Classification Forests and Markov Random Field to Segment Chronic Ischemic Infarcts from Multimodal MRI*

Jhimli Mitra[1], Pierrick Bourgeat[1], Jurgen Fripp[1], Soumya Ghose[1], Stephen Rose[1], Olivier Salvado[1], Alan Connelly[2], Bruce Campbell[3], Susan Palmer[2], Gagan Sharma[3], Soren Christensen[3], and Leeanne Carey[2]

[1] Australian e-Health Research Centre, CSIRO, Australia
[2] The Florey Institute of Neuroscience and Mental Health, Australia
[3] Royal Melbourne Hospital, Australia
www.START.csiro.au

Abstract. Accurate identification of ischemic lesions and brain atrophy is critical in the management of stroke patients and may serve as an important biomarker in studying post-stroke depression. In this paper we present an automated method to identify chronic ischemic infarcts in gray matter and gray/white matter partial volume areas that may be used to measure the amount of tissue loss due to atrophy in the area. The measure of tissue loss may then be used as a potential biomarker in analyzing the relation between stroke and post-stroke depression. The automated segmentation method relies on Markov random field (MRF) and random forest based classifications. The MRF classification identifies the possible lesion areas from the fluid attenuated inversion recovery (FLAIR) magnetic resonance (MR) images. Thereafter, the multimodal (T1-, T2-weighted, FLAIR, and apparent diffusion coefficient (ADC)) MR images of the possible lesion areas are fed into the classification forests along with other context-aware features and probability maps of the gray and white matter regions. The results of classification from the MRF and the classification forests are finally combined using connected component analysis to identify the final lesion area. The accuracy of the method in identifying infarcted regions from multimodal MR images has been validated on 17 patient datasets with a mean accuracy of 99%, a mean positive predictive value (PPV) of 75% and a mean negative predictive value (NPV) of 99% and a volume correlation of $r = 0.98$.

Keywords: Ischemic lesion segmentation, classification forest, MRF, MRI.

1 Introduction

The magnetic resonance imaging (MRI) parameters within an ischemic lesion are time dependent and are heterogeneous; i.e., they vary between acute (less than

* The START Research Team.

L. Shen et al. (Eds.): MBIA 2013, LNCS 8159, pp. 107–118, 2013.

7 days) and chronic (3 months) stages. Therefore, it is unlikely that a single MR parameter can characterize the complexity of the cerebral tissue [1]. In clinical practice, diffusion-, T1-, T2-weighted (DWI, T1W, T2W) and fluid attenuated inversion recovery (FLAIR) images are acquired during the progression of stroke. Chronic ischemic lesions appear as hyperintense regions in FLAIR with heterogeneities around the lesions. In comparison, in T1W images the ischemic lesion intensities are hypointense compared to the normal tissue areas.

The fate of some cells in the ischemic area is normally death at the chronic stage while others are salvageable early. The dead cells are usually removed in the process of phagocytosis by the macrophages [2]. These areas are in turn filled with the cerebrospinal fluid (CSF). Figure 1(a) shows a stroke lesion at the chronic stage where parts within the stroke are filled with CSF and hence are hypointense compared to the remaining lesion. Therefore, the CSF areas around the ischemic lesions are usually enlarged compared to the normal part of the cortex. Measuring the CSF volume difference between the ipsilesional volume-of-interest (VOI) around the ischemic area and the contralesional VOI would provide an estimate of the tissue lost due to stroke. The measure of tissue loss associated with the index stroke, as well evidence of atrophy in putative depression regions may be used as a useful biomarker in the study of post-stroke depression [3]. In most cases, ischemic strokes appear in partial volume areas between the gray matter (GM) and the white matter (WM). Therefore, the macrophagic action and hence, filling of CSF is mostly observed in these areas. Although, CSF filling can also occur for pure WM strokes, in this paper we have limited our automatic detection method to the GM areas and considered the WM strokes only if they are connected to the GM strokes. This is beacuse detection of pure WM strokes in FLAIR images becomes extremely difficult in the presence of WM diseases or hyperintensities around the peri-ventricular regions [4], especially when WM strokes have an increased tendency to fuse with the periventricular region WM hyperintensities. Therefore, DWI images of the acute/sub-acute stages are necessary to isolate the stroke area from other areas of WM hyperintensities. Figure 1(b) shows a stroke region and other similar hyperintensities on a FLAIR image. In the available literature, ischemic lesion detection at the acute and sub-acute stages is predominant, rather than detection at the chronic stage. Dastidar et al. [5] proposed a semi-automatic method for segmenting ischemic lesions at the sub-acute stage from T1W and T2W MRI. The method involved amplitude segmentation of intensity histograms followed by region growing and decision trees [6], but required extensive manual intervention at each step to refine the segmentation. High correlation was reported between the extracted lesion volumes using this method and that by an expert. Jacobs et al. [7] applied the ISODATA (Iterative Self-organizing Data) algorithm [8] on integrated images comprising of T1W, T2W and DWI MR images to segment ischemic lesions at 3 time-points. The lesion volume extracted by the automated method showed high correlation with either the T2W or DWI volume across different time-points. Kabir et al. [9] suggested an automatic segmentation method based on MRF from multimodal MRI images comprising of DWI, T2W and FLAIR. The primary objective of the

