A Modality-Agnostic Patch-Based Technique for Lesion Filling in Multiple Sclerosis

Ferran Prados[1,2], Manuel Jorge Cardoso[1,3], David MacManus[2], Claudia A.M. Wheeler-Kingshott[2], and Sébastien Ourselin[1,3]

[1] Translational Imaging Group, CMIC, UCL, NW1 2HE, UK
[2] NMR Research Unit, UCL Institute of Neurology, WC1N 3BG, UK
[3] Dementia Research Centre, UCL Institute of Neurology, WC1N 3BG, UK

Abstract. Multiple Sclerosis lesions influence the process of image analysis, leading to tissue segmentation problems and biased morphometric estimates. With the aim of reducing this bias, existing techniques fill segmented lesions as normal appearing white matter. However, due to lesion segmentation errors or the presence of neighbouring structures, such as the ventricles and deep grey matter structures, filling all lesions as white matter like intensities is prone to introduce errors and artefacts. In this paper, we present a novel lesion filling strategy based on in-painting techniques for image completion. This technique makes use of a patch-based Non-Local Means algorithm that fills the lesions with the most plausible texture, rather than normal appearing white matter. We demonstrate that this strategy introduces less bias and fewer artefacts and spurious edges than previous techniques. The advantages of the proposed methodology are that it preserves both anatomical structure and signal-to-noise characteristics even when the lesions are neighbouring grey matter and cerebrospinal fluid, and avoids excess blurring or rasterisation due to the choice of segmentation plane, and lesion shape, size and/or position.

1 Introduction

Multiple Sclerosis (MS) is an immune-mediated demyelinating disease affecting both white matter (WM) and grey matter (GM). It is characterised pathologically by areas of inflammation, demyelination, axonal loss, and gliosis scattered throughout the central nervous system. White matter plaques are relatively easy to detect using current conventional MRI techniques, whereas grey matter lesions can be observed on double inversion recovery MRI [6]. MS plaques that correspond to necrotic lesions appear on T1-weighted sequences as areas of low-signal intensity compared with normal appearing white matter (NAWM), and the active lesions are displayed with hyperintense signals [10]. From an image processing perspective, these MS lesions influence tissue segmentation procedures, resulting in the misclassification of GM and WM. Furthermore, other studies have suggested that MS lesions may affect the estimation of segmentation parameters, resulting in a shift of tissue boundaries [4] and influencing subsequent morphometric studies. Thus, there is a clear need to reduce the negative impact that MS lesions have on image analysis procedures.

P. Golland et al. (Eds.): MICCAI 2014, Part II, LNCS 8674, pp. 781–788, 2014.

Various techniques have been developed in recent years based on the concept of in-painting T1 images [11,4,2]. In short, the process of T1 lesion in-painting is based on filling a WM lesion region of interest (ROI) with synthetic estimates of WM-like intensities. Under the assumption that WM lesions are the ones that most affect morphometric studies (i.e. GM lesions have a minor influence), the process of lesion in-painting reduces the overall algorithmic bias. More specifically, Sdika and Pelliter [11] presented three different in-painting algorithms. The first, denoted *basic in-painting* and inspired by [12], consists of filling the lesion ROI in an inner-radial manner using a Gaussian kernel average $3 \times 3 \times 3$ of the neighbouring intensities. The second, denoted *local white matter in-painting* (LWMI), uses *a priori* information obtained from an image segmentation technique to iteratively fill the border of the lesions using a Gaussian kernel. Finally, the *global white matter in-painting* (GWMI) method fills the MS lesions with the mean intensity of the normal WM over the whole brain, meaning that all lesions will have the same intensity regardless of their neighbourhood.

Later, Chard et al. [4] developed the LEAP (LEsion Automated Preprocessing) technique, with the aims of: filling lesions as normal WM; reproducing the WM noise characteristics; and avoiding operator intervention.

Subsequently, Battaglini *et al.*[2] presented a similar method based on replacing the lesion voxel intensities with values that are randomly sampled from an intensity distribution that is measured from the surrounding WM and GM voxels. This method is available as part of FSL [8]. In short, regardless of their approach, the previously presented algorithms have been restricted to images of a specific modality, require accurate lesion segmentations in boundary regions (e.g periventricular lesions), can create shape gradients around the lesion ROI, and are prone to errors in model fit when estimating the WM distribution properties.

