MRIs from different sequences, we use a multivariate Gaussian mixture model. For each point labeled as a particular class c , we consider a 3-components vector Ψ composed with the intensities of each MRI at this point. We then compute the mean vector $\bar{\Psi}_c$ and the covariance matrix Σ_c . The class-membership probability of a voxel v is computed with the multivariate normal distribution formula :

$$P(\Psi_v|c) = \exp -\frac{1}{2}(\Psi_v - \bar{\Psi}_c)^T \cdot \Sigma_c^{-1} \cdot (\Psi_v - \bar{\Psi}_c) \quad (2)$$

As we consider a two (or three) classes object (depending on what we want to achieve), the seeds of each class are differentiated from the beginning and the source-membership probability of a voxel v is then the highest of the two (or three) class-membership probabilities. As the sink is always a single class, the sink-membership probability is equal to the considered class-membership probability.

2.2.2 The “How to...” Of Spectral Gradient

In order to use the color of tissues as an highly discriminative decision criterion, we need to build an invariant color-edge detector. We propose to use the spectral gradient, first introduced by Geusebroek et al.,⁹ which is based on the psycho-visual color theory and on Koenderink’s Gaussian derivative color model.¹⁰

Color can be interpreted by spectral intensity (e) that falls onto the retina ; this intensity depends on the spectral reflection function r of the surface material but also on the light spectrum l - which is a function of the wavelength - falling onto it. In addition, the shading s takes a great part in this seen intensity whereas it is only position-dependent. To sum it up, we can write this equation :

$$e(x, y, z, \lambda) = r(x, y, z, \lambda) \cdot l(\lambda) \cdot s(x, y, z) \quad (3)$$

As we seek colors invariants, we need to get rid of $l(\cdot)$ and $s(\cdot)$ since r is the only “true” color, which does not depend on illumination conditions. This can be achieved quite easily followings these steps : first, we take the derivative with respect to λ and we normalize the expression, thus leading to :

$$\frac{1}{e(x, y, z, \lambda)} \frac{\partial e(x, y, z, \lambda)}{\partial \lambda} = \frac{l_\lambda}{l} + \frac{r_\lambda}{r} \quad (4)$$

Then, a simple differentiation to the spatial variable (x , y or z) makes an expression which suits all of our constraints :

$$\frac{\partial \left(\frac{1}{e(x, y, z, \lambda)} \frac{\partial e(x, y, z, \lambda)}{\partial \lambda} \right)}{\partial x} = 0 \iff \frac{e \cdot e_{x\lambda} - e_x \cdot e_\lambda}{e^2} = 0 \quad (5)$$

Hence, we now have an expression which is fully expressed by spatial and spectral derivative of e , which is the observable spatio-spectral intensity distribution

Geusebroek et al.¹¹ have proven that these terms can be very well approximated by simply multiplying the RGB values (seen as a column vector) by two matrices :

$$\begin{pmatrix} e \\ e_\lambda \\ e_{\lambda\lambda} \end{pmatrix} = \underbrace{\begin{pmatrix} -0.019 & 0.048 & 0.011 \\ 0.019 & 0 & -0.016 \\ 0.047 & -0.052 & 0 \end{pmatrix}}_{XYZ \text{ to } e} \cdot \underbrace{\begin{pmatrix} 0.621 & 0.133 & 0.194 \\ 0.297 & 0.563 & 0.049 \\ -0.009 & 0.027 & 1.105 \end{pmatrix}}_{RGB \text{ to } XYZ} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix} \quad (6)$$

The first matrix transforms the RGB values to the CIE 1964 XYZ basis for colorimetry and the second one gives the best linear transform from the XYZ values to the Gaussian color model. These two matrices can be merged in a 3×3 matrix \mathcal{M} that characterises the transformation from RGB values to e and its derivatives.




Figure 2. From left to right : Color MRI (from T1, T2 and Flair sequences), spectral intensity e , first order derivative of spectral intensity e_λ and second order derivative of spectral intensity $e_{\lambda\lambda}$.