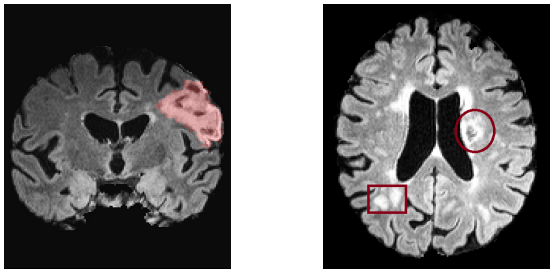

(a) (b)

Fig. 1. Intensity heterogeneities in stroke lesions at the chronic stage and white matter hyperintensities in FLAIR. (a) Heterogenous intensities inside chronic stage stroke lesion due to CSF filling. The light red mask denotes the entire stroke area as marked by an expert, (b) anomalous FLAIR hyperintensities. The circular region in the image depicts a stroke region while the square area contains other white matter hyperintensities.

method was to categorize the stroke sub-types at the acute, sub-acute and later stages by registering with the blood supply territories atlas, therefore no quantitative results were presented related to the accuracy of the stroke detection. Mean-shift procedure on edge-confidence maps was proposed by Hevia-Montiel et al. [10] with subsequent region adjacency analysis and pruning to segment acute ischemic lesions from DWI images. The method showed a high average volume correlation with a high variance when compared to the manually segmented lesion volumes. Shen et al. [11] used extended fuzzy C-means to include the prior GM, WM, and CSF probability maps in the objective function. Finally Bray Curtis distance [12] between the fuzzy membership and prior probability of each voxel was used to isolate lesions from normal voxels. Qualitative results on T1W images showed good agreement of the method with an expert observation.

The methods in the literature show good accuracy in terms of volume correlation and visual agreement, although some methods required manual intervention and some required threshold values to be set at some stages. Most of the available methods deal with DWI images of the acute stage showing hyperintense stroke regions and can be easily targeted to isolate the ischemic stroke from other brain diseases. It appears that the methods in the literature also do not deal with the challenges of ischemic stroke segmentation in the presence of WM hyperintensities in FLAIR. We have designed our automated segmentation method, motivated by some of the works in literature and Zikic et al. [13] to: a) identify the location of the lesion with high confidence (without DWI and avoiding known WM hyperintensities); b) extract the extent of the ischemic lesion and c) find out possible focal atrophies other than known WM hyperintensities that may have significant impact on post-stroke impairment and recovery. These are achieved by: 1) a two-stage MRF classification of the probable lesion areas; 2) using random forests on the intensities of the probable lesion areas; and finally 3) applying an ad-hoc rule to refine the results iterating between random forests and MRF classifications using connected component analysis. To the best of our

knowledge this is the first attempt to apply random forests for stroke lesion segmentation. The paper is organized as follows: Section 2 provides the detail of the multimodal patient data and the preprocessing steps, Section 3 provides the description of the MRF classification using Iterated Conditional Modes (ICM), the random forests method with the feature-space and the ad-hoc rule for refinement of the segmentation in its respective subsections, and finally Section 4 and 5 deal with the results, discussions and conclusions.

