

RESEARCH

Open Access

Multiple sclerosis lesion detection with local multimodal Markovian analysis and cellular automata 'GrowCut'

Daniel Biediger^{1,2}, Christophe Collet^{1*} and Jean-Paul Armspach¹

*Correspondence:

c.collet@unistra.fr

¹iCube Lab - Université de
Strasbourg - CNRS UMR 7357, 300
Bd Sébastien Brant, CS 10413, 67412
Illkirch Cedex, France
Full list of author information is
available at the end of the article

Abstract

Multiple sclerosis affect over 2.5 million people world-wide. This autoimmune disease of the central nervous system causes damage to the insulating myelin sheaths around the axons in the brain. The disease progresses at different rates in different people and can have periods of remission and relapse. A fast and accurate method for evaluating the number and size of MS lesions in the brain is a key component in evaluating the progress of the disease and the efficacy of treatments. Manual segmentation is slow and difficult and the results can be somewhat subjective. It requires a physician to consider several MRI slices across multiple modalities. The power and speed of computer systems provide an obvious avenue to help. While many automated methods exist, they have not reached human-level accuracy of the segmentation results. There exists a need for a robust, fast and accurate method to improve the results of automatic MS lesion segmentation methods. We propose a post-processing stage to improve the segmentation results of an existing system. It uses two different strategies to improve the segmentation results of an automated system based on whole-brain tissue classification and lesion detection. The first strategy leverages the current processing system at a granularity finer than the whole brain to detect lesions at a local level. This reflects the way that a physician considers only a part of the brain at a time. It then combines the series of local results to produce a whole-brain segmentation. This approach better captures the local lesion properties and produces encouraging results, with a general improvement in the detection rate of lesions. The second method dives deeper and looks at the individual voxel level. Just as a physician might look more closely at a lesion, it considers the local neighborhood around a lesion detection. The method selects seed points from the existing results and uses a region growing method based on cellular automata. It grows the lesion areas based on a local neighborhood similarity in intensity. Over the eleven patients examined, some results improved over the base case and show the efficiency of the proposed approach.

Keywords: Image processing; Brain segmentation; MRI; Lesion detection; Multiple sclerosis

Background

Multiple sclerosis (MS) is a disease of the central nervous system that causes damage to the insulating myelin sheaths around the axons in the brain. MS causes the immune system to attack these nerve fibers. The resulting demyelination interferes with the nerve's ability to communicate electrical signals. The damage manifests as a range of cognitive and physical disabilities. The disease progresses at different rates in different people and can have periods of remission and relapse. Even in cases where the myelin can re-grow, there is a permanent degradation in the transmission of electrical impulses in the cell. It is important to be able to detect and evaluate the location of, size of, and changes in MS lesions in the brain. The healthy brain contains white matter (WM), grey matter (GM), and cerebrospinal fluid (CSF). While lesions can be present in both WM and GM, they do not affect the CSF. Magnetic resonance (MR) images provide a non-invasive way to examine the tissues of the brain and to detect MS lesions.

MRI modalities

A number of magnetic resonance imaging (MRI) modalities are available for identifying the tissues and any lesions present in the brain, and each provides different information. The common MRI modalities for the identification of MS lesions are T1-weighted, T2-weighted, PD-weighted, and fluid-attenuated inversion recovery (FLAIR), see Figure 1. For all modalities, it is possible to achieve a high spatial resolution, with voxel cubes of about 1 mm on the side. To save time in the acquisition of MRI sequences, it is also possible to have slice thickness from 3 to 7 mm and even gaps between slices. While the spatial resolution is good, the brain contains complex structures in three dimensions. As a result, the voxels in an MRI can be impacted by several tissue types within the volume. This partial volume (PV) effect can cause problems in the identification of tissues. Fortunately, it is also possible to collect multiple scans of the same patient using different acquisition weights. Taken together, these scans provide anatomical and tissue information that is valuable for identifying MS lesions.

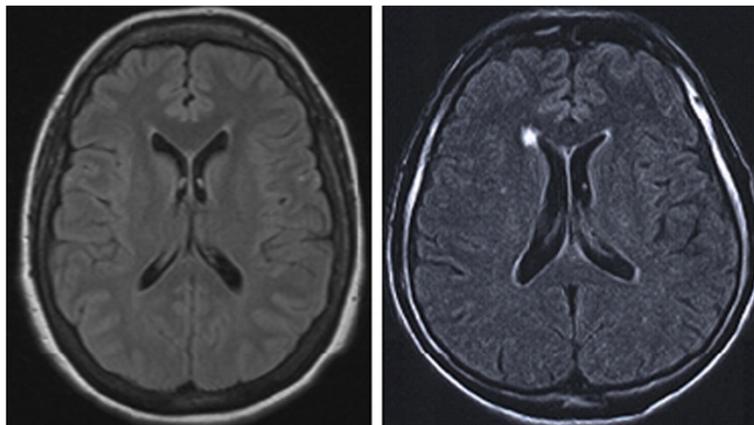


Figure 1 Comparison of normal tissue and MS lesions. The differences between the modalities relate to the method of image acquisition and manifest in the appearance of relative intensity of lesions compared to healthy tissue.