Once the spectral intensity and its derivatives are computed, we can use the following differential properties of the invariant color-edge detector :

$$\varepsilon = \frac{1}{e} \cdot \frac{\partial e}{\partial \lambda} = \frac{e_\lambda}{e} \quad (7)$$

As stated in Geusebroek's work,¹¹ yellow-blue transitions can be found with the first order gradient, which magnitude is:

$$\Gamma = \sqrt{(\partial_x \varepsilon)^2 + (\partial_y \varepsilon)^2 + (\partial_z \varepsilon)^2} \quad (8)$$

The second order gradient detects the purple-green transitions. Its magnitude can be computed as follows :

$$\Upsilon = \sqrt{(\partial_{x,\lambda} \varepsilon)^2 + (\partial_{y,\lambda} \varepsilon)^2 + (\partial_{z,\lambda} \varepsilon)^2} = \sqrt{(\partial_x \varepsilon_\lambda)^2 + (\partial_y \varepsilon_\lambda)^2 + (\partial_z \varepsilon_\lambda)^2} \quad (9)$$

$$\text{with : } \varepsilon_\lambda = \frac{\partial \varepsilon}{\partial \lambda} = \frac{e \cdot e_{\lambda\lambda} - e_\lambda^2}{e^2} \quad (10)$$

Finally, the detection of all color edges can be performed with :

$$\begin{aligned} \aleph &= \sqrt{\Gamma^2 + \Upsilon^2} \\ &= \sqrt{(\partial_x \varepsilon)^2 + (\partial_y \varepsilon)^2 + (\partial_z \varepsilon)^2 + (\partial_x \varepsilon_\lambda)^2 + (\partial_y \varepsilon_\lambda)^2 + (\partial_z \varepsilon_\lambda)^2} \end{aligned} \quad (11)$$




Figure 3. Left : Color MRI - Right : Spectral gradient.

2.3 Introducing Multivariate Mixture Model And Spectral Gradient in the Graph Cut Paradigm

Let's get back to the energy formulation of the Graph Cut :

$$E(\mathcal{V}) = \alpha \cdot \sum_{p \in \mathcal{P}} R_p(V_p) + \sum_{\substack{\{p,q\} \in \mathcal{N} \\ V_p \neq V_q}} B_{\{p,q\}} \quad (12)$$

To form the graph, we need to compute the t -links. Two cases are to be considered : first, the weight W_{so} of the t -link involving the "source" node is :

$$W_{so} = \begin{cases} 0 & \text{if } p \in \mathcal{B} \\ \infty & \text{if } p \in \mathcal{O} \\ \alpha \cdot R_p(\mathcal{B}) & \text{elsewhere} \end{cases} \quad (13)$$

Similar, the weight W_{si} of the t -link involving the “sink” node is computed as follows :

$$W_{si} = \begin{cases} \infty & \text{if } p \in \mathcal{B} \\ 0 & \text{if } p \in \mathcal{O} \\ \alpha \cdot R_p(\mathcal{O}) & \text{elsewhere} \end{cases} \quad (14)$$

According to the scheme by Boykov et al.,^{5,6} the t -link weights of a voxel p are the negative log-likelihoods :

$$R_p(\mathcal{B}) = -\ln P(\Psi_p|\mathcal{B}) \quad \text{and} \quad R_p(\mathcal{O}) = -\ln P(\Psi_p|\mathcal{O}) \quad (15)$$

To compute the n -links, we use an *ad-hoc* function :

$$B_{\{p,q\}} \propto \exp\left(-\frac{(\varepsilon(p) - \varepsilon(q))^2 + (\varepsilon_\lambda(p) - \varepsilon_\lambda(q))^2}{2\sigma^2}\right) \cdot \frac{1}{\text{dist}(p,q)} \quad (16)$$

where ε and ε_λ are the quantities defined in equations 7 and 10. Changing the channels assignment only change the spectral gradient intensity but not the location of the border, so, in this graph cut scheme, the channel assignment does not impact the segmentation.

In order to deal with the quite huge data sets we have which are very resource consuming, we needed a hierarchical scheme that allows us to compute the graph cut in a decent time. Consequently, we choose to implement the graph cut as a multiresolution algorithm. This method, first developed by Lombaert et al.,¹² is inspired by the multilevel graph partition technique¹³ and the narrow band from the level sets.¹⁴ It first computes the graph cut on the coarsest level and then in the successive higher level but only on a narrow band derived from the minimal cut found at the previous coarser level. This pyramidal approach with a Gaussian decimation has proven to be robust, even with high downscaling factor.

3. EXPERIMENTS AND RESULTS

First, we will present the validation experiments run on synthetic data, thus permitting us a comparison with methods from the literature. Then, we will show the results obtained with real data with ground truth from an expert.

3.1 Validation on synthetic data

The validation of the accuracy of our segmentation method was performed on synthetic phantom by using the BrainWeb data.¹⁵ We built the color MRI from simulated T1, T2 and PD sequences. All the images belong to the same subject and are constituted of 217 slices of 181 x 181 isometric 1 mm voxels with 3% noise (relative to the brightest tissue in the images).