2 Data and Preprocessing

A total of 20 patient datasets of size $181 \times 217 \times 181$ voxels are used in our experiments. All patients have stroke lesions and T1W, T2W, FLAIR and ADC (apparent diffusion coefficient) maps of the chronic stage are available. The T1W images are acquired with TE=2.55ms, TR=1900ms, flip angle=9°, voxel spacing=$(1 \times 1 \times 1)$mm; the T2/FLAIR volumes are acquired with TE=388ms, TR=6000ms, T1=2100ms, flip angle=120° with isotropic 1mm voxels. The ADC maps are generated from the diffusion weighted images with B=1000 in 25 directions with isotropic voxel spacing of 2.5mm. The voxels sizes for all the modalities are converted to 1mm isotropic voxels for our experiment. All patient data are rigidly registered to the (Montreal Neurological Institute) MNI atlas. Bias correction for the patient MRIs is done using the method of Van Leemput et al. [14] that also provides the segmented masks of GM, WM and CSF. Skull stripping is done by applying the GM, WM, CSF segmented masks to remove the remaining part of the brain. The probabilistic GM and WM estimates for each patient are obtained by non-rigidly aligning [15] the GM/WM masks of other patients to the target patient and then averaging the aligned masks. This method ensures that the effect of atrophies that are likely to be present in the individual patient tissue masks is minimized.

3 The Proposed Method

This section explains the different steps involved in our method. Figure 2 shows the schematic diagram of our method. The MRF classification from the FLAIR image provides the possible lesion areas. The multimodal MRI images of these areas are then used as inputs to the random forests training and classification, and finally connected component analysis with ad-hoc rules are applied on the random-forests and MRF classifications to refine the segmentation results. The rationale behind using a part of the brain i.e. possible lesion areas as input to the random forest as opposed to the sub-sampled entire brain is to reduce the overhead of computational complexity and to avoid missing smaller lesions in the random forests training phase. Sections 3.1 and 3.2 provide terse theoretical backgrounds of MRF and random forests and how each classification method is applied in our experiment. Finally, Section 3.3 describes the connected component analysis rules used to refine the segmentation results.

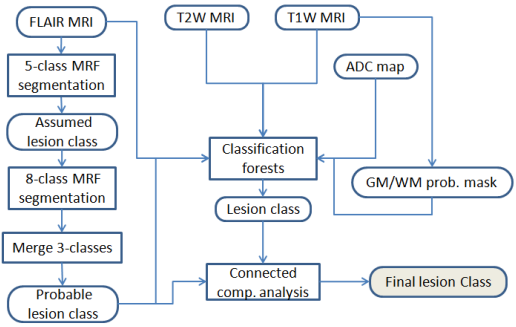

Fig. 2. A schematic diagram of the proposed method for ischemic stroke lesion segmentation

3.1 Markov Random Field (MRF) Segmentation

Given $\mathcal{F} = \{f_s\}_{s \in S}$ a set of image data where f_s denotes the gray value at pixel s. A segmentation problem is to find the labeling $\hat{\omega}$ which maximizes $P(\omega|\mathcal{F})$. Bayes theorem suggests that $P(\omega|\mathcal{F}) = \frac{P(\mathcal{F}|\omega)P(\omega)}{P(\mathcal{F})}$. Actually $P(\mathcal{F})$ does not depend on the labeling ω and we have the assumption from conditional probability that $P(\mathcal{F}|\omega) = \prod_{s \in S} P(f_s|\omega_s)$. The MRF on each of the pixel s is defined by the clique potentials and by a neighborhood-system $\mathcal{G}$. Let C denote a clique of $\mathcal{G}$, and $\mathcal{C}$ the set of all cliques. The restriction of the label ω to the site of a given clique C is denoted by ω_C. The clique potentials $E_C(\omega_C)$ for every C in $\mathcal{C}$ and every $\omega \in \Omega$, where Ω is the set of all possible discrete labellings. Following the definition of MRF and the Bayes theorem, the global labeling which we are trying to find is given by:

$$\hat{\omega} = \max_{\omega \in \Omega} \prod_{s \in S} P(f_s|\omega_s) \prod_{C \in \mathcal{C}} \exp(-E_C(\omega_C)). \tag{1}$$

