



Factorisation-Based Image Labelling

Yu Yan^{1,2}, Yaël Balbastre^{1,3}, Mikael Brudfors^{1,2} and John Ashburner^{1*}

¹ Wellcome Centre for Human Neuroimaging, UCL Queen Square Institute of Neurology, University College London, London, United Kingdom, ² School of Biomedical Engineering & Imaging Sciences, King's College London, London, United Kingdom, ³ Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Boston, MA, United States

Segmentation of brain magnetic resonance images (MRI) into anatomical regions is a useful task in neuroimaging. Manual annotation is time consuming and expensive, so having a fully automated and general purpose brain segmentation algorithm is highly desirable. To this end, we propose a patched-based label propagation approach based on a generative model with latent variables. Once trained, our Factorisation-based Image Labelling (FIL) model is able to label target images with a variety of image contrasts. We compare the effectiveness of our proposed model against the state-of-the-art using data from the *MICCAI 2012 Grand Challenge and Workshop on Multi-Atlas Labelling*. As our approach is intended to be general purpose, we also assess how well it can handle domain shift by labelling images of the same subjects acquired with different MR contrasts.

Keywords: label propagation, atlas, machine learning, latent variables, variational bayes

OPEN ACCESS

Edited by:

Suyash P. Awate,
Indian Institute of Technology Bombay,
India

Reviewed by:

Yiming Xiao,
Concordia University, Canada
Hongzhi Wang,
IBM, Egypt

*Correspondence:

John Ashburner
j.ashburner@ucl.ac.uk

Specialty section:

This article was submitted to
Brain Imaging Methods,
a section of the journal
Frontiers in Neuroscience

Received: 19 November 2021

Accepted: 10 December 2021

Published: 17 January 2022

Citation:

Yan Y, Balbastre Y, Brudfors M and
Ashburner J (2022)
Factorisation-Based Image Labelling.
Front. Neurosci. 15:818604.
doi: 10.3389/fnins.2021.818604

1. INTRODUCTION

Accurate automated labelling of brain structures in MRI scans has many applications in neuroscience. For example, studies of resting state fMRI or diffusion weighted imaging often involve summarising measures of connectivity among a relatively small number of brain regions. Typically, these regions are only very approximately defined, using simple methods such as aligning with a single manually labelled brain (Tzourio-Mazoyer et al., 2002). While more accurate methods of brain parcellation are available, they usually have limitations in terms of the types of MRI scans that can be labelled, and they are often very computationally expensive. For these reasons, they have not yet been widely adopted by the neuroimaging field. We attempt to overcome some of these limitations by presenting a novel brain image labelling toolbox for the widely used SPM software.

In this work, we adopt a multi-atlas labelling approach. Given a training set of scans from N individuals, $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$, along with their corresponding manual annotations $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$, the aim would be to estimate a suitable labelling $\hat{\mathbf{y}}^*$ for a target image \mathbf{x}^* that is not part of the training set. This objective is often achieved by obtaining the single most probable labelling, given by

$$\hat{\mathbf{y}}^* = \arg \max_{\mathbf{y}^*} P(\mathbf{y}^* | \mathbf{x}^*, \mathbf{X}, \mathbf{Y}). \quad (1)$$

Most label propagation methods require alignment between all the atlases in the training data (\mathbf{X}) and the target image (\mathbf{x}^*), which is usually achieved by a series of pairwise registrations. The Symmetric Normalisation algorithm (in the ANTS package) (Avants et al., 2008) is popular for this, although other algorithms are available. This pair-wise strategy allows the label propagation to be performed directly in the space of the target image, which may lead to increased robustness because a small proportion of failures should have a relatively small impact on the results.

Once mappings between the training images and target image have been established, the manually defined labels are warped using the same mappings. Each of these then provides a candidate labelling of the target image. At this stage, some form of machine learning approach is used to predict the labels (\hat{y}^*) for the target image from all the candidate labellings. This procedure is often conceptualised in terms of some form of local weighted voting strategy, the simplest of which is to give each an equal vote. In practice, this system is less effective than one that is weighted according to how well informed the voters are.

Earlier methods were similar to k-nearest neighbour classification, whereby a subset of atlases were chosen to label each brain region (Rohlfing et al., 2004). Others use a non-local patch-based framework (Coupé et al., 2013), with greater weighting for votes based on a more accurate model of the target image. This typically involves a measure of similarity between patches in \mathbf{x}^* and the corresponding patches within each aligned image in \mathbf{X} . Many such approaches can be conceptualised as a joint non-parametric generative model of both image and label data (Sabuncu et al., 2010). An alternative framework is Simultaneous Truth and Performance Level Estimation (STAPLE) (Warfield et al., 2004), which is based on weighting the votes from candidate atlases according to how well they match the consensus over all votes. More local versions of STAPLE have also been devised (Asman and Landman, 2012b; Commowick et al., 2012), as well as hybrids between STAPLE and the non-local approaches (Asman and Landman, 2012a). Other approaches involve assuming that all voxels within each labelled region should have similar intensities, in a similar way to how many domain-adaptive tissue segmentations work. Some methods treat this as a post-processing step in conjunction with a Markov random field (MRF) (Ledig et al., 2013), whereas others have used this assumption to drive the image registration (Tang et al., 2013).

Rather than conceptualise label propagation entirely in terms of optimising vote weightings, more recent work considers it within the more general pattern recognition framework. This includes using Random Forest approaches (Zikic et al., 2014) and, more recently, using convolutional neural networks (de Brebisson and Montana, 2015; Moeskops et al., 2016; Mehta et al., 2017; Roy et al., 2017, 2019; Wachinger et al., 2018; Kushibar et al., 2019; Rashed et al., 2020).

Many of these methods are impacted by the problem of domain shift, which is the situation where images in the training data (\mathbf{X}) have different properties from those that the algorithm is to be applied to (\mathbf{x}^*). Typically, this is due to differences between image acquisition settings, scanner vendors, field strength, and so on. Our aim is to release an automated labelling procedure for general purpose use, which would require overcoming the domain shift problem. We attempt to circumvent it by working with images that have previously been segmented into different tissue types. This can be achieved using one of the many domain-adaptive brain image segmentation approaches that have been developed. Such approaches generally build on the idea of fitting some form of clustering model to the data (Wells et al., 1996; Ashburner and Friston, 1997), often with MRFs (Van Leemput et al., 1999; Zhang et al., 2001) or deformable

tissue priors (Ashburner and Friston, 2005; Pohl et al., 2006; Puonti et al., 2016) built in. While reducing medical images to a few tissue types inevitably leads to some useful information being lost, we consider that this is a price worth paying for the increase in generality of the approach. A related strategy (this time separating segmentation from diagnosis) has been used for increasing the generalisability of deep learning approach for diagnosing retinal disease (De Fauw et al., 2018).

In our proposed Factorisation-based Image Labelling (FIL) method, we consider multi-atlas labelling as a special case of the more general problem of image-to-image translation, but where some of the data are binary or categorical in nature. Hence, our approach differs from previous methods in a number of ways.

Because running many pairwise registrations can be quite time consuming, we propose to label target images using only a single image registration. Training involves first running an image registration approach to warp all training data (\mathbf{X}) into the same atlas space (i.e., spatially normalise the data). After this, only a single image registration step is required to align any target image with the average atlas space. While we acknowledge that this may sacrifice some registration accuracy and robustness, we argue that working in a symmetric space (where no image is a *source* or *target*) facilitates pattern recognition across the set of training scans.

Other methods use purely discriminative methods for the label propagation itself. These methods model \mathbf{y}^* conditional on \mathbf{x}^* , whereas we propose to use a parametric generative model that encodes the joint probability of \mathbf{x}^* and \mathbf{y}^* . In addition to providing a new way of thinking about label propagation, we hope this generative model will open up other avenues of exploration in the future, particularly regarding multi-task and semi-supervised learning (Zhu, 2005).

2. METHODS

The basic idea is to learn a generative latent variable model that can be used to project labels onto images that have been already warped into approximate alignment with each other (i.e., spatially normalised). Alignment across individuals provides prior knowledge about the approximate locations of the various brain structures. Therefore, the fully convolutional machinery of convolutional neural networks (CNNs) is not employed because it may not fully use this prior knowledge.

Rather than use the original pixel/voxel values, the approach aims to achieve generalisability across different types of images by working with categorical maps obtained from one of many tissue classification algorithms. Some confusion may arise from the fact that both our *images* and *labels* are categorical. We will use *categorical images* to denote the original images segmented into tissue types and *categorical labels* to denote manually delineated anatomical labels. The number of anatomical classes is typically larger (and more fine-grained) than the number of tissue classes; although this is not a requirement of the model, which is symmetric with respect to both entries.

The approach is patch-based and applied to spatially normalised tissue maps (i.e., categorical image data). It can

be seen as a multinomial version of principal component analysis, similarly to how linear regression can be generalised to multinomial logistic regression. For each patch, a set of basis functions model both the categorical image data and the corresponding categorical label data, with a common set of latent variables controlling the contributions from the two sets of basis functions. The results are passed through a softmax to encode the means of a multinomial distribution. Training the model at a patch involves making point estimates of the set of spatial basis functions ($\mathbf{W}^{(1)}$ and $\mathbf{M}^{(1)}$) that model the categorical image data ($\mathbf{F}^{(1)}$), along with the basis functions ($\mathbf{W}^{(2)}$ and $\mathbf{M}^{(2)}$) that model the label data ($\mathbf{F}^{(2)}$). For subject n , the contributions of the basis functions to both the image patch and its corresponding label patch are controlled by the same set of latent variables (\mathbf{z}_n). The overall model for a patch is presented in **Figure 1**.

Once trained, the strategy is to determine the distribution of \mathbf{z}^* for each patch within the target image data, which we refer to as the “encoding step”. This is achieved by fitting the learned $\mathbf{W}^{(1)}$ and $\mathbf{M}^{(1)}$ to each patch ($\mathbf{F}^{*(1)}$). Then by using $\mathbf{W}^{(2)}$ and $\mathbf{M}^{(2)}$, and given the estimated distribution of \mathbf{z}^* , it is possible to use a “decoding step” to probabilistically predict the unknown labels ($\mathbf{F}^{*(2)}$) for the patch.

The next section describes a simplified version of the full model shown in **Figure 1**. Rather than model both $\mathbf{F}^{(1)}$ and $\mathbf{F}^{(2)}$, it describes how to encode a single patch \mathbf{F} using an approach similar to a generalisation of principal component analysis (PCA). Generalising this simple approach to jointly model $\mathbf{F}^{(1)}$ and $\mathbf{F}^{(2)}$ is relatively straightforward.

2.1. Multinomial Logistic Principal Component Analysis Model

This section describes a multinomial principal component analysis (PCA), which is based on Khan et al. (2010). For PCA of categorical data, the arrays involved are multi-dimensional, so we represent them as collections of matrices at each of the I voxels

$$\text{Categorical images: } \mathbf{F} = \left\{ \mathbf{F}_i \in [0, 1]^{M \times N} \mid \sum_{m=1}^M f_{mni} \in [0, 1] \right\}_{i=1}^I \quad (2)$$

$$\text{Basis functions: } \mathbf{W} = \{ \mathbf{W}_i \in \mathbb{R}^{M \times K} \}_{i=1}^I \quad (3)$$

$$\text{Means: } \mathbf{M} = \{ \boldsymbol{\mu}_i \in \mathbb{R}^{M \times 1} \}_{i=1}^I. \quad (4)$$

Note that categories lie on the simplex: there are $M + 1$ mutually exclusive categories within the data (i.e., the number of possible labels), but there is no need to represent all categories because the final category is determined by the initial M categories. Dimension N denotes the number of items (i.e., the number of images in the training set). Dimension K denotes the number of basis functions in the model, and therefore also the number of latent variables per item

$$\mathbf{Z} \in \mathbb{R}^{K \times N}. \quad (5)$$

Our notation uses a bold sans serif font to denote 3D tensors (e.g. \mathbf{F}). Each slice is a 2D matrix, shown in bold serif upper case (e.g. \mathbf{F}_i). The next level of indexing extracts column vectors from matrices, which are shown in bold lower case (e.g. \mathbf{f}_{ni}). Dimensions are shown in upper case italic (e.g. N), whereas indices and other scalars are in lower case italic (e.g. n or f_{mni}). This work considers training a generalised PCA model to find the most probable values of \mathbf{M} and \mathbf{W} , while marginalising with respect to \mathbf{Z}

$$\hat{\mathbf{M}}, \hat{\mathbf{W}} = \arg \max_{\mathbf{M}, \mathbf{W}} p(\mathbf{F}, \mathbf{W} | \mathbf{M}) \quad (6)$$

$$p(\mathbf{F}, \mathbf{W} | \mathbf{M}) = \int_{\mathbf{Z}} P(\mathbf{F} | \mathbf{Z}, \mathbf{W}, \mathbf{M}) p(\mathbf{Z}) p(\mathbf{W}) d\mathbf{Z}. \quad (7)$$

