

SPM fMRI Analysis

Authors: D Glahn

Last Modified: 9/1/00

This document contains a set of note for step by step instructions of how to perform fMRI analysis using SPM (www.fil.ion.ucl.ac.uk/spm/). It is intended for those in the Cannon Lab. The notes are based on my knowledge of SPM and are likely wrong.

Table of Contents

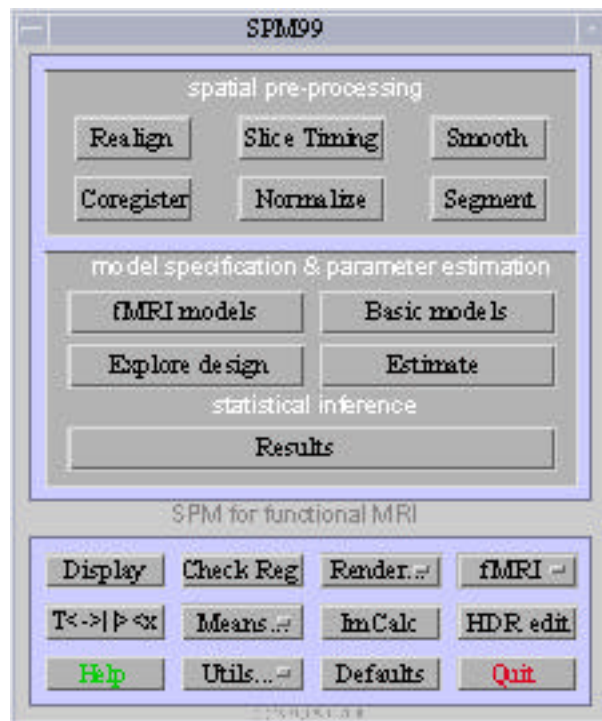
<i>Starting Up</i>	2
<i>Command Window</i>	2
<i>Spatial Pre-processing</i>	2
REALIGN	3
NORMALIZE	4
SMOOTH	5
<i>Model Specification and Parameter Estimation</i>	5
fMRI Models	5
Explore designs	8
Results	8
Basic Models	12

Starting Up

To start SPM, one must begin a MATLAB session. To do so, type 'matlab' in any window on any of our SUN stations. Once MATLAB loads, type 'spm'. A small window will appear asking the type of data you hope to analyze. Press the 'fMRI' button.

Command Window

Three windows will appear on the screen, two smaller windows on the left and a large window on the right. The upper left window is the command window (see below). From this window you will initiate different processes. The window is divided into three sections: spatial pre-processing, model specification & parameter estimation and SPM for functional MRI. The 'help' button on the lower left of the command window is somewhat helpful and is the basis for much of what is to follow.



Directly below the command window is the process window, where information is entered for a process. To the right is the graphics window where results are presented.

Spatial Pre-processing

We use three of the spatial pre-processing tools: realignment, spatial normalization and smoothing. SPM can only read in two or three-dimensional data sets in ANALYZE-7 format.

SPM assumes that images are in the following orientation:

X increases from Left to Right

Y increases from Posterior to Anterior

Z increases from Inferior to Superior

Images must be in this orientation for correct spatial normalization.

REALIGN

Realign is SPM for motion correction. This realignment uses a least squares approach and a 6 parameter (rigid body) spatial translation (Friston et al., 1996). The first image in a time-series is used as a reference scan to which all subsequent scans are aligned. It is good to note that after realignment, SPM creates a text file called "realignment_params_*imagename*.txt" which has estimates of movement (based on center of mass I think) for each time-point in the X, Y and Z directions. These estimates can be used to determine group or subject related movement differences.