The energies of cliques of order 2 and more reflect the probabilistic modelings of labels taking the neighborhood-system into account, i.e. the pixels would be labeled considering the neighborhood labels. Let us assume that $P(f_s|\omega_s)$ is Gaussian, the discrete class λ is represented by its mean value μ_λ and its deviation σ_λ. We get:

$$\hat{\omega} = \min_{\omega \in \Omega} \left(\sum_{s \in S} \left(\log(\sqrt{2\pi\sigma_{\omega_s}}) + \frac{(f_s - \mu_{\omega_s})^2}{2\sigma_{\omega_s}^2} \right) + \sum_{C \in \mathcal{C}} E_C(\omega_C) \right) \tag{2}$$

Using the above equation, the local energy of any labeling ω at site s would be:

$$\varepsilon_s(\omega) = \log(\sqrt{2\pi\sigma_{\omega_s}}) + \frac{(f_s - \mu_{\omega_s})^2}{2\sigma_{\omega_s}^2} + \sum_{C_s \in \mathcal{C}} E_{C_s}(\omega_{C_s}). \tag{3}$$

The estimation of $\hat{\omega}$ is done through the energy minimization using the ICM algorithm [16]. The key to the ICM method is that the probability of the label

at pixel s given the observed image and the current estimates ω_C of all the labels in the neighborhood $\mathcal{G}$ of pixel S. Therefore, the ICM method requires an initial estimate of the labels. This initial estimate is obtained from a simple Bayesian classification of the Gaussian likelihood estimates of the classes based only on the image intensities and uniform probability priors.

The first stage of the MRF classification on the FLAIR MRI considers 5 tissue classes as the background, WM, GM, CSF and a possible lesion class. The segmentation of the lesion class at this stage does not only contain the lesion area but also hyperintense GM areas. To increase the separability of the hyperintense lesion areas from the hyperintense GM areas we further apply a MRF classification on the FLAIR intensities of the lesion class labels from the first stage. This second stage of classification attempts to separate the possible lesion class into 8 classes to maximize the separation between the Gaussian density functions. As mentioned in the introduction, ischemic stroke lesions have intensities dropping around the edges, therefore a merging of the labeled components from the 3 classes with highest mean values ensures that a maximum of the lesion area is included in the segmentation. Nevertheless, hyperintensities in the periventricular regions, basal ganglia and the hippocampal regions similar to the stroke areas, still incorporate errors in the MRF segmentation process. Therefore, the intensities of the merged possible lesion class are further passed into classification forests described in the following section.

3.2 Random Forests (RF) Classification

In this method, standard classification forests based on spatially non-local features combined with probabilistic tissue estimates of GM (P_{GM}) and WM(P_{WM}) are used. The MRI input channels used in the classification are FLAIR, T1W, T2W and ADC maps. Additionally, the labeled probable tissue class (L_{les}) is also used as a channel input. Therefore an input channel comprises of $C_j = (I_{FLAIR}, I_{T1}, I_{T2}, I_{ADC}, P_{GM}, P_{WM}, L_{les})$.

In our method we use 3 spatial and context-aware features [13] along with the voxel position and its respective intensity set. Let us consider $x \in \mathcal{X}$ is a spatial point, to be assigned a class (lesion/background) and I_j is an input channel. $N_j^s(x)$ denotes a 3D neighborhood with edge lengths $s = (s_x, s_y, s_z)$, $\upsilon \in \mathbb{R}^3$ is an offset vector. The contextual features are as follows: 1) $I_{j1}(x) - I_{j2}(x + \upsilon)$, where $I_{j1} = I_{j2}$ is allowed; 2) difference of cuboid means, i.e. $\mu(N_{j1}^{s1}(x)) - \mu(N_{j2}^{s2}(x + \upsilon))$, where $j1 \neq j2$; and 3) this feature assumes that lesions usually appear as partial volumes of GM and WM areas and CSF borders the other side of the lesion, therefore the difference of intensity along a 3D line as $\max_\lambda(I_j(x + \lambda\upsilon)) - \min_\lambda(I_j(x + \lambda\upsilon))$ with $\lambda \in [0, 1]$.