Within our model, the basis functions (\mathbf{W}) are assumed to be drawn from

$$\mathbf{w}_{ki} \sim \mathcal{N} \left(0, \left(\mathbf{I}_M + \frac{1}{M+1} \right)^{-1} \right). \quad (8)$$

No priors are imposed on the mean (\mathbf{M}). The latent variables are assumed to be drawn from an empirically determined (see section 2.2.3) Gaussian distribution that imposes spatial contiguity between neighbouring patches, such that

$$p(\mathbf{Z}) = \prod_{n=1}^N p(\mathbf{z}_n) \quad (9)$$

$$\mathbf{z}_n \sim \mathcal{N}(\mathbf{z}_0, \mathbf{P}_0^{-1}). \quad (10)$$

In linear PCA, the likelihood takes the form of an isotropic Gaussian distribution conditioned on the reconstructed data ($\boldsymbol{\eta}_{ni} = \boldsymbol{\mu}_i + \mathbf{W}_i \mathbf{z}_n$). Here, it is based around a categorical distribution conditioned on the *softmaxed* reconstruction, similar to frameworks used for multinomial logistic regression

$$P(\mathbf{F} | \mathbf{Z}, \mathbf{W}, \mathbf{M}) = \prod_{i=1}^I \prod_{n=1}^N P(\mathbf{f}_{ni} | \mathbf{z}_n, \mathbf{W}_i, \boldsymbol{\mu}_i) \quad (11)$$

$$P(\mathbf{f}_{ni} | \mathbf{z}_n, \mathbf{W}_i, \boldsymbol{\mu}_i) = P(\mathbf{f}_{ni} | \boldsymbol{\rho}_{ni}) = \prod_{m=1}^M \rho_{mni}^{f_{mni}}. \quad (12)$$

The mean ($\boldsymbol{\rho}_{ni}$) of each categorical distribution is computed as the softmax (σ) of the linear combination of basis functions

$$\boldsymbol{\rho}_{ni} = \sigma(\boldsymbol{\eta}_{ni}) = \frac{\exp(\boldsymbol{\eta}_{ni})}{1 + \sum_{m=1}^M \exp(\boldsymbol{\eta}_{mni})} \quad (13)$$

$$\boldsymbol{\eta}_{ni} = \mathbf{W}_i \mathbf{z}_n + \boldsymbol{\mu}_i. \quad (14)$$

The log-likelihood from Equations (12) and (13) is re-written to make use of the log-sum-exp (lse) function

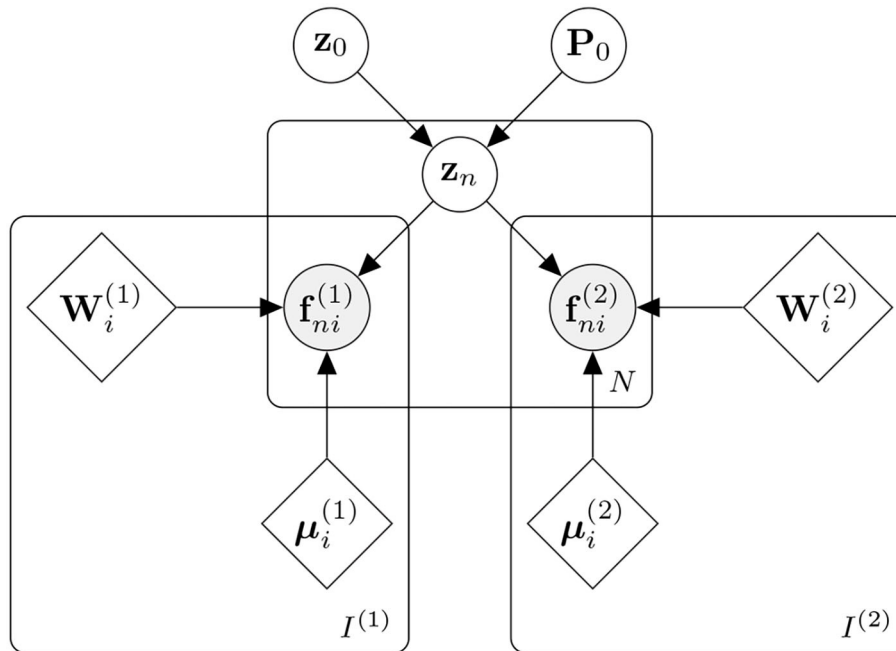


FIGURE 1 | Generative model for a single patch. Data from N subjects are assumed to be patches containing categorical data. Image voxels are denoted by $r_{ni}^{(1)}$ (where $n \in [1..N]$ and $i \in [1..I^{(1)}]$) and label voxels by $r_{ni}^{(2)}$ ($i \in [1..I^{(2)}]$). These are encoded by their means ($\mu_i^{(1)}$ and $\mu_i^{(2)}$, respectively) and a linear combination of basis functions ($W_i^{(1)}$ and $W_i^{(2)}$, respectively). For each subject, the contributions of the two sets of basis functions are jointly controlled by latent variables z_n , which are assumed to be drawn from a normal distribution of mean z_0 and precision P_0 .

$$\ln P(\mathbf{f}_{ni}|\eta_{ni}) = \eta_{ni}^T \mathbf{f}_{ni} - \text{lse}(\eta_{ni}) \tag{15}$$

where

$$\text{lse}(\eta_{ni}) = \ln \left(1 + \sum_{m=1}^M \exp(\eta_{mni}) \right). \tag{16}$$

Numerous computations within this work benefit from having a local quadratic approximation to the lse function around some point $\hat{\eta}_{ni} = \mathbf{W}_i z_n + \mu_i$. This is based on a second order Taylor polynomial that uses the gradient ($\hat{\rho} = \sigma(\hat{\eta}_{ni})$), but replaces the true Hessian with Böhning’s approximation (\mathbf{A}) (Böhning, 1992)

$$\text{lse}(\eta_{ni}) \leq \text{lse}(\hat{\eta}_{ni}) + (\eta_{ni} - \hat{\eta}_{ni})^T \hat{\rho}_{ni} + \frac{1}{2}(\eta_{ni} - \hat{\eta}_{ni})^T \mathbf{A}(\eta_{ni} - \hat{\eta}_{ni}) \tag{17}$$

where

$$\hat{\rho}_{ni} = \sigma(\hat{\eta}_{ni}) = \exp(\hat{\eta}_{ni} - \text{lse}(\hat{\eta}_{ni})) \tag{18}$$

$$\mathbf{A} = \frac{1}{2} \left(\mathbf{I}_M - \frac{1}{M+1} \right). \tag{19}$$

It is worth noting that Böhning’s proposed approximation is more positive definite (in the Loewner ordering sense) than any of the actual Hessians, which guarantees that the approximating function provides an upper bound to the true lse function (Böhning, 1992; Khan et al., 2010) (see **Figure 2** for an illustration where $M = 1$).

2.2. Model Training

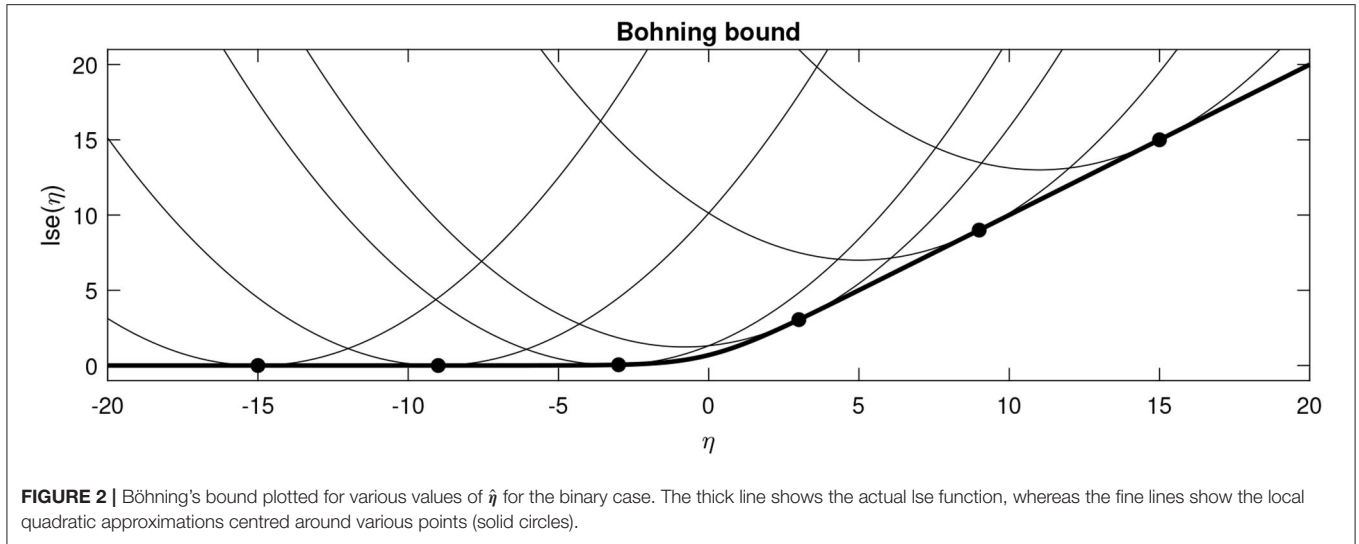
This section describes how the basis functions are estimated for each patch. A variational expectation maximisation (EM) approach is used for fitting the model in such a way that the uncertainty of \mathbf{Z} can be accounted for. Variational Bayesian methods are a family of technique for approximating the types of intractable integrals often encountered in Bayesian inference. Further explanations may be found in textbooks, such as Bishop (2006) and Murphy (2012). In summary, it uses an approximating distribution to enable a lower bound on a desired log-likelihood to be sought. In this work, the approximating distribution is $q(\mathbf{Z})$, which is used to provide a lower bound $\mathcal{L}(q)$ on the likelihood $P(\mathbf{F}|\mathbf{M}, \mathbf{W})$. This bound is tightest when $q(\mathbf{Z})$ most closely approximates $p(\mathbf{Z}|\mathbf{F}, \mathbf{M}, \mathbf{W})$ according to the Kullback-Liebler divergence (i.e. $\text{KL}(q||p)$)

$$\ln p(\mathbf{F}, \hat{\mathbf{W}}|\hat{\mathbf{M}}) = \mathcal{L}(q) + \text{KL}(q||p) \tag{20}$$

$$\mathcal{L}(q) = \int_{\mathbf{Z}} q(\mathbf{Z}) \ln \left(\frac{p(\mathbf{F}, \mathbf{Z}, \hat{\mathbf{W}}|\hat{\mathbf{M}})}{q(\mathbf{Z})} \right) d\mathbf{Z} \tag{21}$$

$$\text{KL}(q||p) = - \int_{\mathbf{Z}} q(\mathbf{Z}) \ln \left(\frac{p(\mathbf{Z}|\mathbf{F}, \hat{\mathbf{W}}, \hat{\mathbf{M}})}{q(\mathbf{Z})} \right) d\mathbf{Z}. \tag{22}$$

Fitting the model by variational EM involves an iterative algorithm that alternates between a variational E-step that updates the approximation to the distribution of the latent variables $q(\mathbf{Z})$ using the current point estimates of \mathbf{M} and \mathbf{W} (i.e.,



$\hat{\mathbf{M}}$ and $\hat{\mathbf{W}}$), and a variational M-step that uses $q(\mathbf{Z})$ to update $\hat{\mathbf{M}}$ and $\hat{\mathbf{W}}$. These two steps are outlined next.

