

# VBM Tutorial

John Ashburner

March 12, 2015

## 1 Getting Started

The data provided are a selection of T1-weighted scans from the freely available IXI dataset<sup>1</sup>. The overall plan will be to

- Start up SPM.
- Check that the images are in a suitable format (*Check Reg* and *Display* buttons).
- Segment the images, to identify grey and white matter (using the *Segment* button). Grey matter will eventually be warped to MNI space. This step also generates “imported” images, which will be used in the next step.
- Estimate the deformations that best align the images together by iteratively registering the imported images with their average (*SPM*→*Tools*→*Dartel Tools*→*Run Dartel (create Templates)*).
- Generate spatially normalised and smoothed Jacobian scaled grey matter images, using the deformations estimated in the previous step (*SPM*→*Tools*→*Dartel Tools*→*Normalise to MNI Space*).
- Do some statistics on the smoothed images (*Basic models*, *Estimate* and *Results* options).

The tutorial will use SPM12, which is available from <http://www.fil.ion.ucl.ac.uk/spm/>, and is the latest SPM software release<sup>2</sup>. SPM12 may already be installed on the system, so there is no need to download it now<sup>3</sup>.

SPM runs in the MATLAB package, which is worth learning how to program a little if you ever plan to do any imaging research. Start MATLAB, and type

---

<sup>1</sup>Now available via <http://www.brain-development.org/>.

<sup>2</sup>Older versions are SPM8, SPM5, SPM2, SPM99, SPM96 and SPM91 (SPMclassic). Anything before SPM5 is considered ancient, and any good manuscript reviewer will look down their noses at studies done using ancient software versions.

<sup>3</sup>Normally, the software would be downloaded and installed, and the patches/fixes then downloaded and unpacked so that they overwrite the buggy files in the original release.

```
editpath
```

This will give you a window that allows you to tell MATLAB where to find the SPM12 software. More advanced users could use the “path” command to do this.

You may then wish to change to the directory where the example data are stored.

SPM is started by typing

```
spm
```

or (eg)

```
spm pet
```

This will pop open a bunch of windows. The top-left one is the one with the buttons.

A manual is available in “man\manual.pdf”. Earlier chapters tell you what the various options within mean, but there are a few chapters describing example analyses later on. The developers of SPM always appreciate bug reports. There are about half a million lines of code in SPM12, so it’s bound to contain a few mistakes.

## 1.1 Check Reg button

SPM requires the image data in a suitable format. Most scanners produce images in DICOM format. The DICOM button of SPM can be used to convert many versions of DICOM into NIfTI format, which SPM (and a number of other packages, eg FSL, MRICron etc) can use. There are two main forms of NIfTI:

- “.hdr” and “.img” files, which appear superficially similar to the older ANALYZE format. The .img file contains the actual image data, and the .hdr contains the information necessary for programs to understand it.
- “.nii” files, which combines all the information into one file.

NIfTI files have some orientation information in them, but sometimes this needs “improving”. Ensure that the images are approximately aligned with MNI space. This makes it easier to align the data with MNI space, using local optimisation procedures such as those in SPM. NIfTI images can be single 3D volumes, but 4D or even 5D images are possible. For SPM5 onwards, the 4D images may have an additional “.mat” file.

Click the button and select one (or more) of the original IXI scans from the “examples” folder, as well as the “canonical\avg152T1.nii” file in the SPM12 release. This shows the relative positions of the images, as understood by SPM (see Fig 1). Note that if you click the right mouse button on one of the images, you will be shown a menu of options, which you may wish to try. If images need to be rotated or translated (shifted), then this can be done using the *Display* button.



Figure 1: *Check Reg and Display.*

Before you begin, each image should be approximately aligned within about 5cm and about 20 degrees of the template data released with SPM.

## 1.2 *Display* button

Click the *Display* button, and select an image. This should show you three slices through the image (see Fig 1). If you have the images in the correct format, then:

- The top-left image is coronal with the top (superior) of the head displayed at the top and the left shown on the left. This is as if the subject is viewed from behind.
- The bottom-left image is axial with the front (anterior) of the head at the top and the left shown on the left. This is as if the subject is viewed from above.
- The top-right image is sagittal with the front (anterior) of the head at the left and the top of the head shown at the top. This is as if the subject is viewed from the left.

You can click on the sections to move the cross-hairs and see a different view of the 3D image. Below the three views are a couple of panels. The one on the left gives you information about the position of the cross-hair and the intensity of the image below it. The positions are reported in both mm and voxels. The

coordinates in mm represent the orientation within some Cartesian coordinate system (eg scanner coordinate system, or MNI space). The vx coordinates indicate which voxel<sup>4</sup> you are at (where the first voxel in the file has coordinate 1,1,1). Clicking on the horizontal bar, above, will move the cursor to the origin (mm coordinate 0,0,0), which should be close(ish) to the anterior commissure (AC). The intensity at the point in the image below the cross-hairs is obtained by interpolation of the neighbouring voxels. This means that for a binary image, the intensity may not be shown as exactly 0 or 1.

Below this are a number of boxes that allow the image to be re-oriented and shifted. If you click on the AC, you can see its mm coordinates. The objective is to move this point to a position close to 0,0,0. To do this, you would translate the images by minus the mm coordinates shown for the AC. These values would be entered into the *right {mm}*, *forward {mm}* and *up {mm}* boxes.

If rotations are needed, they can be entered into rotation boxes. Notice that the angles are specified in radians, so a rotation of 90 degrees would be entered as “1.5708” or as “pi/2”. When entering rotations around different axes, the result can be a bit confusing, so some trial and error may be needed. Note that “pitch” refers to rotations about the x axis, “roll” is for rotations about the y axis and “yaw” is for rotations about the z axis<sup>5</sup>.

To actually change the header of the image, you can use the *Reorient images...* button. This allows you to select multiple images, which would all be rotated and translated in the same way.

The panel on the right shows various pieces of information about the image being displayed. Note that images can be represented by various datatypes. Integer datatypes can only encode a limited range of values, and are usually used to save some disk space. For example, the “uint8” datatype can only encode integer values from 0 to 255, and the “int16” datatype can only encode values from -32768 to 32767. For this reason, there is also a scale-factor (and sometimes a constant term) saved in the headers in order to make the values more quantitative. For example, a uint8 datatype may be used to save an image of probabilities, which can take values from 0.0 to 1.0. This would use a scale-factor of 1/255.

