

Multiple Sclerosis Lesions Segmentation using Spectral Gradient and Graph Cuts

Jérémy Lecoœur^{1,2,3}, Sean Patrick Morissey^{1,2,3,4}, Jean-Christophe Ferré⁵,
Douglas L. Arnold⁶, D. Louis Collins⁶, and Christian Barillot^{1,2,3}

¹ INRIA, VisAGeS Unit/Project, IRISA, Rennes, France

² University of Rennes I, CNRS IRISA, Rennes, France

³ INSERM, U746 Unit/Project, IRISA Rennes, France

⁴ Département of Neurology, University Hospital Pontchaillou, Rennes, France

⁵ Département of Neuroradiology, University Hospital Pontchaillou, Rennes, France

⁶ Montreal Neurological Institute, McGill University, Montreal, Canada

Abstract. We present a new tool for segmenting multiple sclerosis lesions that can take advantage of the complementary modalities we usually use for this purpose. Based on the integration of multi channel information in a scale-space paradigm, its optimization by graph cuts provides a powerful and accurate tool. After presenting the mathematical and physical background on which is based the spectral gradient, we introduce its computational optimization by the graph cuts approach. The validation on synthetic and real data shows both its accuracy and reliability on different sets of standard MRI of sequences.

1 Introduction

Taking advantage of the various protocols that acquire images using multiple modalities 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 in order 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 especially when dealing with pathologies.

Multichannel segmentation usually relies on clustering or classification. In the current work, we propose a new and original scale-space approach for segmenting multiple sclerosis lesions from multidimensional Magnetic Resonance Images. We propose a technique that can perform a joint segmentation of three MRI volumes at a time with a supervised approach.

As the intensity distribution of the interesting tissues follows a Gaussian law in each sequence, 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 MS Lesions from multi-dimensionnal MRI sequences is to use a scale-space color invariant edge detector - *i.e.* the spectral gradient - in a graph cut optimisation 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 Cuts.

In the following sections, we will first present the spectral gradient and its graph cuts optimization that make it possible to exploit the intensity consistencies and the inter-intensities gradient. Then, we will show validation on various sets of MRI sequences for the segmentation of MS lesions. 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). Figure 1 summarizes this framework.

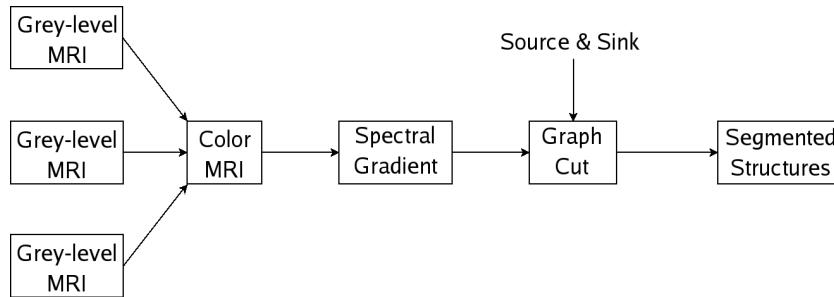


Fig. 1. Our framework

In the following sub-sections, we will explain the two main methodological steps of this framework which are the Spectral Gradient and the Graph Cuts.

2.1 Spectral Gradient

Based on Koenderink's Gaussian derivative color model [5] and the psycho-visual color theory, we propose to use a scale-space approach, first introduced by Geusebroek in [6] to build a color-edge detector. This color edge detector can then be applied to MRI by merging three gray-level MRIs into a single MRI.

The spectral intensity (e) that falls onto the retina depends on the spectral reflection function r of the surface material and the light spectrum l - a function of the wavelength λ - falling onto it. It also depends on the shading s that is only position-dependent. Hence :

$$e(x, y, z, \lambda) = r(x, y, z, \lambda) \cdot l(\lambda) \cdot s(x, y, z) \quad (1)$$

In order to define color invariants, we need an expression that is independent to $l(\cdot)$ and $s(\cdot)$, as r is the only “true” color we are searching for. By taking the derivative with respect to λ and normalizing, we get the following expression :

$$\frac{1}{e(x, y, z, \lambda)} \frac{\partial e(x, y, z, \lambda)}{\partial \lambda} = \frac{l_\lambda}{l} + \frac{r_\lambda}{r} \quad (2)$$

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

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

The last equation is fully expressed by spatial and spectral derivatives of e , the observable spatio-spectral intensity distribution.

Geusebroek et al. [7] have proven that these terms can be very well approximated by simply multiplying the RGB values (seen as a column vector) by two matrices :

$$\begin{pmatrix} e \\ e_\lambda \\ e_{\lambda\lambda} \end{pmatrix} = \underbrace{\begin{pmatrix} -0.019 & 0.048 & 0.011 \\ 0.019 & 0 & -0.016 \\ 0.047 & -0.052 & 0 \end{pmatrix}}_{XYZ \text{ to } e} \cdot \underbrace{\begin{pmatrix} 0.621 & 0.133 & 0.194 \\ 0.297 & 0.563 & 0.049 \\ -0.009 & 0.027 & 1.105 \end{pmatrix}}_{RGB \text{ to } XYZ} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix} \quad (3)$$

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.

