

MS Lesion Segmentation using Markov Random Fields

N. K. Subbanna¹, M. Shah^{1,3}, S. J. Francis^{2,3}, S. Narayanan², D. L. Collins²,
D. L. Arnold^{2,3}, and T. Arbel¹

1 - Centre for Intelligent Machines, McGill University, Montreal, Quebec, Canada

2 - Montreal Neurological Institute, McGill University, Montreal, Quebec, Canada

3 - NeuroRX Research, Montreal, Quebec, Canada

Abstract. We present a fully automated framework for identifying multiple sclerosis (MS) lesions from multispectral human brain magnetic resonance images (MRIs). The brain tissue intensities and lesions are both modeled using Markov Random Fields (MRFs) to incorporate local spatial variations and neighborhood information. In this work, we model all brain tissues, including lesions, as separate classes as opposed to the common approach of modelling the lesions as outliers of the brain tissues. A maximum probability estimate is obtained by arriving at the global convergence of the MRFs using Simulated Annealing. Finally, probability surface discontinuities due to noise and local intensity variations are avoided by incorporating a spline based smoothing function following the MRF modelling. The algorithm is validated on a set of real MRI brain volumes of MS patients with widely varying lesion loads by comparing the results against a silver standard derived from manual expert labellings. The algorithm yields favorable results, including in the posterior fossa where few methods have proved successful. Further, our algorithm yields fewer false negatives than is usual in practice.

Keywords: MS lesion segmentation, MRFs, simulated annealing.

1 Introduction

Multiple sclerosis (MS) is a de-myelinating disease of the central nervous system (CNS) where myelin, the insulation around a nerve fibre (axon), is attacked resulting in focal MS lesions. MS is characterized by episodic, immunological attacks on the CNS resulting in various neurological impairments. MRIs are critical to the process of both studying the disease and developing therapies because it permits visualising lesion formation with a very high sensitivity. Consequently, the ability to detect and monitor lesion volumes and activity in brain MRIs is very crucial to the process of measuring the burden of the disease. Automatic lesion classification can lead to consistency in lesion detection, the elimination of the subjectivity (or intra-rater variability), (inter-rater) variability, and the reduction in the cost associated with manual labelling.

Over the years, various approaches have been proposed to perform automated lesion detection from the brain MRIs with pathology, showing varied degrees of

success. As categorised by Harmouche [13], the principal techniques are based on fuzzy reasoning [2], geometric models to determine lesion boundaries [3], connected component analysis and deterministic annealing [4], artificial neural networks [1], clustering techniques [6], atlas based methods [10] and contouring approaches [11]. A common approach has been to model lesions implicitly as outliers rather than explicitly including them in the tissue intensity model. One of the more successful examples of detecting MS lesions using outlier detection was proposed by van Leemput et. al. [5]. Other semi-automatic or interactive methods have also been suggested. For instance, Lecoer et. al., [12] propose a method based on spectral gradients and graph cuts to detect lesions. A common observation made by most of these approaches is the importance of incorporating local contextual tissue intensity information as very significant in lesion identification. MRFs have also been used to model the neighborhood information before. However, even when MRFs were used, they did not model the lesions as classes explicitly. Some probabilistic approaches such as the one proposed by Harmouche [13] have attempted to obtain a Bayesian model of the MS lesions for different regions of the brain. However, they utilise a Maximum Likelihood estimation using neighborhood information only *after* a labelling solution is obtained. There are approaches that use expert heuristics to incorporate contextual information (e.g. Francis [6]). Finally, very few methods have shown promise in identifying the lesions in the posterior fossa where lesion intensities approach that of the grey matter (GM) and the contrast between healthy tissue and the lesions is greatly reduced.