In classification forests training, each tree t in the forest receives the full training set V, along with the label at the root node and selects a test along randomly chosen dimensions of the feature space to split V into two subsets in order to maximize the information gain. The left and the right child nodes receive their respective subsets V_i and the process is repeated at each child node to grow the next level of the tree. Each node in a decision tree also contains a class predictor

$p_t^i(c|x)$, which is the probability corresponding to the fraction of points with class c in V_i. Growth is terminated when either information gain is minimum or the tree has grown to maximum depth. At testing, a point to be classified is pushed through each tree t by applying the learned split functions. When reaching a leaf node l, posterior probabilities $p_t^l(c|x))$ are gathered in order to compute the final posterior probability of the voxel defined by $p(c|x) = \frac{1}{n}\sum_{t=1}^{n} p_t^l(c|x)$. The actual class estimate is chosen as the most probable class by $\arg\max_c p(c|x)$.

3.3 Ad-hoc Connected Component Analysis

In this step, we firstly remove the periventricular areas and the WM areas from the MRF classification results by applying a non-rigidly registered [15] mask of the automated anatomical labeling (AAL) atlas [17] that models the GM. This means that the lesion areas falling under the zero labels (WM) and the labels 71/72 (periventricular regions) of the registered atlas are ignored. Since the lesion areas are much smaller compared to the whole brain, the input to the RF are probable lesion areas from the MRF classification to avoid unbalanced data. Nonetheless, the probable lesion areas are much larger than the actual lesion area. Therefore, RF provides sparse classification of the true positive areas while managing to eliminate the false positives in most of the cases. Therefore, we applied a pruning step to remove more than 10 connected voxels (analyzing 8-connectivity) and then a morphological closing operator with a circular mask of radius 3. Isolated lesion components from RF may be parts of the same lesion which may be verified from the MRF classification. The components of the RF classification that have more than ϵ percent overlap with the corresponding components in MRF are retained and the others are removed as outliers. In the final output, the component from MRF classification is replaced for these lesions. It may be the case, when the resulting component is less than ϵ percent of the volume of the corresponding RF component, then the larger component from the RF is replaced in the final output. This step further ensures that the WM lesion areas that are connected with the GM areas are also included in the final lesion volume and only the pure WM strokes are eliminated. In summary, the iterative approach ensures that maximum lesion areas either from the MRF or RF are included. Finally, a morphological closing operation of radius 2 ensures that some of the hypointense regions embedded inside the hyperintense regions are mostly included in the lesion volume.

4 Evaluation and Discussions

A total of 20 patient datasets are used in the training and validation procedure. Three patients have pure WM strokes; although we have included these data in the MRF and RF classification stages, they are automatically removed in the connected component analysis step when the WM areas and periventricular areas are removed from the MRF classification using the AAL atlas. Therefore, the final lesion results are evaluated only on 17 patients. The RF training/testing

Table 1. Quantitative results for ischemic lesion detection using MRF and RF. DSC (Dice similarity coefficient) is in fraction, PPV (positive predictive value), NPV (negative predictive value) and ACC (accuracy) are in %. μ is the average and σ is the standard deviation.

Patient#	DSC	PPV	NPV	ACC
1	0.71	87.46	99.92	99.90
2	0.75	97.78	99.90	99.89
3	0.55	51.32	99.98	99.96
4	0.33	31.23	99.99	99.97
5	0.81	94.35	99.86	99.84
6	0.44	73.37	99.98	99.98
7	0.73	88.40	99.89	99.87
8	0.44	40.66	99.99	99.97
9	0.65	79.38	99.83	99.77
10	0.78	88.38	100.00	100.00
11	0.34	35.15	99.99	99.98
12	0.46	93.76	99.97	99.97
13	0.73	94.14	99.76	99.74
14	0.36	53.90	99.93	99.91
15	0.58	81.77	99.96	99.96
16	0.62	85.97	99.98	99.98
17	0.70	95.57	99.88	99.87
μ	0.59	74.86	99.93	99.92
σ	0.16	22.95	0.07	0.08