After pressing the 'realign' button, SPM requests (process window):

- *Number of Subjects* In SPM99 one can do spatial pre-processing on a number of subjects at one time. Here one enters the number of different subjects.
- *Number of sessions* If you scanned one subject on the same task twice and are confident that the subject did not move significantly, you can motion correct both functional runs (sessions) to the first image on the first run. In that case, you enter 2 in this field.
- *Select Scans* A new window will appear and request that you select the images for Subject 1, Session 1. By changing the path and filter, you can navigate to the images you want. SPM groups images with similar names together and reports the number of images in those groups in white to the left of the common name. Once you select the images to be motion corrected, press 'done'.
- *Which Option?* SPM offer the options of Coregister Only, Reslice Only, or Coregister and Reslice (the default). In this case, the tem coregister means to generate parameters necessary for motion correction (these parameters are written to a .mat file). Reslice is the creation of new images that have been registered to the reference image.
If you intend to spatially normalize your images (place them into Talairach space) you should choose Coregister Only, as one should minimize the number of reslices performed on a data set. SPM will use these parameters when performing the spatial normalization.

Realignment will commence as soon as the option is chosen. If you choose to reslice your data, new images will be created with the "r" prefix. For all options, a .mat file will be created.

NORMALIZE

This module spatially (stereotactically) normalizes images into a standard space defined by some ideal model or template image. The template images supplied with SPM conform to the space defined by the ICBM, NIH P-20 project, and approximate that of the space described by the atlas of Talairach and Tournoux (1988).

Generally these algorithms work by minimizing the sum of squares differences between your image and the template. Spatial normalization is performed with affine (12 parameter) and quadratic (3rd order) automated algorithms (Ashburner & Friston, 1997, 1999).

After pressing the ‘normalize’ button, SPM requests (process window):

- *Which Options?* Determine Parameters Only, Write Normalized Only, Determine Parameters and Write Normalized (the default). These are the same as Coregister Only, Reslice Only, of Coregister and Reslice in the realign module. Typically we use the default.
- *Number of Subjects* In SPM99 one can do spatial pre-processing on a number of subjects at one time. Here one enters the number of different subjects
- *Image to Determine Parameters* A new window appears and requests you to choose an image to be used to morphing determine parameters for the first subject. Here people often use a higher resolution T1 or T2 weighted images to generate morphing parameters, assuming that the higher resolution image will provide a better morph. However, I have found that there is often a good deal of movement between the acquisition of the higher resolution images and the functional images. Hence, I advocate the use of the first EPI image in a time-series.
- *Images to Write Normalized* If you chose the Determine Parameters and Write Normalized option, a new window appears and requests you to choose the images to be write normalized. Here you choose the data you want in Talairach. Typically, this includes the functional images for the first subject. If you choose more than one subject, the last two steps are repeated until you have cycled through all of the subjects. If you make a mistake, you have to start again form ‘normalized’ button.
- *Template Image* A new window appears and requests you to choose the template image to be used. The window should automatically be in /spm99/templates. If you do not see any images, see if your filter is set to *.img. For the least squares registration to be most effective (“unbiased”), your image and the template should have similar contrasts. Hence, if you are normalizing a T1 weighted image, use the T1 weighted template or, conversely, if you are normalizing an EPI image, use the EPI template.

- *Interpolation Method?* SPM offer the options of Nearest Neighbor, Bilinear Interpolation (the default) and Sinc Interpolation (9x9x9). We typically use the default. The Sinc Interpolation method should be more accurate, but is much slower and the speed-accuracy tradeoff is poor.

Normalization will commence as soon as the option is chosen. For all options, .mat files are created. If you choose to write normalized images, new files with a “n” prefix will be created.

SMOOTH

This module convolves images with an isotropic Gaussian kernel. Issues concerning smoothing are discussed in detail in Friston et al., 2000.

After pressing the ‘smooth’ button, SPM requests (process window):

- *Smoothing {FWHM in mm}* You must enter the size of the kernel to be used for smoothing. We use a kernel that is more than one but less than the width of two voxels (e.g. 8 mm).
- *Select Scans* A new window appears and requests the scans to be smoothed. Here you can choose as many scans as you like (from as many subjects as you like). Smoothing in SPM takes quite a while.