More recently, Guizard *et al.*[7] calculated the most similar patches using only the surrounding regions after prefilling the lesions with the average of the intensities of the immediately surrounding healthy tissue. This prefilling completely biases the final results because the whole patch is used for calculating the distance between the surrounding patch and the lesion patch.

In the field of computer graphics, structurally aware in-painting algorithms are common, with many of these algorithms permitting a user to simply erase an unwanted portion of an image without any previous knowledge about its composition. These techniques attempt to fill regions by synthesising plausible textural matches from the remainder of the image [5,9,1]. In doing so, these algorithms, commonly used for scratch removal, photo restoration, and object/text removal, are agnostic to the structure of the input image. The most successful techniques for in-painting in computer graphics, here denoted as exemplar-based methods, attempt to fill the unknown ROI by simply copying content from the observed part of the images [9] under some constraints. This class of methods commonly divide the image into a large number of small sub-images, or patches, followed by either a patch-search method [5], or the use of the Non-Local Means algorithm [3]. Finally, the intensities can be synthesised using either pixel- or patch-based textures from the most similar patch.

In this work, we formulate a task-specific patch-search algorithm for the purpose of filling MS lesions. The proposed algorithm presents two main advantages: first, due to its general formulation, the proposed algorithm is able to fill any type of MR image modality that has a non-local structure. Furthermore, due to its contextual nature, the proposed algorithm is also more robust to over-segmentation of the lesion ROI, thus reducing the accuracy requirements when manually or automatically defining the in-painting region of interest.

2 Method

The proposed lesion filling technique can be described in three main steps: (1) estimating the patch with the most similar neighbourhood structure, (2) synthesising the intensity from the best patch, (3) followed by a buffing step through the application of a minimal kernel-based convolution over the filled region.

First, we assume that we have a grayscale-valued image I^*, $X \times Y \times Z$, previously corrected for intensity inhomogeneity, and a well-defined lesion mask $\mathcal{L}$. The filled image I can then be defined as $I(p) = I^*(p)\ \forall p \notin \mathcal{L}(p)$, and as $I(p) = \mathcal{F}(p)\ \forall p \in \mathcal{L}(p)$, where p denotes the voxel location (x, y, z) in the image I and $\mathcal{F}(p)$ is the function that synthesises the intensity of voxel p. We define Ω_p as a search region of size W^3 voxels around voxel p. Within the region Ω, we define a cubic target patch $T(p)$ of size w^3 voxels, centred at a voxel p which is on the boundary of the lesion, and a search patch $S(q)$ of size w^3 voxels, centred in q, with $w < W$ and $q \in \Omega$ and $q \notin \mathcal{L}$.

Given w and W, we propose to replace (or fill) the voxel intensity $I^*(p)$ with the intensity $I^*(q)$ if $S(q)$ is the most similar patch to $T(p)$, under the constraint that q is within the search region Ω, outside the lesion region $\mathcal{L}$ and that $q \neq p$. Formally, a temporary estimate $\hat{I}(p)$ for all $p \in \mathcal{L}$ can be generated by finding $\hat{I}(p) = I(\hat{q})$ with:

$$\hat{q} = \underset{\forall q \in \Omega | (q \neq p) \wedge (q \notin \mathcal{L})}{\arg\min} \mathcal{D}\left(T(p), S(q)\right) \tag{1}$$

where the distance $\mathcal{D}$ between two patches T and S is equal to

$$\mathcal{D}(T(p), S(q)) = \frac{\sum_{i \in T(p) \wedge j \in S(q) | \{i,j\} \notin \mathcal{L}} (I(i) - I(j))^2}{\kappa^2} \tag{2}$$