2.2.1. Variational E-Step

We choose a product of Gaussian distributions for the approximate latent distribution, such that each factor is parameterised by a mean \hat{z}_n and a covariance matrix \mathbf{V}_n

$$q(\mathbf{Z}) = \prod_{n=1}^N \mathcal{N}(\mathbf{z}_n | \hat{z}_n, \mathbf{V}_n). \quad (23)$$

Keeping only terms that depend on $(\hat{z}_n, \mathbf{V}_n)$, the evidence lower bound in Equation (21) can be written as

$$\mathcal{L}(\hat{z}_n, \mathbf{V}_n) = \int_{\mathbf{z}_n} \left(\ln p(\mathbf{F}_n, \mathbf{z}_n, \hat{\mathbf{W}} | \hat{\mathbf{M}}, \mathbf{z}_0, \mathbf{P}_0) - \ln q(\mathbf{z}_n) \right) q(\mathbf{z}_n) d\mathbf{z}_n + \text{const.} \quad (24)$$

Making use of the local approximation in Equation (17) about $\hat{\eta}_{ni} = \hat{\mathbf{W}}_i \hat{z}_n^{\text{prev}} + \hat{\mu}_i$ and of Jensen's inequality, we can devise a quadratic lower bound on the evidence lower bound. Note that, as in classical EM, this quadratic lower-bound touches the variational lower-bound in each estimate; increasing the former therefore ensures to also increase the latter. A closed-form solution to the substitute problem exists, and we find

$$\mathbf{V}_n = \left(\mathbf{P}_0 + \sum_{i=1}^I \hat{\mathbf{W}}_i^T \mathbf{A} \hat{\mathbf{W}}_i \right)^{-1} \quad (25)$$

$$\hat{z}_n = \mathbf{V}_n \left(\mathbf{P}_0 \mathbf{z}_0 + \sum_{i=1}^I \hat{\mathbf{W}}_i^T (\mathbf{f}_{ni} - \hat{\rho}_{ni} + \mathbf{A}(\hat{\eta}_{ni} - \hat{\mu}_i)) \right). \quad (26)$$

This Gaussian approximation has the following expectations, which can be substituted into various other equations when required

$$\mathbb{E}[\mathbf{z}_n] = \hat{z}_n \quad (27)$$

$$\mathbb{E}[\mathbf{z}_n \mathbf{z}_n^T] = \hat{z}_n \hat{z}_n^T + \mathbf{V}_n. \quad (28)$$

2.2.2. Variational M-Step

The M-step uses $q(\mathbf{Z})$ to update the point estimates of \mathbf{M} and \mathbf{W} . For simplicity, our implementation updates \mathbf{M} and \mathbf{W} separately, although it would have been possible to update them simultaneously. The strategy for updating $\hat{\mathbf{M}}$ is similar to a Gauss-Newton update, but we formulate it in a manner that would be familiar to those working with variational Bayesian methods. This involves using the estimates of $\hat{\mathbf{M}}$, $\hat{\mathbf{W}}$ and $\hat{\mathbf{Z}}$ to set the variational parameters to $\hat{\eta}_{ni} = \hat{\mathbf{W}}_i \hat{z}_n + \hat{\mu}_i$. Then the local approximation of Equation (17) is substituted for $\text{lse}(\hat{\mathbf{W}}_i \hat{z}_n + \mu_i)$ in the expectation of $\ln p(\mathbf{F}_i, \mathbf{Z}, \hat{\mathbf{W}}_i | \mu_i)$ with respect to $q(\mathbf{Z})$. Terms that do not involve μ_i are ignored, giving

$$\begin{aligned} \mathbb{E}_{q(\mathbf{Z})}[\log p(\mathbf{F}_i, \mathbf{Z}, \hat{\mathbf{W}}_i | \mu_i)] \leq \\ \mu_i^T \sum_{n=1}^N (\mathbf{f}_{ni} - \hat{\rho}_{ni} + \mathbf{A}(\hat{\eta}_{ni} - \hat{\mathbf{W}}_i \mathbb{E}[\mathbf{z}_n])) - \frac{N}{2} \mu_i^T \mathbf{A} \mu_i \\ + \text{const.} \end{aligned} \quad (29)$$

Completing the square shows that μ_i would be drawn from a multivariate Gaussian distribution, although we are only interested in its mean. Substituting Equation (27) gives the following update for $\hat{\mu}_i$

$$\hat{\mu}_i = (NA)^{-1} \left(\sum_{n=1}^N (\mathbf{f}_{ni} - \hat{\rho}_{ni} + \mathbf{A}(\hat{\eta}_{ni} - \hat{\mathbf{W}}_i \hat{z}_n)) \right). \quad (30)$$

A similar approach is used for updating \mathbf{W}_i , although the vec operator is required to treat the \mathbf{W}_i matrices as vectors of parameters. The Kronecker tensor product (\otimes) is also used to

construct the following upper bound, which is substituted into $\mathbb{E}_{q(\mathbf{Z})}[\log p(\mathbf{F}_i, \mathbf{Z}, \mathbf{W}_i | \hat{\boldsymbol{\mu}}_i)]$

$$\begin{aligned} & \mathbb{E}_{\mathbf{z}_n}[\mathbf{f}_{ni}^T(\mathbf{W}_i \mathbf{z}_n + \hat{\boldsymbol{\mu}}_i) - \text{lse}(\mathbf{W}_i \mathbf{z}_n + \hat{\boldsymbol{\mu}}_i)] \\ & \leq -\frac{1}{2} \text{vec}(\mathbf{W}_i)^T (\mathbb{E}[\mathbf{z}_n \mathbf{z}_n^T] \otimes \mathbf{A}) \text{vec}(\mathbf{W}_i) \\ & \quad + \text{vec}(\mathbf{W}_i)^T (\mathbb{E}[\mathbf{z}_n] \otimes (\mathbf{f}_{ni} - \hat{\boldsymbol{\rho}}_{ni} + \mathbf{A}(\hat{\boldsymbol{\eta}}_{ni} - \hat{\boldsymbol{\mu}}_i))) \\ & \quad + \text{const.} \end{aligned} \quad (31)$$

Substituting (Equations 27, 28) into $\mathbb{E}_{q(\mathbf{Z})}[\log p(\mathbf{F}_i, \mathbf{Z}, \mathbf{W}_i | \hat{\boldsymbol{\mu}}_i)]$ and completing the square reveals the following Gaussian distribution, and update for \mathbf{W}_i

$$\text{vec}(\mathbf{W}_i) \sim \mathcal{N}(\text{vec}(\hat{\mathbf{W}}_i), \mathbf{H}_i^{-1}) \quad (32)$$

$$\mathbf{H}_i = \sum_{n=1}^N (\hat{\mathbf{z}}_n \hat{\mathbf{z}}_n^T + \mathbf{V}_n) \otimes \mathbf{A} + \mathbf{I}_K \otimes (\mathbf{I}_M + \frac{1}{1+M}) \quad (33)$$

$$\text{vec}(\hat{\mathbf{W}}_i) = \mathbf{H}_i^{-1} \left(\sum_{n=1}^N \hat{\mathbf{z}}_n \otimes (\mathbf{f}_{ni} - \hat{\boldsymbol{\rho}}_{ni} + \mathbf{A}(\hat{\boldsymbol{\eta}}_{ni} - \hat{\boldsymbol{\mu}}_i)) \right). \quad (34)$$

2.2.3. Conditional Random Field

This section describes how spatial contiguity is achieved. Rather than treat the data as a collection of independent patches, the proposed approach attempts to model the relationship between the latent variables encoding each patch and those encoding the six immediately neighbouring patches (or four neighbouring patches in 2D). For a valid mean-field approximation, updating patches is done via a “red-black” (checkerboard) ordering scheme (see **Figure 3**), such that one pass over the data updates the “red” patches, while making use of the six neighbouring patches of each (which would correspond to ‘black’ patches). The next pass would update the “black” patches, while making use of the six (“red”) neighbouring patches of each. The remainder of this section explains how information from neighbouring patches is used to provide empirical priors (\mathbf{z}_0 and \mathbf{P}_0) for the latent variables of a central patch. We note that this approach is related to work by Zheng et al. (2015) and Brudfors et al. (2019).

Recall that the posterior means and covariances of each latent variable are denoted $\hat{\mathbf{z}}_n$ and \mathbf{V}_n . Here, we refer to the concatenated means and covariances of the latent variables in all adjacent patches as $\hat{\mathbf{y}}_n$ and \mathbf{U}_n , respectively, where \mathbf{U}_n is block diagonal. We assume that

$$\begin{bmatrix} \mathbf{z}_n \\ \mathbf{y}_n \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{P}^{-1}). \quad (35)$$

Using K to denote the order of \mathbf{P} , the model assumes a Wishart prior on \mathbf{P} .

$$\begin{aligned} \ln p(\mathbf{P}) &= \ln \mathcal{W}(\boldsymbol{\Psi}_0, \nu_0) \\ &= \frac{\nu_0 - K - 1}{2} \ln \det |\mathbf{P}| - \frac{1}{2} \text{Tr}(\boldsymbol{\Psi}_0^{-1} \mathbf{P}) \\ & \quad - \frac{\nu_0}{2} \ln \det |\boldsymbol{\Psi}_0| - \frac{\nu_0 K}{2} \ln 2 - \ln \Gamma_K \left(\frac{\nu_0}{2} \right) \end{aligned} \quad (36)$$

This leads to the following approximating distribution for \mathbf{P} .

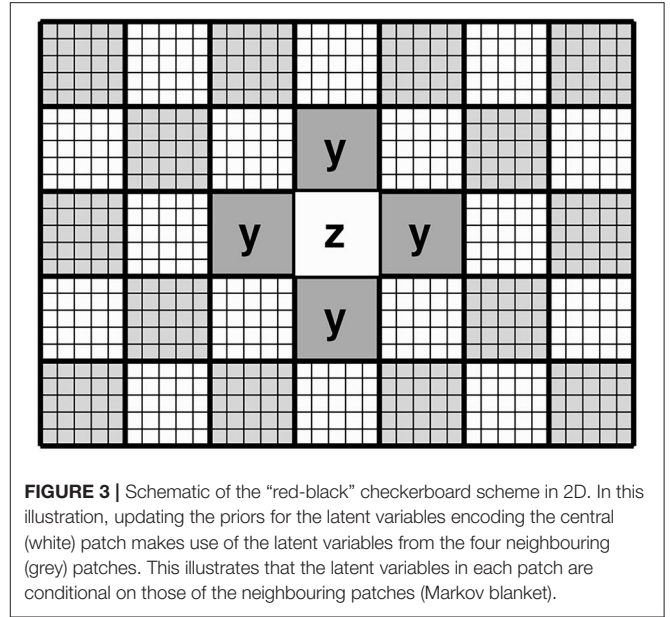


FIGURE 3 | Schematic of the “red-black” checkerboard scheme in 2D. In this illustration, updating the priors for the latent variables encoding the central (white) patch makes use of the latent variables from the four neighbouring (grey) patches. This illustrates that the latent variables in each patch are conditional on those of the neighbouring patches (Markov blanket).

$$\begin{aligned} \ln q(\mathbf{P}) &= \mathbb{E}_{q(\mathbf{Z}), q(\mathbf{Y})}[\ln p(\mathbf{Z}, \mathbf{Y} | \mathbf{P}) + \ln p(\mathbf{P})] \\ &= \frac{N}{2} \ln \det |\mathbf{P}| - \frac{1}{2} \text{Tr} \left(\mathbb{E} \left[\sum_{n=1}^N [\mathbf{z}_n^T \ \mathbf{y}_n^T]^T [\mathbf{z}_n^T \ \mathbf{y}_n^T] \right] \mathbf{P} \right) \\ & \quad + \frac{\nu_0 - K - 1}{2} \ln \det |\mathbf{P}| - \frac{1}{2} \text{Tr}(\boldsymbol{\Psi}_0^{-1} \mathbf{P}) + \text{const} \end{aligned} \quad (37)$$

By substituting expectations from Equations (27) and (28), we can represent this as the Wishart distribution

$$q(\mathbf{P}) = \mathcal{W}(\mathbf{P} | \boldsymbol{\Psi}, \nu) \quad (38)$$

where

$$\boldsymbol{\Psi} = \left(\sum_{n=1}^N \begin{bmatrix} \hat{\mathbf{z}}_n \hat{\mathbf{z}}_n^T + \mathbf{V}_n & \hat{\mathbf{z}}_n \hat{\mathbf{y}}_n^T \\ \hat{\mathbf{y}}_n \hat{\mathbf{z}}_n^T & \hat{\mathbf{y}}_n \hat{\mathbf{y}}_n^T + \mathbf{U}_n \end{bmatrix} + \boldsymbol{\Psi}_0^{-1} \right)^{-1} \quad (39)$$

$$\nu = N + \nu_0. \quad (40)$$

From the properties of Wishart distributions, we have

$$\mathbb{E}[\mathbf{P}] = \nu \boldsymbol{\Psi}. \quad (41)$$

For the next step, the matrices \mathbf{P} and $\boldsymbol{\Psi}$ are conformably decomposed into

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_{zz} & \mathbf{P}_{zy} \\ \mathbf{P}_{yz} & \mathbf{P}_{yy} \end{bmatrix} \text{ and } \boldsymbol{\Psi} = \begin{bmatrix} \boldsymbol{\Psi}_{zz} & \boldsymbol{\Psi}_{zy} \\ \boldsymbol{\Psi}_{yz} & \boldsymbol{\Psi}_{yy} \end{bmatrix}. \quad (42)$$

Now, we wish to compute priors for the latent variables of each patch, conditional on the latent variables of the neighbours. This can be derived from