After smoothing is complete, new files with a “s” prefix will be written.

Model Specification and Parameter Estimation

We typically use three of the tools in this section: fMRI models, Basic models and Results. Here I will describe how one would do an analysis on a single subject and then, using a random effects model, how one might do group analysis. It should be noted that I am only describing a simple on-off analysis and will not give complete explanations for all of my choice of options (some because I don’t know what they do and others for the sake of shortness). That being said, I will attempt to be thorough.

In general, one analyzes data in SPM in three steps. First one generates a design matrix or model made up of column vectors, which predict physiological responses to changing task conditions. Hence, the design matrix defines the experimental design and the nature of hypothesis testing. Second, one estimates “fit” parameters for these predictor vectors using a least squares approach to (multi) linear regression. Finally, one uses these parameter estimates for statistical inferences about single subject or group level hypotheses.

For the purposes of illustration, we will use a mock data set collected from a simple on-off paradigm (25 scan blocks) over 150 time-points.

fMRI Models

This module allows one to (i) specify a statistical model in terms of the design matrix, (ii) review that design, (iii) link data to a pre-existing design and (iv) specify and link data.

After pressing the ‘fMRI models’ button, SPM requests (process window):

- *What would you like to do?* Here you choose to specify a model, review a specified model, estimate a specified model, or specify and estimate a model. These options will be discussed below
- **SPECIFY A MODEL** a.k.a. create a design matrix. After choosing to specify a model, SPM requires the following information:
 - *Interscan Interval* Here one should input the time (in seconds) required to acquire a set of images for a single time point. This is often termed “TR”. We typically use a TR of 2.5 or 3.
 - *Scans Per Session* This is a request for the total number of time-points per subject in the run. For example, if 150 images were acquired during a task and you intend to analyze 3 subjects, enter 150 150 150. Alternately, if you have one subject performing the task three times, you would enter 150 150 150.
 - *Number of conditions or trials* Here one enters the number of conditions present in experiment. Although SPM99 does not require explicit modeling of the baseline condition, I prefer to do so. Thus, from our example, we would choose 2.
 - *Name of condition/trial1* At various points in the analysis, SPM will use this name to refer to this condition. However, at other points, it will not.... If you have more than one condition/trial, you will have to name each of them.
 - *Stochastic Design* This is SPM for single-trial design. Given that the design we propose is blocked, we will choose the ‘no’ button. I may give an explanation of single-trial designs below.
 - *SOA* SPM’s use to the term SOA (Stimulus Onset Asynchrony) is idiosyncratic. What they mean is the number of scans from between the beginning of a condition in one block and the next occurrence of that condition. Here SPM allows you to use a ‘fixed’ or ‘variable’ SOA. Our example task (ABABAB) uses a fixed SOA because the order of the conditions does not change over the experiment. However, if constructed our experiment in any other way, we would use a ‘variable’ SOA (e.g. ABBA, ABCBCA).
 - *SOA for trial1* If you choose a fixed SOA, SPM requires you to enter the fixed SOA for condition XX. In our example, this is 50.
 - *Time to 1st trial* Here SPM wants the number of the first scan that this condition appears (i.e. 0 or 25).
 - *Vector of onsets for trial1* If you choose a variable SOA, SPM requires you to enter a vector (by hand) of each time-point where the condition begins.
 - *Variable duration* of block length for a single condition over the task. If ‘yes’ you will need to enter an array of the length of the block for each occurrence of that condition.
 - *Session 1/1 Parametric specification* Here SPM will create a new single predictor, based on one of the task conditions, that will be parametrically varied across the experiment. This option is a bit

misleading, in that the SPM will not allow you to use this predictor independently and I think they intend its use for partitioning out (error) variance. You are given the choice of 'none', 'time' or 'other'. If you choose 'time', you must decide the type of expansion (linear, exponential, or polynomial) and which condition to model. If you choose 'other', you must enter a name for the parameter, the type of expansion, which condition to model, and input a vector of weights for each block of that condition.