In this paper, we aim at constructing an MRF model of multispectral MRI (T1-weighted, T2-weighted and the proton density (PD) weighted image modalities) tissue intensities for the whole brain including the posterior fossa (PF) by incorporating voxel level spatial information in a standard anatomical space. In contrast to the previous approaches, we explicitly model all the tissues of the brain, including the lesions as classes. Further, simulated annealing is used to obtain an optimal labelling solution. Although simulated annealing requires a higher computational time, it finds the mode of the distribution, so that samples from the mode are from the global optimum. We then validate our results using both the κ coefficient and the false positive and false negative rates. In particular, we obtain a low rate of false negatives. Our results are compared against the published values of Shiee et. al [17] and van Leemput et. al. [5]. The technique we suggest performs better than the others reviewed here on the data sets we worked on.

Our principal contributions are constructing a MRF model of the multispectral MRI tissue intensities for the whole brain, including the posterior fossa. We explicitly model the lesions also as tissue classes, in contrast to most approaches which consider them outliers of the healthy brain tissues. Further, we use gradient information in addition to the intensity histogram information, which allows us to detect the smaller sub-cortical lesions better.

The rest of the paper is organized as follows: Section 2 briefly details the MRF approach used in this work along with the learning algorithm utilised. Section

3 then presents qualitative and quantitative empirical results on ten brain MRI volumes. Section 4 highlights the main observations and concludes with a brief discussion on the future work.

2 MRF Tissue Classification

Our goal is to classify each MRI voxel as belonging to one of the following six classes: Background (Bk), White Matter (WM), Grey Matter (GM), Cerebral Spinal Fluid (CSF), T1-hypointense lesions ($T1_{les}$) and T2-hyperintense lesions ($T2_{les}$). These two lesion types can be characterized as follows: $T2_{les}$ are hyperintense as compared to the healthy WM tissues in the $T2$ modality and $T1_{les}$ are hypointense as compared to the surrounding healthy white matter (WM) tissues in the $T1$ image modality. $T1$ lesions are a subset of the $T2$ hyperintense lesions, and are presumed to be the tissue destructive ones.

Our approach is divided into two parts. In the first part, distributions needed for MRF modelling are obtained from pre-labeled training volumes. Next, distributions learned from the training volumes are used in classifying new volumes.

2.1 Training

Given the MRI image modalities, the intensity of voxel v_i in the MRI volume is denoted as a three dimensional vector $\mathbf{I}_i = (I_i^{T1}, I_i^{T2}, I_i^{PD})$ with the elements corresponding to the voxel intensities in T1-weighted, T2-weighted and PD-weighted modalities respectively. Let each voxel v_i be modelled as a random variable f_i and $\mathcal{S}$ be a collection of sites, where each site is a voxel. Each f_i can take one of the six labels corresponding to the classes above.




Fig. 1. (left) The histograms of the white matter voxels in $T2$ image and (middle) the grey matter voxels and (right) the CSF voxels. We see that the histograms fit the multivariate Gaussian model quite well.

In the training phase, given the expert labelled set of R images, $\mathbf{I}^{(k)}$, $k \in 0, \dots, R-1$, we plot the intensities of the voxels belonging to the six classes for all three modalities. Our observations as shown in Fig. 1 confirms that the histogram of the intensities of the healthy brain tissues can be approximated as

multivariate Gaussians. Further, the distribution of intensities of lesions, though a little different (they are bimodal) can be made to fit the multivariate Gaussian distribution by dividing the lesions that are visible only on $T2$ volumes and those observable in both $T1$ and $T2$ volumes into $T2_{les}$ and $T1_{les}$ respectively. In Fig. 2, we see that there is perceptible difference between the different Gaussians and their peaks, allowing the algorithm to distinguish between various classes based on the intensities. Consequently, we fit multivariate Gaussians to each tissue class in the histograms and then compute the mean, μ_{f_i} , of the intensities of classes in the multispectral intensity space, along with covariance matrices Σ_{f_i} .




Fig. 2. The Gaussians fitted to the histograms of the volumes in the $T1$ (left), $T2$ (middle) and PD (right) modalities. We can see that there are considerable differences between the histograms of the classes [13].