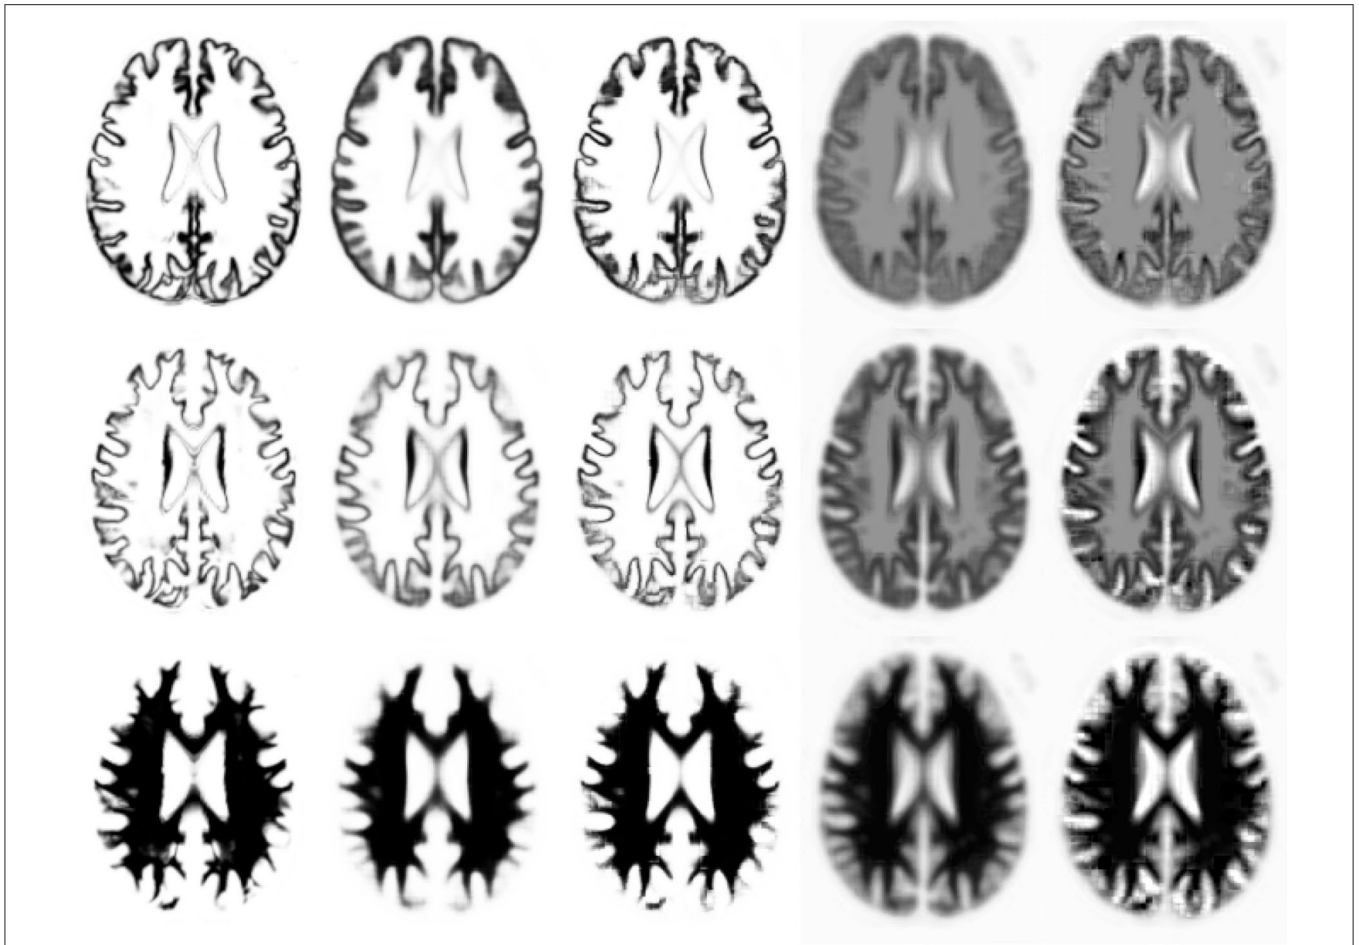


FIGURE 4 | Illustration of (encoding) model fit to an example image, where each row shows a different tissue class. The first column shows the warped categorical image data for one subject. The second and fourth columns shows the mean (parameter M within each patch), with and without a softmax, respectively. The third and fifth columns show the model fit (mean plus linear combination of basis functions within each patch) to the categorical image data, with and without a softmax.

$$\begin{aligned} \mathbb{E}_{q(y_n), q(\mathbf{P})} [\ln p(\mathbf{z}_n, y_n)] &= -\frac{1}{2} \mathbb{E}_{q(y_n), q(\mathbf{P})} \left[[\mathbf{z}_n^T y_n^T] \mathbf{P} [\mathbf{z}_n^T y_n^T]^T \right] \\ &\quad + \text{const} \\ &= -\frac{1}{2} \mathbf{z}_n^T \mathbb{E}[\mathbf{P}_{zz}] \mathbf{z}_n - \mathbf{z}_n^T \mathbb{E}[\mathbf{P}_{zy}] \mathbb{E}[y_n] \\ &\quad + \text{const}. \end{aligned} \tag{43}$$

By completing the square and substituting expectations from Equations (27) and (41), we can derive suitable priors for use in Equation (10) from section 2.1.

$$\mathbf{z}_n \sim \mathcal{N}(\mathbf{z}_0, \mathbf{P}_0^{-1}) \tag{44}$$

where

$$\mathbf{P}_0 = \mathbb{E}[\mathbf{P}_{zz}] = \nu \Psi_{zz}. \tag{45}$$

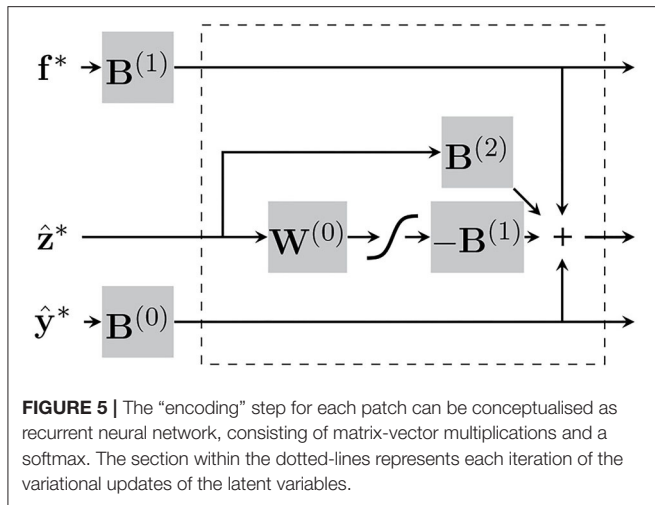
$$\mathbf{z}_0 = -\mathbb{E}[\mathbf{P}_{zz}^{-1} \mathbf{P}_{zy} y_n] = -\Psi_{zz}^{-1} \Psi_{zy} \hat{y}_n. \tag{46}$$

2.2.4. Implementation Details

Training the model is an iterative procedure. Our implementation consists of a number of outer iterations, each involving a “red” and “black” sweep through the data. During each outer iteration, patches are updated by first determining a new \mathbf{P}_0 and a new prior expectation \mathbf{z}_0 for each \mathbf{z}_n . Then from these priors, the variational EM steps are repeated five times within each patch. For each of these sub-iterations, the E-step is run five times, as is the M-step.

Unlike most other methods, our proposed approach performs label propagation on spatially normalised versions of the images. To account for this, our implementation considers the expansions and contractions involved in warping the images using the Jacobian determinants to weight the data appropriately, which is effected by a slight modification to the likelihood term in Equation 15. This weighting essentially is an integration by substitution, and is used both during the training and testing phases.

Because the method is patch-based, not every patch needs to encode all possible brain structure labels. To save memory and computation, the model is set up so that each patch only encodes



the categories that it requires. This is determined by whether that category exists in the corresponding patch in all training scans, and results in dimension $M^{(2)}$ varying across patches.

The Bayesian formulation of the model tends towards an automatic relevance determination solution, whereby the distributions of some latent variables approach a delta function at zero. To speed up the computations, the model is “pruned” after every second iteration of the training. This involves using PCA to make $\hat{\mathbf{Z}}^T \hat{\mathbf{Z}}$ orthonormal within each patch (along with applying the corresponding rotations to each $\hat{\mathbf{W}}_i^{(1)}$, $\hat{\mathbf{W}}_i^{(2)}$ and \mathbf{V}_n). Any latent variables that contributed a negligible amount to the model fit were then removed, along with the basis functions they controlled.

As a simple attempt to make better use of the limited number of labelled images, the training data are augmented by also using versions translated by integer numbers of voxels along all three directions, up to some maximum radius. We weight the contribution of each presentation of the data according to the amount of translation used. This weighting is based on Gaussian function of distance ($\exp(-\frac{1}{2}d^2/s^2)$), and is parameterised by a standard deviation (s). The weights are re-normalised to sum to one over all possible amounts of translation. Weighting enters into the algorithm as a modification to the updates in the variational M-step, as well as when updating the parameters of the conditional random field.

2.3. Labelling a Target Image

In this section, we explain the computations that take place during deployment of a trained model, such that the model can be used for labelling new and unseen images. To re-iterate, labelling a new image involves an encoding step, where the distribution of \mathbf{z}^* is estimated for each target image patch ($\mathbf{F}^{*(1)}$) by fitting $\mathbf{M}^{(1)}$ and $\mathbf{W}^{(1)}$ to it (see Figure 4). This is followed by a decoding step, whereby the label probabilities ($\mathbf{F}^{*(2)}$) are reconstructed from the distribution of \mathbf{z}^* using $\mathbf{M}^{(2)}$ and $\mathbf{W}^{(2)}$. We show that encoding a patch in a new image can be achieved by expressing Equations (26) and (46) as a type of recurrent ResNet (He et al., 2016).

2.3.1. Encoding

Latent variables are all initialised to zero, before the “red-black” scheme is used to update them. One cycle updates the estimates of the latent variables ($\hat{\mathbf{z}}$) for the “red” patches, based on values of the latent variables in the neighbouring “black” patches ($\hat{\mathbf{y}}$). The next cycle updates the latent variables of the “black” patches, using the recently updated neighbouring “red” latent variables. Our implementation repeats this procedure for a fixed number of iterations, although it would be possible to terminate based on some convergence criterion. For each patch, the parameters computed during training that are required during encoding are $\hat{\mathbf{W}}^{(1)}$, $\mathbf{M}^{(1)}$, ν and Ψ . The procedure for computing the distribution of the latent variables \mathbf{z} for a patch in a target image $\mathbf{F}^{*(1)} \in \{0, 1\}^{M \times I}$ can be made more efficient by first pre-computing some new matrices. The following covariance matrix is required, which is obtained by combining 25 and 45.

$$\mathbf{V} = \left(\Psi_{zz} \nu + \sum_{i=1}^I \hat{\mathbf{W}}_i^T \mathbf{A} \hat{\mathbf{W}}_i \right)^{-1} \quad (47)$$

The spatial basis functions ($\hat{\mathbf{W}}^{(1)}$ and $\mathbf{M}^{(1)}$) are reshaped to make them easier to work with.

$$\mathbf{W} = \begin{bmatrix} \hat{\mathbf{w}}_{11}^{(1)} & \hat{\mathbf{w}}_{21}^{(1)} & \dots & \hat{\mathbf{w}}_{K1}^{(1)} \\ \hat{\mathbf{w}}_{12}^{(1)} & \hat{\mathbf{w}}_{22}^{(1)} & \dots & \hat{\mathbf{w}}_{K2}^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{\mathbf{w}}_{1I}^{(1)} & \hat{\mathbf{w}}_{2I}^{(1)} & \dots & \hat{\mathbf{w}}_{KI}^{(1)} \end{bmatrix} \text{ and } \boldsymbol{\mu} = \begin{bmatrix} \mu_1^{(1)} \\ \mu_2^{(1)} \\ \vdots \\ \mu_I^{(1)} \end{bmatrix} \quad (48)$$

The additional matrices that are pre-computed to speed up the updates in 26 are

$$\mathbf{B}^{(0)} = -\mathbf{V} \Psi_{zy} \nu \quad (49)$$

$$\mathbf{B}^{(1)} = \mathbf{V} \mathbf{W}^T \quad (50)$$

$$\mathbf{B}^{(2)} = \mathbf{V} \mathbf{W}^T (\mathbf{I}_I \otimes \mathbf{A}) \mathbf{W}. \quad (51)$$

Each patch in the target image is reshaped to a column vector.

$$\mathbf{f}^* = \begin{bmatrix} \mathbf{f}_1^{*(1)} \\ \mathbf{f}_2^{*(1)} \\ \vdots \\ \mathbf{f}_I^{*(1)} \end{bmatrix} \quad (52)$$

Computing the distribution of the latent variables can then be achieved by iterating the following.

$$\hat{\mathbf{z}}^* \leftarrow \mathbf{B}^{(1)} \mathbf{f}^* + \mathbf{B}^{(0)} \hat{\mathbf{y}}^* + \mathbf{B}^{(2)} \hat{\mathbf{z}}^* - \mathbf{B}^{(1)} \sigma(\mathbf{W} \hat{\mathbf{z}}^* + \boldsymbol{\mu}) \quad (53)$$

In practice, this is iterated five times per patch for every full sweep through the image. This specific number was chosen based on