Below this lot is some positional and voxel size information. Note that the first element of the voxel sizes may be negative. This indicates whether there is a reflection between the coordinate system of the voxels and the mm coordinate system. One of them is right-handed, whereas the other is left-handed (see Fig 2).

Another option allows the image to be zoomed to a region around the cross-hair. The effects of changing the interpolation are most apparent when the image is zoomed up. Also, if the original image was collected in a different orientation,

---

<sup>4</sup>A voxel is a like a pixel, but in 3D.

<sup>5</sup>In aeronautics, pitch, roll and yaw would be defined differently, but that is because they use a different labelling of their axes.

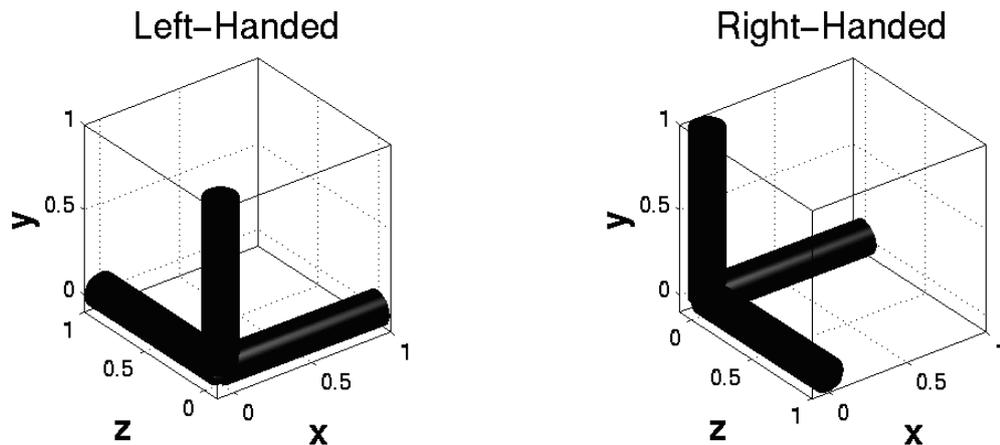


Figure 2: Left- and right-handed coordinate systems.

then the image can be viewed in *Native Space* (as opposed to *World Space*), to show it this way.

## 2 Image processing for VBM

At this point, the images should all be in a suitable format for SPM to work with them. The following procedures will be used. For the tutorial, you should specify them one at a time. In practice though, it is much easier to use the batching system, by clicking the *Batch* button of the top-left figure. The sequence of jobs in the batching system would be:

- **Module List**

- **SPM→Spatial→Segment**: To generate the roughly (via a rigid-body) aligned grey and white matter images of the subjects.
- **SPM→Tools→Dartel Tools→Run Dartel (create Template)**: Determine the nonlinear deformations for warping all the grey and white matter images so that they match each other.
- **SPM→Tools→Dartel Tools→Normalise to MNI Space**: Actually generate the smoothed “modulated” warped grey and white matter images.

### 2.1 Using Spatial→Segment

The objective now is to automatically identify different tissue types within the images, using *Segment* [Ashburner and Friston, 2005]. The output of the segmentation will be used for achieving more accurate inter-subject alignment using Dartel (to be explained later). Segmentation can be found within *SPM→Spatial→Segment*

in the batching system. It is suggested that *Native Space* versions of the tissues in which you are interested are generated, along with “Dartel imported” versions of grey and white matter. It is useful to also have either a native or “imported” version of the CSF, so that intra-cranial volumes may be computed (see later).

For VBM, the native space images are usually the `c1*.nii` files, as it is these images that will eventually be warped to MNI space. Both the imported and native tissue class image sets can be specified via the Native Space options of the user interface.

Segmentation in SPM can work with images collected using a variety of sequences, but the accuracy of the resulting segmentation will depend on the particular properties of the images. Although multiple scans of each subject were available, the dataset to be used only includes the T1-weighted scans. There won't be time to segment all scans, so the plan is to demonstrate how one or two scans would be segmented, and then continue with data that was segmented previously. If you know how to segment one image with SPM, then doing lots of them is pretty trivial. The *Segment* module should be set up as follows:

- *Data*: Clicking here allows more channels of images to be defined. This is useful for multi-spectral segmentation (eg if there are T2-weighted and PD-weighted images of the same subjects), but as we will just be working with a single image per subject, we just need one channel.
  - *Channel*
    - \* *Volumes*: Here you specify all the IXI scans to be segmented.
    - \* *Bias regularisation*: Leave this as it is. It works reasonably well for most images.
    - \* *Bias FWHM*: Again, leave this as it is.
    - \* *Save Bias Corrected*: This gives the option to save intensity inhomogeneity corrected version of the images, or a field that encodes the inhomogeneity. Leave this at *Save nothing* because we don't have a use for them here.
- *Tissues*: This is a list of the tissues to identify.
  - *Tissue*: The first tissue usually corresponds to grey matter.
    - \* *Tissue probability map*: Leave this at the default setting, which points to a volume of grey matter tissue probability in one of the images released with SPM12.
    - \* *Num. Gaussians*: This can usually be left as it is.
    - \* *Native Tissue*: We want to save *Native + Dartel imported*. This gives images of grey matter at the resolution of the original scans, along with some lower resolution “imported” versions that can be used for the Dartel registration.

- \* *Warped Tissue*: Leave this at *None*, as grey matter images will be aligned together with Dartel to give closer alignment.
- *Tissue*: The second tissue is usually white matter.
  - \* *Tissue probability map*: Leave alone, so it points to a white matter tissue probability map.
  - \* *Num. Gaussians*: Leave alone.
  - \* *Native Tissue*: We want *Native + Dartel imported*.
  - \* *Warped Tissue*: Leave at *None*.
- *Tissue*: The third tissue is usually CSF.
  - \* *Tissue probability map*
  - \* *Num. Gaussians*
  - \* *Native Tissue*: Just chose *Native Space*. This will give a map of CSF, which can be useful for computing total intra-cranial volume.
  - \* *Warped Tissue*: Leave at *None*.
- *Tissue*: Usually skull.
  - \* *Tissue probability map*
  - \* *Num. Gaussians*
  - \* *Native Tissue*: Leave at *None*.
  - \* *Warped Tissue*: Leave at *None*.
- *Tissue*: Usually soft tissue outside the brain.
  - \* *Tissue probability map*
  - \* *Num. Gaussians*
  - \* *Native Tissue*: Leave at *None*.
  - \* *Warped Tissue*: Leave at *None*.
- *Tissue*: Usually air and other stuff outside the head.
  - \* *Tissue probability map*
  - \* *Num. Gaussians*
  - \* *Native Tissue*: Leave at *None*.
  - \* *Warped Tissue*: Leave at *None*.
- *Warping & MRF*
  - *MRF Parameter*: This tries to remove isolated mis-classified voxels, and generally tidy up the tissue classes. It's probably best to leave this at the default setting of 1.

