NeuroImage 139 (2016) 376–384

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

A multi-time-point modality-agnostic patch-based method for lesion filling in multiple sclerosis

Ferran Prados [a,b,*], Manuel Jorge Cardoso [a], Baris Kanber [a], Olga Ciccarelli [b], Raju Kapoor [c], Claudia A.M. Gandini Wheeler-Kingshott [b,d], Sebastien Ourselin [a,e]

[a] Translational Imaging Group, Centre for Medical Image Computing (CMIC), Department of Medical Physics and Bioengineering, University College London, Malet Place Engineering Building, Gower Street, London, WC1E 6BT, UK
[b] NMR Research Unit, Queen Square MS Centre, Department of Neuroinflammation, UCL Institute of Neurology, 1st Floor, Russell Square House, 10-12 Russell Square, London WC1B 5EH, UK
[c] Department of Neuroinflammation, National Hospital for Neurology and Neurosurgery, Queen Square, London WC1N 3BG, UK
[d] Brain MRI 3T Center, C. Mondino National Neurological Institute, Pavia, Italy
[e] Dementia Research Centre, Department of Neurodegenerative Disease, UCL Institute of Neurology, Queen Square, London WC1N 3BG, UK

ARTICLE INFO

Article history:
Received 9 March 2016
Revised 10 June 2016
Accepted 27 June 2016
Available online 1 July 2016

Keywords:
Error correction
Multiple sclerosis
Lesions
Segmentation errors
Artefacts
MRI

ABSTRACT

Multiple sclerosis lesions influence the process of image analysis, leading to tissue segmentation problems and biased morphometric estimates. Existing techniques try to reduce this bias by filling all lesions as normal-appearing white matter on T1-weighted images, considering each time-point separately. However, due to lesion segmentation errors and the presence of structures adjacent to the lesions, such as the ventricles and deep grey matter nuclei, filling all lesions with white matter-like intensities introduces errors and artefacts. In this paper, we present a novel lesion filling strategy inspired by in-painting techniques used in computer graphics applications for image completion. The proposed technique uses a five-dimensional (5D), patch-based (multi-modality and multi-time-point), Non-Local Means algorithm that fills lesions with the most plausible texture. We demonstrate that this strategy introduces less bias, fewer artefacts and spurious edges than the current, publicly available techniques. The proposed method is modality-agnostic and can be applied to multiple time-points simultaneously. In addition, it preserves anatomical structures and signal-to-noise characteristics even when the lesions are neighbouring grey matter or cerebrospinal fluid, and avoids excess of blurring or rasterisation due to the choice of the segmentation plane, shape of the lesions, and their size and/or location.

Introduction

Multiple sclerosis (MS) is an immune-mediated demyelinating disease that affects 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. These pathological processes affect several quantitative MRI indices, and therefore can be indirectly measured with advanced imaging methods. Among these, tissue volume and, in particular, brain/tissue-specific atrophy, are very sensitive to subtle changes over a scale of a few months, making *in vivo* MRI measurements of these indices very appealing for studying the mechanisms of disease and for clinical trials. White matter plaques are relatively easy to detect using conventional MRI techniques, whereas grey matter lesions can be observed using specialised sequences, such as double inversion recovery (DIR) (Geurts et al., 2012) or phase sensitive inversion recovery (PSIR) (Sethi et al., 2012). MS plaques appear as areas of low-signal intensity and high-signal intensity compared with normal-appearing white matter (NAWM) on T1-weighted and T2-weighted sequences respectively. On the other hand, active lesions exhibit hyper-intense signals on gadolinium-enhanced scans (Lladó et al., 2012). Lesions and atrophy are two interconnected aspects of the disease, linked to different disease mechanisms, and both are extremely important for MS studies.

From an image processing perspective, MS lesions influence tissue segmentation, resulting in the misclassification of the GM and the WM. It has been suggested that MS lesions may affect the estimation of segmentation parameters, resulting in a shift of tissue boundaries (Chard et al., 2010), thus influencing the subsequent morphometric studies, including atrophy measurements. Thus, there is a clear need to reduce the negative impact that MS lesions may have on image analysis in order to improve tissue segmentation and longitudinal registration, increasing sensitivity to subtle changes, reducing the time-intervals and sample sizes needed for longitudinal studies and treatment trials.

* Corresponding author at: Translational Imaging Group, Centre for Medical Image Computing (CMIC), Department of Medical Physics and Bioengineering, University College London, Malet Place Engineering Building, Gower Street, London, WC1E 6BT, UK.
E-mail address: f.carrasco@ucl.ac.uk (F. Prados).

http://dx.doi.org/10.1016/j.neuroimage.2016.06.053

Various techniques have been developed in recent years based on the concept of in-painting T1-weighted MRI images: Sdika and Pelletier (2008), Chard et al. (2010), Battaglini et al. (2012), Magon et al. (2014), Valverde et al. (2014), and Guizard et al. (2015).

In short, the process of lesion in-painting is based on filling a WM lesion with synthetic estimates of WM-like image intensities. The process of WM lesion in-painting is expected to reduce the overall algorithmic bias. Sdika and Pelletier (2008) presented three different in-painting algorithms. The first, denoted *basic in-painting* and inspired by Telea (2004), consists of filling the lesion ROI in an inner-radial manner using a Gaussian kernel average $3\times3\times3$ of the neighbouring intensities. The second, called *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.