are carried out with 100 full-grown trees in a leave-one-patient-out manner. The offset vector υ in Section 3.2 is empirically chosen as 8mm and the 3D neighborhood is restricted to $3 \times 3 \times 3$ voxels. The value of ϵ for the connected component analysis of Section 3.3 is empirically learned from the datasets as 20. An expert neurologist segmented the ischemic lesions on the FLAIR images. Table 1 shows the quantitative results in terms of Dice similarity coefficient (DSC)= $2 \times TP/(FP + 2 \times TP + FN)$, i.e. the fraction of overlap between the manual segmentation and the automated segmentation method showing a mean of 0.59 ± 0.16; positive predictive value or the precision rate (PPV)=$TP/(TP + FP)$ denotes the proportion of true positives i.e. a high PPV would indicate that a patient identified with a lesion does actually have the lesion. Negative predictive value (NPV)=$TN/(TN + FN)$ denotes the proportion of negative results in the test, i.e. a high NPV indicates that the method rarely misses the lesion; while accuracy (ACC)=$(TP + TN)/(TP + TN + FP + FN)$ measures the degree of closeness to the actual measurements. Here TP ='true positives', TN ='true negatives', FP ='false positives' and FN ='false negatives'. The volume correlation between the manually segmented volume and that by the automated method is $r = 0.98$. The DSC values show a maximum high value of 0.81 and a low value of 0.33. This variation is due to several reasons, for e.g., the expert neurologist marked the lesions which were deemed to be stroke,

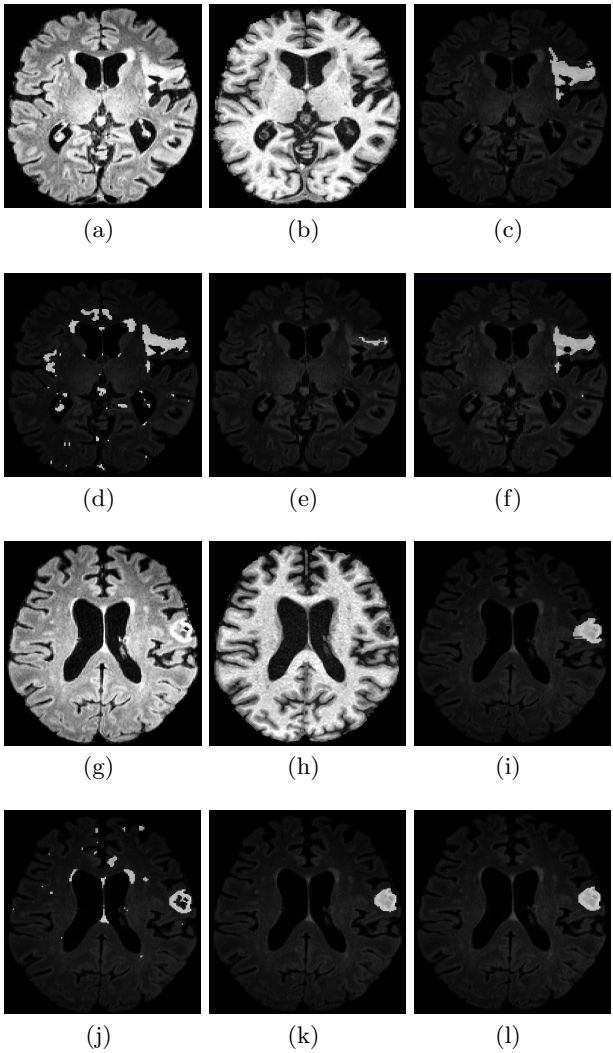

Fig. 3. Classification results from MRF, RF and connected component analysis. (a) to (f) are the results corresponding to patient 1 and (g) to (l) are the results corresponding to patient 2. (a) & (g) are the FLAIR images that show hyperintensities for stroke lesions and (b) & (h) are the T1W images showing hypointense regions of stroke. The manual segmentations by the expert radiologist are shown in (c) & (i); segmentations from MRF (merged 3-classes) are shown in (d) & (j); segmentations from RF classifications are shown in (e) & (k) and the final lesions after ad-hoc connected component analysis are shown in (f) & (l) for patients 1 and 2 respectively.

but the patients might actually contain other hyperintense areas other than WM diseases which are extracted by our algorithm and may eventually cause atrophies that are important for the study of post-stroke recovery and depression.