Lesion detection

In T2-weighted images, water appears with a high intensity of images. In T2 images, lesions appear with higher (hyper-) intensities when compared to the surrounding WM. Unfortunately, regions containing CSF also appear with high image intensities. As a result, it becomes difficult to segment lesions in areas near the CSF-filled ventricles. Figure 1 shows an example of a normal brain and a brain with hyper-intense lesions visible. In FLAIR images, a special weighting effectively removes the CSF from T2-weighted images. The result is an image where CSF appears with low intensity and lesions appear with a high intensity. FLAIR images can be useful for identifying lesions in WM, even in the periventricular region. Unfortunately, FLAIR imagery has a long acquisition time and may not be available for all patients. Proton-density-weighted images also show reduced intensity of CSF compared to T2-weighted, though not as much as FLAIR. The advantage of proton-density-weighted series is that they can be acquired together with T2 images in the same sequence.

The task of identifying MS lesions in MRI is difficult and normally requires an expert physician. The identification process is also time consuming and includes some subjectivity in interpreting the voxel data. It requires the fusion of intensity data from different MRI modalities, anatomical understanding, and spatial thinking. Lesions can take on many different appearances from ovoid to irregular and fuzzy to distinct. As a result, the lesion segmentations provided by different physicians can vary in the number and size of lesions identified. Fortunately, computerized methods for lesion identification have been proposed to tackle this problem. These methods promise consistency and repeatability, though they do not always agree with the expert segmentations. The methods take different approaches to the problem of lesion segmentation and include a number of steps, including pre- and post-processing. This article introduces a two-step approach to improve the results of an existing automated segmentation method. It includes an analysis of the existing results which explains the rationale behind the proposed changes. It presents the base segmentation results and improved results for the two modifications performed on a set of 11 patients. It compares all the results to the expert segmentations for each patient. This report concludes with a discussion of the techniques and results as well as the perspectives on future improvements.

Methods

Segmentation methods: state of the art

The detection of lesions in MRI scans generally requires a significant time investment from a skilled physician. A number of automated methods seek to lessen this burden by providing fast, accurate, and repeatable segmentation results. The goal of this project was to improve upon the results of an existing segmentation method [1]. The current system represents a significant investment of time and produces good results for a range of real cases. Unfortunately, the process falls short of the abilities of an expert physician.

A number of general categories exist for automated segmentation of MS lesions in MRI scans of the brain. The methods can be divided based on the approach and grouped based on their implementation [2]. There are three main types of segmentation approaches: manual, semi-automatic, and automatic. Manual segmentation is the base method for lesion segmentation. An expert physician examines different modalities to select the lesion voxels. Unfortunately, the manual process is time consuming and somewhat

subjective. Different experts can report different results and the same expert can provide different results for the same data on subsequent evaluations. Even so, manual segmentations are considered the best results available and serve as the baseline for evaluating other methods. The expert segmentations can be considered as a 'silver standard' since they are not perfect representations of the ground truths but provide the best *in vivo* estimates available. Computer-aided methods do provide some benefit to MS lesion segmentation. Where experts can have difficulty combining information from multiple MRI modalities and from multiple adjacent slices, well-designed algorithms can efficiently blend this data. As a result, it is interesting to pursue the development of semi-automated and automated lesion segmentation methods.

Semi-automatic methods require some human input as the starting point for an automated processing step. This information could be a region of interest or a coarse selection of lesion voxels. While semi-automatic methods can relieve some of the work from physicians, they do require some input. The input can be as simple as a region of interest, with suspected lesions. The user provides a rectangle around suspected lesions to narrow the focus of the algorithm. The required input might be as detailed as a coarse painting of lesion and non-lesion tissues. An algorithm could then use information based on the appearance and features of these selections to grow the two regions without any other knowledge. In any case, the automated portion of the segmentation is sensitive to the quality of the input. Because they require some level of user input, the semi-automatic methods may be unsuitable for large patient studies. Fully automatic methods require no human interaction and can be grouped, at several levels, based on the method used to perform the segmentation. In general, there are three main types of fully automated segmentation schemes: data-driven methods, intelligent methods, and statistical methods. The data-driven methods use thresholding and region growing to segment the lesions in an image, like the watershed and grow-cut methods. The learning-based methods require a training set and some feature extraction. These methods learn the characteristics of lesions and then classify based on fuzzy rules or decision forests. The statistical methods involve estimations of probability density functions. These methods are based on inference methods with some neighborhood or classification examples and include Markov models and support vector machines. All have advantages and disadvantages in their use and the results they provide (see [3] for more details).

Graph cuts (GC) is a method for finding the *maximum a posteriori* (MAP) estimate of a binary image [4]. The method treats the image like a flow graph with two nodes, the 'source' and the 'sink'. The source represents the object class in the image, in this case the lesions. The sink represents the background, the non-lesion tissue. The other nodes of the graph are the image voxel. A network of weighted and directed edges connect the nodes in the graph. The GC makes use of regional and voxel-neighborhood information to differentiate between the two classes.

The MAP estimate corresponds to the maximum flow through the node network. Essentially, the method removes the inter-label connections in favor of intra-label connections. The result is two sets of strongly connected nodes that correspond to the foreground and background image elements.

The *fuzzy c-mean* seeks to cluster pixels into a number of groups that maximize inter-cluster variability while minimizing intra-cluster variability [5,6]. Rather than a crisp or hard classification, the fuzzy approach specifies the degree to which a pixel belongs to a

given cluster. In this way, a pixel can belong to more than one cluster with some degree of probability. Manual segmentation results varying from expert to expert and for repeated evaluations by the same expert can be taken into account in this way.