2.2 Markov Random Field Tissue Classification

The choice of MRFs is motivated by the fact that tissues tend to occur and agglomerate in locally contiguous patterns. Consequently, local tissue information is of very high importance in the detection and location of lesions and local tissue information influences the accurate classification of images. Given this predilection, MRFs, that depend on local neighbourhood information to cluster voxels to their classes, are a natural choice.

Given a new image to classify, we model the entire volume as an instance of an MRF and the goal is to obtain the best configuration of labels over all the voxels in the image. The set of classes is denoted by $L = \{l_1, \dots, l_M\}$. In our case, we have $M = 6$, with $L = \{Bk, WM, GM, CSF, T1_{les}, T2_{les}\}$. Let the vector $\mathbf{f} = (f_1, \dots, f_n)$ denote a label configuration on the brain volume, where n is the number of voxels in the MRIs. A random instantiation of this vector hence induces a random field on $\mathcal{S}$. Let $\mathcal{F} = \{(f_1, \dots, f_n) | f_i \in L, i \in \mathcal{S}\}$ be the set of all such possible configurations. We define a homogeneous and isotropic

neighbourhood¹ system $\mathcal{N} = \{\mathcal{N}_i, i \in \mathcal{S}\}$ such that $\mathcal{N}_i$ denotes neighbours of the voxel at i , and $i \in \mathcal{N}_j \iff j \in \mathcal{N}_i$. Then by standard definitions [14] a random field $\mathbf{f}$ on $\mathcal{S}$ is an MRF on $\mathcal{S}$ with respect to $\mathcal{N}$ iff $P(\mathbf{f}) > 0$, $\forall \mathbf{f} \in \mathcal{F}$ and

$$P(f_i | f_{\mathcal{S}-\{i\}}) = P(f_i | f_{\mathcal{N}_i}). \quad (1)$$

The tissue class membership probability of voxels depends on the neighborhood.




Fig. 3. The 7-voxel clique containing the voxel of interest. The clique contains its four in-slice neighbours and the two neighbours from preceeding and succeeding slices.

According to the Hammersley-Clifford theorem [14], this MRF is equivalent to a Gibbs distribution. The probability of any configuration $\mathbf{f}$ takes the form:

$$P(\mathbf{f}) = \frac{1}{\sum_{\mathbf{f} \in \mathcal{F}} \exp(-\frac{1}{T}U(\mathbf{f}))} \exp(-\frac{1}{T}U(\mathbf{f})) \quad (2)$$

where T , the temperature, is an optimisation parameter and $U(\mathbf{f}) = \sum_s V_s(\mathbf{f})$, the sum of clique potentials $s \in \mathcal{G}$. A clique G is a subset of sites in $\mathcal{S}$ and can consist of a single site, or a pair of neighbouring sites or a triple of neighbouring sites and so on. We choose the one voxel clique and the seven voxel cliques, (shown in Fig. 3) because the energy of these two cliques dominates the 2 voxel, 3 voxel and 5 voxel cliques used by Harmouche [13]. The energy $U(\mathbf{f})$ is given by

$$U(\mathbf{f}) = \sum_{i=1}^n \left[(\mathbf{I}_i - \mu_{f_i})^T \Sigma_{f_i}^{-1} (\mathbf{I}_i - \mu_{f_i}) + \sum_{\forall j, j \in \mathcal{N}_i} \alpha (\mu_{f_i} - \mathbf{I}_j)^2 + \beta m(f_i) \right], \quad (3)$$

where μ_{f_i} is the mean intensity value of the voxels belonging to the class f_i in all the three modalities in the training set, Σ_{f_i} is the covariance matrix of the class given by f_i , $m(f_i)$ is the number of voxels that belong to the class of f_i in the 7 neighbourhood of voxel i and α and β are MRF parameters that determine the attraction between like intensities and similar clusters respectively. α is the weight of the gradient in $3D$. The dimensions of the voxels are not uniform (in

¹ The homogeneity implies positional independence in the neighbourhood and isotropism implies orientation independence.