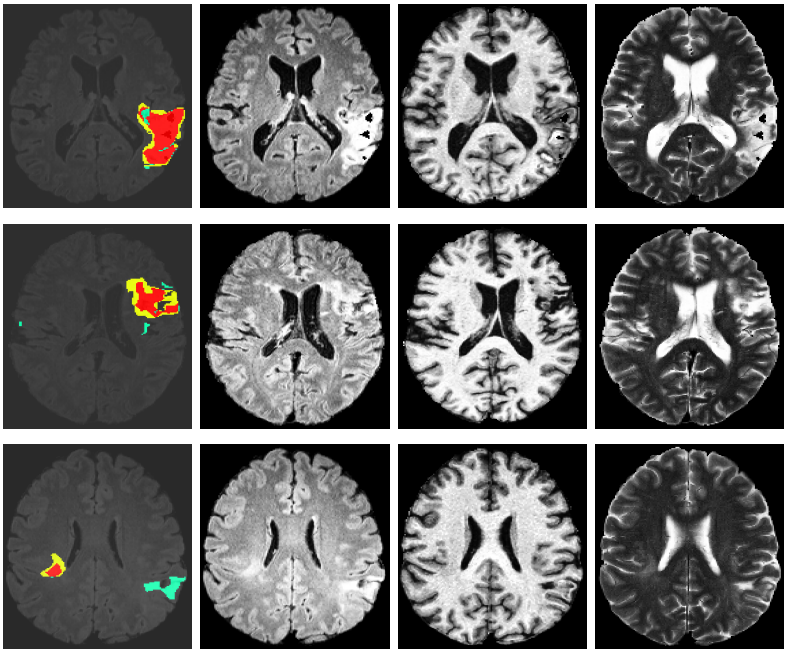

Fig. 4. Qualitative results of ischemic lesion segmentation. Rows correspond to patients 5, 9 and 14. The 1^{st} column shows the overlap image and the 2^{nd}, 3^{rd} and 4^{th} represent the corresponding FLAIR, T1W and T2W images. 'Red' denotes TP, 'yellow' denotes FN and 'green' denotes FP.

The morphological operation helps to fill-up the CSF areas within the stroke in most cases, while if the areas are too large then the uniform morphological operator mask radius used is not good enough to approximate the entire stroke area. Therefore, our algorithm provides an underestimation of the stroke area and hence poor overlap measures for those patients. The post-processing of the lesion areas with connected-component analysis and morphological processing is data dependant and therefore in future we would like to train our model of classification forest with more patient datasets and would be able to remove the ad-hoc step. Nevertheless, at this stage with a limited number of datasets, the ad-hoc step is required to obtain clinically meaningful results.

Figure 3 shows the classification results for patients 1 and 2 at each stage of the proposed method. Figures 3(d) shows the result after 3-classes with highest means are merged from the 8-class MRF classification, while 3(e) shows the results from the classification forest. It is observed that a significant part of the lesion is missing from the RF classification while in the final output (Figure 3(f)), the connected component analysis manages to extract a larger part of the lesion that is similar to the manual segmentation as in Figure 3(c). In the case of patient 2, Figure 3(j) shows a missing part of the lesion from MRF classification (being hypointense in FLAIR), while Figure 3(k) shows the complete lesion after

RF classification followed by morphological closing operation, therefore the final output in Figure 3(l) from connected component analysis shows a lesion area conforming to the manual segmentation of Figure 3(i).

Figure 4 shows the qualitative results of patients 5, 9 and 14 which show very high, medium and low overlap areas with the manual segmentation. The TP areas are shown in red, the FN areas are shown in yellow and the FP areas are shown in green. Observing Figure 4 it is evident that the results of our method show high overlap for patient 5 (row 1). The last row in the figure shows false positive areas in green which are hyperintense in FLAIR and T2W MRI and hypointense in T1W MRI. This suggests that our method is capable of extracting ischemic stroke lesions and some added lesions which might have been caused by some other disease but may have significant structural and functional impact on the patient brain. Therefore identifying such lesions is also important for the post-stroke depression study.

All our implementations are done in C++ and ITK platform on a 6-core CPU of 3.2GHz with 23.5GB of memory. The MRF classification requires 6min−7min on an average for each patient, the RF training requires an average of 2 hours to train in leave-one-patient-out manner, while each testing requires 5min. The connected component analysis also takes 10s−15s on average.