- *Are these trials* Here you choose between 'events', 'epochs' or 'mixed'. Events is SPM for a single trial and should be used with the stochastic option above. Epochs are blocks and Mixed is a time series with both single-trials and blocked trials. Here we choose epochs. I may give an explanation of single-trial designs below.
- *Select type of response* Here SPM requests a function to model the predictor vectors. Choices are 'bias functions (Discrete Cosine Set)', 'bias function (Mean & exponential decay)', 'fixed response (Half-sine)', and 'fixed response (Box-car)'. Your choice of function will be based on assumptions of how BOLD signal varies over the course of the experiment. Without evidence to the contrary, the 'fixed response (Box-car)' option seems to be the most obvious.
- *Convolve with HRF* Here SPM gives the option to convolve the predictor vectors with a hemodynamic-response function. We always choose yes (see Cohen, 1997; Aguirre, Zarahn & D'Esposito, 1998).
- *Add temporal derivatives* No XXX
- *Epoch length for trial 1* The length (in number of scans) for a single block in trial XX. I think this option may only be present if you choose a fixed SOA.
- *Interactions among trials (Volterra)* As far as I can tell, this uses a variance co-variance matrix based on the design to predict interactions among the different conditions and then models that interaction. The math behind this is somewhat difficult (at least for me (Friston, Josephs, Rees & Turner, 1998)) and choosing this option greatly increases the time needed to estimate parameters.
- *User specified regressors* Here SPM would like the number of predictors (regressors) that you would like to enter on your own. For each of these, you will need to enter the weight of that predictor at each time point of the experiment. Typically, I choose 0. Once this option is completed, the design matrix will appear in the Graphics window.
- **REVIEW A SPECIFIED MODEL** With this option you can review a design matrix that was created previously. After choosing this option, a window will appear asking you to select a 'SPM_fMRIDesMtx.mat' file. Choosing one of these and pressing the

‘done’ button will cause a design matrix to appear in the Graphics window. You can use the ‘Explore fMRI design’ pull-down menu to view different conditions.

- **ESTIMATE A SPECIFIED MODEL** With this option you can use a previously created design matrix to create parameter estimates. After choosing this option, a window will appear asking you to select a ‘SPM_fMRIDesMtx.mat’ file, followed by a window asking you to choose fMRI data to be analyzed. After choosing this data, SPM requires the following information:
 - *Remove Global Effects* This is a global scaling placed on the data to have the mean intensity of the data be 1000. Reasons for global scaling include better comparability across subjects and for meaningful units to beta estimates. However, global scaling does add artifacts to your data. Nonetheless, we typically do chose the ‘scale’ button.
 - *High-pass filter* Here SPM will use a filter to remove the effects of high-frequency variations within your data set (mostly slow alterations in the magnetic field). If you choose this option, you must specify the exact frequencies to be discarded.
 - *Low-pass filter* SPM will use a filter to remove the effects of low-frequency variations within your data set (typically variance associated with physiologic processes such as breathing and heart-rate). Here your options are ‘none’, ‘Gaussian’ and ‘hrf’. For the hrf option, any signal with a frequency lower than the hemodynamic response function will be removed. We typically choose this option.
 - *Model intrinsic correlations?* I have yet to find this option described in any real detail and hence choose ‘none’.
 - *Setup trial-specific F-contrasts* This option is new to SPM99 and allows one to save a bit of time in the next phase of analysis. However, I have never used it.
 - *Estimate?* Choose ‘now’ or ‘later’.
- **SPECIFY AND ESTIMATE A MODEL** With this option you can create a design matrix and estimate parameters based on that model. Basically, it is a combination of specifying a model and estimating one. The information requested is the same as when these stages are done separately, just in a slightly different order.