Here, where κ is the cardinality of the set $\{i \in T(p) \wedge j \in S(q) \mid \{i, j\} \notin \mathcal{L}\}$, i.e. the number of voxels within the patches $T(p)$ and $S(q)$ that are not in the lesion region. Note that while the denominator κ^2 favours patches with more information, an extra hard constraint is necessary to avoid situations where only a few voxels have matching intensities. Thus, by defining α as the minimum required percentage of the patch size we can formally define the hard constraint $\kappa > \alpha w^3$ (i.e. the cardinality of the set $\{i \in T(p) \wedge j \in S(q) \mid \{i, j\} \notin \mathcal{L}\}$ has to be more than $\alpha\%$ of the patch size). If this constraint is satisfied, then the p is removed from the set $\mathcal{L}$, otherwise, p remains in $\mathcal{L}$. This process is repeated

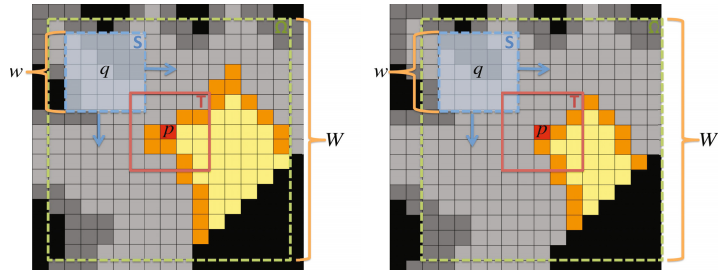

Fig. 1. Algorithm schema that illustrates two different iterations, with the voxels in orange denoting the ones respecting the hard constraint $\kappa > \alpha w^3$. In short, the proposed method moves the search patch S, centred at q, within Ω, until it finds the location $\hat{q}$ where $S(\hat{q})$ and $T(p)$ are the most similar.

until $\mathcal{L} = \emptyset$, i.e., every voxel p initially in $\mathcal{L}$ has an estimate $\hat{I}(p)$. This iterative process results in an inwards filling of the voxels in $\mathcal{L}$ as depicted in Fig. 1.

Finally, when $\mathcal{L} = \emptyset$, $\mathcal{F}(p)$ can then be estimated by buffing the estimates of $\hat{I}(p)$ using a convolution operation $\mathcal{F} = \mathcal{C} * \hat{I}$, where $\mathcal{C}$ is a minimal 6-neighbourhood clique (cross-shape) kernel with its centre voxel set to 1 and all other voxels set to $\mathcal{K}$. The kernel $\mathcal{C}$ is then normalised so that the kernel density sums to 1.

3 Validation

The proposed method has been evaluated both qualitatively and quantitatively. All our experiments use the following untuned empirically defined parameters: $W = 21$ for the search region Ω, $w = 5$ for the T and S patch size, $\alpha = 0.1$, and $\mathcal{K} = 0.4$. For the sake of comparison, two publicly available lesion filling algorithms, the method by Chard *et al.* [4], here denoted as ION, and the method by Battaglini *et al.* [2], denoted as FSL, were used for comparison purposes.

The proposed method has been applied over two datasets. The first dataset comes from the public database *BrainWeb* (`http://www.bic.mni.mcgill.ca/brainweb/`) and it is used for a qualitative analysis only. The second dataset is composed of 104 patients with secondary progressive MS (age range: $30 - 61$ years) and it is used for a qualitative and quantitative analysis. Each patient was scanned at baseline and at 24 months, resulting in a total of 208 scans. The MRI data used here was collected using a single 1.5-T MRI scanner (General Electric, Milwaukee, WI, USA) and analysed (quality control and manual lesion segmentation) by two trained raters. Appropriate quality assurance procedures, involving regular scanning of control subjects with no known neurological deficit and phantoms, were undertaken in keeping with departmental policy. The following sequences were acquired: 2D T1W Spin Echo (SE) (TE=15ms,TR=550ms, in-plane pixel spacing: 0.9375×0.9375 mm, out-of-plane: 3 mm), T2W Dual Fast SE (TE=20ms and 80ms,TR=2500ms,voxel size: $0.9375 \times 0.9375 \times 3$ mm) and 3D T1WGE (TE=5ms,TR=15ms,TI=450ms, $0.976 \times 0.976 \times 1.5$ mm).

