Multi-Brain Manual

Current and former members of the FIL Computational Anatomy Group (and honorary members)

Mikael Brudfors
Yael Balbastre
John Ashburner

Functional Imaging Laboratory
Wellcome Centre for Human Neuroimaging
UCL Queen Square Institute of Neurology
12 Queen Square, London WC1N 3BG, UK
September 19, 2020
https://www.fil.ion.ucl.ac.uk/spm/

Requirements

The algorithm is developed using MATLAB and relies on external functionality from the SPM12 software, which can be download from https://www.fil.ion.ucl.ac.uk/spm/software/spm12/.

It is easiest to run the software as a toolbox of SPM12. To do this, simply move or copy the Multi-Brain software directory into the spm/toolbox directory, which will allow Multi-Brain to be run as an SPM toolbox. This manual relates to its use within this framework.

Alternatively, experts may run the Multi-Brain software separately, but need to ensure that the MATLABPATH includes the main SPM12 source code directory, as well as the directories containing

- *Shoot toolbox*: Add Shoot folder from the toolbox directory of the SPM source code.
- *Longitudinal toolbox*: Add Longitudinal folder from the toolbox directory of the SPM source code.

Chapter 1

Multi-Brain toolbox

Contents

The Multi-Brain (MB) toolbox has the general aim of integrating a number of disparate image analysis components within a single unified generative modelling framework (segmentation, nonlinear registration, image translation, etc.). The model [10] builds on a number of previous works and has the objective of achieving diffeomorphic alignment of a wide variety of medical image modalities into a common anatomical space. This involves the ability to construct a "tissue probability template" from one or more populations of scans through group-wise alignment [3, 8]. Diffeomorphic deformations are computed within a geodesic shooting framework [4], which is optimised with a Gauss-Newton strategy that uses a multi-grid approach to solve the system of linear equations [1]. Variability among image contrasts is modelled using a more sophisticated

version of the Gaussian mixture model with bias correction framework originally proposed in the "Unified Segmentation" paper [2], and which has been extended to account for known variability of the intensity distributions of different tissues [7]. This model has been shown to provide a good model of the intensity distributions of different imaging modalities [9].

This work was funded by the EU Human Brain Project's Grant Agreement No 785907 (SGA2).

1.1 Fit Multi-Brain model

This is where the model fitting is actually done. The outputs from the model fitting are an initial velocity field and deformation field for each subject. In addition, fitting the model may also generate a representation of a set of average shaped tissue probability maps that can serve as a template for future model fitting.

1.1.1 Template

The model can be run using a pre-computed template, or it can implicitly create an average shaped template from the population(s) of scan data. Here, the user gets to choose whether to create a template or use an existing one. Templates are named "mu_*.nii".

Create template

A tissue probability template will be constructed from all the aligned images. The algorithm alternates between re-computing the template and re-aligning all the images with this template. The user gets to choose the voxel size of the template and the number of tissue types it encodes.

Number of classes Specify K, the number of tissue classes encoded by the template. This value is ignored if it is incompatible with the specified data.

Voxel size Specify the voxel size of the template to be created (mm). The algorithm will automatically attempt to determine suitable settings for its orientation and field of view.

Existing template

The model can be fit using a previously computed template, which is not updated. Note that the template contains K-1 volumes within it, and that K should be compatible with various aspects of the data to which the model is fit. The template does not actually encode the tissue probabilities, but rather these probabilities can be generated from the template using a Softmax function ($\mathbf{p} = \frac{\exp \mathbf{p}}{1+\sum_k \exp p_k}$).

1.1.2 Affine

Type of affine transform to use in the model, which may be either none, translations only (T(3)) or rigid body (SE(3)). The fitting begins with affine registration, before continuing by interleaving affine and diffeomorphic registrations over multiple spatial scales.

1.1.3 Shape regularisation

Specify the regularisation settings for the diffeomorphic registration. These consist of a vector of five values, which penalise different aspects of the warps:

- Absolute displacements need to be penalised by a tiny amount. The first element encodes the amount of penalty on these. Ideally, absolute displacements should not be penalised, but it is usually necessary for technical reasons.
- The 'membrane energy' of the deformation is penalised, usually by a relatively small amount. This penalises the sum of squares of the derivatives of the velocity field (i.e., the sum of squares of the elements of the Jacobian tensors).