Explore designs

This module allows one to review a design matrix and seems to be identical to the ‘review a specified model’ option under the fMRI models module.

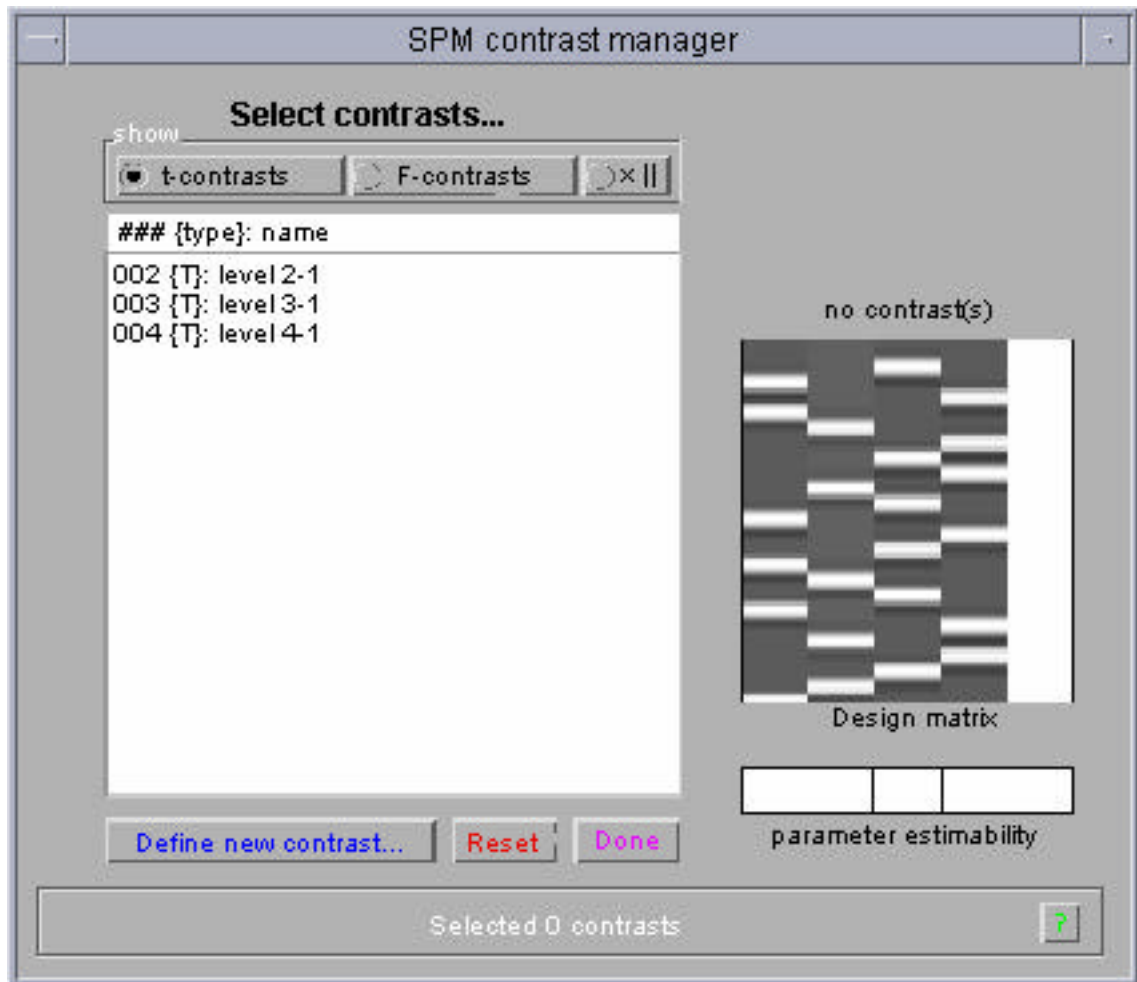
Results

This module is for the interactive exploration and characterization of the results of a statistical analysis. It requires an SPM.mat file which houses the parameter