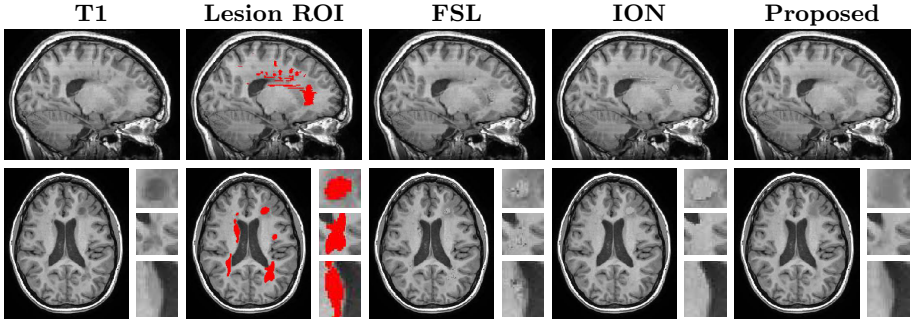

Fig. 2. The coronal, axial and zoomed views of the original T1 image and the lesions mask (LROI), followed by the results with FSL, ION and the proposed methodology. Note the introduction of WM-like intensity in the missegmented caudate region (3^{rd} zoomed view) using the FSL and ION methods.

3.1 Qualitative Analysis

In this evaluation, we compare the different methods for filling the lesions on a MS patient in two situations: with $\mathcal{L}$ defined as the manually-segmented lesion mask (LROI), and with a dilated version of the same mask (DROI). Fig. 2 and 3 show the results obtained using the original mask and the dilated mask respectively. The results show that the proposed method not only preserves better the boundaries of the underlying neighbouring structures (ventricles and WM/GM boundary), but also reduces artefacts and spurious rasterisation due to lesion shape, size and position and due to the choice of imaging plane for manual segmentation. Furthermore, as the proposed method is context aware, it is also able to cope with situations when the human rater erroneously segments a non-pathological region of interest, e.g. the third zoomed region in both Fig. 2 and 3 shows that the caudate nucleus was mislabelled as an MS lesion. This structure

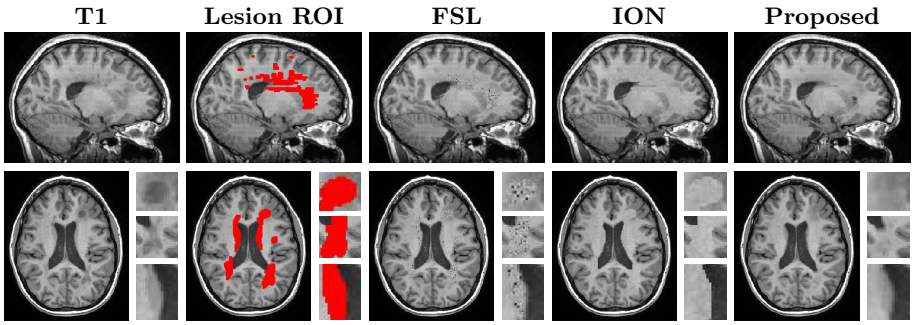

Fig. 3. The coronal, axial and zoomed views of the original T1 image and the dilated lesions mask (DROI), followed by the results with FSL, ION and the proposed methodology. Note the introduction of noisy samples using the FSL method and sharp contrast boundaries with the ION method.

was correctly preserved using the proposed technique but not using both the ION and FSL techniques. Furthermore, the proposed method was also applied to images from different modalities (see Fig. 4), demonstrating the generalisability and agnosticism to the type of image acquisition. The same parameters were used for all modalities.

