Diffusion Tools in the FMRIB Software Library (FSL 4.1)

Saad Jbabdi, Steve Smith, Jesper Andersson & Tim Behrens - and the rest of the FMRIB Analysis Group

FSL is a freely-available comprehensive library of analysis tools for FMRI, MRI and DTI brain imaging data. FSL includes a variety of tools for analysing diffusion MRI data, including pre-processing (registration, eddy-current correction), model fitting (diffusion tensor, crossing fibres, uncertainty modelling), connectivity modelling (probabilistic tractography) and voxelwise analysis of multi-subject data (TBSS). FSL runs on Apple and PCs (Linux), and is very easy to install. Most of the tools can be run both from the command line and as GUIs ("point-and-click" graphical user interfaces).

www.fmrib.ox.ac.uk/fsl

FDT: FMRIB Diffusion Toolbox

Eddy-current correction

Eddy currents in the gradient coils induce (approximate) stretches and shears in the diffusion weighted images. These distortions are different for different gradient directions. Eddy Current Correction corrects for these distortions, and for simple head motion, using affine registration to a reference volume.

DTIFIT: Diffusion tensor fitting

DTIFIT estimates a diffusion tensor at each voxel of the brain, and outputs a set of images including fractional anisotropy (FA), mode (MO), mean diffusivity (MD), tensor eigenvalues & vectors. The ouput of DTIFIT can be used directly within TBSS (below).

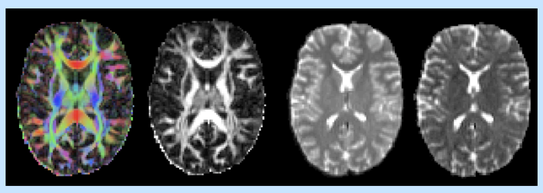

Various output of DTIFIT. From left to right: principal eigenvector encoded in RGB, FA, baseline non-diffusion-weighted signal, mean diffusivity.

BEDPOSTX: Crossing fibres & uncertainty

Bedpostx fits a multi-compartment model to the diffusion data. It uses a partial volume model where the diffusion MR signal is split into an infinitely anisotropic component for each fibre orientation, and a single isotropic component parameter. The Bayesian modelling, with MCMC inference, automatically estimates the number of fibres present in a voxel, as well as the uncertainties in partial volume fractions and tract directions.

Tractography using a single-fibre model

Tractography using a two-fibre model

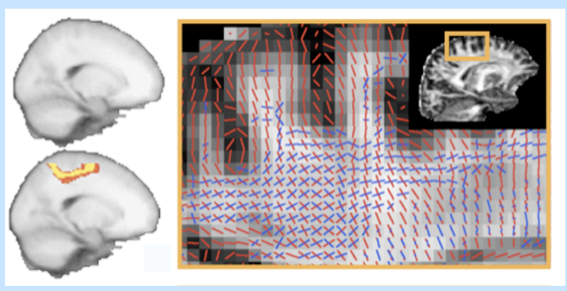

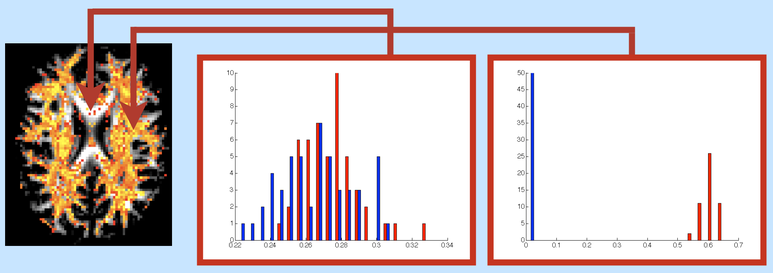

Bedpostx also allows automatic determination of the number of compartments (crossing fibres) in each voxel. Above: samples from the partial volume of the secondary fibre population (blue) are close to zero in the corpus callosum, but not in voxels of the the white matter that support a complex fibre architecture.

PROBTRACKX: Probabilistic tractography

Probabilistic tractography allows us to propagate the uncertainty of the local principal diffusion direction (PDD) into uncertainty on global connections between brain regions. Probtrackx's various options include:
- setting exclusion/inclusion/termination masks
- connectivity-based classification (e.g. thalamus segmentation below)
- surface-based tractography (based on FreeSurfer surfaces)

Figure shows the log probability of connection to the Splenium of the Corpus Callosum (seeded 5mm above the current slice). Probability distributions for the principle diffusion direction are calculated at each voxel, and these probabilities are integrated over all possible connecting paths.

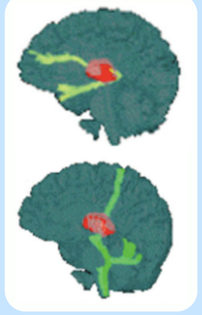

Seeding from different parts of the thalamus results in distinct probabilistic tracts en route to the cortex (left). Using this observation, we can segment the thalamus into voxels connecting to different parts of the cortex (a). This segmentation (c,d) resembles the map of thalamic nuclei proposed from tracer studies in monkeys (b).