The *mean-shift* is an unsupervised non-parametric clustering algorithm for image segmentation [7]. The main idea of the mean-shift algorithm is to treat image points as vectors in a probability density function. The dense regions in this space represent the local maxima of some underlying distribution. The method performs a gradient ascent optimization at each image point until convergence. The mean-shift vector gradually decreases in length as it approaches the maximum. The resulting points are the modes of the distribution. Nearby data points, within some window size, are considered members of the same cluster. The clustering process depends on the selection of a kernel (local neighborhood) and the specification of a window size and not some prior specification of the number of clusters. The correct selection of a window size is the key for obtaining good results. If the window size is too large, the image will be under-segmented and regions will be lumped together. This can remove the fine details of small structures like MS lesions. If the window size is too small, a significant amount of over-segmentation can occur.

The *k-nearest neighbor* (k-NN) is a learning-based approach that attempts to classify voxels based on the consensus of nearby examples [8]. A number of features can be extracted for a voxel, including its appearance, location in the brain, and relation to its neighbors. A labeled training set provides examples in feature space against which a test voxel is compared. The advantage of this approach is that it needs to only locally estimate the probability densities. The classification is based on the agreement of test examples with similar features in some small neighborhood of examples. The method requires and depends on good examples for good classification. Because MS lesions vary in size, shape, and appearance, they will have widely different feature sets. Without a sufficient number of examples, it could be difficult to correctly classify lesions. Even in patients with MS lesions, the actual number of lesion voxels may be far less than the number of voxels representing healthy tissue, perhaps one in a thousand. As with other learning-based method, the large fraction of non-lesion voxels can bias the examples and hurt the lesion detection rate.

Support vector machines (SVM) is a popular and widely used supervised learning algorithm and has been applied to the MS lesion segmentation problem [9,10]. The method extracts some features from examples of lesion and non-lesion voxels. It then attempts to divide the two classes by a hyper-plane in the feature space. While there are many possible dividing planes, the method seeks the plane with the widest margin. Some methods can employ kernels to re-map the feature space and allow for a non-linear division of the classes. One problem with the SVM approach in MS lesion detection is the imbalance between class representations. In general, the number of voxels that represent normal brain tissue far exceeds the number of voxels that represent MS lesions. This can lead to the over representation of non-lesions in the training. Unfortunately, it is difficult to just exclude non-lesion examples, since any given example might represent important information.

Statistical models generally focus on some estimation of the probability of a lesion based on some mixture-models for normal tissue [11,12]. Generally, normal brain tissue is divided into three classes, WM, GM, and CSF. Lesions are generally treated as an

outlier to the normal tissue, although in some cases they can be treated as a separate class. The statistical methods try to assign a classification based on the likelihood that a given voxel is a lesion based on these models. The methods include neighborhood information through Markov random field (MRF) or conditional random field (CRF). In these cases, the nearby voxels contribute to the classification. These methods usually include some probability parameter or threshold, beyond which an outlier is considered a lesion. The existing method uses a hidden Markov chain (HMC) to incorporate neighborhood information into the segmentation process. The main drawback is that we segment the whole brain whereas the physician works locally. To be more efficient, we propose in this paper to combine local and global approaches.

Algorithm evaluation

To evaluate the effectiveness of any improvements, it was necessary to specify a quantitative metric for evaluating the progress of any proposed solutions. A number of metrics exist for comparing the computerized and expert results. The metric of interest for these comparisons is the similarity-index (SI).

The SI represents the amount of overlap in the identification regions provided by the experts and the method of interest. It is computed as the ratio of twice the area of intersection of the regions to the sum of the areas of the regions. It reflects the relative number of correctly segmented voxels to the false-positives and false-negatives in a single metric.

$$SI = \frac{2(A \cap B)}{A + B} \quad (1)$$

Values for the SI will fall between 1 and 0, with values closer to 1.0 representing better results and closer to 0.0 being the worse results. The goal is to improve the segmentation approach in a way that would better match the physician's segmentations as measured by the level of correspondence between the results and penalized by the difference. As the automated method is intended to assist the physicians in identifying lesions, under-segmentation or false-negatives were a more important concern than over-segmentation or false-positives. Physicians would have to search the whole brain for any missed lesions, but could more easily reject the incorrectly identified lesions. It was expected that the improvements of SI scores could include an increase in the number of true-positives as well as an increase both false-positives and false-negatives. The addition of false-negatives would be slightly more troubling than the addition of false-positives, but an increase in the detection of true-positives would supersede the two. Beyond the overall scores, it was interesting to consider a per-lesion evaluation of the efficacy of a given system. By comparing the labeled expert's results to a set of automated results, it was possible to evaluate over- and under-segmentation for each lesion. The per-lesion results provided a quantitative way to evaluate the power of a given method to discriminate the boundaries of lesions as well as its propensity for detecting spurious lesions.

Our development uses hidden Markov chain model [1,13] as a starting point because this approach obtained a very good score at *grand challenge: 3D segmentation in the clinic* in MICCAI'08. This Markovian method computes lesion segmentations that generally agree with expert results but under-estimates lesions in some places and over-estimates in others. After some investigation, we propose a two-stage process (global on the whole

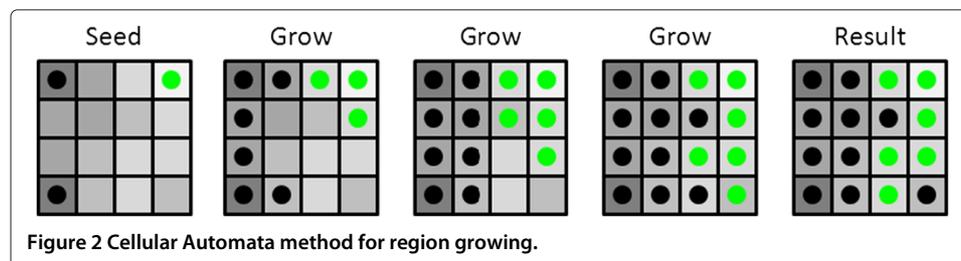
brain then local on reduced areas) for improving the segmentation results. The two-stage method proposed a change to the way the images were processed and then applied a post-processing step to grow the lesions. A discussion of the theoretical basis for the existing system and proposed improvements is provided in the Section ‘New local detection model.’ The specific implementation details are presented in Section ‘Proposed algorithm,’ whereas Section ‘Results’ will validate the approach for a set of patients. A conclusion ends this paper, with some observations and potential avenues for future work.