3.2 Quantitative Analysis

In this second experiment, we want to quantitatively assess which method produces the most realistic patch. In order to do so, the 208 images were filled with the corresponding lesion mask. We computed the normalised entropy inside the filled lesions and the gradient magnitude at the edge of the filled lesions ($|\nabla I|$). The normalised entropy provides a measure of the tissue homogeneity inside the filled lesions, where a small entropy means that all lesions are filled with clustered intensity values. On the other side, $|\nabla I|$ provides information on the presence of discontinuities at the boundary of the lesion, with small values meaning a soft transition between real and synthetic intensities and high values show the presence of spurious image gradients at the edge of the region. Furthermore, in order to test the robustness of the filling procedure to the oversegmentation of the MS lesions, the manually segmented lesions were dilated with a cubic (26 connection) kernel. A boundary ROI (BROI) was obtained by an XOR operation between the lesion mask and the dilated lesion mask. As this boundary ROI is outside the lesion area, i.e. it contains only non-pathological tissues (WM, GM or CSF), it will be used to assess the error in the synthesis process. We used the mean square error (MSE) to measure the difference between the synthetic values and the real values within the BROI region, with a smaller MSE meaning more realistic synthetic intensities in the BROI.

In short, Table 1 shows the results of two experiments: first, using the LROI as filling ROI, we estimate the entropy in the LROI region after lesion filling and also the gradient magnitude $|\nabla I|$ at the edge of the LROI. Second, we fill the lesions using the DROI as a filling ROI, and calculate again the gradient at the edge of the DROI and also the MSE between the synthetic values and the real values within the BROI region.

Table 1. Results from the quantitative analysis, with the mean (std) and t-test (against the proposed method) over all the 208 subjects of the LROI region entropy and the gradient magnitude $|\nabla I|$ at the edge of the LROI, followed by the mean squared error (MSE) in the BROI region and the gradient at the boundary of the DROI region.

		FSL	**ION**	**Proposed**
LROI	Entropy	0.2 (0.03) p=0.13	0.09 (0.02) p<0.001	0.19 (0.03)
	$\|\nabla I\|_{LROI}$ ($\times 10^{-2}$)	6.1 (0.9) p=0.47	7.9 (1.3) p<0.001	6.1 (0.9)
DROI	MSE_{BROI} ($\times 10^{-4}$)	15.9 (6.1) p<0.001	69.3 (27.4) p<0.001	5.9 (1.7)
	$\|\nabla I\|_{DROI}$ ($\times 10^{-2}$)	6.5 (0.8) p<0.001	9.9 (1.4) p<0.001	6.0 (0.7)

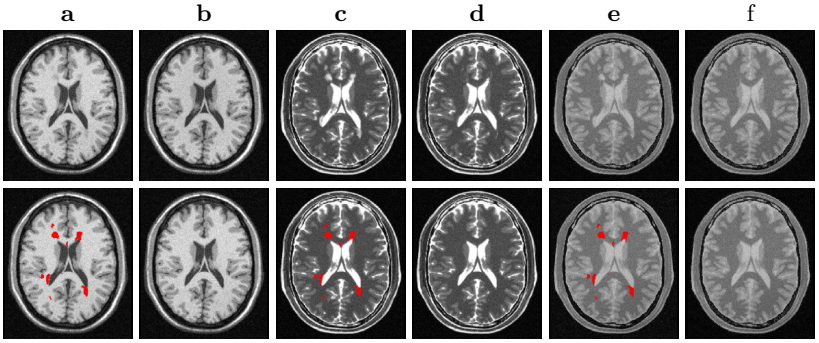

Fig. 4. Lesion filling results on different modalities using the "severe" MS phantom from the *BrainWeb* dataset. First row, the T1 and the lesion filled T1 ($a - b$), the T2 and the lesion filled T2 ($c - d$), and the PD and the lesion filled PD ($e - f$). Second row, the lesions mask over each different modality image (a, c and e), and the ground truth for each different modality image (b, d and f).

4 Discussion and Conclusion

In this paper, we propose a new and robust multi-modality lesion filling technique that relies on a non-local patch match strategy. The method shows improved results compared to previously published publicly available methods.

The presented method requires less prior information than previous methods, as it only needs a roughly defined lesion ROI mask. Conversely, the FSL method requires a precisely drawn lesion mask and the ION method requires an accurate skull stripping methodology.

Furthermore, the proposed method is not only less affected by the lack of contrast between tissues, as it fills the lesion ROI with the most similar non-local patches and not according to a class-specific intensity model, but at the same time more robust to the location of the lesions, i.e. previous algorithms have problems with lesions located close to non white matter regions.