5 Conclusions

The method described in this paper deals with a two-stage classification based on MRF and classification forests to identify chronic ischemic stroke lesions in the human brain. The first stage results from the MRF are used as inputs to the classification forests as well as help reduce misclassified and missing lesion areas obtained from the classification forests. The method shows good accuracy in identifying the lesion areas which is useful for the post-stroke depression study. Since, our method is also capable of identifying most of the hyperintense areas on FLAIR images, we would like to use the DWI images of the acute/sub-acute stages to discriminate between strokes (both GM and WM lesions) and other WM hyperintensities which would provide additional biomarkers consistent with the hypotheses in the post-stroke depression study [3]. We would like to validate our method with more stroke patients in future and would probably not require the ad-hoc connected component analysis step if the RF is trained with more patients which would increase the generalization power of classification forests.

Acknowledgments. We would like to acknowledge the Stroke Imaging Prevention and Treatment (START) program of research which is supported in part by the CSIRO of Australia through the Preventative Health Flagship Cluster, the National Health and Medical Reserch Council of Australia, and a Victorian Government Operational Infrastructure Support Grant. In particular, we wish to acknowledge the stroke patients, radiologists and START researchers who contributed to the data collected for this study. Leeanne Carey is supported by an Australian Research Council Future Fellowship [number FT0992299]. The funding sources had no role in conduct of the study or writing of the report.

References

1. Baird, A.E., Warach, S.: Magnetic resonance imaging of acute stroke. J. Cereb. Blood Flow Metab. 18, 583–609 (1998)
2. Mosser, D.M., Edwards, J.P.: Exploring the full spectrum of macrophage activation. Nature Reviews Immunology 8, 958–969 (2008)
3. Carey, L., et al.: START-PrePARE PREdiction and Prevention to Achieve Optimal Recovery Endpoints after stroke: Study rationale and protocol. In: 7th World Cong. of Neurorehab. (2012)
4. Campbell, B.C., et al.: Assessing response to stroke thrombolysis: validation of 24-hour multimodal magnetic resonance imaging. Arch. Neurol. 69(1), 46–50 (2012)
5. Dastidar, P., et al.: Volumetric measurements of right cerebral hemisphere infarction: use of a semiautomatic MRI segmentation technique. Comput. in Biol. & Med. 30, 41–54 (2000)
6. Heinonen, T., et al.: Semiautomatic tool for segmentation and volumetric analysis of medical images. Med. Biol. Eng. Comput. 36, 291–296 (1998)
7. Jacobs, M.A., et al.: Multiparametric MRI tissue characterization in clinical stroke with correlation to clinical outcome: Part 2. Stroke 30, 950–957 (2001)
8. Jacobs, M.A., et al.: A model for multiparametric MRI tissue characterization in experimental cerebral ischemia with histological validation in rat: Part 1. Stroke 32, 943–949 (2001)
9. Kabir, Y., et al.: Multimodal MRI segmentation of ischemic stroke lesions. In: IEEE EMBS, pp. 1595–1598 (2007)
10. Hevia-Montiel, N., et al.: Robust nonparametric segmentation of infarct lesion from diffusion-weighted MR images. In: IEEE EMBS, pp. 2102–2105 (August 2007)
11. Shen, S., et al.: An improved lesion detection approach based on similarity measurement between fuzzy intensity segmentation and spatial probability maps. Mag. Reson. Imag. 28, 245–254 (2010)
12. Bray, J.R., Curtis, J.T.: An ordination of the upland forest communities of southern wisconsin. Ecol. Monogr. 27, 325–349 (1957)
13. Zikic, D., et al.: Decision forests for tissue-specific segmentation of high-grade gliomas in multi-channel MR. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part III. LNCS, vol. 7512, pp. 369–376. Springer, Heidelberg (2012)
14. Van Leemput, K., et al.: Automated model-based tissue classification of MR images of the brain. IEEE Trans. in Med. Imag. 18(10), 897–908 (1999)
15. Modat, M., et al.: Fast free-form deformation using graphics processing units. Comput. Meth. Prog. Biomed. 98(3), 278–284 (2010)
16. Besag, J.: On the statistical analysis of dirty pictures. J. R. Statist. Soc. Ser. B 48, 259–302 (1986)
17. Tzourio-Mazoyer, N., et al.: Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brains. NeuroImage 15(1), 273–289 (2002)