- The 'bending energy' is penalised (3rd element). This penalises the sum of squares of the 2nd derivatives of the velocity.
- Linear elasticity regularisation is also included. This parameter (μ) is similar to that for linear elasticity, except it penalises the sum of squares of the Jacobian tensors after they have been made symmetric (by averaging with the transpose). This term essentially penalises length changes, without penalising rotations.
- The final term also relates to linear elasticity, and is the weight that denotes how much to penalise changes to the divergence of the velocities (λ). This divergence is a measure of the rate of volumetric expansion or contraction.

The default settings work reasonably well for most cases.

1.1.4 Output name

Specify a key string for inclusion within all the output file names.

1.1.5 Output directory

All output is written to the specified directory. If this is not specified, the current working directory will used by default.

1.1.6 Classes

Images might have been segmented previously into a number of tissue classes. This framework allows such pre-segmented images to be included in the model fitting, in a similar way to the old Dartel toolbox for SPM. The user sets up a series of tissue class types (e.g., grey matter and white matter).

Class

For each of the tissue class types, the user should specify the data to be included within the model fitting by selecting the files. It is important that the subject ordering of the files is the same across all classes.

1.1.7 Populations

Multiple populations of subjects may be combined. For example, there may be T1-weighted scans and manually defined labels for one population, whereas another population may have T2-weighted and PD-weighted scans without labels. Yet another population might have CT scans. All subject's data would be subdivided into the same tissue classes, although the intensity distributions of these tissues is likely to differ across populations.

Pop. of scans

Information about a population of subjects that all have the same set of scans.

Channels Multiple image channels may be specified. For example, two channels may be used to contain the T2-weighted and PD-weighted scans of the subjects.

Channel There may be multiple scans of different modalities for each subject. These would be entered into different channels. Note that all scans within a channel should be the same modality.

Scans Select one NIfTI format scan for each subject. Subjects must be in the same order if there are multiple channels. Image dimensions can differ over subjects, but (if there are multiple channels) the scans of each subject must all have the same dimensions and orientations.

Intensity nonuniformity Specify the intensity nonuniformity (INU) settings for the current channel, which consist of a regularisation setting and a cutoff.

Regularisation Specify the bending energy penalty on the estimated intensity nonuniformity (INU) fields (bias fields). Larger values give smoother INU fields.

Cut off Specify the cutoff (mm) of the intensity nonuniformity (INU) correction (bias correction). Larger values use fewer parameters to encode the INU field. Note that a global intensity rescaling correction, without INU correction, can also be specified. For quantitative images, it may be better not to use any correction.

Modality Specify the modality of the scans in this channel. The main reason this is done is so that CT files can have a constant value of 1000 added to them to account for the way Hounsfield units are defined.

Labels? Specify whether or not there are any pre-defined label maps for (all) the subjects in the current population.

Has labels If subjects have corresponding label maps to guide the segmentation, these need to be specified along with a confusion matrix that relates values in the label maps to which tissue classes they correspond with.

Label maps Label maps are NIfTI images containing integer values, which must have the same dimensions and orientations as the scans of the corresponding subjects. Voxels of each value in the label map may be included in one or more tissue classes. For example, a label map showing the location of brain will include voxels that can be in grey or white matter classes. This information is specified in the confusion matrix.

Confusion matrix Specify rows of a confusion matrix, where each row corresponds to label values of 1, 2, ..., L + 1, etc in a label map.

L are the number of labels in the label map. The last row (L + 1) needs to specify what classes unlabeled voxels can take.

Row For this value in the label map, specify which tissue classes it can correspond to (including the K+1 implicit background class).

Intensity prior Intensity distributions of each tissue class are modelled by a Gaussian distribution. Prior knowledge about these distributions can make the model fitting more robust.

Definition Knowledge of Gaussian-Wishart priors for the intensity distributions of each cluster can help to inform the segmentation. When available, this information is specified in MATLAB prior*.m files. These files currently need to be hand-crafted. Unless you understand what you are doing, it is advised that you do not specify and intensity prior definition.

Optimise Specify whether the Gaussian-Wishart priors should be updated at each iteration. Enabling this can slow down convergence if there are small numbers of subjects. If only one subject is to be modelled (using a pre-computed template), then definitely turn off this option.