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

$$\varepsilon = \frac{1}{e} \cdot \frac{\partial e}{\partial \lambda} = \frac{e_\lambda}{e} \quad (4)$$

As stated in [7], yellow-blue transitions can be found with the first order gradient, which magnitude is:

$$\Gamma = \sqrt{(\partial_x \varepsilon)^2 + (\partial_y \varepsilon)^2 + (\partial_z \varepsilon)^2} \quad (5)$$

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

$$\Upsilon = \sqrt{(\partial_{x,\lambda} \varepsilon)^2 + (\partial_{y,\lambda} \varepsilon)^2 + (\partial_{z,\lambda} \varepsilon)^2} = \sqrt{(\partial_x \varepsilon_\lambda)^2 + (\partial_y \varepsilon_\lambda)^2 + (\partial_z \varepsilon_\lambda)^2} \quad (6)$$

$$\text{with : } \varepsilon_\lambda = \frac{\partial \varepsilon}{\partial \lambda} = \frac{e \cdot e_{\lambda\lambda} - e_\lambda^2}{e^2} \quad (7)$$




Fig. 2. Upper row : Color MRI (from T1, T2 and Flair sequences) and spectral intensity e . Lower row : first order derivative of spectral intensity e_λ and second order derivative of spectral intensity $e_{\lambda\lambda}$.

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

$$\begin{aligned} \aleph &= \sqrt{\Gamma^2 + \Upsilon^2} \\ &= \sqrt{(\partial_x \varepsilon)^2 + (\partial_y \varepsilon)^2 + (\partial_z \varepsilon)^2 + (\partial_x \varepsilon_\lambda)^2 + (\partial_y \varepsilon_\lambda)^2 + (\partial_z \varepsilon_\lambda)^2} \end{aligned} \quad (8)$$

2.2 Graph Cuts Optimization of Spectral Gradients

In the Graph Cuts framework, the image is represented by a weighted graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ where each image voxel p is represented by a node and each edge links the voxel p to each of its neighboring voxels q . The binary graph cut associates each node to one of two special nodes, called the 'source' node \mathcal{S} and the 'sink' node \mathcal{T} . These two nodes (the terminal nodes) represent the labels (*i.e.* "object" or "background").




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

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 :

$$E(\mathcal{V}) = \alpha \cdot \sum_{p \in \mathcal{P}} R_p(V_p) + \sum_{\substack{\{p, q\} \in \mathcal{N} \\ V_p \neq V_q}} B_{\{p, q\}} \quad (9)$$

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, in [8] these models are the intensity distribution models (*e.g.* histograms) of the source and sink. 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. It is often based on local intensity gradient or Laplacian zero-crossing. The coefficient α is used to adjust the importance of the region and boundary terms.

In the construction of the graph, the regional term is used to compute the edge weight between each voxel and the terminal nodes (*t*-links) whereas the edge weight between neighboring pixels (*n*-links) is calculated using the boundary term.

The object (\mathcal{O}) and background (\mathcal{B}) seeds placed by the user at the beginning of the process are used to compute spectral intensity distribution models. In order to build this spectral intensity of the source (*resp.* sink), we consider a 3-components vector $\Psi = (e, e_\lambda, e_{\lambda\lambda})$ in each voxel labeled as source (*resp.* sink) and we compute the mean vector $\bar{\Psi}$ and the covariance matrix Σ . The probability for a voxel v to be in the source is computed with the multivariate normal distribution formula :

$$P(\Psi_v|\mathcal{O}) = \exp -\frac{1}{2}(\Psi_v - \bar{\Psi})^T \cdot \Sigma^{-1} \cdot (\Psi_v - \bar{\Psi}) \quad (10)$$

Thus, it allows us to compute the t -links and n -links as follows :

Edge	Case	Weight
$\{p, q\}$	$\{p, q\} \in \mathcal{N}$	$B_{\{p, q\}}$
	$p \in \mathcal{B}$	0
$\{p, \mathcal{S}\}$	$p \in \mathcal{O}$	∞
	$p \notin \mathcal{O} \cup \mathcal{B}$	$\alpha \cdot R_p(\mathcal{B})$
	$p \in \mathcal{B}$	∞
$\{p, \mathcal{T}\}$	$p \in \mathcal{O}$	0
	$p \notin \mathcal{O} \cup \mathcal{B}$	$\alpha \cdot R_p(\mathcal{O})$

The t -link weight of a voxel p is then the negative log-likelihoods:

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

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

$$B_{\{p, q\}} \propto \exp \left(-\frac{(\varepsilon(p) - \varepsilon(q))^2 + (\varepsilon_\lambda(p) - \varepsilon_\lambda(q))^2}{2\sigma^2} \right) \cdot \frac{1}{dist(p, q)} \quad (11)$$

where ε and ε_λ are the quantities defined in equations (4) and (7). The two terms in the exponential (11) refer respectively to the yellow/green edge local detector and the purple/green edge local detector described in the previous section.