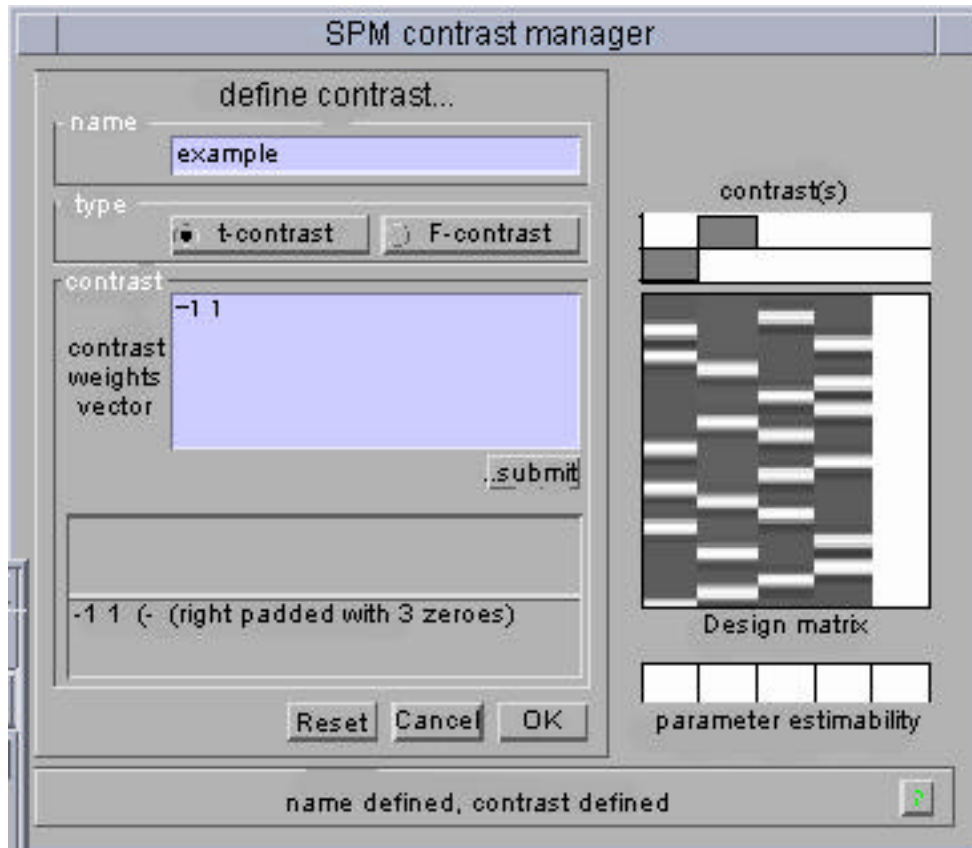
estimates generated in the previous analyses. Once the 'results' button is pushed, a window will appear asking for the SPM.mat file.

After choosing this file, the SPM contrast manager will appear (see below). This somewhat clunky tool allows you to create contrasts based on the parameter estimates of the predictor vectors created in the design matrix. The first choice to be made is that of statistical test (t or f). The basic difference for this context is that t-tests are uni-directional while f tests are bi-directional. That is, t-tests either look at positive or negative differences between parameter estimates where f-tests look simply at differences in parameter estimates. Given that the authors of SPM argue that only t-tests can be used in latter random effects models, we will focus on the t-test.

- **USING THE CONTRAST MANAGER** To create a contrast, press the 'Define new contrast' button on the lower left portion of the contrast manager window.



The window will change to the contrast “maker” orientation (see below).