our case, the voxels measure $1mm \times 1mm \times 3mm$). Since we are computing the gradient for non-uniform voxels, it is useful to weight the gradient accordingly, to eliminate errors due to differences in voxel dimensions. $\beta m(\mathbf{f}_i)$ may be thought of as the regularisation of the gradient term. Both α and β are determined experimentally from the training set. The term $(\mathbf{I}_i - \mu_{f_i})^T \Sigma_{f_i}^{-1} (\mathbf{I}_i - \mu_{f_i})$ can be thought of as the energy of the voxel itself (or the one voxel clique energy) and the term $\sum_{j,j \in N_i} \alpha (\mu_{f_i} - \mathbf{I}_j)^2 + \beta m$ the energy contributed to the voxel by its neighbours in the clique (seven voxel clique energy). In order to maximise $P(\mathbf{f})$, we use simulated annealing [16] to minimise $U(\mathbf{f})$, where

$$\mathbf{f}_{min} = \operatorname{argmin}_{\mathbf{f} \in \mathcal{F}} U(\mathbf{f}). \quad (4)$$

3 Experiments and Results

3.1 Image Preprocessing and Training

Before the images can be used for training, the brain MRI volumes need to be corrected for acquisition artefacts and co-registered into a common space. Image preprocessing includes bias-field inhomogeneity correction using N3 [7], intra-subject registration of the multispectral volumes [8], extraction of non-brain regions from the MRI (brain parenchyma) [9] and intensity range normalization. The intra-subject registration of the volumes involves registering the three modalities to the $T2$ space before any further processing is done, either manually or automatically. Finally, for each patient, the lesions were manually segmented by 5 expert raters resulting in a silver standard such that any lesion voxel is the result of a consensus among the experts.

The algorithm was trained on 14 real MRI brain volumes acquired from different MS patients. For all subjects, brain MRI scans were acquired on a 1.5T Philips Gyroscan ACS II scanner using a body coil transmitter and a quadrature head-coil receiver. Fifty contiguous 3mm thick $T2$ and PD weighted images were acquired parallel to the callosal line using a dual turbo-spin-echo sequence with repetition time (TR) = 2075 ms, echo time (TE) = 31.6 and 90 ms, 256×256 mm field of view. $T1$ weighted images were acquired with the same matrix using a 3D gradient echo sequence (TR = 35ms, TE = 10.2 ms). Each MRI volume was corrected for image inhomogeneity and $T2$ and PD weighted volumes were registered to the $T1$ -weighted volumes using a mutual information based approach. To ensure sufficient diversity in our training set, we chose subjects having with different amounts of lesion load, with 5 volumes having lesions less than 5cc, 5 volumes having lesion load between 5cc and 15cc, and 4 volumes with lesion load between 15cc and 70cc.

3.2 Classification and Post-processing

Our classification algorithm was tested on 10 MRI brain volumes of real patients with MS. The validation set had lesion loads varying from 6cc to 71cc.




Fig. 4. (a) (b) (c) (d)
 In all the above figures, we have (a) Central Slice of a T2-weighted MRI image of a person suffering from MS (b) the same slice in the T1-weighted image (c) the lesions labelled by experts and (d) lesions labelled by our algorithm. In the top row of figures, we have low lesion load. The algorithm manages to localise lesions accurately with $\kappa = 0.69$. In the central row of figures, there are a large number of sub-cortical lesions. However, the algorithm did manage to locate them reasonably correctly with $\kappa = 0.72$. Finally, in the bottom row of figures, there is heavy lesion load in the slice, but the algorithm detects them fairly accurately. It also managed to localise the smaller lesions away from the ventricles accurately. The κ value here is 0.76.