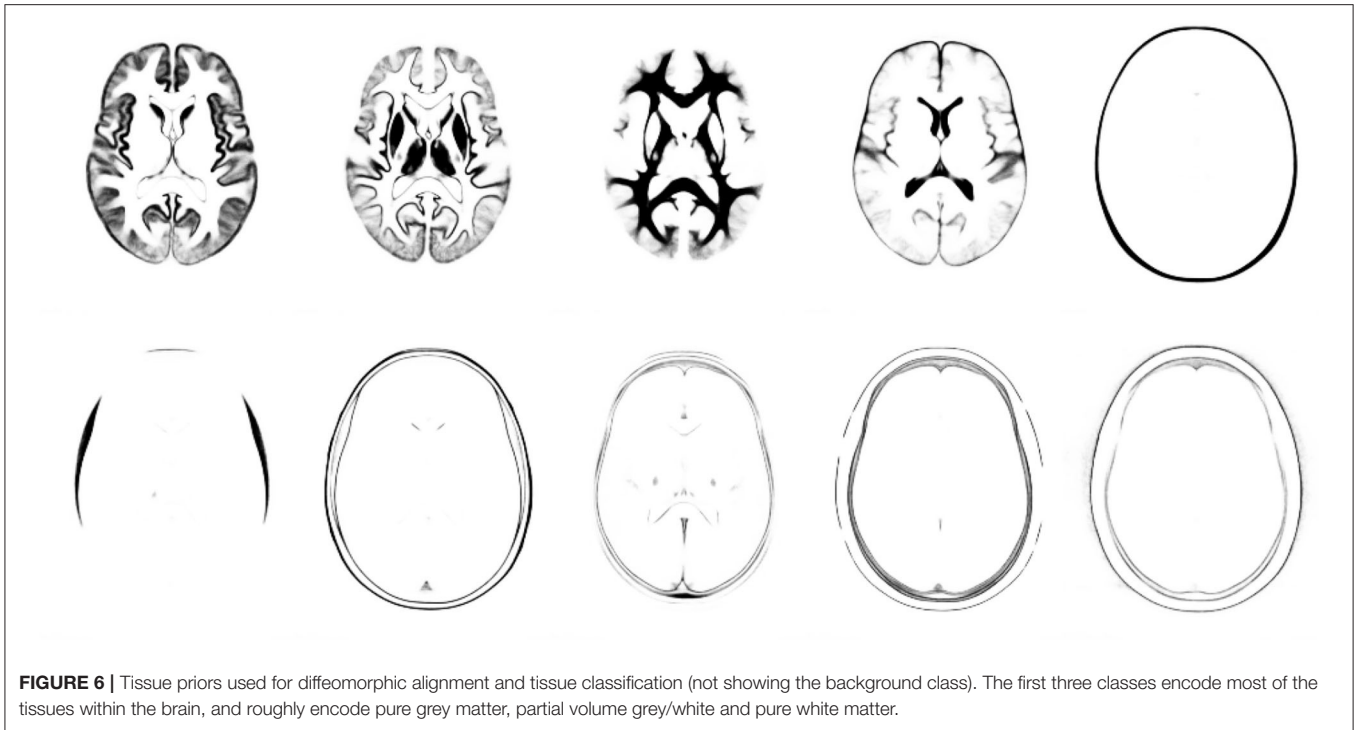


FIGURE 6 | Tissue priors used for diffeomorphic alignment and tissue classification (not showing the background class). The first three classes encode most of the tissues within the brain, and roughly encode pure grey matter, partial volume grey/white and pure white matter.

what was reported in Khan et al. (2010). It may be worth noting that these variational updates of the latent variables consist only of matrix-vector multiplications, additions and a softmax. The procedure can be conceptualised as a sort of ResNet, which we have attempted to illustrate in Figure 5.

2.3.2. Decoding

Once the expectation of the latent variables has been computed for the patches, the probabilistic label map can then be generated using $\hat{\mathbf{W}}^{(2)}$ and $\hat{\mathbf{M}}^{(2)}$. Slightly more accurate probabilities could be achieved by repeated sampling from $\mathbf{z}^* \sim \mathcal{N}(\hat{\mathbf{z}}^*, \mathbf{V}^*)$, but our approach simply reconstructs voxel probabilities using $\hat{\mathbf{z}}^*$.

$$\mathbf{f}_i^{*(2)} \sim \text{Cat} \left(\sigma \left(\hat{\mathbf{W}}_i^{(2)} \hat{\mathbf{z}}^* + \hat{\boldsymbol{\mu}}_i^{(2)} \right) \right) \quad (54)$$

2.3.3. Registration With Trained Model

Labellings achieved from the simple encoding-decoding model, while fast to compute, are of limited accuracy. We note that our proposed model also allows a subject-specific template (see Figure 4) to be generated, such that

$$\mathbf{f}_i^{*(1)} \sim \text{Cat} \left(\sigma \left(\hat{\mathbf{W}}_i^{(1)} \hat{\mathbf{z}}^* + \hat{\boldsymbol{\mu}}_i^{(1)} \right) \right). \quad (55)$$

Because the trained model is able to generate synthetic template images with which new images can be aligned, this leads to a strategy that allows the alignment between any new image and the training data to be improved. Higher labelling accuracies can be achieved by finessing the warps that align the images to label and the trained model. For each subject's image data, this involves

alternating between running the encoding and decoding model to generate a subject-specific template, and using this to refine the diffeomorphic alignment to achieve a closer match with the training data.

3. EXPERIMENTS AND RESULTS

Two experiments were performed. The first involved assessing labelling accuracy compared with a ground truth based on manual annotations. The second involved assessing the replicability of the labelling using different image contrasts.

3.1. Datasets Used

We used the dataset from the *MICCAI 2012 Grand Challenge and Workshop on Multi-Atlas Labelling*, which is available through http://www.neuromorphometrics.com/2012_MICCAI_Challenge_Data.html. These data were provided for use in the MICCAI 2012 Grand Challenge and Workshop on Multi-Atlas Labelling (Landman and Warfield, 2012). The data is released under the Creative Commons Attribution-NonCommercial license (CC BY-NC) with no end date. Original MRI scans are from OASIS (<https://www.oasis-brains.org/>). Labellings were provided by Neuromorphometrics, Inc. (<http://Neuromorphometrics.com/>) under academic subscription. The dataset consists of 35 manually labeled volumetric T1-weighted (T1w) MRI brain scans (1 mm isotropic resolution) from 30 unique subjects, with five of the subjects scanned twice. The dataset is split into 15 training images and 20 testing images, where the testing images included those subjects who were scanned twice (the test-retest subjects). We note that each of the images is an average of several rigidly-aligned face-stripped

raw MRI scans. Because missing voxels in the face-stripped scans were coded with a value of zero, it was necessary to erode the images to remove voxels containing an average of present and missing values.

We also used T1-, T2- and PD-weighted scans from the IXI dataset (EPSRC GR/S21533/02), which is available from <https://brain-development.org/ixi-dataset/>. This dataset is a collection of MR images from almost 600 normal, healthy subjects, which was collected on 1.5T and 3T MRI scanners at three different London hospitals.

3.2. Tissue Probability Atlas

Alignment and tissue segmentation of images prior to training and applying our proposed method used the approach described by Brudfors et al. (2020), which extends ideas presented in Blaiotta et al. (2018). These works combine diffeomorphic image registration and Gaussian mixture model based tissue classification within the same generative model, and also allow average shaped tissue probability maps to be computed. The software is incorporated into SPM12 as the Multi-Brain (MB) toolbox. Default settings were used throughout, except for the regularisation for the diffeomorphic registration, which was set to be higher than the default settings (“Shape Regularisation” on the user interface was set to [0.0001 0.5 0.5 0.0 1.0]).

First of all, the Brudfors et al. (2020) approach was used to construct a tissue probability map from T2-weighted and PD-weighted scans of the first 64 subjects from the IXI dataset, along with the T1-weighted scans of the next 64 IXI subjects. The 15 training subjects’ scans from the MICCAI Challenge Dataset were also included in the template construction. After merging several of the automatically identified tissue classes, the tissue probability map has 1 mm isotropic resolution, dimensions of $191 \times 243 \times 229$ voxels and consists of 11 tissue types, three of which approximately corresponded with brain tissues. This atlas of tissue priors (illustrated in **Figure 6**) was used for all subsequent image alignment and tissue classification.

3.3. Tuning Settings on the MICCAI Challenge Training Dataset

Scans from all 15 training subjects from the *MICCAI 2012 Grand Challenge and Workshop on Multi-Atlas Labelling* dataset were aligned with the tissue prior atlas (described previously) and the brains segmented into three tissue classes using the method of Brudfors et al. (2020). To facilitate training, warped versions of the tissue classes and labels were generated for the training subjects at 1 mm isotropic resolution. Warped tissue class images were saved in NIfTI format, whereas a custom sparse matrix format was used to encode the warped labels.

Subsequently, various settings for training the proposed method were then tuned using data from these 15 subjects, whereby the first 10 were used for model training, with the remaining five used for validation.

Following training, labelling of the five validation scans was done with either:

- No additional diffeomorphic registration.

- Four Gauss-Newton updates of additional diffeomorphic registration, which involves alternating between estimating the latent variables and updating the alignment. The registration regularisation was one quarter of that used for the initial registration/segmentation of the data.

Because label predictions were made in normalised space, the label probabilities were warped back to match the original image volumes using partial volume interpolation (Maes et al., 1997), where they were converted into categorical data by assigning the most probable label at each voxel. Accuracy was assessed by the average Dice-Sorensen Coefficient (DSC), measuring the overlap between the ground truth and our predictions.

It was not feasible to explore all possible model settings via a full grid search, so we selected a few choice settings and examined the DSC that these achieved on the validation set. Because training is slower when augmentation is used, the initial tuning was done without augmentation. Four outer iterations were used during each training attempt.

Because relatively small regions had proven successful in previous label propagation works, we chose to use patch sizes of $4 \times 4 \times 4$ voxels. The model accounts for the covariance among latent variables in neighbouring patches, so it is able to induce locally correlated behaviour across neighbouring patches. If the model can benefit from weights in neighbouring patches always varying in unison, then this will be exploited. Therefore, we did not explore patch sizes and instead focussed on those settings that control the behaviour of the conditional random field.

Initially, the main settings that were varied were those controlling the Wishart prior on the conditional random field (ν_0 and ν_0 , where $\Psi_0^{-1} = \mathbf{I}\nu_0\nu_0$). Three values for ν_0 were used: the first was a value of 1.0, which encodes an improper prior; the second used the least informative proper prior in a way that varied according to how the model was pruned (named “var” in **Table 1**); the third was a relatively uninformative proper prior based on $\nu_0 = 7K - 0.9$. Three different values for ν_0 were also explored. For these experiments, the maximum number of basis functions K was set to nine, although one run involved training with $K = 0$ to serve as a majority-voting baseline.

Once suitable Wishart prior settings were identified, the next step was to continue the tuning using data that have been augmented by translating by up to 1.5 voxels. This scaled the amount of training data by a factor of 19, leading to a concomitant increase in training times. We considered that $K = 16$ would be a reasonable maximum number of latent variables to use for each patch. These experiments varied the standard deviation of the Gaussian weighting used for augmentation. A final run involved augmenting by translating by up to 3 voxels during training, which increased the training time by about a factor of 123.

DSC scores are presented in **Table 1**. Varying the settings of the Wishart prior made relatively little difference to the DSC, although the variable setting for ν_0 with $\nu_0 = 1.0$ was the most effective without additional registration. Augmentation and using additional registration gave greater improvements. Although using the 3 voxel radius augmentation led to less accurate labelling of non-cortical structures, it gave the best overall DSC.

TABLE 1 | Training/validation DSC with different settings.

Settings					DSC-no registration			DSC-with registration		
v_0	v_0	K	r	sd	Overall	Non-cort.	Cortical	Overall	Non-cort.	Cortical
N/A	N/A	0	0.0	0.0	0.7008	0.8215	0.6564	0.7260	0.8341	0.6862
1.0	0.3	9	0.0	0.0	0.7264	0.8370	0.6858	0.7423	0.8429	0.7053
1.0	1.0	9	0.0	0.0	0.7277	0.8374	0.6874	0.7427	0.8427	0.7059
1.0	3.0	9	0.0	0.0	0.7284	0.8373	0.6883	0.7426	0.8428	0.7058
var	0.3	9	0.0	0.0	0.7294	0.8385	0.6894	0.7440	0.8442	0.7072
var	1.0	9	0.0	0.0	0.7302	0.8385	0.6904	0.7435	0.8440	0.7066
var	3.0	9	0.0	0.0	0.7301	0.8386	0.6902	0.7424	0.8436	0.7052
62.1	0.3	9	0.0	0.0	0.7285	0.8366	0.6889	0.7435	0.8426	0.7071
62.1	1.0	9	0.0	0.0	0.7295	0.8369	0.6900	0.7432	0.8428	0.7065
62.1	3.0	9	0.0	0.0	0.7294	0.8374	0.6898	0.7421	0.8430	0.7051
var	1.0	16	1.5	1.0	0.7405	0.8419	0.7033	0.7504	0.8453	0.7156
var	1.0	16	1.5	3.0	0.7408	0.8416	0.7037	0.7506	0.8451	0.7158
var	1.0	16	1.5	10.0	0.7408	0.8416	0.7037	0.7506	0.8450	0.7159
var	1.0	16	3.0	3.0	0.7438	0.8392	0.7087	0.7526	0.8429	0.7194