New local detection model

The ‘GrowCut’ approach to image segmentation using a Cellular Automaton seemed to be a good match for the local-region-growing idea to improve segmentation results [13]. Cellular Automata (CA) are discrete models both in space and time that govern the evolution over time of a grid of cells. The cells exist in some finite set of states and have simple deterministic rules that govern the change of these states at each time step. In the case of image segmentation, it is possible to construct an update rule to grow regions given some example seeds. The method is generally applied to semi-automatic segmentation, with a user providing the seed points for the fore- and background. Since the existing method provides a base segmentation, these results could be adapted to provide seed points in a fully automatic way. Figure 2 shows a schematic representation of the way the method progresses, from seed selection through growth iterations to the final convergence.

The CA method operates on a set of voxels V in the MRI image. These voxels become cells p in the lattice of the automaton. This relationship between voxels in the 3D image and the corresponding cells is summarized as $p \in V \subseteq \mathbb{Z}^3$.

The cellular automaton A is composed of the triplet $A = (S, N, \delta)$. S represents the non-empty set of states within the automaton. N represents the neighborhood of points, in this case the, six-neighborhood in three dimensions around a given cell. The transition rule, δ , is the function which updates the states of the cells at each time step t . The state of a given voxel, S_V , is composed of three pieces of information and given by $(l_V, \theta_V, \vec{C}_V)$. The label of the current cell is given by l_V and can have the integer values $[-1, 1]$. The values $-1, 0$, and 1 correspond to the labels unknown, WM, and lesion, respectively. The unknown label represents those cells which have yet been assigned a value or that are specifically excluded from the analysis because, for example, they belong to a different tissue class. The strength of the current label is given by θ_V with values given by $\theta_V \in [0, 1]$. The feature vector \vec{C} represents the properties of each voxel, in this case, a scalar intensity.



Region-growing algorithm

The following are the steps for region growing:

1. Select some seed points as representatives of the WM and lesions and assign them the corresponding labels, strengths of 1.0, and their respective intensities.
2. Assign all other points an initial label, strength in the range $[0, 1]$, and their respective intensities.
3. Assign the non-WM and non-lesion points the unknown label and strengths of 1.0. These points, identified as GM and CSF by the atlas, are ignored in the iterations.
4. Iterate over the points in the image using the evolution rule presented above to update the strengths and labels of the cells.
5. Terminate the iteration when there are no additional changes (no label changes) or after some maximum number of iterations.

In the first three steps, the values for each cell are initialized to a suitable value. Seed points are normally selected by the user, but in a fully automated system the main processing step provides these points. For these seed points, the lesion label is applied and the corresponding label strength is set to the maximum value of 1.0. In a similar way, points can be set as anti-seeds: those cells that could not be part of the lesion class because they have very low intensities. These anti-seed points are given then non-lesion label with the maximum label strength of 1.0. Some cells represent other tissue types and must be ignored by the region-growing method. These cells are excluded by giving them maximum strengths and a label that indicates they should not be considered by the algorithm. The remaining cells represent viable growth regions. The algorithm will attempt to consume them for one label or the other. They must be given initial strength values in the range $[0, 1]$. With the cells of the automaton initialized, the iteration process of step 4 begins. The new values at time $t + 1$ are determined based on an evolution rule for the cells. The states for all cells in p at time t are given by S_p^t . These values are updated to the next time-step $t + 1$ and given by S_p^{t+1} . In this way, the label at each time step, l_V^t and its strength θ_V^t are updated to for each cell to l_V^{t+1} and θ_V^{t+1} , respectively.

Evolution rule for cellular automaton

As mentioned above, the states of the cells in the Cellular Automaton evolve over time. Through the application of the evolution rule, the cells move through a number of intermediate states to a final stable set of states. The evolution rule describes how the states are updated at each time step. At each step, the evolution rule iterates over the cells p in the image. The cell label, l_v , and the strength of that label, θ_v , are copied from the current time step t to the next time step $t + 1$. The neighbors around the cell of interest are then considered in turn. Each neighbor is compared to the cell of interest and *attacks* the cell, attempting to modify its label and strength. The attack strength for the comparison and the subsequent update in strength depend on a function $g(x)$. This function can take a number of forms but must be monotonic decreasing and restricted to the range $[0, 1]$. The monotonic decreasing requirement ensures that there are no local extrema in the function that will cause repeated skirmishes. The decreasing requirement also ensures that larger input differences result in smaller strength attacks. This is a necessary behavior, since only neighbors with some close resemblance should have a strong influence on a given cell. Restricting the range to $[0, 1]$ ensures that the strength values will always be bounded and

will not increase unexpectedly or uncontrollably. As seen in the algorithm, the input value x is determined at each point by the L2 norm of the difference of the feature vectors C . In the simplest case, the feature vectors are simply the voxel intensity values (Figure 3), but it is easy to imagine this method applied to multiple modalities taken together to form a collective feature vector. This evolution algorithm is then applied at each time step in the overall method. The process of updating continues until convergence or some pre-determined number of iterations.

