Chapter 2

Materials and Methods

Part II Statistical Analysis of Functional MRI Time Series

2.4 Introduction

Functional neuroimaging techniques provide a means for making inferences about differences in regional brain activity between different conditions or states. To localise a function to a specific anatomical region it is critical that the experimental design allows one to unambiguously consider only the effect of the appropriate cognitive manipulation. Identifying functionally specialised brain regions generally proceeds using statistical parametric mapping (SPM). Functional specialisation is a fundamental principle of brain organisation and this thesis focuses on functional specialisation for learning within the human medial temporal lobe. It is important, however, to acknowledge that the brain is also organised on the basis of functional integration, where integration within and among specialised areas is mediated by effective connectivity. Advances in SPM have enabled characterisation of functional integration, an example of which is given in Chapter 5 part I.

In order to assign an observed response to a particular brain structure or cortical area, the data must conform to a known anatomical space. Hence, the analysis of fMRI data starts with a series of spatial transformations of brain images aimed at reducing artifactual variance components in each voxel time series. The imaging time series is first realigned to a common reference frame to correct for subject movement during scanning. After realignment the data are transformed using linear and nonlinear mappings into a standard anatomical space (Friston *et al*., 1995a). This normalisation procedure allows averaging data across subjects and permits data reporting within a standardised reference co-ordinate system. Following this preprocessing, a statistical model must be created in order to draw inferences

about differences in regional brain activity between different conditions (Friston *et al*., 1995b). Finally, these inferences must be corrected to guard against Type I error.

2.5 Spatial preprocessing

2.51 Realignment

2.51a Spatial

Head motion during fMRI can give rise to artifactual change in signal intensity. Despite subjects being firmly immobilised with soft head pads, even the best subjects show movement up to a millimetre or so. Realignment removes variance from a time series which would otherwise be attributable to error (hence decreased sensitivity) or to evoked effects i.e. if movement is correlated with the cognitive task.

Realignment involves the following steps:

1. estimating the 6 parameters of the affine 'rigid body' transformation that minimises the [sum of squared] differences between each successive scan and the first. In three dimensions, a rigid body transformation can be defined by 6 parameters, typically three translations and three rotations about orthogonal axes.
2. Applying the transformation by resampling the data using sinc or trilinear interpolation.

2.51b Temporal

In multi-slice acquisitions, different slices will be acquired at slightly different times. In fMRI analysis, stimulus onset times are specified in scans, hence

posing the problem for event-related fMRI studies (see later) that certain slices will be more sensitive to a particular model of haemodynamic responses. Temporal realignment ensures that the data from all slices within a given volume correspond to the same time point. This is achieved using sinc interpolation over time (when temporal dynamics of the evoked responses are important, i.e. event-related fMRI studies, and the TR sufficiently small to permit interpolation).

2.52 Spatial normalisation

Realignment produces a mean image for the time series. This mean image is used to estimate the warping parameters that map this image onto a template (in fMRI this is a template EPI image) that already conforms to a standard anatomical space (in the case of SPM this is the space defined by the atlas of Talairach and Tournoux, 1988). The estimation is achieved using:

1. A 12-parameter affine transformation where the parameters constitute a spatial transformation matrix (the affine transformation is similar to that used during realignment but also includes zooms and shears).
2. Low-frequency basis functions (a discrete cosine set) where the parameters to be estimated are the coefficients of the basis functions employed.

A Bayesian framework is used to estimate the parameters of both models, where one wants to find the deformation that is most likely given the data. The deformation is updated iteratively to minimise the sum of squared differences between the template and the deformed image and reflects the probability of actually getting that image if the transformation was correct. Prior information about the likelihood of a

given transformation is incorporated by weighting the least squares (Ashburner *et al*., 1997).

This procedure can be extended to allow normalisation across image modalities, allowing, for example, the SPM obtained from a fMRI time series to be overlaid on an individual subject's structural T1 image. First, that subject's T1 image must be mapped into the same space as the EPI images. Here, the difference between the EPI template and T1 image can be minimised using a combination of templates depicting gray, white, CSF, and skull tissue partitions. This approach was adopted in the first experiment reported in this thesis (chapter 3).

2.53 Spatial smoothing

After normalisation, the fMRI data are smoothed by applying a Gaussian kernel (point spread function), of known width, to each voxel. The motivations for smoothing are:

1. Smoothing the data render them more parametric in their distribution and ensures the validity of parametric statistical tests.
2. Smooth data is one of the assumptions of Gaussian Field Theory (see later).
3. In order to average across subjects it is necessary to smooth so that regional effects are expressed at a spatial scale where homologies in functional anatomy exist over subjects.
4. The matched filter theorem states that the optimum smoothing kernel corresponds to the size of the effect anticipated. According to optical imaging experiments, the spatial scale of the haemodynamic response is about 2-5 mm.