In order to classify correctly the different classes of tissues, we followed a hierarchical segmentation scheme. A first Graph Cut with source seeds mixing all the tissues and sink seeds on the background gives us the brain mask ; then inside this mask, we perform one graph cut per tissue with only its own source seeds - the sink seeds being the background seeds and the other tissues seeds. The whole computation time is between 50 to 80 seconds on a laptop (Dual core at 2.16 Ghz and 2 GB of RAM running Linux).

We tested our method on different non-uniformity levels and compared the obtained segmentation to the ground truth available, using the Dice Similarity Coefficient with five components being considered : cerebrospinal fluid, white matter, grey matter, the whole brain and the lateral ventricles. All the tests were performed on data with 3% of noise, relative to the brightest tissue.

The following table shows that the accuracy of the segmentation is not significantly influenced by bias as the DSC is quite similar for all non-uniformity values. More over, the segmentation is more or less equally accurate for little or thin structures (ventricles, CSF) or bigger ones (grey and white matter).




Figure 4. Validation on Brainweb - Left : our segmentation (DSC = 0.983) - Right : BrainWeb ground truth

	Non Uniformity		
	0%	20%	40%
Cerebro Spinal Fluid	0.892	0.891	0.892
Grey Matter	0.932	0.924	0.927
White matter	0.961	0.954	0.958
Whole brain	0.983	0.981	0.982
Ventricles	0.946	0.944	0.944

Those results are better or similar to those found in the literature on similar data, such as Ashburner et al.¹⁶ with DSC running from 0.932 for grey matter to 0.978 for the whole brain ; the LOCUS approach by Scherrer et al.¹⁷ scores 0.83 for CSF, 0.94 for grey matter and 0.96 for white matter, or Bricq et al.¹⁸ with a DSC around 0.96 for each class of tissue.

An important data in the evaluation of this tool is the amount of seeds needed to achieve a correct segmentation on pathological tissue. We ran experiments in order to quantify the similarity between the obtained segmentation and the ground truth. As input to our algorithm, we used the ground truth randomly decimated.

In order to assess this algorithm, we used synthetic data from BrainWeb¹⁵ with simulated multiple sclerosis lesions. Results of this test are presented on figure 5.




Figure 5. Dice Similarity Coefficient versus relative number of seeds (in percentage of the ground truth). Blue line : DSC obtained after running our algorithm. Black line : DSC from initialisation seeds only. Red Line : difference of the preceding two curves, showing the efficiency of the proposed method

As explained by Zijdenbos,¹⁹ a DSC score above 0.7 is generally considered as very good, especially when the segmented structures are small. Here, this threshold is reached when the input seeds are around 5% of the ground truth. The performance of our algorithm was compared to Van-Leemput's,²⁰ Freifeld's²¹ and Rousseau's

algorithm²² on moderate MS lesions. For the optimal value (10% of relative number of seed), our DSC is 0.93 when Van-Leemput scores 0.80 (calculated by Freifeld in²¹), Freifeld 0.77, Rousseau only 0.63 and Bricq 0.79.

3.2 Quantification of the accuracy on various sequences

Usually simulated data don't address the same issues than clinical data. That's why it is important to run experiments on this type of data which rarely allows the use of automatic or generic approaches. As we aim to use this new tool in a clinical context, we evaluated the results on different clinical data sets. Those sets are the ones that are most likely to be used for diagnosis purpose. Three combinations of sequences were considered relevant with our goal : T1-weighted, T2-weighted and Flair (which will be referred as TTF), T1-w, T2-w and PD (which will be referred as TTP) and T1-w, T1-w injected with Gadolinium and Flair (which will be referred as TGF). The data covers a large range of lesion load and clinical grades (from RR to SP).

The following table summarizes the different parameters of the color MRIs we built in order to assess our tool:

Type of MRI	Number of Subjects	Number of slices	Size of slices	Size of voxels
TTF	6	138	256 × 256	isometric 1 mm
TTP	8	217	181 × 181	isometric 1 mm
TGF	5	160	256 × 256	isometric 1 mm

We run similar experiments than those for BrainWeb validation with an addition : we also used a ground truth decimated by successive erosions as input, hence simulating the classical behavior of the users which would preferably take seeds in the center of the bigger lesions and forget smaller lesions.