The solution proposed by simulated annealing led to a labelling of each voxel in each brain volume. In some cases, due to noise and intensity aberrations, it was found that voxels in the middle of lesions were not correctly labelled as lesions. In order to correct these errors, we use smoothing by b-splines since they are good surface approximators. The technique has been applied as in [15]. The general form of a B-spline surface of a given set of data points p_k is given by

$$x(u, v) = \sum_{i=0}^{N-1} B_i(u, v) d_i, \quad (5)$$

where $B_i(u, v)$ are the spline basis vectors and d_i are the coefficients. We need to minimise the functional

$$F = \sum_k \|x(u_k, v_k) - p_k\|^2 \quad (6)$$

to get the surface. We have implemented this technique as suggested in [15]. In our case, let us assume that a voxel V is flagged as a lesion. We choose square grids of 7×7 voxels centred at V in the axial plane and a grid of 3×3 in the coronal and sagittal planes since the voxel size is $1 \times 1 \times 3mm$. We assume that the lesion probabilities are themselves samples of the surfaces to be created, and generate control points at the corners of the squares. Then we smooth over the entire plane, approximating the other probabilities inside the squares as best as possible using the surface. At the location of the voxels, new probabilities are the resampled values of the continuous surface generated by the spline.




Fig. 5. (a) (b) (c) (d)
(a) Posterior fossa in the T2-weighted image of a patient with MS, (b) the same slice in the T1-weighted image, (c) the lesions labelled by the experts and (d) the lesions labelled by our algorithm image. While detection of lesions is difficult in the posterior fossa, our algorithm captured at least some lesions accurately with $\kappa = 0.49$.

3.3 Qualitative Results

Figs. 4, 5 and 6 present some qualitative results from five different slices, from different brain volumes with varying types of lesion loads. These include: (a) a

slice of a volume with heavy lesion load (bottom row Fig. 4), (b) a slice with low lesion load (top row Fig. 4), (c) a slice where there are sub-cortical lesions present (middle row Fig. 4), (d) a supra-ventricular slice with considerable lesion load (Fig. 6), and (e) a slice with lesions in the posterior fossa (Fig. 5).

The algorithm manages to find lesions even when they are sparse and only a few voxels in size as seen in top row of Fig. 4. Our detection of lesions is in consonance with the silver standard, even if there is a slight overestimate of the lesions at the top of the right ventricle in comparison with the silver standard.

In the bottom row of Fig. 4, one can observe that there is a considerable amount of lesion in the subject. In such cases, smaller lesions tend to be ignored by many commonly employed algorithms. In our case, it may be seen in Fig. 4, that not only are the larger lesions localised reasonably well, but also the two smaller lesions on either side of the ventricles have been identified correctly.

In Fig. 6, we have chosen to highlight a supra ventricular slice. Algorithms generally perform well in the central slices where the lesions often tend to aggregate around the ventricles. Our algorithm also detects these lesions accurately.

As mentioned earlier, proper detection of lesions in the posterior fossa is difficult due to differences in lesion intensities and contrast, compared to the cerebrum. Fig. 5 shows the lesions detected by our algorithm in the posterior fossa and we observe that they coincide with those detected by experts.

Sub-cortical lesions are often difficult to detect since they are very similar other brain tissue, apart from being very small. In the middle row of Fig. 4, we show that our algorithm is able to detect the cortical lesions as well, with our results are in agreement with the silver standard. In addition, our algorithm is able to detect peri-ventricular lesions accurately; however, in this case, it may be observed that our technique does overestimate the lesions near the boundaries.




Fig. 6. (a) (b) (c) (d)
(a) Upper brain slice of a T2-weighted MRI image of a person suffering from MS (b) the same slice in the T1-weighted image (c) the lesions labelled by experts and (d) lesions labelled by our algorithm. Although we choose a slice near the top of the ventricles, the algorithm still functions well as may be noticed with $\kappa = 0.72$.

3.4 Quantitative Results

Having described our results qualitatively, we compute the accuracy of our results using the traditional κ statistic (or Dice Similarity Coefficient) $\kappa = \frac{2(A \cap B)}{A + B}$ where A is the set of voxels labelled by experts as lesions, and B is the set of voxels labelled as lesion by our algorithm. Hence, the κ factor may be thought as the degree of agreement between the algorithm's and raters' lesion label assignment.