2.6 Characterising haemodynamic responses using the General Linear Model

Following spatial preprocessing, the data are ready for statistical analysis. In the absence of prior information regarding the physical location of a particular function, statistical analysis of evoked haemodynamc responses must test for experimentally-induced effects at each intracerebral voxel individually and simultaneously. This involves two steps. Firstly, statistics indicating evidence against the null hypothesis (i.e. no experimentally-induced effect) are computed at each voxel. This statistic is usually a t or F statistic and is based on the parameter estimates calculated using the general linear model. This procedure results in an 'image' of statistics (i.e. SPM). Secondly, an inference must be drawn from this SPM, reliably locating voxels where an effect is present whilst controlling against the probability of false positives.

Data analysis as implemented in SPM is parametric. Statistics with a known null distribution are used, such that under the null hypothesis, the probability of obtaining a statistic greater than, or equal to, that observed can be computed. The statistical model used is a special case of the general linear model.

2.61 Parameter estimation using the General Linear Model

Commonly used parametric models, such as linear regression, t-tests and analysis of variance (ANOVA) are special cases of the general linear model. This model explains variation in the data, Y, in terms of a linear combination of the explanatory variables (x), plus an error term:

$$Y_j = x_{j1}\beta_1 + \ldots + x_{jl}\beta_l + \ldots + xj_L\beta_L + \varepsilon_j.$$

The β_l are unknown parameters, corresponding to each of the L explanatory variables for the *j*th observation of Y. The errors ε are assumed to be identically and normally distributed.

For *J* observations of Y, the general linear model can be expressed in matrix formulation:

$$\mathbf{Y} = \mathbf{X\beta} + \boldsymbol{\varepsilon}$$

for the column vector of observations **Y**, the column vector of error terms **ε** and the column vector of parameters **β; β** = $[\beta_1 \ldots \beta_j \ldots \beta_J]^T$. Matrix **X**, of size *J* x *L*, is the design matrix. This matrix has one row per observation and one column per model parameter. The number of parameters *L* is (usually) less than the number of observations *J* hence the simultaneous equations implied by the general linear model (obtained by expanding the matrix formulation with **ε=0**) cannot be solved (it is overdetermined). Therefore, some method is required for estimating parameters that "best fit" the data. The method adopted is that of least squares.

Each column in the design matrix (**X**) corresponds to some effect that one has built into the experiment, such as the alternating 'boxcar' function modelling alternating activation and control epochs in the experiment presented in chapter 3, or effects that may confound the results. The latter include a series of terms that are

designed to remove or model low-frequency variations in signal due to artifacts such as aliased biorhythms or scanner drift. The relative contribution of each of these columns to the experimental variance (i.e. the parameter estimate for each column) is assessed using generalised least squares estimators.

Inferences about the parameter estimates are made using their estimated variances. This allows for two types of statistical test. One can test the null hypothesis that all the estimates are zero using the *F* statistic to give a SPM{*F*}or, alternatively, that some particular linear combination or a "contrast" (e.g. a subtraction) of the estimates is zero using a SPM{*T*}. The T statistic is calculated by dividing the contrast (specified by contrast weights) of parameter estimates by the standard error of that contrast. This error term is estimated using the variance of the residuals about the least squares fit. An example of contrast weights could be [1 –1 0 0 …] to compare the differential responses evoked by two conditions that have been modelled by the first two condition-specific regressors (columns) in the design matrix.

As stated above, an important assumption in the analysis of time-series is that the residuals are identically and normally distributed. However, the haemodynamic response is of longer duration than the typical scan acquisition time, which leads to serial correlations among the error terms. The general linear model accounts for these autocorrelations by imposing a known temporal smoothing function on the time-series and adjusting the estimators and degrees of freedom accordingly (Worsley and Friston, 1995).

2.62 Statistical inference and the Theory of Gaussian Fields

As already described, SPM calculates the significance of brain activity in a voxel-wise manner. However, in many instances, one cannot have an *a priori* prediction for the precise voxel which is going to be activated in a particular condition. Inferences can be of two sorts depending on whether one has an *a priori* hypothesis about the particular region of the brain engaged. In classical statistics, multiple tests require a correction (such as Bonferroni correction) for the number of tests performed on the data. If the hypothesis is not constrained by anatomy (i.e. the null hypothesis states there is no effect anywhere in the brain), the p values of the ensuing maxima must be corrected for multiple comparisons. This correction is based on Gaussian field theory, which takes into account the fact that neighbouring voxels are not independent by virtue of spatial smoothing in the original EPI images. This correction is similar to a Bonferroni correction for multiple comparisons but less severe provided that the data are sufficiently smooth. Alternatively, if the hypothesis is anatomically constrained to effects in a particular brain region, the uncorrected p value associated with the height (magnitude) of the local maximum in that region of the SPM can be used to reject the null hypothesis.

2.62a Anatomically constrained hypotheses

If one has predicted activation in a particular brain region, a correction of the p value in this region for the entire search volume is inappropriate. As the *a priori* hypothesis is regional, and not voxel-specific, some form of correction is still required. This can be achieved by prespecifying an appropriate search volume and making the appropriate correction (Worsley *et al*., 1996).