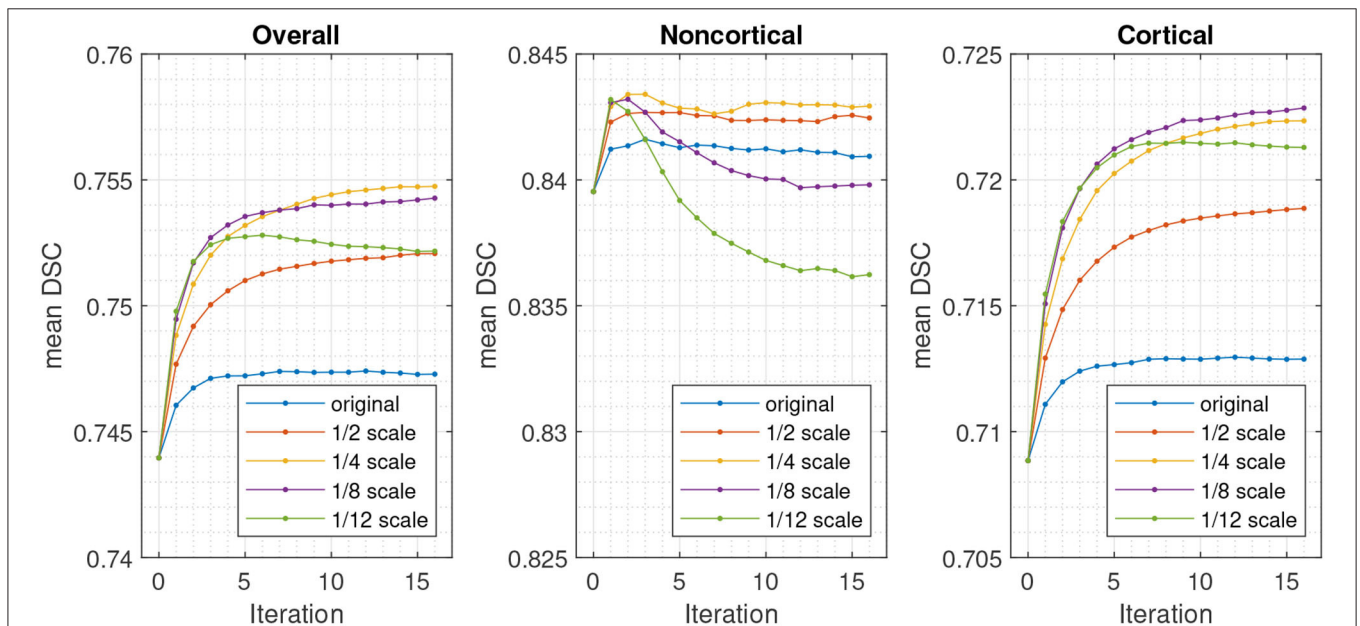


FIGURE 7 | DSC for validation set using different numbers of registration iterations and degrees of regularisation, relative to that used for the initial registration and segmentation.

The effects of the number of Gauss-Newton iterations used for the registration refinement was assessed, as well as the amount of regularisation. This used the model trained with 3 voxel radius augmentation, and involved scaling the regularisation used for the initial registration by various different amounts. For each registration iteration, re-estimation of latent variables used 10 outer iterations (sweeps over the patches) and 16 inner iterations (recomputing latent variables at each patch), which was more than was used for the results in **Table 1** (9 and 5, respectively). The resulting DSC are plotted in **Figure 7**.

3.4. Accuracy on MICCAI Challenge Test Dataset

The FIL model was then re-trained on all 15 training subjects, using a minimally informative, but proper Wishart prior, with $v_0 = 1.0$. An augmentation search radius of 3 voxels was used with a Gaussian weighting standard deviation of 2.0 voxels. Patch sizes were $4 \times 4 \times 4$ voxels, and four outer iterations were used for model training. Up to $K = 24$ basis functions were available to encode each patch, although no more than 15 were needed for any patch after the automatic pruning.

TABLE 2 | Average Dice-Sorenson coefficients (DSC) across testing subjects using the proposed method in comparison with published methods, including the top three approaches from the MICCAI workshop.

Method	Overall	Non-cortical	Cortical
MICCAI-2012: 1. PICSL-BC (Wang et al., 2012)	0.765	0.838	0.739
MICCAI-2012: 2. NonLocalSTAPLE (Asman and Landman, 2012a)	0.758	0.830	0.732
MICCAI-2012: 3. MALP-EM (Ledig et al., 2012)	0.758	0.825	0.733
MICCAI-2012: 4. PICSL-Joint (Wang et al., 2012)	0.750	0.827	0.722
MICCAI-2012: 5. MAPER (Heckemann et al., 2012)	0.741	0.814	0.714
MICCAI-2012: 25. Last-place entry	0.711	0.786	0.683
Joint label fusion + corrective learning (Wang and Yushkevich, 2013)	0.771	0.836	0.747
Random Forest (Zikic et al., 2014)	0.728	0.805	0.699
CNN (de Brebisson and Montana, 2015)	0.725	—	—
CNN (Moeskops et al., 2016)	0.735	0.785	0.717
CNN (Mehta et al., 2017)	0.743	0.805	0.720
Transfer learning FCN (Huo et al., 2019)	0.776	—	—
Majority voting (proposed)	0.717	0.804	0.684
FIL (proposed)	0.744	0.819	0.716

The test subjects were automatically labelled, using six Gauss-Newton iterations for the additional registration. As for the MICCAI challenge, the DSC was computed over all 20 target scans, with results presented as average DSC for cortical labels, non-cortical labels and the overall average. **Table 2** shows the accuracies from the proposed method, alongside previous results from the literature where the same MICCAI challenge data were used. We also trained the model without any basis functions (i.e., $K = 0$) or augmentation, which gave majority voting predictions for the spatially normalised data. These results are also presented in **Table 2** and show that the proposed FIL method increases the overall overlap from a majority voting baseline by about 2.7% (1.5 and 3.2% for non-cortical and cortical regions, respectively). For comparison, the PICSL joint label fusion method (Wang and Yushkevich, 2013) gave similar improvements over majority voting (3.1% overall, 2.9% non-cortical and 3.1% cortical), although the authors achieved additional DSC increases by including their corrective learning step (additional 1.4, 1.1, and 1.5%, respectively). We note that the regularisation used for the registration was quite high, and a higher DSC baseline may have been achieved if this regularisation was lower.

There were 25 entries to the original MICCAI challenge, and the top five entries are also shown in **Table 2**. These were PICSL-BC (Wang et al., 2012), NonLocalSTAPLE (Asman and Landman, 2012a), MALP-EM (Ledig et al., 2012), PICSL-Joint (Wang et al., 2012) (same as PICSL-BC, but without

the corrective learning Wang et al., 2011) and MAPER (Heckemann et al., 2012). Since then, a number of other papers have reported accuracies on these data that were obtained using other methods, so the table also includes several of those results.

While the average DSC from our proposed FIL method were not as high as those from the top performing methods, they would still have achieved fifth position on the leaderboard of the MICCAI challenge. Most of the challenge entries required each test scan to have been registered pairwise with all of the 15 training scans, but our proposed method used a single nonlinear registration for each test scan to the tissue probability template, followed by iterative refinement of the registration, which saves a considerable amount of time. After a single tissue classification and registration (taking about 23 min per subject), the labelling itself took about 24 min (3.25 min without the additional registration) to label each volumetric T1w scan. The laptop computer used in this work is an ASUS ZenBook 14 UX434 with 8 GB of RAM and an AMD Ryzen 5 3500U processor. A GPU implementation would likely lead to much better performance.

Figure 8 illustrates the predicted and ground truth labellings, along with their tissue classifications, for the scans with the highest and lowest average DSC. As can be seen, the white matter hyperintensities in the scan with the lowest average DSC led to less accurate tissue classification, which in turn resulted in less accurate FIL labellings. The DSC (all individuals and mean) for non-cortical and cortical brain regions are shown in **Figures 9, 10**, respectively.

To better understand the upper limit of the accuracies that may be achieved, the test-retest subjects were coregistered together using SPM12's implementation of normalised mutual information coregistration (Studholme et al., 1999) and the label maps that had been manually defined on the second scans were resliced to match those of the first scans using partial volume interpolation. DSC was computed between the first scan labels and the resliced second scan labels and the overall average was found to be 0.816 (0.846 for non-cortical and 0.805 for cortical). Similarly, the hard labels generated by the FIL method for the second scans were resliced, and the overall average DSC was 0.933 (0.932 for non-cortical and 0.934 for cortical). This shows higher test-retest reliability for the FIL method compared with manual labelling.

3.5. Replicability Under Domain Shift

This section assesses the replicability of the proposed label propagation method, by computing DSC between labellings computed from T1w scans, versus those obtained from jointly using T2w and PDw scans of the same subjects. The last 10 subjects for each of the three scanning sites within the IXI dataset were used for this work (Guys Hospital: IXI639 – IXI662; Hammersmith Hospital: IXI632 – IXI646; Institute of Psychiatry: IXI553 – IXI596). The T2w and PDw scans were rigidly aligned with the T1w scans of each subject using normalised mutual information (Studholme et al., 1999).

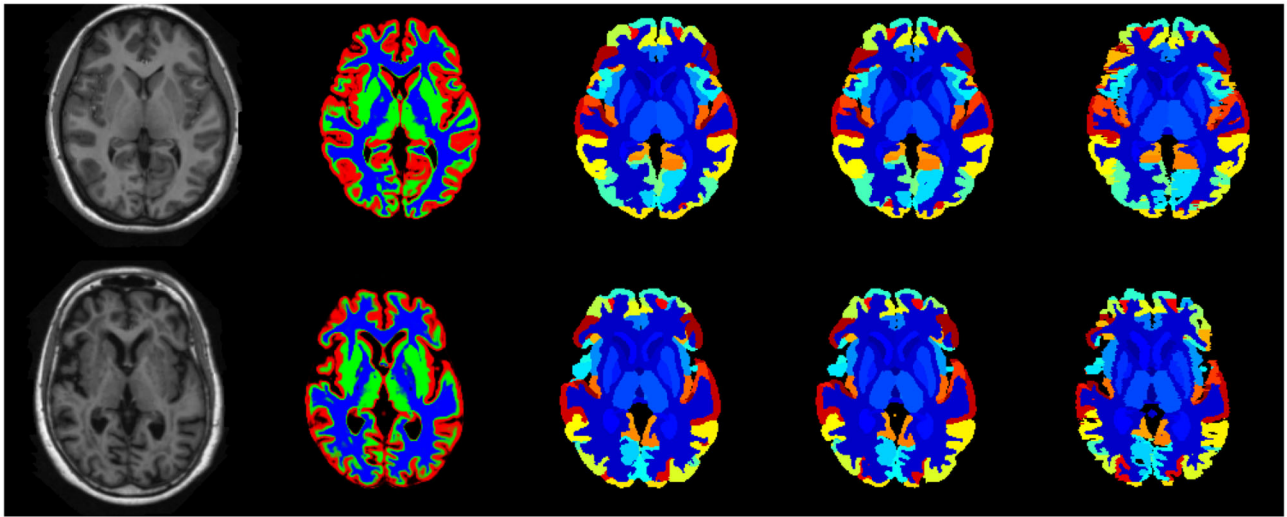


FIGURE 8 | Best case (scan 1,038, top row) and worst case (scan 1128, bottom row) labellings from the fusion challenge dataset. The figure shows the original T1w scan (first column), the tissue classes (second column) identified using the method of Brudfors et al. (2020), the majority voting labelling (third column), the predicted labels using the proposed FIL method (fourth column) and the manually defined ground truth labels (fifth column).