3 Application to Multiple Sclerosis Lesions segmentation

In order to correctly classify the Multiple sclerosis lesions, 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 a graph cut with source seeds inside lesions - 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 for 3 MRI volumetric sequences). Since we can observe that changing the channels assignment doesn't change the segmentation results, no information about which sequence is used for which channel are necessary.

3.1 Validation on BrainWeb

An important data in the evaluation of this tool is the amount of seeds needed to correctly segment the lesions. We run 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. We then computed the *Dice Similarity Coefficient* (DSC) as follows :

$$S = 2 \cdot \frac{Card(R \cap V)}{Card(R) + Card(V)} \quad (12)$$

R being the segmentation result and V the ground truth.

In order to assess this algorithm, we used synthetic data from BrainWeb [9]. We built the color MRI from simulated T1-weighted, T2-weighted 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) and 20% non-uniformity field. Results of this test are presented on figure 4.




Fig. 4. Dice Similarity Coefficient versus relative number of seeds (in percentage of the ground truth). Blue line : DSC obtained after running our algorithm according to the number of seeds used at the initialisation step. Black line : DSC from initialisation seeds only. Red Line : difference of the preceding two curve, showing the efficiency of the proposed method

As stated in [10], 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 [11], Freifeld's [12] and Rousseau's algorithm [13] on moderate MS lesions. For the optimal value (8% of relative number of seed), our DSC is 0.83 when Van-Leemput scores 0.80 (calculated by Freifeld in [12]), Freifeld 0.77 and Rousseau only 0.63.

3.2 Quantification of the accuracy on various sequences

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.




Fig. 5. 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 without applying our algorithm ; solid line : randomly decimated ground truth as input, dash dotted line : ground truth decimated by successive erosions as input.

The figure 5 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 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.




Fig. 6. 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.

The figure 6 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 6, 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 figures 7, 8 and 9.




Fig. 7. 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)




Fig. 8. Results on TTP images. Upper row from left to right : T1-w, T2-w and PD Sequence. Lower left initialisation from randomly decimated ground truth. Lower center : initialisation from ground truth decimated by erosion. Lower right : segmentation results.




Fig. 9. Results on TGF images. Upper row from left to right : T1-w, Flair and T1-w injected with Gadolinium Sequence. Lower left initialisation from randomly decimated ground truth. Lower center : initialisation from ground truth decimated by erosion. Lower right : segmentation results.

4 Conclusion

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 intensity processing and scale-space approach in a new way. With such little computational time, it allows the user to interactively correct the results.

We've built our work on an analogy between RGB channels and multi-modal MRIs. This method may be not optimal, especially the transform matrix \mathcal{M} (*cf* eq. 4). 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. [14]

The next steps of this work will be to run more advanced validation studies and particularly on longitudinal data sets in order 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., Seheult, A.: Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society* **51**(2) (1989) 271–279
2. Ford, L., Fulkerson, D.: Maximal flow through a network. *Canadian Journal of Mathematics* **8** (1956) 399–404
3. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(11) (2001) 1222–1239
4. Boykov, Y., Jolly, M.P.: Interactive organ segmentation using graph cuts. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer-Verlag (2000) 276–286
5. Koenderink, J.: *Color Space*. Utrecht University (1998)
6. Geusebroek, J.M., Dev, A., van den Boomgaard, R., Smeulders, A., Cornelissen, F., Geerts, H.: Color invariant edge detection. In: *Scale-Space Theories in Computer Vision*. Volume 1252 of *Lecture Notes in Computer Science*. (1999) 459–464
7. Geusebroek, J.M., van den Boomgaard, R., Smeulders, A., Dev, A.: Color and space : The spatial structure of color images. In: *IEEE European Conference on Computer Vision*. Volume 1842 of *Lecture Notes in Computer Science*. (2000) 331–341
8. Boykov, Y., Jolly, M.P.: Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images. In: *International Conference on Computer Vision*. Volume 1. (July 2001) 105–112
9. Cocosco, C., Kwan, K., Evans, R.: Brainweb: online interface to a 3-d mri simulated brain database. *Neuroimage* **5**(4) (1997) part 2/4, S425
10. Zijdenbos, A.P., Dawant, B.M., Margolin, R.A., Palmer, A.C.: Morphometric analysis of white matter lesions in mr images: Method and validation. *IEEE Transactions on Medical Imaging* **13**(4) (1994) 716–724
11. Van Leemput, K., Maes, F., Vandermeulen, D., Colchester, A., Suetens, P.: Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE Transactions on Medical Imaging* **20**(8) (2001) 677–688
12. Freifeld, O., Greenspan, H., 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*. (2007) 596–599
13. Rousseau, F., Blanc, F., de Sèze, J., Rumbac, L., Armspach, J.: An a contrario approach for outliers segmentation : application to multiple sclerosis in MRI. In: *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. (2008) 9–12
14. Aït-Ali, L., Prima, S., Hellier, P., Carsin, B., Edan, G., 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*. Volume 3749 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg (October 2005) 409–416