Figure 6. Dice Similarity Coefficient versus relative number of seeds (in percentage of the ground truth). Red lines : TTF MRIs, green lines: TTP MRIs, blue lines : TGF MRIs, black line : DSC from input data alone ; solid line : randomly decimated ground truth as input, dash dotted line : ground truth decimated by successive erosions as input.

Figure 6 presents the results of this experiment. One can notice that the behavior of the curves are quite similar, however, the TTF sets seem to be more suitable for segmenting MS lesions, with an average 0.972 DSC score for a number of seeds around 6% of the ground truth; TTP being second best with a 0.777 average DSC score for the same relative number of seeds and TGF only scoring 0.593.

The other important fact showed by this experiment is that a random decimation performs better than the successive erosions as the DSC is lower by about 10% when the erosions are used. We can interpret that from two hypothesis : this may come from keeping small lesions in the initial stage in the random decimation (those small lesions being the first to disappear in the decimation method) or the Gaussian multivariate law computed




Figure 7. Dice Similarity Coefficient Enhancement versus relative number of seeds (in percentage of the ground truth). Red lines : TTF MRIs, green lines: TTP MRIs, blue lines : TGF MRIs ; solid line : randomly decimated ground truth as input, dash dotted line : ground truth decimated by successive erosions as input. The variation is shown as an error bar on the optimal point of each curve and is about $\pm 2\%$ for TTF, $\pm 1.5\%$ for TTP and $\pm 1\%$ for TGF.

from the seeds has a lower variance than the real one as we only keep the centers of lesions whose intensities are not fully representative of the normal intensity distribution of all lesions.

Figure 7 shows the difference between the DSC score of the initial portion of the ground truth retained from initialisation alone and the DSC score obtained after running our algorithm from the same initialisation. It somehow computes the enhancement given by our tool with respect to the DSC. From top to bottom of the curves in figure 7, the maximum enhancement is obtained respectively for 6, 8, 11, 14, 16 and 18% of relative number of seeds. This gives our optimal range for initialisation constraints with respect to the expected lesion load.

Results and seeds of the algorithm are shown on figure 8 for TTF sequences.




Figure 8. Results on TTF images. Upper row from left to right : T1-w, T2-w and Flair Sequence. Lower left : initialisation from randomly decimated ground truth. Lower center : initialisation from ground truth decimated by erosion. Lower right : segmentation results (green : correctly classified ; blue : false positive ; yellow : false negative)

We then applied the spectral gradient Graph Cuts on MRIs of patients with tumors. The aim of these tests were to be able to segment accurately the peri-tumoral edema and the tumor itself. The MR sequences used for this purpose are T1 and T1 with Gadolinium injection images composed of 182 slices of 256×256 isometric 1 mm voxels and an interpolated FLAIR image (original size $704 \times 704 \times 34$). As in the preceding applications, the first pass gives us the brain mask, then we apply the graph cut for each class of interest (white matter, grey matter, edema and tumor). On the following figure, we can see that the tumor (in blue) and edema (in dark grey) regions are well segmented. This result was obtained in roughly 90 seconds. Visual assesment have been performed by an expert.




Figure 9. Our segmentation (right), T1 image (middle) and Flair image (left).

4. CONCLUSION AND FUTURE WORK

In this paper, we have introduced the spectral gradient in the field of 3D medical images. This new, fast and robust multiple images segmentation framework allows us to process multichannel data from scale-space derivatives.

Its optimization by a hierarchical graph cuts has proven to be accurate with very effective results. It combines region-based intensity processing and contour-based scale-space approach in a new way. With such little computational time, it allows the user to interactively correct the results from a very limited input effort.

We've built our work on an analogy between RGB channels and multi-modal MRIs. This method may not be optimal, especially the transform matrix \mathcal{M} (cf eq. 6). To find the optimal matrix, we could learn its parameters from a supervised procedure.

We propose here a semi-automatic but supervising the intensity model of lesions could give an automatic initialisation. Such models could be given by the STREM algorithm from Ait-Ali et al.²³

The next steps of this work will be to run more advanced validation studies and particularly on longitudinal data sets in order to be able, for instance, to quantify the evolution of multiple sclerosis lesions and to study the influence of the parameter of the \mathcal{M} matrix parameters.