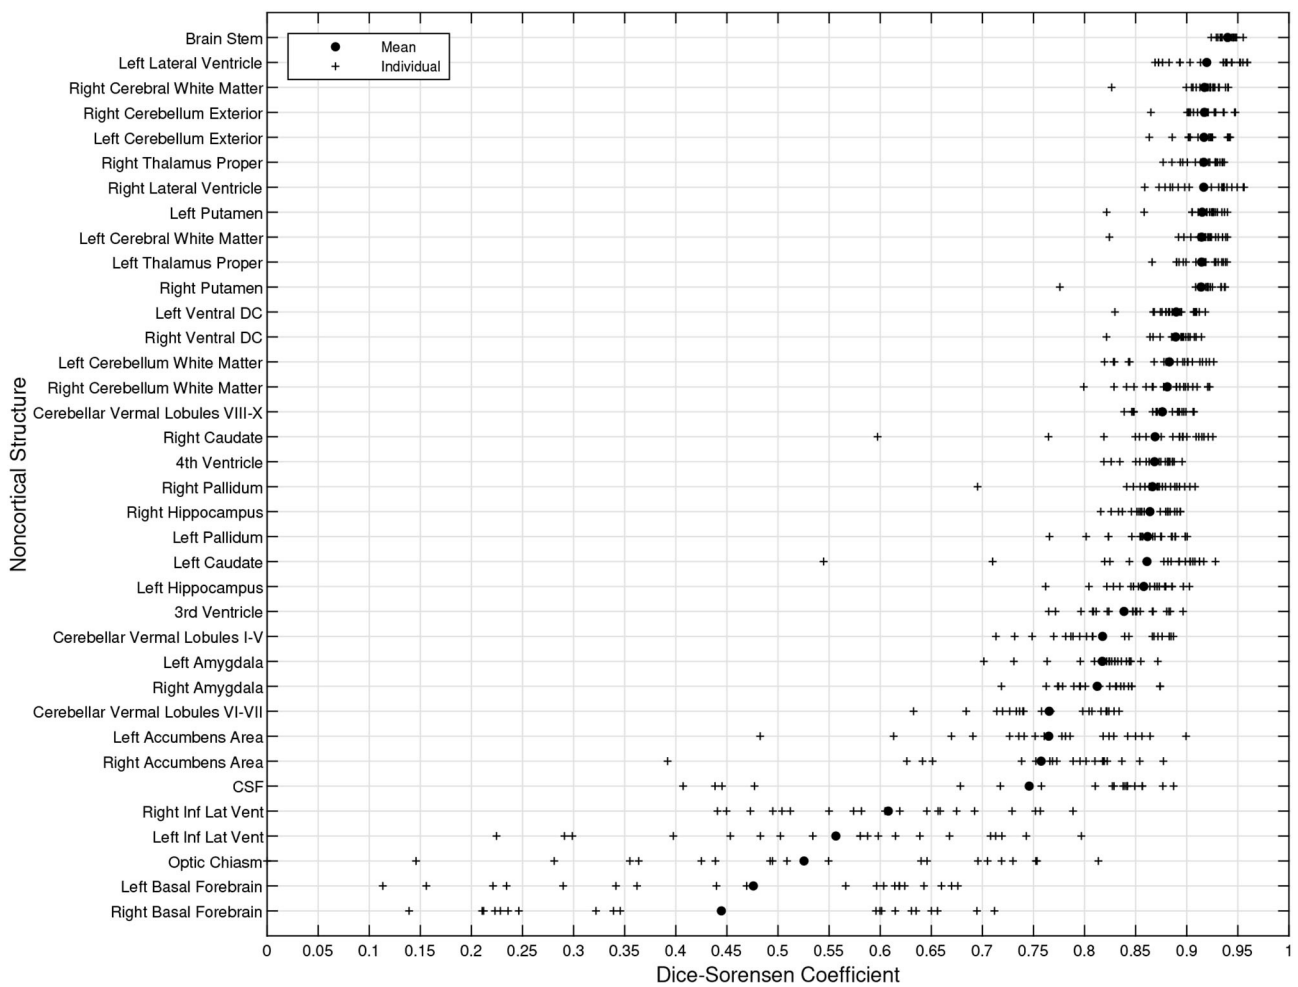
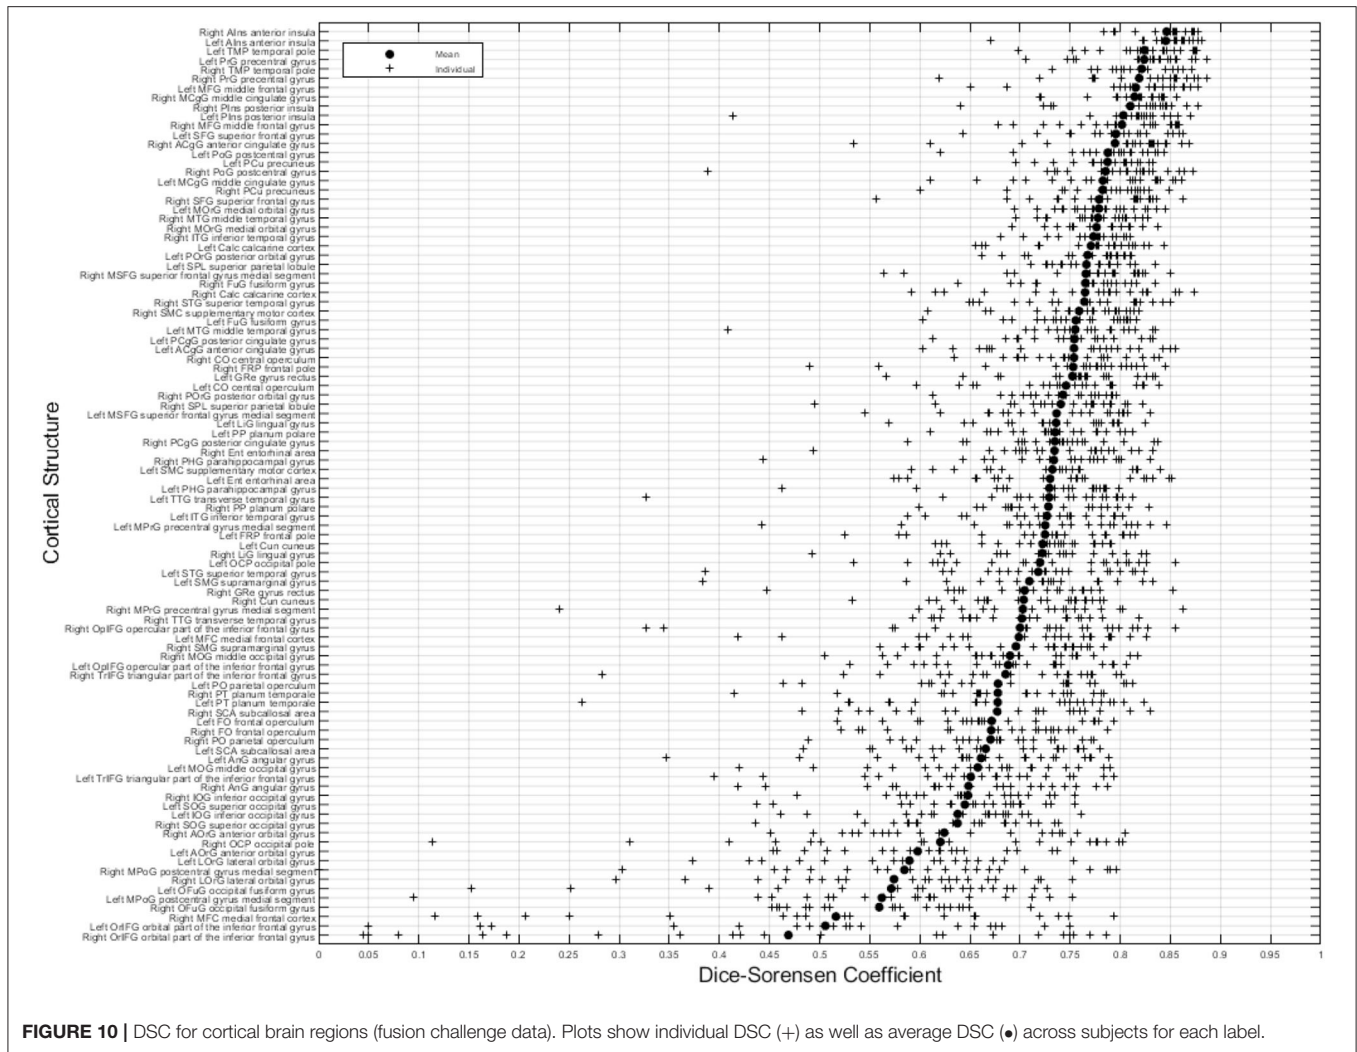


FIGURE 9 | DSC for non-cortical brain regions (fusion challenge data). Plots show individual DSC (+) as well as average DSC (•) across subjects for each brain region.



The scans were subsequently processed as described previously, and labelled using the FIL model trained on the *MICCAI 2012 Workshop on Multi-Atlas Labelling* dataset, as described in the previous section. No ground truth is available for these data so we simply assess the DSC between the two sets of labellings, so that the DSC measures how similar a labelling obtained from a subject's T1w scan is to the labelling obtained from their T2w and PDw scans. Example labellings are shown in **Figure 11** and average DSC from scans from the three different sites are presented in **Table 3**.

As can be seen from **Table 3**, the overall DSC of 0.785 indicates that the two sets of labellings are similar, although there was considerable systematic variability in this similarity across the different sites. Visual inspection of the *IXI* scans (not shown) suggests that better grey-white contrast in the T2w and PDw scans in the Guys Hospital data led to more consistent tissue segmentations, which in turn led to greater labelling consistency. These results suggest that the FIL method is able to generalise well to scans that have sufficient contrast to achieve good tissue classification.

4. DISCUSSION

We have proposed a patch-based, variational Bayesian model for label map prediction using aligned tissue segmented MR scans from SPM12. We computed the DSC for overall brain, cortical and non-cortical regions, which we compared with the *MICCAI 2012* challenge leader table. We also assessed the replicability of labelling MRI data acquired with different image contrasts.

Applying the basic method to a new scan only uses a single image registration, which is subsequently refined, to align with an average shaped template, rather than several separate registration steps to align with all the training scans. We also note that the proposed method is applied to automatically identified tissue classes, which we hope will enable it to generalise better to a broader range of new image types.

In its current form, our proposed labelling approach is not as accurate as some other approaches when applied to scans with the same MRI contrast as the training images. One likely reason for this is that the tissue classification itself may be the limiting factor, as it uses an atlas of tissue priors based only on

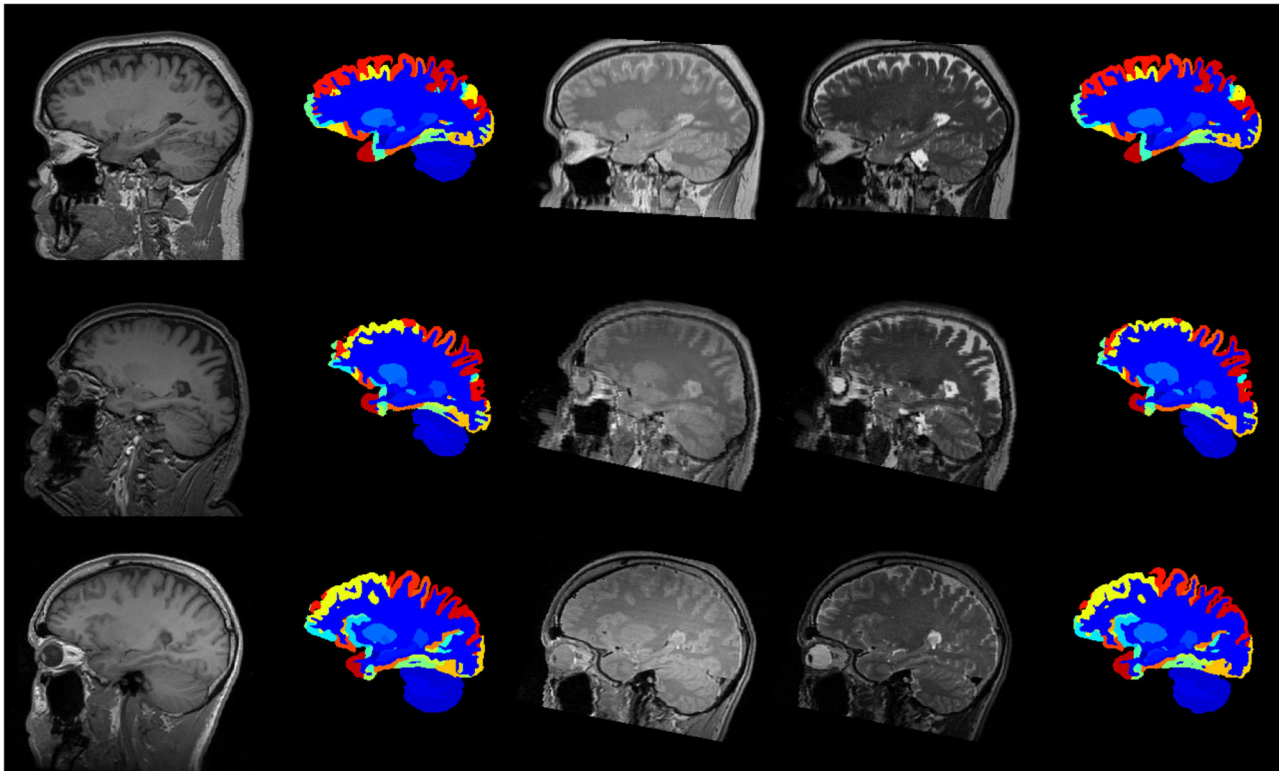


FIGURE 11 | Example labellings of different modalities from three different scanners. First column: T1-weighted scan. Second column: Labelled T1-weighted scan. Third and fourth columns: PD- and T2-weighted scans. Fifth column: Labelled PD- and T2-weighted scans. First row: Guys Hospital; Second row: Hammersmith Hospital; Third row: Institute of Psychiatry.

spatial warping. Because our proposed model is generative, it may be better able to encode priors that could be used for tissue classification, overcoming the limitations of purely warping-based priors. This would also increase compatibility between the tissue classification method and the labelling method, as both would use the same underlying model. However, this may come with the cost of making it more difficult to formulate the training so that the results generalise to scans with a wide variety of MR contrasts. We leave these explorations to future work.