- *Clean Up*: This is a bit of an ad hoc procedure that tries even harder to eliminate mis-classified tissues outside the brain<sup>6</sup>. Again the default settings should work reasonably well.
- *Warping regularisation*: This is a penalty term to keep deformations smooth. Leave alone.
- *Affine regularisation*: Another penalty term. Leave alone.
- *Sampling distance*: A speed/accuracy balance. Sampling every few voxels will speed up the segmentation, but may reduce the accuracy. Leave alone.
- *Deformation fields*: Not needed here, so leave at *None*.

Once everything is set up (and there are no “<-” symbols, which indicate that more information is needed), then you could click the *Run* button (the green triangle) - and wait for a while as it runs. This is a good time for questions. If there are hundreds of images, then it is chance to spend a couple of days away from the computer.

After the segmentation is complete, there should be a bunch of new image files generated. Files containing “c1” in their name are what the algorithm identifies as grey matter. If they have a “c2” then they are supposed to be white matter. The “c3” images, are CSF. The file names beginning with “r” (as in “rc1”) are the Dartel imported versions of the tissue class images, which will be aligned together next.

I suggest that you click the *Check Reg* button, and take a look at some of the resulting images (see Fig 4). For one of the subjects, select the original, the c1, c2 and c3. This should give an idea about which voxels the algorithm identifies as the different tissue types. Also try this for some of the other subjects.

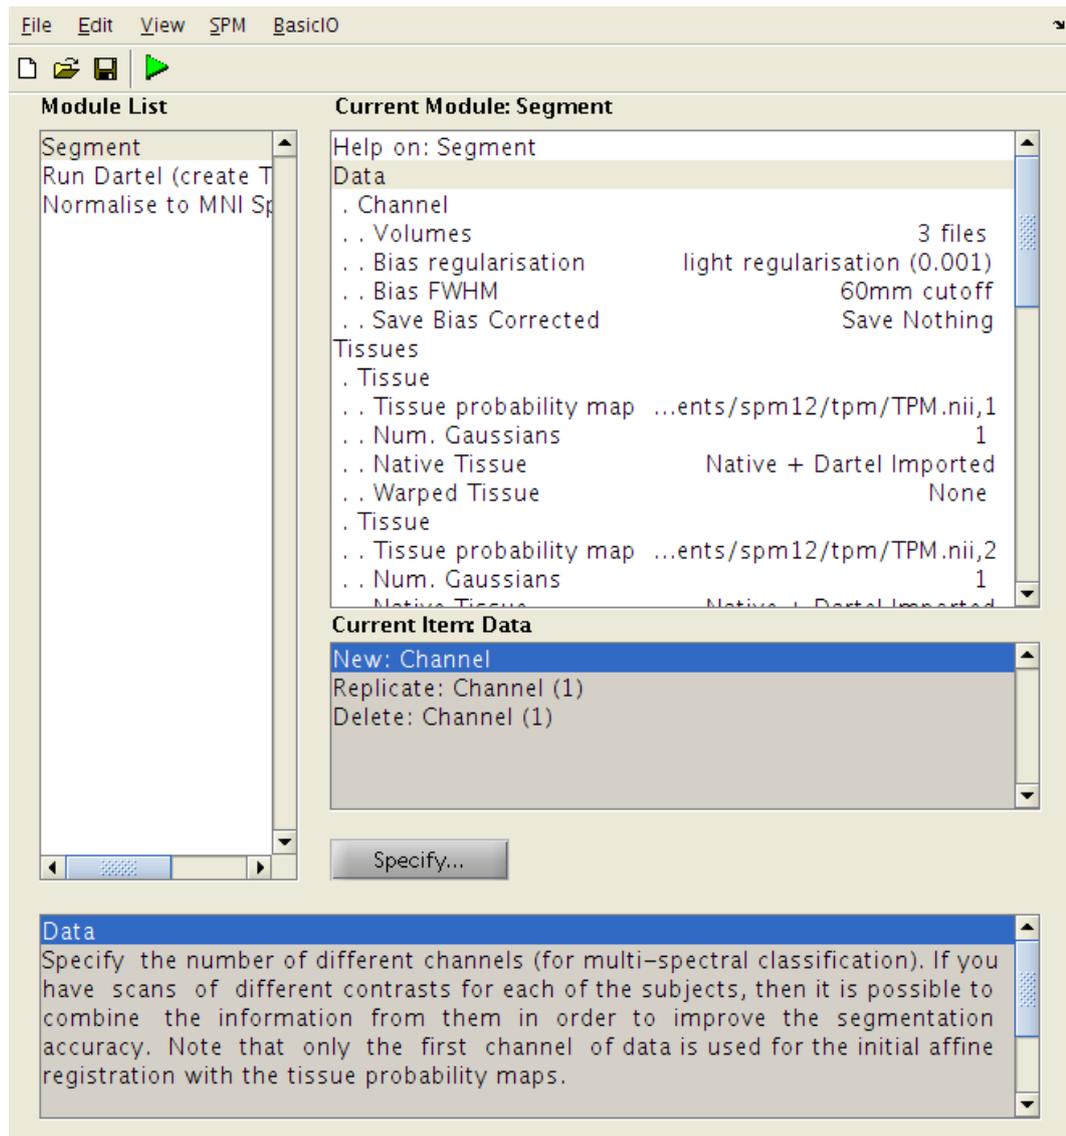
## 2.2 Run Dartel (create Templates)

The idea behind Dartel [Ashburner, 2007] is to increase the accuracy of inter-subject alignment by modelling the shape of each brain using millions of parameters (three parameters for each voxel). Dartel works by aligning grey matter among the images, while simultaneously aligning white matter. This is achieved by generating its own increasingly crisp average template data (see Fig 5), to which the data are iteratively aligned [Ashburner and Friston, 2008]. This uses the imported “rc1” and “rc2” images, and generates “u\_rc1” files, as well as a series of template images. The Dartel module would be set up as follows:

- *Images*: Two channels of images need to be created.