1.2 Merge tissues

This option is for merging template tissues together and extracting intensity priors for later use when running "Fit Multi-Brain model". This is typically used for re-ordering tissue classes or combining multiple classes (e.g. from air, which has a non-Gaussian intensity distribution that is often has multiple "tissues" fitted to it) into one. Generated templates also usually have a large field of view, so it is often desirable to trim them down so the field of view covers a smaller region of anatomy.

1.2.1 MB results file

Specify the results file obtained from running "Fit Multi-Brain model". This will be named mb_fit_*.mat and contain a link to where any resulting template may be found.

1.2.2 Indices

Specify indices. For example, if the original model had K=9 and you wish to combine the final three classes, then enter 1 2 3 4 5 6 7 7 7. Note that K refers to the total number of tissue maps – including the implicit background.

1.2.3 Bounding box

The bounding box (in voxels) of the merged template volume. Non-finite values indicate to use original template dimensions.

1.2.4 Output name

Specify a key string for inclusion within all the output file names.

1.2.5 Output directory

All output is written to the specified directory. If this is not specified, the current working directory will used by default.

1.3 Output

When "Fit Multi-Brain model' is run, the resulting model fit contains information that allows a lot of other derived images to be generated. For example, the results file encodes the INU fields, which allows images to be INU corrected. It contains information about intensity distributions, which (when combined with other extracted information) enables tissue segmentation to be achieved. And of course, the estimated deformations allow spatially normalised versions of these results to be generated.

The "Output" functionality allows a range of derived data to be generated from these parameter estimates. When it is executed, it generates the derived data from the images originally entered into "Fit Multi-Brain model". In order to do this, it needs to assume that the images, as well as the results files, have not been moved from their original locations. If they can not be fund, then "Output' will crash out in a not very elegant way.

1.3.1 MB results file

Specify the results file obtained from previously running "Fit Multi-Brain model".

1.3.2 Images

Specify whether versions of the original images, but with missing values filled in, should be written out.

1.3.3 INU corrected

Specify whether INU corrected versions of the original images (with missing values filled in) should be written out.

1.3.4 Warped images

Specify whether spatially normalised versions of the INU corrected images (missing values filled in) should be written out.

1.3.5 Warped INU corrected

1.3.6 INU

Specify whether the estimated INU fields should be written out.

1.3.7 Tissues

Specify the indices of any native-space tissue class images to be written.

1.3.8 Warped tissues

Specify the indices of any spatially normalised tissue class images to be written.

1.3.9 Warped mod. tissues

1.3.10 MRF Parameter

When tissue class images are written out, a few iterations of a simple Markov random field (MRF) cleanup procedure are run. This parameter controls the strength of the MRF. Setting the value to zero will disable the cleanup.

Bibliography

[1] J. Ashburner. A fast diffeomorphic image registration algorithm. *NeuroImage*, 38(1):95–113, 2007.

[2] J. Ashburner and K.J. Friston. Unified segmentation. *NeuroImage*, 26:839–851, 2005.

[3] J. Ashburner and K.J. Friston. Computing average shaped tissue probability templates. *NeuroImage*, 45(2):333–341, 2008.

[4] J. Ashburner and K.J. Friston. Diffeomorphic registration using geodesic shooting and Gauss–Newton optimisation. *NeuroImage*, 55(3):954–967, 2011.

[5] J. Ashburner and G. Ridgway. Symmetric diffeomorphic modeling of longitudinal structural mri. *Frontiers in Neuroscience*, 6(197), 02 2013.

[6] Y. Balbastre, M. Brudfors, K. Bronik, and J. Ashburner. Diffeomorphic brain shape modelling using gauss-newton optimisation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 862–870. Springer, 2018.

[7] Claudia Blaiotta, M.J. Cardoso, and J. Ashburner. Variational inference for medical image segmentation. *Computer Vision and Image Understanding*, 151:14–28, 2016.

[8] Claudia Blaiotta, P. Freund, J.M. Cardoso, and J. Ashburner. Generative diffeomorphic modelling of large mri data sets for probabilistic template construction. *NeuroImage*, 166:117–134, 2018.

[9] M. Brudfors, J. Ashburner, P. Nachev, and Y. Balbastre. Empirical bayesian mixture models for medical image translation. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 1–12. Springer, 2019.

[10] Mikael Brudfors, Yaël Balbastre, Guillaume Flandin, Parashkev Nachev, and John Ashburner. Flexible bayesian modelling for nonlinear image registration. *arXiv preprint arXiv:2006.02338*, 2020.