The NL-Means models usually use a weighted averaging of the best-matched patches, rather than using the best match. In our testing, we found that weighted averaging can introduce edge blurring and an artificially low SNR.

Lastly, manual lesion editing is still the gold standard for lesion masking in MS, with the accuracy of the rater and the choice of segmentation plane being sources of bias. By exploiting contextual information, the proposed algorithm has been shown to be more robust to lesion over-segmentation than previously published techniques. Thus, it would be interesting to see if the proposed method can be used in conjunction with a highly sensitive automatic lesion detection methodology, thus removing rater bias from the analysis process. If the lesions are under-segmented, there is the possibility for our method to fill the segmented lesion ROI with lesion-like intensities, but a mask dilation can avoid this problem.

Although, Guizard's method [7] also used NL-means strategy, the two algorithms are different. Guizard's method searches the similar patches in

surrounding regions, having fewer potential patches to fill the lesion. It uses the whole patch and prefills the lesions with an average intensity, introducing bias and make the method less robust in the presence of under or over-segmentation.

Future work will explore a scaling/rotation invariant extension of the patch search technique, multi-time-point filling and model parameter optimisation.

Acknowledgments. FP is funded by the National Institute for Health Research University College London Hospitals Biomedical Research Centre (NIHR BRC UCLH/UCL High Impact Initiative). MJC is funded by the EPSRC (EP/H046410/1). SO receives funding from the EPSRC (EP/H046410/1, EP/J020990/1, EP/K005278), the MRC (MR/J01107X/1), the NIHR Biomedical Research Unit (Dementia) at UCL and the NIHR BRC UCLH/UCL. This work was also supported by the Medical Research Council, the Multiple Sclerosis Society in UK and the Brain Research Trust.

References

1. Barnes, C., Shechtman, E., Goldman, D., Finkelstein, A.: The generalized patchmatch correspondence algorithm. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010, Part III. LNCS, vol. 6313, pp. 29–43. Springer, Heidelberg (2010)
2. Battaglini, M., Jenkinson, M., De Stefano, N.: Evaluating and reducing the impact of white matter lesions on brain volume measurements. Human brain mapping 33(9), 2062–2071 (2012)
3. Buades, A., Coll, B., Morel, J.M.: A Review of Image Denoising Algorithms, with a New One. Multiscale Modeling & Simulation 4(2), 490–530 (2005)
4. Chard, D.T., Jackson, J.S., Miller, D.H., Wheeler-Kingshott, C.: Reducing the impact of white matter lesions on automated measures of brain gray and white matter volumes. Journal of Magnetic Resonance Imaging 32(1), 223–228 (2010)
5. Criminisi, A., Perez, P., Toyama, K.: Object removal by exemplar-based inpainting. In: CVPR 2003, vol. 2, pp. 721–728. IEEE Comput. Soc. (2003)
6. Geurts, J.J.G., Calabrese, M., Fisher, E., Rudick, R.A.: Measurement and clinical effect of grey matter pathology in multiple sclerosis. Lancet Neurology 11(12), 1082–1092 (2012)
7. Guizard, N., Nakamura, K., Coupe, P., Arnold, D.L., Collins, D.L.: Non-local MS MRI lesion inpainting method for image processing. In: The endMS Conf. (2013)
8. Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M.: Fsl. NeuroImage 62(2), 782–790 (2012)
9. Komodakis, N., Tziritas, G.: Via Priority Scheduling and Dynamic Pruning. IEEE Transactions on Image Processing 16(11), 2649–2661 (2007)
10. Lladó, X., Oliver, A., Cabezas, M., Freixenet, J., Vilanova, J.C., Quiles, A., Valls, L., Ramió, L., Rovira, A.: Segmentation of multiple sclerosis lesions in brain MRI: A review of automated approaches. Information Sciences 186(1), 164–185 (2012)
11. Sdika, M., Pelletier, D.: Nonrigid registration of multiple sclerosis brain images using lesion inpainting for morphometry or lesion mapping. Human Brain Mapping 30(4), 1060–1067 (2008)
12. Telea, A.: An Image Inpainting Technique Based on. Journal of Graphics Tools 9(1), 25–36 (2004)