In addition to experimenting with different model settings, there are a number of other areas that could lead to potential improvement. As suggested in Sabuncu et al. (2010), the simple data augmentation approach could probably be improved by augmenting based on the uncertainty of the image registration (Simpson et al., 2012; Iglesias et al., 2013; Wang et al., 2018), which would effectively “integrate out” this source of uncertainty. Such an approach would also need to consider the expected uncertainty with which target images could be aligned. While it benefited cortical segmentation, we found that our simple augmentation strategy often decreased DSC for many non-cortical structures, where registration uncertainty was lower. This supports the idea that it would be more effective to augment according to this uncertainty.

The proposed method is generative, so we speculate that this may make it easier to extend to do semi-supervised and

TABLE 3 | Average DSC between FIL labelling of T1w scans versus jointly labelling T2w and PDw scans.

Method	Overall	Non-cortical	Cortical
FIL: Guys Hospital (Philips 1.5T)	0.835	0.848	0.830
FIL: Hammersmith Hospital (Philips 3T)	0.776	0.754	0.784
FIL: Institute of Psychiatry (GE 1.5T)	0.745	0.743	0.745
FIL: Overall	0.785	0.782	0.786

The average DSC was computed for data from each site, as well as the overall average DSC.

multi-task learning. For example, it might be possible to achieve greater accuracy by training using additional data that does not have manually defined labels associated with it. Although we know that this would allow the model parameters used for encoding (i.e., $\hat{\mathbf{W}}^{(1)}$, $\mathbf{M}^{(1)}$, ν and Ψ) to be more accurately characterised, we can currently only speculate on whether this would translate into more accurate labelling. Similarly, the model could be learned by combining sets of annotations defined using different labelling protocols. While this could be another approach for estimating more accurate encoding parameters, we do not yet know whether encoding many different aspects of images using the same sets of latent variables would improve overall performance.

DATA AVAILABILITY STATEMENT

The code for the Factorisation-based Image Labelling and Multi-Brain toolbox is available at the following link: <https://github.com/WCHN/Label-Training/> and at the following link: <https://github.com/WTCN-computational-anatomy-group/mb>, will be released with the next set of updates, available at the following link: <https://www.fil.ion.ucl.ac.uk/spm/software/spm12/>.

AUTHOR CONTRIBUTIONS

YY contributed to the code, performed the evaluations, and helped draft the manuscript. YB and MB provided important

advice and helped draft the manuscript. JA formulated the equations, contributed to the code, ran some evaluations, and helped draft the manuscript. All authors contributed to the article and approved the submitted version.

FUNDING

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No. 785907 (Human Brain Project SGA2). This research was funded in whole, or in part, by the Wellcome Trust [203147/Z/16/Z].

REFERENCES

- Ashburner, J., and Friston, K. (1997). Multimodal image coregistration and partitioning—a unified framework. *Neuroimage* 6, 209–217. doi: 10.1006/nimg.1997.0290
- Ashburner, J., and Friston, K. J. (2005). Unified segmentation. *Neuroimage* 26, 839–851. doi: 10.1016/j.neuroimage.2005.02.018
- Asman, A., and Landman, B. (2012a). “Multi-atlas segmentation using non-local STAPLE,” in *MICCAI Workshop on Multi-Atlas Labeling France*, 87–90.
- Asman, A. J., and Landman, B. A. (2012b). Formulating spatially varying performance in the statistical fusion framework. *IEEE Trans. Med. Imaging* 31, 1326–1336. doi: 10.1109/TMI.2012.2190992
- Avants, B. B., Epstein, C. L., Grossman, M., and Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* 12, 26–41. doi: 10.1016/j.media.2007.06.004
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York, NY: Springer.
- Blaiziot, C., Freund, P., Cardoso, M. J., and Ashburner, J. (2018). Generative diffeomorphic modelling of large MRI data sets for probabilistic template construction. *Neuroimage* 166, 117–134. doi: 10.1016/j.neuroimage.2017.10.060
- Böhning, D. (1992). Multinomial logistic regression algorithm. *Ann. Inst. Stat. Math.* 44, 197–200. doi: 10.1007/BF00048682
- Brudfors, M., Balbastre, Y., and Ashburner, J. (2019). “Flexible bayesian modelling for nonlinear image registration,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (New York, NY: Springer), 253–263.
- Brudfors, M., Balbastre, Y., Flandin, G., Nachev, P., and Ashburner, J. (2020). Flexible Bayesian modelling for nonlinear image registration. *arXiv preprint arXiv:2006.02338*. doi: 10.1007/978-3-030-59716-0_25
- Commowick, O., Akhondi-Asl, A., and Warfield, S. K. (2012). Estimating a reference standard segmentation with spatially varying performance parameters: Local MAP STAPLE. *IEEE Trans. Med. Imaging* 31, 1593–1606. doi: 10.1109/TMI.2012.2197406
- Coupé, P., Manjon, J. V., Fonov, V., Pruessner, J., Robles, M., and Collins, D. L. (2013). Nonlocal patch-based label fusion for hippocampus segmenta. *Med. Image Comput. Comput. Assist. Interv.* 13, 129–136. doi: 10.1007/978-3-642-15711-0_17
- de Brebisson, A., and Montana, G. (2015). “Deep neural networks for anatomical brain segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (New York, NY: IEEE), 20–28.
- De Fauw, J., Ledsam, J. R., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., et al. (2018). Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat. Med.* 24, 1342–1350. doi: 10.1038/s41591-018-0107-6
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). “Deep residual learning for image recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV: IEEE), 770–778.
- Heckemann, R. A., Keihaninejad, S., Ledig, C., Aljabar, P., Rueckert, D., Hajnal, J. V., et al. (2012). “Multi-atlas propagation with enhanced registration-MAPER,” in *MICCAI 2012 Workshop on Multi-Atlas Labeling France*, 83.
- Huo, Y., Xu, Z., Xiong, Y., Aboud, K., Parvathaneni, P., Bao, S., et al. (2019). 3D whole brain segmentation using spatially localized atlas network tiles. *Neuroimage* 194, 105–119. doi: 10.1016/j.neuroimage.2019.03.041
- Iglesias, J. E., Sabuncu, M. R., and Van Leemput, K. (2013). Improved inference in Bayesian segmentation using monte carlo sampling: application to hippocampal subfield volumetry. *Med. Image Anal.* 17, 766–778. doi: 10.1016/j.media.2013.04.005
- Khan, M. E., Bouchard, G., Murphy, K. P., and Marlin, B. M. (2010). “Variational bounds for mixed-data factor analysis,” in *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates), 1108–1116.
- Kushibar, K., Valverde, S., González-Villá, S., Bernal, J., Cabezas, M., Oliver, A., et al. (2019). Supervised domain adaptation for automatic sub-cortical brain structure segmentation with minimal user interaction. *Sci. Rep.* 9, 1–15. doi: 10.1038/s41598-019-43299-z
- Landman, B., and Warfield, S. (2012). Landman, B., and Warfield, S. (2012). *MICCAI 2012 Workshop on Multi-Atlas Labeling*. Virginia Beach, VA: CreateSpace Independent Publishing Platform.
- Ledig, C., Heckemann, R. A., Aljabar, P., Wolz, R., Hajnal, J. V., Hammers, A., et al. (2012). “Segmentation of MRI brain scans using MALP-EM,” in *MICCAI Workshop on Multi-Atlas Labeling France*, 79–82.
- Ledig, C., Heckemann, R. A., Hammers, A., and Rueckert, D. (2013). “Improving whole-brain segmentations through incorporating regional image intensity statistics,” in *Medical Imaging 2013: Image Processing*, eds S. Ourselin, and D. R. Haynor (New York, NY: SPIE), 442–448. doi: 10.1117/12.2006966
- Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., and Suetens, P. (1997). Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imaging* 16, 187–198. doi: 10.1109/42.563664
- Mehta, R., Majumdar, A., and Sivaswamy, J. (2017). BrainSegNet: a convolutional neural network architecture for automated segmentation of human brain structures. *J. Med. Imaging* 4:024003. doi: 10.1117/1.JMI.4.2.024003
- Moeskops, P., Viergever, M. A., Mendrik, A. M., De Vries, L. S., Benders, M. J., and Išgum, I. (2016). Automatic segmentation of MR brain images with a convolutional neural network. *IEEE Trans. Med. Imaging* 35, 1252–1261. doi: 10.1109/TMI.2016.2548501
- Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. Cambridge, CB: MIT Press. p.359.
- Pohl, K. M., Fisher, J., Grimson, W. E. L., Kikinis, R., and Wells, W. M. (2006). A Bayesian model for joint segmentation and registration. *Neuroimage* 31, 228–239. doi: 10.1016/j.neuroimage.2005.11.044
- Puonti, O., Iglesias, J. E., and Van Leemput, K. (2016). Fast and sequence-adaptive whole-brain segmentation using parametric Bayesian modeling. *Neuroimage* 143, 235–249. doi: 10.1016/j.neuroimage.2016.09.011

- Rashed, E. A., Gomez-Tames, J., and Hirata, A. (2020). End-to-end semantic segmentation of personalized deep brain structures for non-invasive brain stimulation. *Neural Netw.* 125, 233–244. doi: 10.1016/j.neunet.2020.02.006
- Rohlfing, T., Brandt, R., Menzel, R., and Maurer C. R. Jr. (2004). Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains. *Neuroimage* 21, 1428–1442. doi: 10.1016/j.neuroimage.2003.11.010
- Roy, A. G., Conjeti, S., Navab, N., and Wachinger, C. (2019). QuickNAT: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *Neuroimage* 186, 713–727. doi: 10.1016/j.neuroimage.2018.11.042
- Roy, A. G., Conjeti, S., Sheet, D., Katouzian, A., Navab, N., and Wachinger, C. (2017). “Error corrective boosting for learning fully convolutional networks with limited data,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Springer), 231–239.
- Sabuncu, M. R., Yeo, B. T., Van Leemput, K., Fischl, B., and Golland, P. (2010). A generative model for image segmentation based on label fusion. *IEEE Trans. Med. Imaging* 29, 1714–1729. doi: 10.1109/TMI.2010.2050897
- Simpson, I. J., Woolrich, M. W., Andersson, J. R., Groves, A. R., and Schnabel, J. (2012). Ensemble learning incorporating uncertain registration. *IEEE Trans. Med. Imaging* 32, 748–756. doi: 10.1109/TMI.2012.2236651
- Studholme, C., Hill, D. L., and Hawkes, D. J. (1999). An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognit.* 32, 71–86. doi: 10.1016/S0031-3203(98)00091-0
- Tang, X., Oishi, K., Faria, A. V., Hillis, A. E., Albert, M. S., Mori, S., et al. (2013). Bayesian parameter estimation and segmentation in the multi-atlas random orbit model. *PLoS ONE* 8:e65591. doi: 10.1371/journal.pone.0065591
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., et al. (2002). Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 15, 273–289. doi: 10.1006/nimg.2001.0978
- Van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P. (1999). Automated model-based tissue classification of MR images of the brain. *IEEE Trans. Med. Imaging* 18, 897–908. doi: 10.1109/42.811270
- Wachinger, C., Reuter, M., and Klein, T. (2018). DeepNAT: Deep convolutional neural network for segmenting neuroanatomy. *Neuroimage* 170:434–445. doi: 10.1016/j.neuroimage.2017.02.035
- Wang, H., Avants, B., and Yushkevich, P. (2012). “A combined joint label fusion and correctivelearning approach,” in *MICCAI Workshop on Multi-Atlas Labeling* France, 91–94.
- Wang, H., Das, S. R., Suh, J. W., Altinay, M., Pluta, J., Craige, C., et al. (2011). A learning-based wrapper method to correct systematic errors in automatic image segmentation: consistently improved performance in hippocampus, cortex and brain segmentation. *Neuroimage* 55, 968–985. doi: 10.1016/j.neuroimage.2011.01.006
- Wang, H., and Yushkevich, P. (2013). Multi-atlas segmentation with joint label fusion and corrective learning—an open source implementation. *Front. Neuroinform.* 7:27. doi: 10.3389/fninf.2013.00027
- Wang, J., Wells, W. M., Golland, P., and Zhang, M. (2018). “Efficient Laplace approximation for Bayesian registration uncertainty quantification,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* ((New York, NY: Springer), 880–888.
- Warfield, S. K., Zou, K. H., and Wells, W. M. (2004). Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Trans. Med. Imaging* 23, 903–921. doi: 10.1109/TMI.2004.828354
- Wells, W. M., Grimson, W. E. L., Kikinis, R., and Jolesz, F. A. (1996). Adaptive segmentation of MRI data. *IEEE Trans. Med. Imaging* 15, 429–442. doi: 10.1109/42.511747
- Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* 20, 45–57. doi: 10.1109/42.906424
- Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., et al. (2015). “Conditional random fields as recurrent neural networks,” in *Proceedings of the IEEE International Conference on Computer Vision* (Santiago: IEEE), 1529–1537.
- Zhu, X. J. (2005). *Semi-supervised learning literature survey*. Technical report, University of Wisconsin-Madison Department of Computer Sciences.
- Zikic, D., Glocker, B., and Criminisi, A. (2014). Encoding atlases by randomized classification forests for efficient multi-atlas label propagation. *Med. Image Anal.* 18, 1262–1273. doi: 10.1016/j.media.2014.06.010

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Yan, Balbastre, Brudfors and Ashburner. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.