The results in Fig. 4 demonstrate qualitatively that the lesions identified by the algorithm coincide with the lesions labelled by the experts. Quantitatively, Table 1 shows the lesion κ scores for ten patients from the MNI database. These patients exhibited significant variation in lesion load. The mean value of $\kappa = 0.69$ indicates considerable relative agreement with the silver standard. Van Leemput et. al., [5] work with 50 real cases, which were rated by two manual experts. Their algorithm yielded a κ value of 0.47 with the first expert and a value of κ value of 0.51 with the second. On the other hand, our mean κ value compares favourably against the results obtained by both van Leemput et. al., [5] and Harmouche [13], the latter obtaining a mean κ of 0.61. Shiee et. al., [17] used both real and simulated data. With simulated data, their κ value for all classes is 0.677. They also used 10 real volumes (modalities were FLAIR, $T1$ and $T2$) and obtained an average $\kappa = 0.51$ with the expert raters for the entire brain. It is important to note that the κ values reported in the studies reflect the accuracy of segmentation of *all* the classes, and not merely lesions. Shiee. et. al., also report a κ value of 0.53 for lesions in 8 volumes, where the output of their algorithm was compared against the lesions detected by a single expert rater. Our κ value refers only to the segmentation of MS lesions, where there is often a lack of consensus even among experts, particularly in the vicinity of the boundaries.

Patient	1	2	3	4	5	6	7	8	9	10	Mean
κ with PF	0.64	0.76	0.77	0.72	0.69	0.64	0.74	0.63	0.62	0.67	0.69
κ without PF	0.67	0.77	0.78	0.74	0.71	0.66	0.75	0.64	0.64	0.69	0.71

Table 1. The κ values comparing our results to the experts' consensus (the silver standard with and without the posterior fossa (PF)).

While the κ statistic is a reasonable measure of similarity between the algorithm and the experts, it provides no information about the kind of errors generated by the algorithm, nor does it possess any information about the localisation of the errors. Consequently, we also compute the false positive and the false negative rates, so that we may measure the accuracy of the algorithm.

The false positive rate is given by $f_p = \frac{\# \text{ false positives}}{\# (\text{false positives} + \text{true negatives})}$ and the false negative rate is given by $f_n = \frac{\# \text{ false negatives}}{\# (\text{true positives} + \text{false negatives})}$. Table 2 presents the false negative rate and the false positive rate scores for all

the ten volumes whose κ value we computed in Table 1. This improves the utility of the statistics shown in the paper.

Patient	1	2	3	4	5	6	7	8	9	10	Mean
False +ve Rate	0.003	0.002	0.001	0.004	0.003	0.0004	0.0006	0.0003	0.001	0.005	0.002
False -ve Rate	0.13	0.06	0.05	0.06	0.09	0.19	0.11	0.14	0.12	0.06	0.10

Table 2. The false positive and false negative rates for lesions over the ten volumes.

As we can see from Table 2, the rate of false negatives is rather small, never exceeding 20% of the total number of lesion voxels. This shows that we do manage to capture most lesions accurately. The false positive statistic is quite low as well, however this statistic is subject to misinterpretation. In general, the number of lesions is seldom greater than 1% of the brain volume (in terms of voxel). Consequently, the number of true negatives is far greater than the number of false positives. As a result, the false positive rate is consistently low, regardless of the absolute number of false positives. The proposed approach does, in fact, produce some false positive classification results relative to the silver standard. However, recall that our silver standard reflects the consensus among experts. As such, the silver standard tends to underestimate the number of lesions and generally reflects the most conservative classification result. Our false positives are generally small as in Fig. 4 or occur at the borders of the lesions, where there is considerable disagreement even among experts.

4 Conclusions and Future Work

In this paper, we present an automatic tissue classification scheme based on Markov Random Fields that probabilistically models the local spatial relationships between voxels and their neighbours. Our approach explicitly builds distributions for lesions as separate tissue classes, as opposed to considering them as outliers. We adapt simulated annealing techniques to obtain the required MRF parameters. Our framework is evaluated on a dataset of 10 full MRI volumes of brain images acquired from real patients suffering from MS. The data reflects patients at various stages of the disease, thus exhibiting various degrees of lesion loads. This lies in contrast to many approaches in the literature, which post results on simulated images or focus their findings on specific slices in brain volumes. Our results indicate that the algorithm produces high lesion kappas relative to a silver standard, and relatively low false positive and negative rates. In addition, the results are consistent with expert labellings in the posterior fossa, a brain region generally considered “difficult” to correctly segment by experts.