Algorithm 1 Evolution rule for cellular automaton

| | |
|--|--|
| for $\forall p \in P$ do | ▷ For all cells p in the image |
| $l_V^{t+1} \leftarrow l_V^t$ | ▷ Preserve the cell label |
| $\theta_V^{t+1} \leftarrow \theta_V^t$ | ▷ and the label strength |
| for $\forall q \in N(p)$ do | ▷ For all cells in the neighborhood of p |
| if $g(\ \vec{C}_p - \vec{C}_q\ _2) \cdot \theta_q^t > \theta_p^t$ then | ▷ If the attacker is stronger |
| $l_p^{t+1} \leftarrow l_q^t$ | ▷ Preserve the cell label |
| $\theta_p^{t+1} \leftarrow g(\ \vec{C}_p - \vec{C}_q\ _2) \cdot \theta_q^t$ | ▷ and the label strength |
| end if | |
| end for | |
| end for | |

Proposed algorithm

Because the region-growing method based on CA operates at the local voxel level, it should mimic the approach that a physician would take in segmenting lesions. The method depends on a number of details that were important in the implementation. One key detail was the selection of the weighting function for the automaton growth based on the difference in appearance between the lesion voxels. The second major consideration was the initialization method and selection of appropriate seed points for the grown region. The third and related question was the scale at which to perform the processing. These three factors contribute directly to the performance of the method and are weighed for their potential advantages and disadvantages.

Region-growing method

The region-growing method is summarized below:

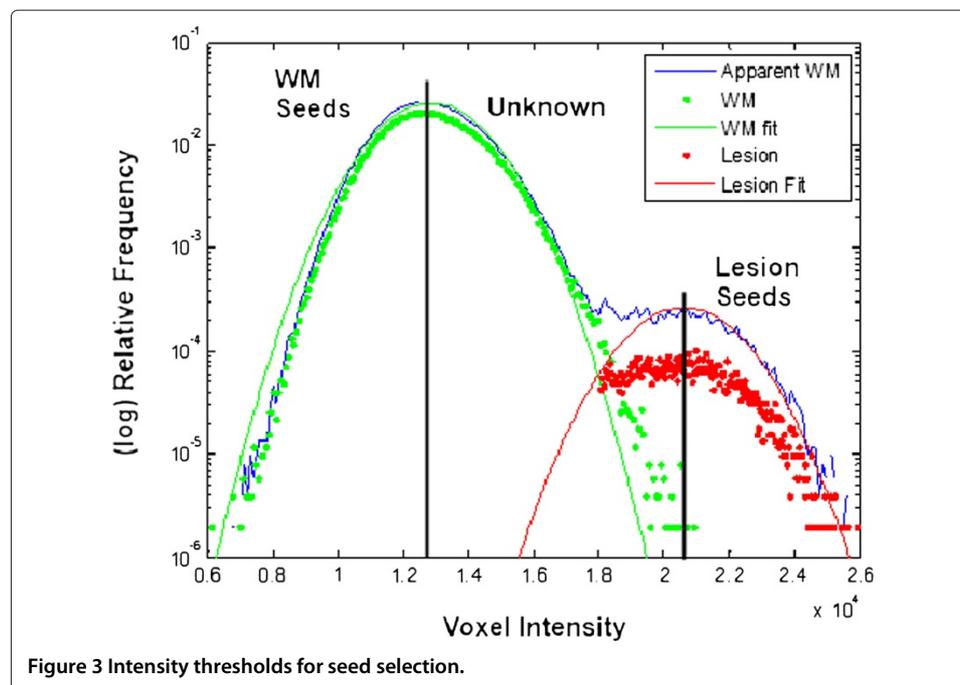
1. Perform the normal analysis using an automated method to identify potential lesions.
2. Compute region statistics for the WM and lesion voxels identified by the analysis.
3. Select WM seeds from the WM voxels with intensities below the regional mean for the WM. Assign them the WM label (value 0) and strengths of 1.0.
4. Select the lesion seeds by accepting all lesion voxels in the region with intensities above the lesion median. Assign them the lesion label (value 1) and strengths of 1.0.
5. Assign all other WM and lesion voxels their corresponding labels. Assign the WM voxels strengths proportional to their Gaussian probability in the range [0, 1]. Give the lesion voxels strengths of 0.0.

- Assign all non-WM and non-lesion voxels the unknown label (-1) and strengths of 1.0. These are excluded from computation.
- Perform the region-growing method by iterating over the cells as outlined.
- Combine the regional results to recover the whole-brain results.

Because the method requires some seed points, it is necessary to perform some automated analysis before applying the region-growing method. The first step in the method is to apply the HMC method [1] to sub-cubes in the brain and collect the results. The sub-cube provide a more local region for making estimates of the seed appearances in the brain. The second step computes the relevant statistics for the selecting the seeds. The selection of suitable seed points is critical for good segmentation results. While this method usually relies on a user to select these points, we made use of the existing segmentation results and attempted to grow the regions from there. To select a suitable pool of points, we considered the results provided by each sub-cube. By taking the results from a cube, it was hoped that a sufficient number of lesion voxels would be present for a good local estimate of the lesion class. Seed points were selected from the upper half of the lesion intensities and potentially unwanted lesions might be excluded. By operating at a regional rather than global level, the method hoped to avoid the regional differences in tissue and lesion appearances across the brain. The seed points were selected from an analysis of the intensities of the lesions and WM voxels within the sub-cube region (Figure 2).