---

<sup>6</sup>Classification of fluid in eyeballs is very different with and without the cleanup. Without it, eyeballs are in the same class as CSF, because they are both a sort of fluid. With the cleanup, eyeballs fall into the non-brain soft tissue class (making it easier to compute intra-cranial volumes).

Figure 3: The form of a *Segment* job.

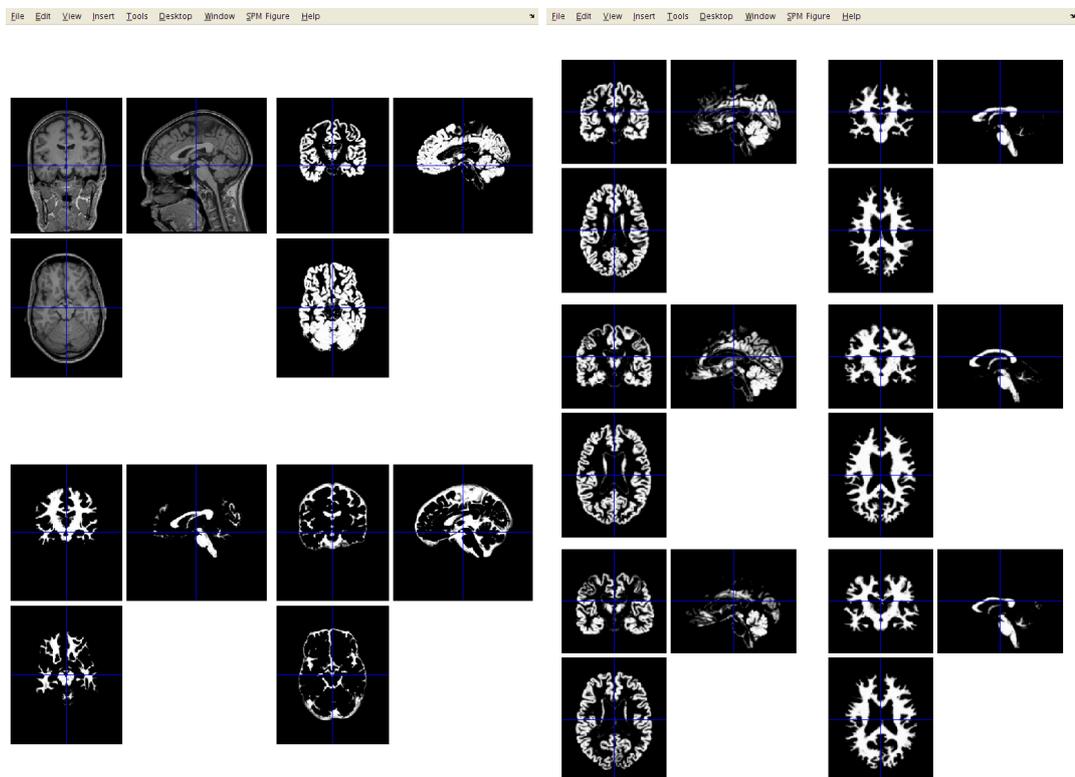


Figure 4: Left: An image, along with grey matter (c1), white matter (c2) and CSF (c3) identified by *Segment*. Right: Imported grey (rc1) and white matter (rc2) for three subjects.

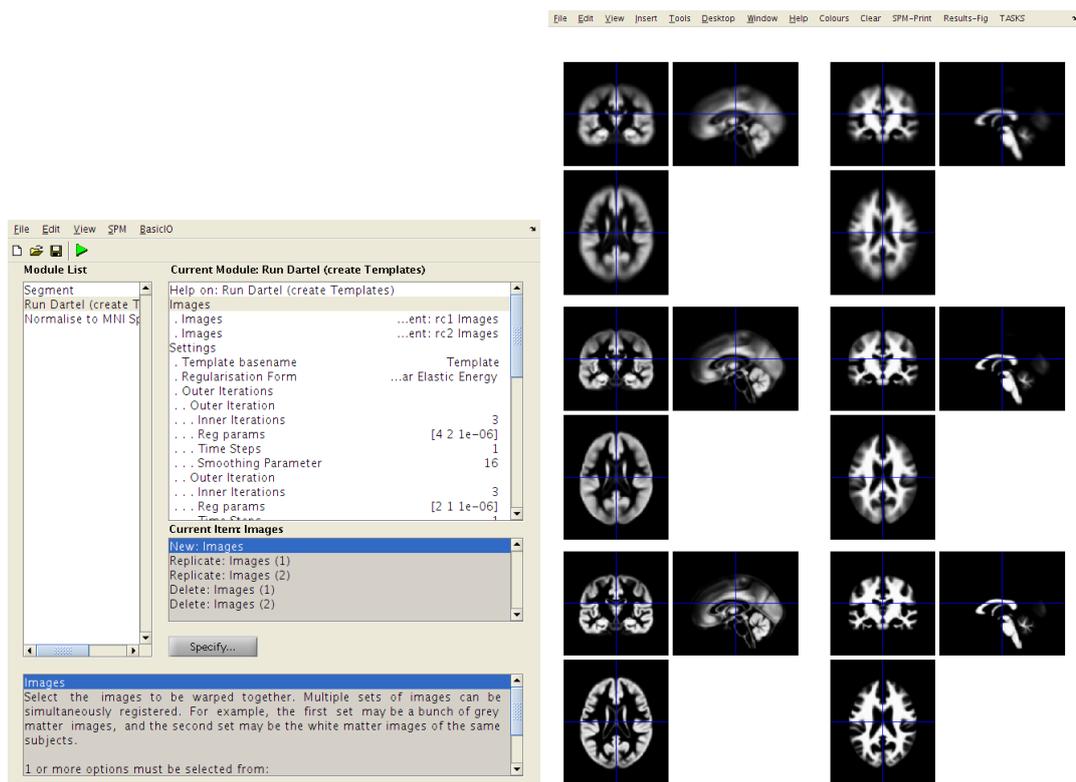


Figure 5: The form of a *Dartel* job (left) and the template data after 0, 3 and 18 iterations (right).

- *Images*: Select the imported grey matter images (rc1\*.nii).
- *Images*: Select the imported white matter images (rc2\*.nii). These should be specified in the same order as the grey matter, so that the grey matter image for any subject corresponds with the appropriate white matter image.
- *Settings*: There are lots of options here, but they are set at reasonable default values. Best to just leave them as they are.

Dartel takes a long time to run. If you were to hit the *Run* button, then the job would be executed. This would take a long time to finish, so I suggest you don't do it now. If you have actually just clicked the *Run* button, then find the main MATLAB window and type Ctrl-C to stop the job. This will bring up a long error message<sup>7</sup>, and there may be some partially generated files to remove.