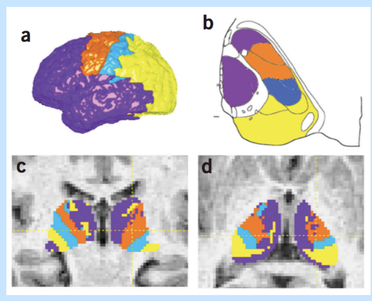

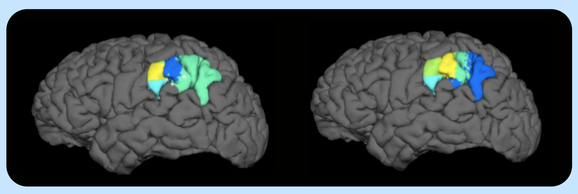

Seeding from the cortical surface (left), with knowledge of the orientation of this surface, leads to a more robust tractography-based segmentation than seeding from voxels (right).

TBSS: Tract-Based Spatial Statistics

Overview of TBSS

Many multi-subject voxelwise statistical analyses of white matter use fractional anisotropy (FA) images in order to localise brain changes related to development, plasticity or disease.

The strengths of VBM-style analyses are that they are fully automated, simple to apply, investigate the whole brain, and do not require pre-specifying regions/features of interest. Limitations include problems caused by alignment inaccuracies, and the lack of a principled way to choose smoothing extent. Tractography-based approaches have opposite advantages and disadvantages. They can overcome alignment problems by working in the space of individual subjects' tractography results, and for similar reasons do not necessarily require pre-smoothing. However, such approaches do not allow the whole brain to be investigated, and generally require user intervention in order to define the tracts.

In TBSS, we attempt to bring together the strengths of each approach. We aim to solve the alignment and smoothing issues, while being fully automated, investigating the "whole" brain - not requiring pre-specification of tracts of interest. This is achieved by estimating a group mean FA skeleton, which represents the centres of all fibre bundles that are generally common to the subjects involved in a study. Each subject's FA data is then projected onto the mean FA skeleton in such a way that each skeleton voxel takes the FA value from the local centre of the nearest relevant tract, thus hopefully resolving issues of alignment and correspondence.

TBSS in summary:

1. Identify a common registration target and align all subjects' FA images to this target using non-linear registration. At this stage, perfect alignment is not expected or required.
2. Create the mean of all aligned FA images and apply thinning (non-maximum-suppression perpendicular to the local tract structure), to create a skeletonised mean FA image.
3. Threshold this mean FA (typically between 0.2 and 0.3) to suppress areas of low mean FA and/or high inter-subject variability.
4. Project each subject's (aligned) FA image onto the skeleton, by filling the skeleton with FA values from the nearest relevant tract centre. This is achieved, for each skeleton voxel, by searching perpendicular to the local skeleton structure for the maximum value in the subject's FA image.
5. Carry out voxelwise statistics across subjects on the skeleton-space FA data

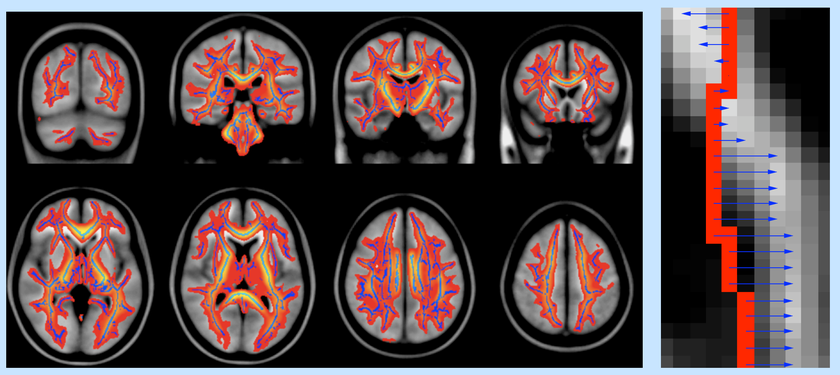

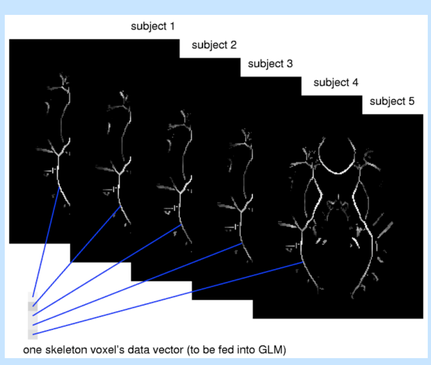

Example - Multiple Sclerosis

The figure shows the mean lesion probability distribution in blue. Red voxels show where FA correlates negatively across subjects with total lesion volume. There is strong negative correlation in left superior cingulum and many parts of the corpus callosum, including midline parts of the CC, well away from areas of lesion. This suggests that FA is reduced even in normal appearing white matter as disease progresses.