- *Name* Type the name of the contrast
- *Type* Chose either ‘t-contrast’ or ‘f-contrast’
- *Contrast* Here one enters the weights to be assigned the various predictor variables in questions. It may be a good idea to keep these weights orthogonal (i.e. sum to 0), but that is based mostly on unsubstantiated ideas. Using our simple example experiment, one might enter $-1\ 1$. This contrast asks the question, what voxels showed increases in the second condition relative to the first. Pressing the ‘submit’ button below the contrast box will provide a graphical representation of the contrast (found just above the graphical representation of the design matrix on the right half of the screen). It is important to note that if your design matrix includes multiple subjects and you intend to perform a fixed effect model, you must enter in contrast weights for each predictor for each subject in your model. Thus if you had three subjects data on our experimental task and you wanted to know how these subjects (as a whole) performed on condition 2 as compared to condition 1, you would enter the following contrast $-1\ 1\ -1\ 1\ -1\ 1$. Press the ‘ok’ button on the lower right when the contrast meets your specifications.

Once a contrast has been specified (by clicking on it), press the ‘done’ button. This will return you to the Results module. In the process window, you will be asked to provide the following information:

- *Mask with other contrast* Here SPM will mask the current contrast with another contrast. If you choose to do so, SPM will ask you to choose the additional contrast (via the contrast manager), choose an *uncorrected mask t-value* for that mask, and specify the nature of your mask, ‘inclusive’ or ‘exclusive’. The idea here is that you may want to only query the voxels that were found “active” in a previous contrast. By doing so, you decrease the number of voxels used in the current analysis, decreasing the number of comparisons and minimizing the correction for these comparisons.
- *Title for comparison* This is the title that will appear above the output in the graphics window. Once a title is chosen, the process window will show the % of the computation completed. After these computations are complete:
- *Corrected height threshold* This is a correction for the height of the statistic in question. That is, in the case of a t-test, what is the probability of that large of a t-test occurring by chance. If you choose ‘yes’, you will be asked to enter a corrected p-value (e.g. 0.05). This p-value is based on random Gaussian field theory and is one of the “wonders” of SPM. If you choose ‘no’, SPM will ask for a threshold for display purposes.
- *Extent threshold* This is a correction for the extent of activation. Here you specify how large (in voxels) you feel an area of activation should be to consider it “real” (see Forman et al., 1995). I typically use 0. With this choice made, the results of your analysis should appear on the Graphics window.
- **USING THE RESULTS WINDOW** The results window (located on the bottom half of the process window) provides tools to review your analysis (see below). I will give a brief description of each.



- *Volume* After pressing this button, SPM will report (in a table on the lower half of the Graphics window) each of the clusters of activation above the chosen threshold. On the far left are two columns titled “set-level”. This gives the p-value and ‘c’ statistic (no idea what this is) for the entire volume. The next set of columns are marked “cluster-level” and provide statistics relevant to clusters of activation within your data set. These statistics include a corrected p-value (the probability that a cluster

of this size would appear by chance, corrected using Gaussian field theory), the k statistic (indicating in some way the size of that cluster) and the raw p-value (the probability that a cluster of this size would appear by chance). The next set of columns give voxel-level statistics. These are local maxima within the relevant cluster. If more than one entry is presented for a cluster, it suggests that this cluster has more than one maximum separated by some amount of space. The statistics provided here include the corrected p-value (the probability that a t-value of this amplitude would be present due to chance, corrected for the number of observations using Gaussian field theory), the t-value, the Z-value (the t-value represented as a Z statistic) and the uncorrected p-value. The final set of columns are the coordinates (x, y, z) of the reported voxels. If the data was normalized into Talairach space, these are Talairach coordinates.