Results

In the manual segmentation process, experts identify areas that appear brighter in the FLAIR MRI than the surrounding white matter. The locally bright points serve as a starting point for lesion identification, and the expert examines a close neighborhood around these points to complete the segmentation. They grow the lesions to some boundary with



a sufficient local intensity difference. Physicians incorporate anatomical and other information to select the starting points; this method relies on the bright points identified by the HMC method. The key element that this approach brings is the strategy of selecting regions of interest and then growing them based on local information. In this way, the method duplicates the strategy employed by expert physicians. To measure the efficacy of the method, results were computed for 11 patients using the base method, the sub-cube method, and the region-growing method. The patients were selected at random from a set of adult-patient data. Co-registered T1 and FLAIR MRI images with 1-mm voxel size were used. Table 1 shows the number of lesion voxels identified by an expert physician and show a range of 319 to 41,169 lesion voxels corresponding to a range of 0.319 to 41.169 cm³. These patients exhibited a range of lesions: from low to high loading and from uniform to non-uniform, which are located throughout the brain and with different sizes and shapes.

The baseline results were based on the HMC method using T1-weighted MRI and a robust atlas for tissue classification. The lesions were detected as outliers of the WM distribution in the FLAIR image. Table 2 shows a summary of the results for the patients in the study.

The common volume column indicates the number of voxels in common between the expert and the automated segmentations. These voxels are the true-positives that have been correctly detected by the system. The false-positive errors represent the voxels that have been identified by the automated system but not by the expert, over-segmenting the lesions. The false-negatives represent lesion voxels that were identified by the expert but missed by the automated method, i.e., the false-negatives. As explained before, the SI score represents the quality of overlap between the segmentation results of the two methods. While the scores were generally good, some of the results stand out as poor. Patients 4 and 11 had very high over-segmentation values that hurt their overall score. In both cases, image artifacts have caused excessive false detections. In one case, the patient had some form of tumor that was identified by the automated method as lesions. In the other case, artifacts from the brain extraction left some non-brain voxels in the image. These voxels were incorrectly identified as lesions. These artifacts were recognized and understood by the expert and not included in the manual segmentation. Unfortunately, the method itself cannot filter these results since they require actual understanding of the artifacts in these images. In all other cases, the under-segmentation posed a more significant problem than

Table 1 Expert segmentation results

| Patient | Lesion voxels | Lesion volume cm ³ |
|---------|---------------|-------------------------------|
| 1 | 2,421 | 2.421 |
| 2 | 7,719 | 7.719 |
| 3 | 998 | 0.998 |
| 4 | 3,026 | 3.026 |
| 5 | 319 | 0.319 |
| 6 | 5,449 | 5.449 |
| 7 | 3,539 | 3.539 |
| 8 | 5,218 | 5.218 |
| 9 | 918 | 0.918 |
| 10 | 41,169 | 41.169 |
| 11 | 619 | 0.619 |

Table 2 Base segmentation results using HMC method

| Patient | Correct | False-positive | False-negative | SI score |
|---------|---------|----------------|----------------|----------|
| 1 | 2,421 | 763 | 1,821 | 0.652 |
| 2 | 7,719 | 778 | 5,621 | 0.707 |
| 3 | 998 | 373 | 1,070 | 0.580 |
| 4 | 3,026 | 5,248 | 941 | 0.494 |
| 5 | 319 | 100 | 285 | 0.624 |
| 6 | 5,449 | 644 | 2,587 | 0.771 |
| 7 | 3,539 | 605 | 1,744 | 0.751 |
| 8 | 5,218 | 70 | 1,178 | 0.893 |
| 9 | 918 | 393 | 1,380 | 0.509 |
| 10 | 41,169 | 50 | 65,953 | 0.555 |
| 11 | 619 | 1,311 | 327 | 0.430 |

the over-segmentation. Patient 10 exhibited an exceptionally high lesion load and showed a significant amount of under-segmentation in the automated results. This could be, in part, due to the assumption that the lesions represent only the most significant outliers to a Gaussian distribution. Because the patients represent a wide range of lesion loads and appearances, they represent an interesting test group for the method.

Using a set of smaller sub-regions based on the HMC method, the sub-cube method produced an improvement over the base case. As can be seen in Table 3, the results while generally better were mixed. The approach performs much better in some cases and much worse in others. In nearly all cases, the method provided a better detection rate of the individual lesions and, in some cases, added a significant number of spurious detection points as well.

Using the seed points selected from the results of the HMC method, the region-growing method produced an interesting set of results. As can be seen in Table 4, the results were somewhat mixed, performing much better in some cases and much worse in others. In nearly all cases, the method provided a better hit rate on the lesions, even if it also added a significant number of spurious detection points as well.