<sup>7</sup>If you become a regular SPM user, and it crashes for some reason, then you may want to ask about why it crashed. Such error messages (as well as the MATLAB and SPM version you use, and something about the computer platform) are helpful for diagnosing the cause of the problem. You are unlikely to receive much help if you just ask why it crashed, without providing useful information.

## 2.3 Normalise to MNI Space

This step uses the resulting ‘u\_rc1” files (which encode the shapes), to generate smoothed, spatially normalised and Jacobian scaled grey matter images in MNI space [Mechelli et al., 2005, Ashburner, 2009].

- *Dartel Template*: Select the final template image created in the previous step. This is usually called Template\_6.nii. This template is registered to MNI space (affine transform), allowing the transformations to be combined so that all the individual spatially normalised scans can also be brought into MNI space.
- *Select according to*: Choose *Many Subjects*, as this allows all flow fields to be selected at once, and then all grey matter images to be selected at once.
  - *Many Subjects*
    - \* *Flow fields*: Select all the flow fields created by the previous step (u\_\*.nii).
    - \* *Images*: Need one channel of images if only analysing grey matter.
      - *Images*: Select all the grey matter images (c1\*.nii), in the same order as the flow fields.
- *Voxel Sizes*: Specify voxel sizes for spatially normalised images. Leave as is (NaN NaN NaN), to have 1.5mm voxels.
- *Bounding box*: The field of view to be included in the spatially normalised images can be specified here. For now though, just leave at the default settings.
- *Preserve*: For VBM, this should be set to *Preserve Amount (“modulation”)*, so that tissue volumes are compared.
- *Gaussian FWHM*: Specify the size of the Gaussian (in mm) for smoothing the processed data by. This is typically between about 4mm and 12mm. Use 10mm for now.

The optimal value for FWHM depends on many things, but one of the main ones is the accuracy with which the images can be registered. If spatial normalisation (inter-subject alignment to some reference space) is done using a simple model with fewer than a few thousand parameters, then more smoothing (eg about 12mm FWHM) would be suggested. For more accurately aligned images, the amount of smoothing can be decreased. About 8mm may be suitable, but I don’t much have empirical evidence to suggest appropriate values. The optimal value is likely to vary from region to region: lower for subcortical regions with less variability, and more for the highly variable cortex.

The smoothed images represent the regional volume of tissue (see Fig 6). Statistical analysis is done using these data, so one would hope that differences