Recently, Chard et al. (2010) developed the LEAP (LEsion Automated Preprocessing) technique with the aim of filling lesions as normal WM, reproducing the WM noise characteristics and avoiding operator intervention. This method starts by skull stripping the brain and applying a non-uniformity intensity correction algorithm. The normal tissue intensity distribution is modelled numerically as the sum of four Gaussian components representing GM, WM, CSF, and partial-volume voxels. Finally, the lesion ROI is filled with random samples from a Gaussian distribution with mean equal to the most probable WM intensity and a standard deviation equal to the WM full-width half maximum noise characteristics. This method is available at: http://www.nmrgroup.ion.ucl.ac.uk/. Valverde et al. (2014) and Magon et al. (2014) have then introduced variations to the LEAP method. Valverde et al. (2014) suggested to fill the volume in a slice-wise manner, whilst Magon et al. (2014) filled the lesions using the mean intensity of two voxel expanded neighbouring over the normal-appearing WM.

Similarly, Battaglini et al. (2012) presented a 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. The surrounding normal-appearing tissue volume is taken as the extra volume obtained by dilating the lesion ROI twice. Lesions are then filled with samples taken from the neighbouring histogram using a uniform random value passed through an interpolated version of the empirical cumulative distribution function of the neighbouring histogram. Both GM and WM voxels are included in the neighbouring histogram in order to represent the surrounding tissue and allow the filled lesions to best visually blend into its environment. This method is available as part of FSL (Jenkinson et al., 2012) at http://fsl.fmrib.ox.ac.uk.

Regardless of their approach, all these algorithms have been restricted to images of a specific MRI modality (*e.g.* T1-weighted scans), and require accurate lesion segmentation, especially when lesions are periventricular or when the methods are based on filling the lesions with values from the surrounding areas. They can also create shape gradients around the lesion ROIs, and are prone to errors coming from estimating WM distribution properties.

More recently, Guizard et al. (2013, 2015) calculated the most similar patches using only the surrounding regions after pre-filling the lesions with the median of the image intensities of the surrounding of healthy tissues (Guizard et al., 2013). The same authors later introduced an hierarchical, concentric filling strategy, where distances between patches are computed over the full patch, the filling process is repeated with different weighting values and at multiple increasing resolutions (Guizard et al., 2015).

In the field of computer graphics, structurally aware in-painting algorithms used for scratch/object/text removal and photo restoration are common, with many of these algorithms permitting a user to simply erase an unwanted portion of an image without any prior knowledge about its composition. These techniques attempt to fill regions by synthesising plausible texture matches from the remainder of the image (Criminisi et al., 2003; Komodakis and Tziritas, 2007; Barnes et al., 2010). In doing so these algorithms are completely 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 (Komodakis and Tziritas, 2007), under some constraints. This class of methods commonly divides the image into a large number of sub-images, or patches, followed by either a patch-search method Criminisi et al. (2003), or the use of the Non-Local Means algorithm (Buades et al., 2005). Finally, the intensities can be synthesised using either pixel values or patch-based textures from the most similar patch.

In this work, we formulate a multi-time-point, task-specific, patch-search algorithm for the purpose of filling MS lesions. This novel work offers three main advantages: first, due to its general formulation, the proposed algorithm should in theory be able to inpaint most types of MR images with repetitive patterns. Secondly, due to its contextual nature, the proposed algorithm is also more robust to over-segmentation of the lesion ROIs, thus reducing accuracy requirements when manually or automatically defining the in-painting region of interest. Thirdly, and finally, it allows filling lesions at different time-points (as in longitudinal studies) at once taking advantage of the information of the lesion evolution.

Material and methods

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

First, we assume that we have a greyscale-valued 5D volume I^*, composed of n different modalities or MRI sequences, over t time points, with each individual 3D volume being of size $X\times Y\times Z$. Each time point and modality has an associated lesion mask, here denoted as $\mathcal{L}$.

Let the filled image I be defined as $I(p)=I^*(p)\ \forall p\notin\mathcal{L}$, and as $I(p)=\mathcal{F}(p)\ \forall p\in\mathcal{L}$, where p denotes the voxel location (x,y,z,n,t) in the image I and $\mathcal{F}(p)$ is the function that synthesises the intensity of voxel p. We define Ω as a search region of size W^3 voxels around voxel p (where W denotes the spatial search region size in voxels in each spatial direction). Within the region Ω, we define a 5D target patch $T(p)$ of size ntw^3 voxels (where w denotes the patch size in voxels), centred at a voxel p, and a search patch $S(q)$ of size ntw^3 voxels, centred in q, with $q\in\Omega$ 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:

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

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

$$\mathcal{D}(T(p),S(q))=\frac{\sum_{\forall i\in\{ntw^3\}|\{T(p+i),S(q+i)\}\notin\mathcal{L}}\left(I(p+i)-I(q+i)\right)^2}{\kappa^c} \tag{2}$$

Here, κ is the cardinality of the set $i\in\{ntw^3\}|\{T(p+i),S(q+i)\}\notin\mathcal{L}$, *i.e.* the number of voxels within the patches $T(p)$ and $S(q)$ that are not in the lesion region. Note that when $c>1$, the denominator κ^c favours patches with more information. A further hard constraint can be added by defining α as the minimum required percentage of non-lesion voxels in a patch. This hard constraint can be formally defined as $\kappa>\alpha ntw^3$, *i.e.* the cardinality of the set $\forall i\in\{ntw^3\}|\{T(p+i),S(q+i)\}\notin\mathcal{L}$ has to be more than α% of the patch size. If this constraint is satisfied, then p is marked as solve and is removed from the set $\mathcal{L}$, otherwise, p remains in