As is apparent in the table, there are some shortcomings with the region-growing process. The method provides less common segmentation and correspondingly greater under-segmentation than the Markovian sub-cube method. Fortunately, it captures the lesion voxels better than the base method. But it also provides more over-segmentation:

Table 3 Sub-cube method results

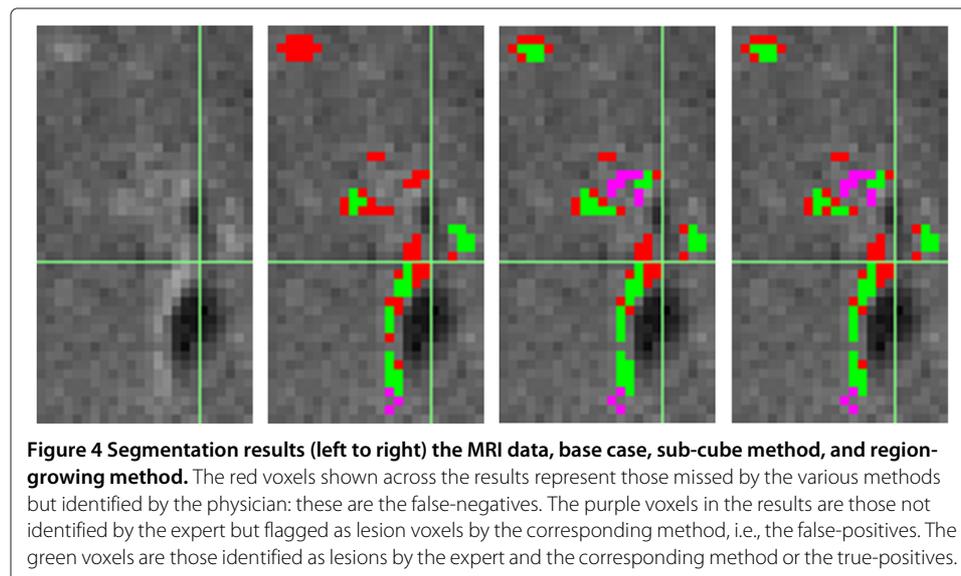
| Patient | Correct | False-positive | False-negative | SI score | Percentage of change from base |
|---------|---------|----------------|----------------|----------|--------------------------------|
| 1 | 2,966 | 1,569 | 1,276 | 0.676 | 3.7 |
| 2 | 8,909 | 1,095 | 4,431 | 0.763 | 8.0 |
| 3 | 897 | 483 | 1,171 | 0.520 | -10.4 |
| 4 | 3,387 | 4,154 | 580 | 0.589 | 19.1 |
| 5 | 371 | 225 | 233 | 0.618 | -0.9 |
| 6 | 6,409 | 1,497 | 1,627 | 0.804 | 4.2 |
| 7 | 4,459 | 2,101 | 824 | 0.753 | 0.3 |
| 8 | 5,866 | 397 | 530 | 0.927 | 3.8 |
| 9 | 1,514 | 1,107 | 784 | 0.616 | 21.0 |
| 10 | 44,944 | 238 | 62,178 | 0.590 | 6.3 |
| 11 | 730 | 2,811 | 216 | 0.325 | -24.4 |

Table 4 Region-growing method results

| Patient | Correct | False-positive | False-negative | SI score | Percentage of change from base |
|---------|---------|----------------|----------------|----------|--------------------------------|
| 1 | 2,779 | 1,487 | 1,463 | 0.653 | 0.2 |
| 2 | 8,301 | 1,044 | 5,039 | 0.732 | 3.5 |
| 3 | 1,194 | 1,370 | 874 | 0.516 | -11.2 |
| 4 | 3,205 | 3,797 | 762 | 0.584 | 18.2 |
| 5 | 352 | 213 | 252 | 0.602 | -3.4 |
| 6 | 6,001 | 1,446 | 2,035 | 0.775 | 0.5 |
| 7 | 4,247 | 1,885 | 1,036 | 0.744 | -0.9 |
| 8 | 5,319 | 374 | 1,177 | 0.880 | -1.5 |
| 9 | 1,446 | 1,002 | 852 | 0.609 | 19.8 |
| 10 | 39,212 | 229 | 67,910 | 0.535 | -3.6 |
| 11 | 703 | 2,657 | 243 | 0.327 | -24.1 |

less than the sub-cubes but more than the base. The results of this method fall between those of the base and sub-cube methods. We might expect the large over-segmentation, since the algorithm greedily accepts voxels of similar appearance. We would also expect the results to depend on the selection of the seed points. The details of the implementation also introduced some additional unforeseen impacts on the quality of the segmentation results

Figure 4 shows an example of the segmentation results generated using the different approaches. The first panel shows FLAIR data with some light lesion regions visible. The next three panels show the results from the various methods evaluated against the expert segmentation: the base case, the sub-cube method, and the region-growing method. The red voxels shown across the results represent those missed by the various methods, but identified by the physician: these are the false-negatives. The purple voxels in the results are those not identified by the expert, but flagged as lesion voxels by the corresponding method, i.e., the false-positives. The green voxels are those identified as lesions by the expert and the corresponding method, or the true-positives. The second panel shows the segmentation results using the base processing method. The lesion in the upper left



corner of the image was missed by the base analysis and appears red in this slice. The red voxels reflect the considerable amount of false-negatives visible in this panel. The green shows some correctly segmented voxels, while the few purple spots denote a small amount of false-positives. The third panel in the set shows the evaluation of the results of the sub-cube method for the same patient. The results show a improvement in the coverage of the lesions with some false-negatives in the base case correctly identified by the sub-cube method. The lesion in the upper left, which was not previously identified, has been detected with the sub-cube method. There are some false-negatives shown near the middle of the panel. The combination of increased true-positives and decreased false-negative offsets the additional false-positives, the sub-cube method scores a higher SI compared to the base case. The method also identifies a previously overlooked lesion, representing a significant improvement. The fourth panel shows the results of the region-growing method. Compared to the base case, it shows some of the same improvements as the sub-cube method. Unfortunately, it shows slightly worse performance than the sub-cube method. Some voxels capture by the sub-cubes are missed, giving it a lower SI score. This places the results of the region-growing method slightly lower than the sub-cube method.