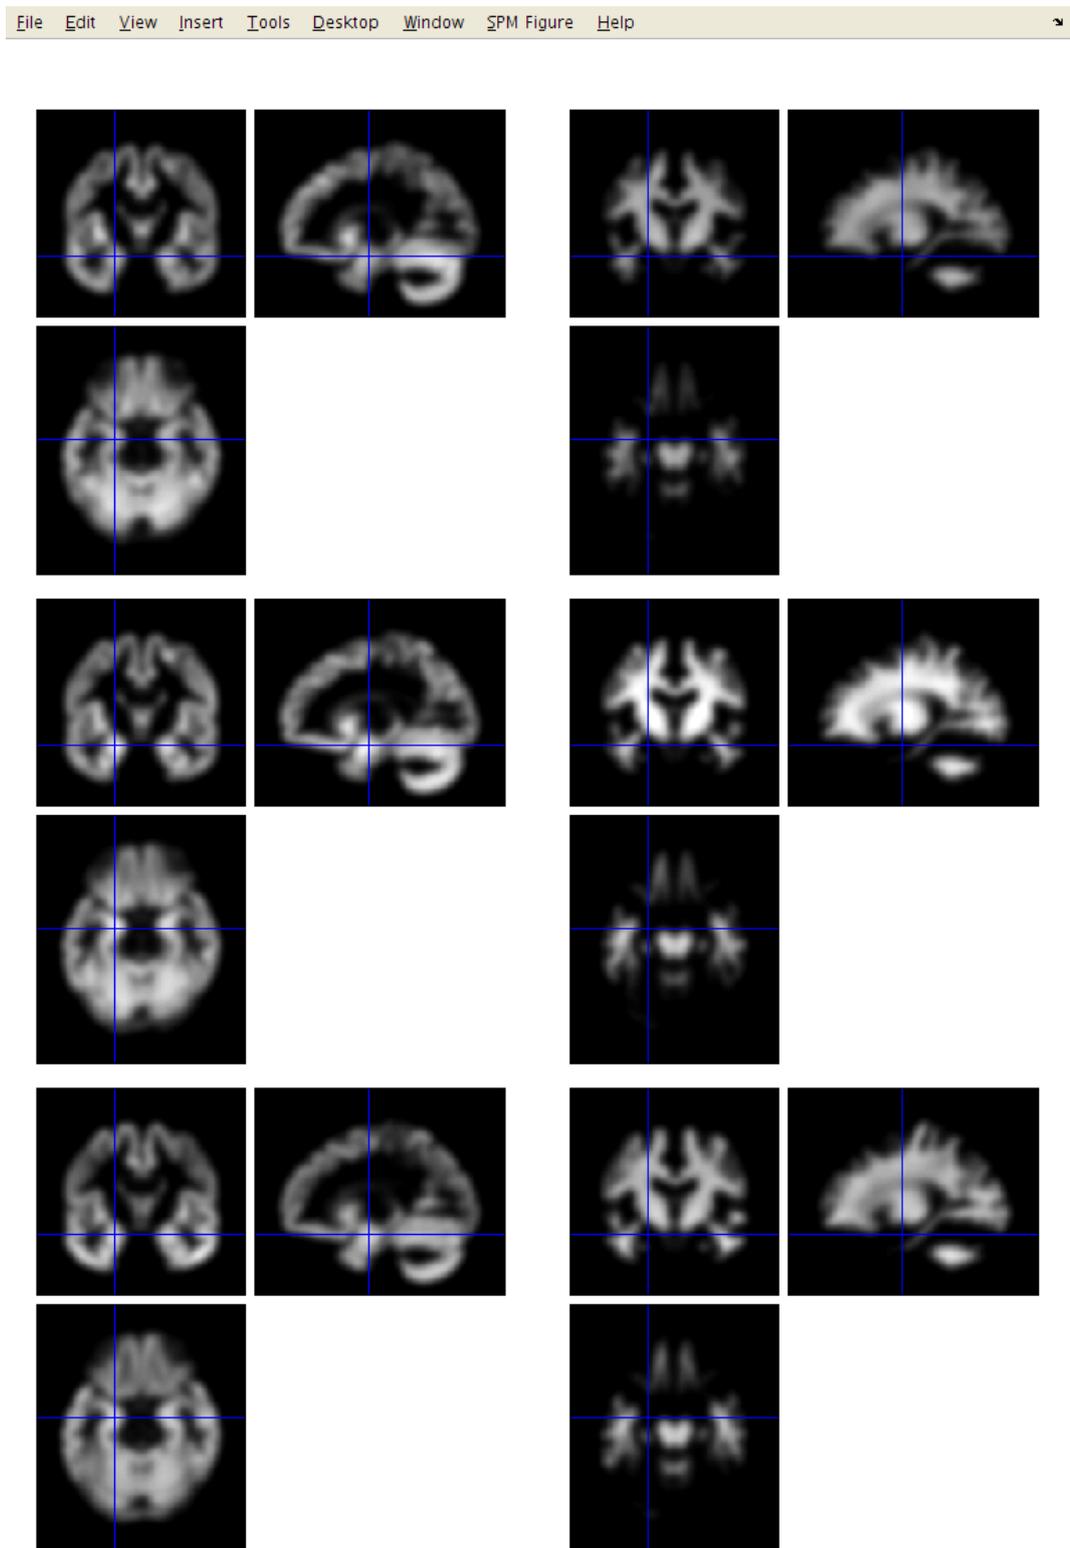


Figure 6: Processed grey and white matter data for three subjects, which are smoothed by 8 mm FWHM. Note that all data are shown scaled the same, which illustrates the effect of different global brain sizes (darker for smaller brains because structures are smaller).

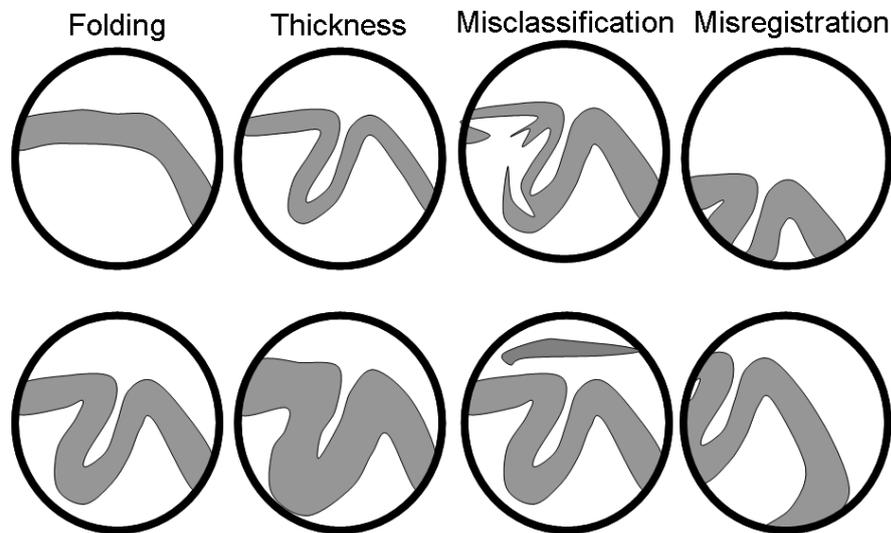


Figure 7: We generally hope that the results of VBM analyses can be interpreted as systematic volumetric differences (such as folding or thickness), rather than artifacts (such as misclassification or misregistration). Because of this, it is essential that the processing be as accurate as possible. Note that exactly the same argument can be made of the results of manual volumetry, which depend on the accuracy with which the regions are defined, and on whether there is any systematic error made.

among the processed data actually reflect differences among the regional volumes of grey matter (see Fig 7).

After specifying all those files, you may wish to keep a record of what you do. This can be achieved by clicking the *Save* button (with a floppy disk<sup>8</sup> icon) and specifying a filename. Filenames can either end with a “.mat” suffix or a “.m”. Such jobs can be re-loaded at a later time (hint: *Load* button).

### 3 Statistical Analysis

Now we are ready to fit a general linear model through the processed data [Friston et al., 1995], and make inferences [Worsley et al., 1996] about where there are systematic differences among the processed data.

The idea now is to set up a linear model in which to interpret any differences among the processed data. A different model will give different results, so in

<sup>8</sup>Floppy disks are from a time long ago when people felt that “640K ought to be enough for anybody”.

publications it is important to say exactly what was included<sup>9</sup>. Similarly, the models used for processing also define how findings are interpreted. Grey matter is defined as what the segmentation model says it as. Homologous structures are defined by what the registration model considers homologous. The amount of smoothing also defines something about what the scale at which the data is examined. Changes to any of these components will generate slightly different processed data, and therefore a different pattern of results.

Returning to the analysis at hand. The strategy is:

- Use *Basic Models* to specify a design matrix that explains how the processed data are caused.
- Use *Review* models, to see what is in the design matrix, and generally poke about.
- Use *Estimate* to actually fit the GLM to the data, and estimate the smoothness of the residuals for the subsequent random field theory corrections.
- Use *Results* to characterise differences.

### 3.1 Basic Models

Click the *Basic Models* button on the top-left figure with the buttons on it (or select *Stats* and *Factorial design specification* via the batching system). This will give a whole list of options about *Directory*, *Design*, *Covariates*, *Multiple covariates*, *Masking*, *Global calculation* and *Global normalisation*.

#### 3.1.1 Directory

The statistics part of SPM generally involves producing a bunch of results that are all saved in some particular directory. If you run more than one analysis in the same directory, then the results of the later analyses will overwrite the earlier results (although you do get a pop-up that asks whether or not you wish to continue). The *Directory* option is where you specify a directory in which to run the analysis (so create a suitable one first).

#### 3.1.2 Global Calculation

For VBM, the “global normalisation” is about dealing with brains of different sizes. Because VBM is a voxel-wise assessment of volumetric differences, then it is useful to consider how the regional volumes are likely to vary as a function of whole brain volume. Bigger brains are likely to have larger brain structures. In a comparison of two populations of subjects, if one population has larger

---

<sup>9</sup>Unless you intend to submit to one of those pop science journals that does not allow a proper Methods section.

brains than the other, then should we be surprised if we find lots of volumetric differences? By including some form of correction for the “global” brain volume, then we may obtain more focused results. So what measure should be used to do the correction? Favourite measures are total grey matter volume, whole brain volume (with or without including the cerebellum) and total intra-cranial volume. Within the aging and dementia research fields, total intra-cranial volume is the most accepted measure, but other researchers may favour different ones.