Several improvements are in the horizon for this technique. Different models for the different regions of the brain as suggested by Harmouche [13] might help in capturing local variations better. We also expect to integrate the smoothing

function completely into the classification scheme, something which should help us overcome the drawbacks of over-estimation of lesions.

5 Acknowledgements

The authors are grateful to the Montreal Neurological Institute for the use of its database of MS patients. The authors would also like to express their gratitude to Mr. Zografos Caramanos for his help during the experiments.

References

1. F. Admasu et. al., "Segmentation of multiple Sclerosis lesions from MR Brain images using the principles of fuzzy connectedness and artificial neuron networks, ICIP, pp. 1081-0184, 2003.
2. J. K. Udupa, et. al "Multiple Sclerosis Lesion Quantification using Fuzzy Connectedness Principles", IEEE TMI, Vol. 16(5), pp. 598-609, 1997.
3. F. Yang, et. al., "White Matter lesion segmentation from volumetric MR Images", MIAR, pp. 113-120, 2004.
4. Z. Y. Ge et. al., "A Statistical 3-d segmentation algorithm for classifying brain tissues in MS", in Proc. CBMS, 14th IEEE Symposium, pp. 455-460, 2001.
5. L. van Leemput et. al, "Automated Segmentation of Multiple Sclerosis Lesions by Model Outlier Detection", IEEE TMI, Vol. 20(8), pp. 677-688, 2001.
6. S. Francis, "Automatic lesion identification in MRI of MS patients", Master's Thesis, McGill University, 2004.
7. J. G. Sled and G. B. Pike, "Correction for b(1) and b(0) variations in quantitative t(2) measurements using MRI", Mag. Reson. Med., Vol. 43(4), pp. 589-593, 2000.
8. D. L. Collins et. al., "Automatic 3d intersubject registration of MR volumetric data in standardised Talairach space", J. Comp. Ass. Tom., Vol. 18, pp. 192-205.
9. S. M. Smity, "Fast Robust automated brain extraction", Hum. Brain Mapp. Vol. 17(3), pp. 143-155.
10. S. Al-Zubi et. al, "Fusing Markov random fields with anatomical knowledge and shape based analysis to multiple sclerosis white matter lesions in magnetic resonance images of the brain", in Proc. SPIE (Medical Imaging), Vol. 4684, pp. 206-215, February 2002.
11. J. Grimaud et. al., "Quantification of MRI lesion load in MS: A comparison of three computer assisted techniques", J. MRI, Vol. 14, pp. 495-505, 1996.
12. J. Lecoer et. al, "Multiple Sclerosis Lesions Segmentation using Spectral Gradient and Graph Cuts", in Proc. MICCAI Workshop (Validation and Methodological Issues), 2008.
13. R. Harmouche, "Bayesian MS lesion classification modelling using Regional and Local Spatial Information", Master's Thesis, McGill University, 2006.
14. J. Besag, "Spatial Interaction and Statistical Analysis of Lattice Systems", J. Royal Stat. Soc., Ser. B, Vol. 36(2), pp. 192-236, 1974.
15. H. Pottmann et. al., "Approximation with Active B-spline Curves and Surfaces", in Proc. 10th Pacific Conference on Computer Graphics and Applications, 2002.
16. R. O. Duda, P. E. Hart, and D. G. Stork, "Pattern Classification", John Wiley and Sons, 2000.
17. N. Shiee et. al., "Topologically Constrained Segmentation of Brain Images with Multiple Sclerosis Lesions", in MICCAI Workshop on. Med. Imag. Analysis on Multiple Sclerosis, pp. 72-81, 2008.