Conclusions

In this paper, we propose a new method based simultaneously on local multimodal Markovian segmentation coupled with cellular automata to detect lesion in MRI with multiple modalities, as close as possible as the physician does. We propose a post-processing stage to improve the segmentation results of a process based on Markovian approach. We used two different strategies to improve the whole-brain tissue classification and lesion detection. The first strategy leverages the current processing system at a granularity finer than the whole brain to detect lesions at a local level : the brain MRI cube (1024^3 voxels) is divided into smaller ones (64^3 voxels), where a Markovian segmentation is processed. This reflects the way that a physician considers only a part of the brain at a time. We observed that this approach better captures the local lesion properties and produces encouraging results, with a general improvement in the detection rate of lesions. The second method dives deeper and looks at the individual voxel level. In this case, we are forced to give up Bayesian detection based on Markov assumption due to the weak number of samples observed locally in a narrow neighborhood around the lesion. Just as a physician might look more closely at a lesion, it considers the local neighborhood around a lesion detection. We select seed points from the existing results and use a region-growing method based on cellular automata to grow the lesion areas based on a local neighborhood similarity in intensity. The method does hold promise in finding more accurately identifying the size and shape of lesion detections as a physician does and improving the results of an automated system to expect the lesion load of the whole brain.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

D.B. carried out the software implementation DB and CC participated together to the mathematical design and image processing methods. DB, CC and JPA worked together to interpret the results obtained on the set of patients. All authors read and approved the final manuscript.

Authors' information

<http://miv.u-strasbg.fr/collet>

Acknowledgements

We would like to thank the European Union and program Atlantis for funding mobility between the University of Houston (Computer Sciences Dept), TX, USA and the University of Strasbourg (Télécom Physique Strasbourg - iCube Lab and CNRS), France.

Author details

¹iCube Lab - Université de Strasbourg - CNRS UMR 7357, 300 Bd Sébastien Brant, CS 10413, 67412 Illkirch Cedex, France.

²University of Houston, Computer Sciences Department, PGH 550E, Houston, TX 77204-3010, USA.

Received: 7 February 2013 Accepted: 29 November 2013

Published: 10 January 2014

References

1. Bricq S, Collet C, Armspach J-P: **Unifying framework for multimodal brain MRI segmentation based on hidden Markov chains.** *Med Image Anal* 2008, **2**(6):639–652.
2. Mortazavi D, Abbas Z, Soltanian-Zadeh H, 299-320: **Segmentation of multiple sclerosis lesions in MRI images: a review.** *Neuroradiology* 2012, **54**(4). doi:10.1007/s00234-011-0886-7.
3. Soltanian-zadeh H: **Segmentation of multiple sclerosis lesions in MR images: a review.** *Neuroradiology* 2011, **54**(4):299–320.
4. Vezhnevets V, Konouchine V: **Growcut - interactive multi-label N-D, image segmentation by cellular automata.** *Proc Graphicon*, **2005**:150–156.
5. Aymerich F, Montseny E, Sobrevilla P, Rovira A: **A fuzzy local-based approach for detecting cerebrospinal fluid regions in presence of MS lesions.** In *Proceedings of the 13th International Conference, Part II, IPMU: 2010; Dortmund*.
6. Aymerich F, Sobrevilla P, Montseny E, Rovira A: **Filtering false detections of small multiple sclerosis lesions using fuzzy regional analysis.** In *2010 IEEE International Conference on Fuzzy Systems (FUZZ); July 18–23 2010; Barcelona*:pp. 1–8.
7. Mayer H, Greenspan A: **An adaptive mean-shift framework for MRI brain segmentation.** *IEEE Trans Med Imaging* 2009, **28**(8):1238–1250.
8. de Boer R, Vrooman H, van der Lijn F, Vernooij M, Ikram M, van der Lugt A, Breteler M, Niessen W: **White matter lesion extension to automatic brain tissue segmentation on MRI.** *NeuroImage* 2009, **45**:1151–1161.
9. Abdullah B, Younis PPA, Saraf-Lavi E: **Textural based SVM for MS lesion segmentation in FLAIR MRIs.** *Open J Med Imaging* 2011, **1**:26–42.
10. Fiot J-B, Cohen L, Raniga P, Fripp J: **Efficient lesion segmentation using support vector machines.** In *ViplImage 2011 – III ECCOMAS Thematic Conference on Computational Vision and Medical Image Processing, Paris: May 14-17 2008*.
11. Bricq S, Collet C, Armspach J-P: **Lesion detection on 3D brain MRI using trimmed likelihood estimator and probabilistic atlas.** *IEEE Int Symp Biomed Imaging* 2008:93–96.
12. Bricq S, Collet C, Armspach J-P: **Markovian segmentation of 3d brain MRI to detect multiple sclerosis.** In *15th IEEE International Conference on Image Process. San Diego, CA; 2008*:733–736.
13. Bricq S, Collet C, Armspach J-P: **MS lesion segmentation based on Hidden Markov chains.** In *MICCAI 11th International Conference on Medical Image Computing and Computer Assisted Intervention - paper selected for a grand challenge: 3D segmentation in the clinic: September 6-10 2008; New-York*. http://grand-challenge2008.bigr.nl/proceedings.

doi:10.1186/2194-3990-1-3

Cite this article as: Biediger et al.: Multiple sclerosis lesion detection with local multimodal Markovian analysis and cellular automata 'GrowCut'. *Journal of Computational Surgery* 2014 **1**:3.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com