

Unified Segmentation Revisited

Mikael Brudfors, Yaël Balbastre, Guillaume Flandin, and John Ashburner

The Wellcome Centre for Human Neuroimaging, UCL, London, UK
{mikael.brudfors.15,y.balbastre,g.flandin,j.ashburner}@ucl.ac.uk

In this paper we describe a brain segmentation method defined by a generative model over a population of MR (magnetic resonance) images. Here, this population consists of the seven sets of original brain MR scans¹, with their respective manual labels, provided by the MRBrainS18 challenge organisers. Our model assumes two forms of prior knowledge: (1) a probabilistic tissue atlas, which is deformed to a subject’s brain; and (2), known hyper-parameters over the model’s intensity distribution. During training, these two priors are learnt from the population by inverting the model using Bayes’ rule (Figure 1). During testing, these learnt priors are used to segment new, unseen subjects. Defining an accurate generative model can enable using less training data to infer on unseen data. Here, our aim is therefore to use no additional data to enrich the training set. The model can however use data from other populations, if provided.




Fig. 1: Class estimates from the $K = 10$ tissue atlas learnt using the generative model presented in this paper. The colors represent: 1. Background, 2. Cortical gray matter, 3. Basal ganglia, 4. White matter, 5. White matter lesions, 6. Cerebrospinal fluid in the extracerebral space, 7. Ventricles, 8. Cerebellum, 9. Brain stem, 10. Other.

Our model builds on the unified segmentation routine developed for the Statistical Parametric Mapping (SPM) software [1]. SPM is used by neuroimaging researchers around the world to test hypotheses about functional MR imaging data. This is the reason we want to work with the data in its original form — in order to keep the segmentation routine as general as possible. Unified segmentation uses a Gaussian Mixture Model (GMM) to model the intensity distribution

¹ Using the data in its original form means the MR images have not been: resliced to the same size, rigidly aligned, bias-field corrected, or skull-stripped.

of an MR image, and registers a probabilistic atlas into alignment with the image, encoded by a multinomial distribution. Furthermore, it also estimates a correction for the bias-field of the MR image.

Two recent publications [2, 3] extended the hierarchical structure of the unified segmentation model by: (1) placing priors on the parameters of the GMM and then using Variational Bayes (VB) to infer the GMM parameters; (2) extending the small deformation model used by the image registration to instead combining affine and diffeomorphic registration; (3) estimating the probabilistic atlas and the hyper-parameters of the VB GMM, from a population of subjects; (4) ability to infer missing voxels between different MR contrasts; and (5), made the unsupervised model semi-supervised by making it possible to include manual labels into the generative model. Although these extensions, in principle, should make for a powerful brain segmentation tool, it has not yet gained momentum. This is because of difficulties in implementing the algorithm and improvements that were necessary to the model.

Our segmentation model aims to improve on the one presented in the previous paragraph — making it work in practice. In order to do so, we have made some key improvements: (1) the non-uniformity of MR image intensities makes it impossible to learn the hyper-parameters of the intensity distribution, we therefore now normalise the intensities between subjects inside the model; (2) an atlas generated from a small number of subjects may be a poor representation of the anatomy of the general population, we therefore now introduce spatial blurring to regularise this average by a zero mean Gaussian prior on the atlas; (3) the labelled information was not propagated optimally across different tissue classes, we have now improved how manual labels are included in the generative model; and (4), careful initialisation of the model parameters is necessary for good convergence and we now initialise all parameters by first fitting the model to histogram representations of the input images.

Additionally, a super-resolution algorithm was used to provide the input data to the segmentation routine, which reconstructed 1 mm isotropic versions of each set of MR scans [4]. To produce the final 8-label segmentation we make maximum-likelihood estimates of the 8-classes of the probabilistic segmentations compromising the tissues of interest (cortical gray matter, basal ganglia, white matter, white matter lesions, cerebrospinal fluid in the extracerebral space, ventricles, cerebellum and brain stem).

References

1. J. Ashburner and K. J. Friston, “Unified segmentation,” *Neuroimage*, vol. 26, no. 3, pp. 839–851, 2005.
2. C. Blaiotta, M. J. Cardoso, and J. Ashburner, “Variational inference for medical image segmentation,” *Comput Vis Image Understand*, vol. 151, pp. 14–28, 2016.
3. C. Blaiotta, P. Freund, M. J. Cardoso, *et al.*, “Generative diffeomorphic modelling of large MRI data sets for probabilistic template construction,” *NeuroImage*, 2017.
4. M. Brudfors, Y. Balbastre, P. Nachev, and J. Ashburner, “MRI Super-Resolution using Multi-Channel Total Variation,” in *MIUA*, Springer, 2018.