“Globals” can be computed in a number of ways. For example, obtaining the total volume of grey matter could be achieved by adding up the values in a “c1” or “mwc1” image, and multiplying by the volume of each voxel. Similarly, whole brain volume can be obtained by adding together the volume of grey matter and white matter. Total intra-cranial volume can be obtained by adding up grey matter, white matter and CSF volumes. The *SPM→Utils→Tissue Volumes* option from the batching system can be used to compute these.

A set of values that can be used as “globals” have been provided, which were obtained by summing up the volumes of grey matter, white matter and CSF. These values can be loaded into MATLAB by loading the “info.txt” file, which is a simple text file. If it is in your current directory, you can do this by typing the following into MATLAB:

```
load info.txt
```

This will load the contents of the file into a variable called “info” (same variable name as file name). In MATLAB, you can see the values of this variable by typing:

```
info
```

Among other things, this variable encodes the values that we’ll use as globals. Note that it contains 55 rows, which is the same as the number of scans to be entered into the statistical analysis. Click on *Global calculation*, and choose *User*, which indicates user-specified values. A *User* option should now have appeared, which has a field that says *Global values*. Select this, so a box should appear, in which you can enter the values of the globals. To do this, enter

```
info(:,3)+info(:,4)+info(:,5)
```

This indicates that the sum of columns 3, 4 and 5 of the “info” variable should be used.

### 3.1.3 Global Normalisation

This part indicates how the “globals” should be used. *Overall grand mean scaling* should be set to *No*. Then click *Normalisation*, to reveal options about how our “globals” are to be used.

- *None* indicates that no “global” correction should be done.

- *Proportional* scaling indicates that the data should be divided by the “globals”. If the “globals” indicate whole brain volume, then this scales the values in the processed data so that they are proportional to the fraction of brain volume accounted for by that piece of grey matter. Similar principles apply when other values are used for “globals”.
- *ANCOVA* correction simply treats the “globals” as covariates in the GLM. The results will show differences that can not be explained by the globals.

For this analysis, I suggest using proportional scaling.

### 3.1.4 Masking

*Masking* includes a number of options for determining which voxels are to be included in the analysis. The corrections for multiple comparisons mean that analyses are more sensitive if fewer voxels are included. For VBM studies, there are also instabilities that occur if the background is included in the analysis. In the background, there is little or no signal, so the estimated variance is close to zero. The t statistics are proportional to the magnitudes of the differences, divided by the square root of the residual variance. Division by zero, or numbers close to zero is problematic. Also, it is possible for blobs to extend out a long way into the background (see [Ridgway et al. \[2012\]](#)), but this is prevented by the masking.

The parametric statistical tests within SPM assume that distributions are Gaussian. Another reason for masking is that values close to zero can not have Gaussian distributions, if they are constrained to be positive. If they were truly Gaussian, then at least some negative values should be possible.

It is possible to supply an image of ones and zeros, which acts as an *Explicit Mask*. The other masking strategies involve some form of intensity thresholding of the data. *Implicit Masking* involves assuming that a value of zero at some voxel, in any of the images, indicates an unknown value. Such voxels are excluded from the analysis.

The main way of masking VBM data is to use *Threshold masking*. At each voxel, if a value in any of the images falls below the threshold, then that voxel is excluded from the analysis. Note that with *Proportional* scaling correction, the values are divided by the globals prior to generating the mask.

For these data, I suggest using *Absolute* masking, with a threshold of 0.2.

### 3.1.5 Covariates

It is possible to include additional covariates in the model, but these will be specified elsewhere. I therefore suggest not adding any here.

### 3.1.6 Design

This is the part where the actual design matrix and data are specified. For the current data, the analysis can be specified as a *Multiple Regression*, so select this *Design* option. There are then three fields to specify: *Scans*, *Covariates* and *Intercept*.

Click *Scans* <-, and select the 55 previously computed “smwc1” files (from the “processed” folder).

For *Intercept*, *Specify Menu Item*, *Include Intercept*. This will model a constant term in the design matrix.

Click *Covariates*, and then *New “Covariate”* two times. The plan will be to model sex<sup>10</sup> and age as covariates. Enter the covariates as follows.

1. Enter “info(:,1)” as the *Vector*, “Age” as the *Name* and *No centering for Centering*.
2. Enter “info(:,2)” as the *Vector*, “Sex” as the *Name* and *No centering for Centering*. Values of 0 indicate a female subject, and 1 indicates male.

All information should now be entered. You can save the job if you like (*Save* button). Hit the *Run* button, and SPM will set up the necessary information to actually perform the statistical analysis. This information is saved in a “SPM.mat” file.

## 3.2 Review

The *Review* button gives the chance to see what has been specified. If you notice any mistakes, then you can set up the analysis again. This is why it could be useful to save job files as you go.

## 3.3 Estimate

The *Estimate* button actually runs the analysis. Click *Estimate*, select the “SPM.mat” file that has just been created, and then wait a while.

The first step is to fit the data at each voxel to some linear combination of the columns of the design matrix. This generates a series of “beta” images, where the first indicates the contribution of the first column, the second is the contribution of the second column etc. There is also a “ResMS” image generated, which provides the necessary standard deviations for computing the t statistics.

---

<sup>10</sup>Sex is usually treated as a categorical variable, but can be modelled as a covariate of ones and zeros.

The second step is to compute residual images, from which the smoothness of the data is computed. By default, these images are temporary, and are automatically deleted as soon as they are no longer needed.

That's it. Pretty simple really.

### 3.4 Results

Once the GLM has been fitted, it is possible to use the information from the resulting “beta” and “ResMS” images (as well as the “SPM.mat” file) to generate images of t statistics. Contrast vectors are specified, which indicate the linear combination of “beta” images to test. These linear combinations are saved as “con” images, and the idea is to identify any regions that differ from zero in some statistical sense. The t statistics themselves are saved in “spmT” images. Because lots of voxels are examined, some correction for the number of tests is needed. The processed data are spatially smooth, and voxels are correlated with their neighbours, so Gaussian random field theory is used to correct for multiple dependent comparisons [Worsley et al., 1996].