REFERENCES

- [1] Greig, D., Porteous, B., and Seheult, A., "Exact maximum a posteriori estimation for binary images," *Journal of the Royal Statistical Society* **51**(2), 271–279 (1989).
- [2] Ford, L. and Fulkerson, D., "Maximal flow through a network," *Canadian Journal of Mathematics* **8**, 399–404 (1956).
- [3] Boykov, Y., Veksler, O., and Zabih, R., "Fast approximate energy minimization via graph cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(11), 1222–1239 (2001).
- [4] Boykov, Y. and Jolly, M.-P., "Interactive organ segmentation using graph cuts," in [*International Conference on Medical Image Computing and Computer-Assisted Intervention*], 276–286, Springer-Verlag (2000).
- [5] Boykov, Y. and Jolly, M.-P., "Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images," in [*International Conference on Computer Vision*], **1**, 105–112 (July 2001).
- [6] Boykov, Y. and Funka-Lea, G., "Graph cuts and efficient N-D images segmentation," *International Journal of Computer Vision* **70**, 109–131 (Nov. 2006).
- [7] Goldberg, A. and Tarjan, A., "A new approach to the maximum flow problem," *Journal of the association for computing Machinery* **35**, 921–940 (10 1988).
- [8] Boykov, Y. and Kolmogorov, V., "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26**, 1124–1137 (Sept. 2004).

- [9] Geusebroek, J.-M., Dev, A., van den Boomgaard, R., Smeulders, A., Cornelissen, F., and Geerts, H., "Color invariant edge detection," in [*Scale-Space Theories in Computer Vision*], *Lecture Notes in Computer Science* **1252**, 459–464 (1999).
- [10] Koenderink, J., [*Color Space*], Utrecht University (1998).
- [11] Geusebroek, J.-M., van den Boomgaard, R., Smeulders, A., and Dev, A., "Color and space : The spatial structure of color images," in [*IEEE European Conference on Computer Vision*], *Lecture Notes in Computer Science* **1842**, 331–341 (2000).
- [12] Lombaert, H., Sun, Y., Grady, L., and Xu, C., "A multilevel banded graph cuts method for fast image segmentation," in [*International Conference on Computer Vision*], **1**, 259–265 (Oct. 2005).
- [13] Karypis, G. and Kumar, V., "Multilevel k-way partitioning scheme for irregular graphs," *Journal of Parallel and Distributed Computing* **48**(1), 96–129 (1998).
- [14] Adalsteinsson, D. and Sethian, J.-A., "A fast level set method for propagating interfaces," *Journal of Computational Physics* **118**, 269–277 (1995).
- [15] Cocosco, C., Kwan, K., and Evans, R., "Brainweb: online interface to a 3-d mri simulated brain database," *Neuroimage* **5**(4), part 2/4, S425 (1997).
- [16] Ashburner, J. and Friston, K., "Unified segmentation," *Neuroimage* **26**, 839–851 (2005).
- [17] Scherrer, B., Dojat, M., F., F., and Garbay, C., "LOCUS: LOcal cooperative Unified Segmentation of MRI Brain Scans," in [*International Conference on Medical Image Computing and Computer-Assisted Intervention*], 219–227 (2007).
- [18] Bricq, S., *Segmentation d'images IRM anatomiques par inférence bayésienne multimodale et détection de lésions*, PhD thesis, Université Louis Pasteur de Strasbourg (2008).
- [19] Zijdenbos, A. P., Dawant, B. M., Margolin, R. A., and Palmer, A. C., "Morphometric analysis of white matter lesions in mr images: Method and validation," *IEEE Transactions on Medical Imaging* **13**(4), 716–724 (1994).
- [20] Van Leemput, K., Maes, F., Vandermeulen, D., Colchester, A., and Suetens, P., "Automated segmentation of multiple sclerosis lesions by model outlier detection," *IEEE Transactions on Medical Imaging* **20**(8), 677–688 (2001).
- [21] Freifeld, O., Greenspan, H., and Goldberger, J., "Lesion detection in noisy MR brain images using constrained GMM and active contours," in [*IEEE International Symposium on Biomedical Imaging: From Nano to Macro*], 596–599 (2007).
- [22] Rousseau, F., Blanc, F., de Sèze, J., Rumbac, L., and Armspach, J., "An a contratio approach for outliers segmentation : application to multiple sclerosis in MRI," in [*IEEE International Symposium on Biomedical Imaging: From Nano to Macro*], 9–12 (2008).
- [23] Aït-Ali, L., Prima, S., Hellier, P., Carsin, B., Edan, G., and Barillot, C., "STREM: a robust multidimensional parametric method to segment MS lesions in MRI," in [*International Conference on Medical Image Computing and Computer-Assisted Intervention*], *Lecture Notes in Computer Science* **3749**, 409–416, Springer Berlin / Heidelberg (Oct. 2005).