- *Cluster* After placing the indicator on an area of activation represented on the “glass” brain on the top of the Graphics window, pressing this button will give you the statistics for that cluster. The statistics are reported in the same way as above.
- *S.V.C. and V.O.I.* buttons are for looking at small volumes within the brain. I think they stand for volume of interest and small volume correction. I have never used this option and don't know much about them.
- *Plot* This button lets one plot parameter estimates with various things. Again, I have never used this option and don't know much about them.
- *Overlays* This button allows the option of ‘slices’, ‘sections’ or ‘rendered’. The slices option will overlay the three slices of activation over an image of your choice. These pictures are color-coded. The sections option will overlay the activation onto a 3d representation of a brain. This option allows you to move around in 3d space and I think is the easiest way to get a feeling for your data. Finally, the render option will overlay your data onto a rendered file. If you choose this option, a window will appear asking you to choose a rendered template (I like the `render_single_subj.mat`). You will then be asked to choose between the new or old style and how bright you want your blobs. After some computation time, the rendered brain will appear on the bottom half of the Graphics window.

Basic Models

Although one can perform standard and even somewhat complex analysis of time-series data using the fMRI models module, SPM99 offers the Basic Models module to perform somewhat less complex analyses that closer to those used with PET data. Each of these models rests on a different set of assumption and here we will only consider one of these models (a two-way t-test) in the context of using SPM99 to perform random effects analyses.

RANDOM EFFECTS ANALYSES These notes thus far have focused upon experimental designs that attempt to address hypotheses regarding neural activity

in individual subjects. Here we turn to considering one of the issues of multi-subject comparisons, that of fixed versus random effects statistical models. Most PET and fMRI studies reported to date in which repeated observations have been obtained from each subject, a fixed-effect statistical model has been employed. This type of model makes the assumption that observations from different subjects are the same as observations from within a single subject (no subject \times task interaction; Holmes et al., 1998). If one wishes to make inferences about an effect that extends to the population from which the subjects were drawn, then it is necessary to employ a statistical model that explicitly accounts for a subject \times task interaction. Such random effects models are actually rather simple to perform in SPM, but require two stages. The first is an analysis of individual subjects (dependent on intra-subject variability) and second is an analysis of the group (dependent on inter-subject variability).

The individual subjects are analyzed at a within-subject level using balanced standard SPM models (such as those found in the fMRI Models module). For each subject, the contrast of the parameter estimates of interest is computed and written out as a “contrast” image (conXX.img). Note this is not the SPM (t or F image), rather it is the linear combination of the parameter estimates). These contrast images surmise the response for each subject and are then used as input for a between-subject level analysis. The between-subject analysis then may simply be a two-way t-test contrasting subjects from two groups. Assuming that we have performed single-subject analysis outlined above, we will now turn to the between-subject analysis.

After pressing the ‘Basic Models’ button, SPM requests (process window):

- *Select design type?* Here you choose the type of statistical model to be performed. Here we will illustrate a two sample t-test.
- *Select images* For a random effects model, one would choose the conXX.img images from subjects from different groups. For example, if we had 10 subjects from two groups (5 in each) we would select the same contrast images for each subject.
- *Group?* Here one enters a vector (0’s and 1’s) specifying each input image to each group. In our example, this may be 0 0 0 0 0 1 1 1 1 1.
- *Gmsca: grand mean scaling* Since we used global normalization for the subject-level analysis, we should choose the default of ‘no grand mean scaling’.
- *Threshold masking* SPM will apply a threshold mask to remove non-brain voxels from the analysis. However, SPM masks all images used in any fMRI Models analysis and hence our images have been previously masked. We choose the ‘none’ option.
- *Implicit masking (ignore zero’s)?* Here SPM is asking if you would like to include voxels that were found in only some of your subjects. Here we typically choose ‘yes’.
- *Explicitly mask images?* As discussed above, our images have been masked previously and we choose ‘no’ here.

SPM notes

- *Global calculation...omit*
- *Estimate? Now or latter*
- specify a model, review a specified model, estimate a specified model, or specify and estimate a model. These options will be discussed bellow

Disclaimer: Whenever possible, I stole the text for these notes from the SPM help or other sources. I hope to include a section on single subject designs soon....