Click the *Results* button, and select the “SPM.mat” file. This brings up the *SPM contrast manager*. Click the *t-contrasts* button, and then *Define new contrast*. Enter “-Age” (or some other name of your choice) where the *SPM contrast manager* says *name*. The idea is to test for regions with proportionally less grey matter in older subjects (ie more grey matter in younger subjects). The t contrast for this is 0, -1, 0, 0. This would be entered into the field that says *contrast*. Note that the trailing zero(s) can be omitted, as SPM will automatically fill these in. Click *OK*, then click *Done*.

The lower-left SPM figure is now the one to focus on. This will prompt you for answers to a series of questions, so here goes.

- *mask with other contrast(s)*: click *no*.
- *title for comparison*: leave this as it is. You can hit return on the keyboard, or click the edge of the text box with your mouse to do this.
- *p value adjustment to control*: Click *FWE* (family-wise error correction).
- *threshold {T or p value}*: Enter 0.05 for regions that are “statistically significant” after corrections for multiple comparisons ( $p < 0.05$ )<sup>11</sup>.
- *& extent threshold {voxels}*: lets see if there are any blobs at all, so go for 0.

---

<sup>11</sup>Note that 0.05 is an arbitrary threshold to reach, with no mathematical or empirical justification, but it does seem to be one of those things that few seem to question (*Nullius in verba* as they say in the Royal Society). Some consider the 0.05 threshold to be much too lenient, and prefer 0.001. Some entire fields use thresholds of less than  $10^{-6}$  (the  $5\sigma$  criterion used by physicists – see <http://www.dcs.science.net/?p=6518>). Laplace (who developed the Bayesian interpretation of probability) said “The weight of the evidence should be proportioned to the strangeness of the facts”, and some findings are certainly more surprising (“high-impact”) than others.

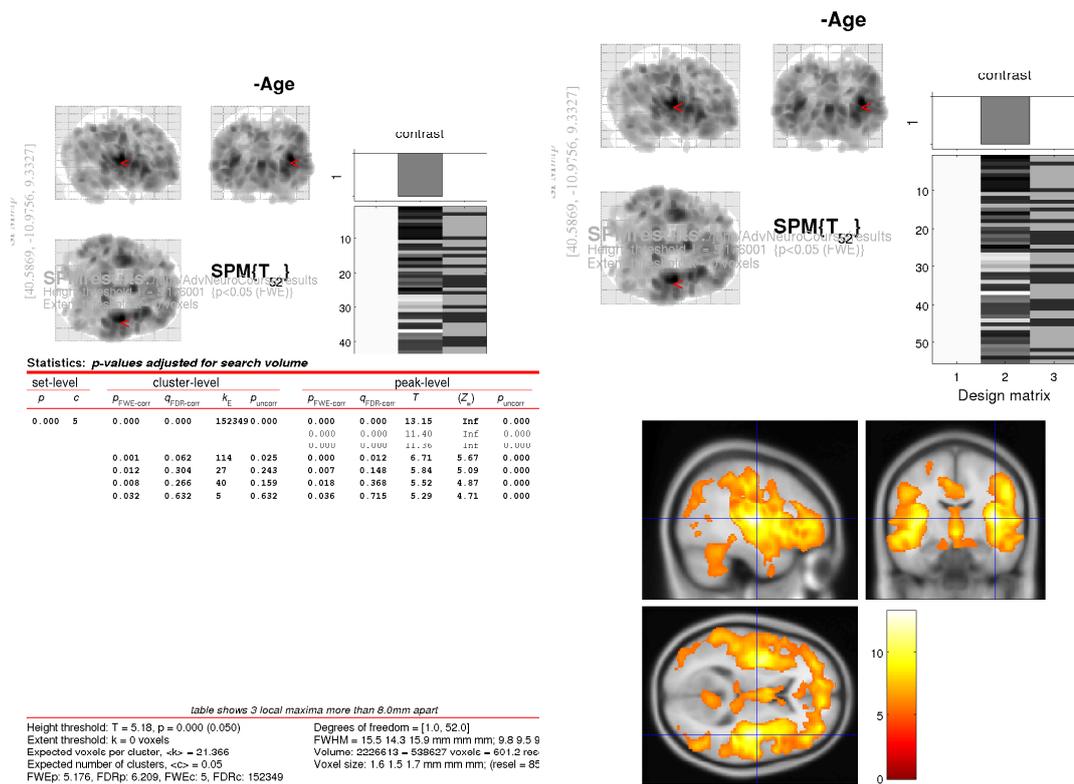


Figure 8: A table of results (left) and the blobs superimposed on one of the average p brains released with SPM (right).

After waiting some time for the results to appear, you should see some blobs that show what is likely to happen to your brain in future.

To better understand the results, it is a good idea to examine the image files generated by the analysis. *Check Reg* can be used for this.

## References

J. Ashburner. A fast diffeomorphic image registration algorithm. *NeuroImage*, 38(1):95–113, 2007. doi: <http://dx.doi.org/10.1016/j.neuroimage.2007.07.007>.

J. Ashburner. Computational anatomy with the SPM software. *Magnetic Resonance Imaging*, 2009. doi: 10.1016/j.mri.2009.01.006.

J. Ashburner and K.J. Friston. Unified segmentation. *NeuroImage*, 26:839–851, 2005. doi: doi:10.1016/j.neuroimage.2005.02.018.

J. Ashburner and K.J. Friston. Computing average shaped tissue probability templates. *NeuroImage*, 45(2):333–341, 2008. doi: 10.1016/j.neuroimage.2008.12.008.

K.J. Friston, A.P. Holmes, K.J. Worsley, J.B. Poline, C. Frith, and R.S.J. Frackowiak. Statistical parametric maps in functional imaging: A general linear

- approach. *Human Brain Mapping*, 2:189–210, 1995. URL [/spm/doc/papers/SPM\\_3/welcome.html](/spm/doc/papers/SPM_3/welcome.html).
- A. Mechelli, C.J. Price, K.J. Friston, and J. Ashburner. Voxel-based morphometry of the human brain: Methods and applications. *Current Medical Imaging Reviews*, pages 105–113, 2005.
- G. Ridgway, V. Litvak, G. Flandin, K.J. Friston, and W.D. Penny. The problem of low variance voxels in statistical parametric mapping; a new hat avoids a 'haircut'. *NeuroImage*, 59(3):2131–2141, 2012. doi: 10.1016/j.neuroimage.2011.10.027.
- K.J. Worsley, S. Marrett, P. Neelin, A. C. Vandal, K.J. Friston, and A. C. Evans. A unified statistical approach for determining significant voxels in images of cerebral activation. *Human Brain Mapping*, 4:58–73, 1996.