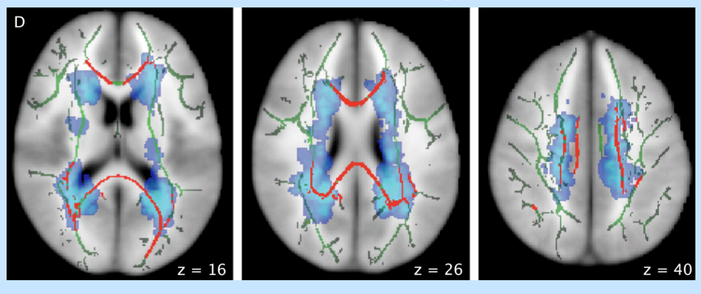

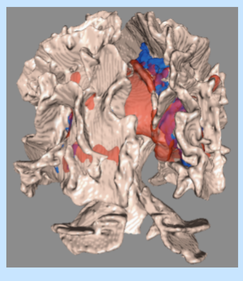

Example - Schizophrenia

TBSS found reduced FA in patients in right-superior, medial and anterior corpus callosum, superior and right-inferior fornix and in long assocation fibres near the junction of the right superior and inferior longitudinal fasciculi. In the majority of these areas the VBM-style analysis also found a difference, though with reduced localisation precision. However, in addition, several spurious results were generated by the VBM, for example, just below the ventricles. It is clear from inspecting the mean FA images for the controls and schizophrenics that while the corpus callosum is well aligned between the two groups, the lower edge of the ventricles is not, due to larger ventricles in the patient group.
This has caused a result that could easily be misinterpreted as a group difference in FA in the VBM-style analysis. TBSS did not show this spurious effect.

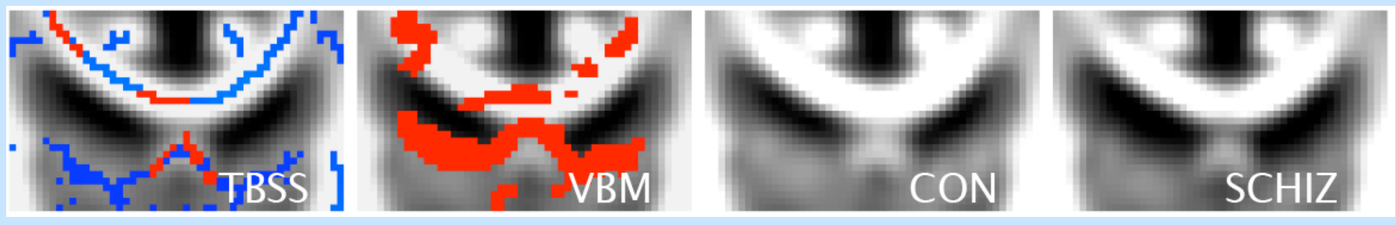

Current developments

Rician noise modelling

The parameters of the tensor model are often estimated from log-transformed data using a linear model. There are several disadvantages with that model: (i) It is based on the assumption of log-Normal distributed errors, (ii) it does not put any constraints on the parameters, allowing e.g. diffusion to become negative; (iii) for poor SNR data (e.g. with high b-values) the errors follow a Rician distribution, and by assuming Normal (or log-Normal) errors, one will underestimate diffusion and FA. We have implemented and tested a method for non-linear estimation of maximum a posteriori estimates of (positivity constrained) diffusion parameters using a Rician noise-model.

Image distortions

EPI images are plagued by distortions caused by sensitivity to (even small) inhomogeneities of the magnetic field. These inhomogeneities may be caused by the object itself (susceptibility-induced distortions), and will hence (as a first approximation) remain constant for that subject. Another source of disturbance of the field is lingering after-effects of the diffusion gradients (eddy current-induced distortions) which will cause unique distortions in each diffusion weighted image. We are working on a model that simultaneously estimates and corrects for all sources of distortions (and subject movements). It is based on manipulating the acquisitions so that a given inhomogeneity manifests itself differently in different images. One can then use a model to make predictions about what the images should look like given knowledge of the different fields, and by inverting this model we estimate the fields and corrected data.

Clustering

Connectivity-based clustering of brain regions is a powerful method for determining functionally distinct brain regions. We are extending the current parcellation method in order to address three outstanding issues: (i) choosing the number of clusters that are present in a region; (ii) combining data from different subjects; and (iii) determining the connections that drive a parcellation. Using infinite mixture modelling, we can address these three issues by using a so-called Dirichlet process prior on the mixture parameters of a Gaussian mixture model. A hierarchical extension of this model allows us to address issue (ii).