2.62b Anatomically open hypotheses: levels of inference

Inference can be drawn at a number of different levels. SPM derives p values pertaining to

1. *set-level inference*: the number of activated regions (i.e. number of clusters above a height and volume threshold)
2. *cluster-level inference*: the number of activated voxels comprising a particular region (i.e. volume of cluster)
3. *voxel-level inference*: the p-value for each voxel within that cluster. Significance testing in this thesis is limited to testing significance at the voxel level.

These p values can be corrected for the multiple dependent comparisons based on *c* or more clusters with *k* or more voxels above a threshold *t* in an SPM of known smoothness. The probability P that a maximum value in a cluster would be greater than that observed under the null hypothesis (when no activation is present), can then be tested. To approximate P, SPM uses the expected Euler characteristic (a topological measure that is effectively the number of clusters minus the number of holes). At high thresholds, the expected Euler characteristic simply counts the number of regions above *t*. The expected Euler characteristic for a given threshold gives the probability of the maximum exceeding that threshold, indicating the test level that would just reject the null hypothesis at that voxel. This corresponds to the adjusted p value at that voxel – the probability that the maxima in a SPM would be greater than the voxel value.

The expected Euler characteristic depends on the smoothness of the data. The smoothness of the statistical images is calculated from the component fields. These

are the residual fields (what is left over after the modelled response at each voxel is subtracted from the data) after they have been normalised by the variance at each individual voxel. Although the smoothness is parameterised by a variance-covariance matrix of the spatial partial derivatives, it is represented as the FWHM of an equivalent Gaussian point spread function (PSF).

The expected number of clusters, the expected Euler characteristic, depends on the expected number of clusters of any size and the probability that the cluster will be bigger than k resels. The resel (resolution element) parameterises the number of independent tests. The total number of resels is equal to the volume of the search region divided by the product of the FWHMs of the smoothness in each dimension. As the smoothness increases, the number of resels decreases yielding fewer independent tests, hence the probability of obtaining a maxima as large as that observed gets smaller.

In this thesis, I report activations surviving a corrected threshold of $p<0.05$. In addition, I report medial temporal activations with an uncorrected $p<0.001$ and correct these for the volume of the medial temporal region of interest. In the case of functional segregation within the hippocampus, the volume of the hippocampus was divided into anterior and posterior portions and the p values adjusted accordingly.

2.7 Event-related fMRI

Event-related fMRI can be defined as the use of fMRI to detect transient haemodynamic responses to brief stimuli or tasks (Josephs *et al.*, 1997). Event-

related, or trial-based measurement is already standard in electrophysiology, namely stimulus-locked event-related potentials. Previous functional imaging methods, such as PET, have limited temporal resolution necessitating measurement of prolonged states of brain activity. Such state-based designs were initially adopted in fMRI and referred to as epoch-related designs, the first experiment presented in this thesis being an example. Improvements in sensitivity and temporal resolution of fMRI have allowed an event-related approach. The event-related approach offers several advantages (Josephs and Henson, 1999).

1. The order of trials can be randomised hence the response to a trial is neither confounded by a subject's cognitive set nor systematically influenced by previous trials (Johnson *et al*., 1997).
2. Individual trials can be categorised or parameterised post-hoc according to a subject's performance. An example is the categorisation of event related responses to presented words according to whether a particular word was subsequently remembered or forgotten in a test of free recall (chapter 6).
3. Some experiments involve events that cannot be blocked, such as 'oddball' paradigms where the event of interest is a stimulus that violates the prevailing context. An oddball experiment is presented in chapter 5.
4. Some events can occur unpredictably and can only be indicated by the subject (such as the spontaneous perceptual transitions measured by Portas *et al*., 1999).
5. Event-related fMRI allows more direct comparison with other techniques such as ERP or psychophysics.
6. Extensions to epoch-related designs. A state can be modelled, to first order, as a continuous train of events, each representing one trial within an epoch. This

method also enables stimulus or response parameters to be modelled within an epoch. For example, the second experiment presented in this thesis was epoch-based, but analysed in an event-related manner to enable correct and incorrect trials to be modelled separately, and trial-by-trial within-epoch performance to be modelled parametrically.

In analysing event-related data, the explanatory variables are created by convolving a set of delta functions, indicating the onset times of a particular event, with a small set of basis functions that model the haemodynamic responses to those events (Josephs *et al*., 1997). The approach adopted in this thesis was to employ a multivariate Taylor expansion of a mixture of gamma-functions which approximate to a canonical haemodynamic response function (HRF; Friston *et al*., 1998a). The higher order basis functions in this expansion include the partial derivative of the HRF with respect to time. This approach has the advantage that the parameter for each covariate is interpretable in terms of response magnitude (canonical HRF) and latency (temporal derivative). The canonical form and its derivative can be tested separately by means of univariate t-tests and together by multivariate F-tests.

2.8 Optimising experimental design

Optimisation involves maximising the sensitivity (signal:noise ratio) for particular contrasts (hypotheses) as a function of stimulus ordering and stimulus onset asynchrony (SOA). Any experimental design can be characterised by the minimal SOA, SOA_{min}, and the probability, p, of an event occurring each SOA_{min} (Friston *et al*., 1999). All experiments in this thesis contain multiple event types, for

which the efficiency of the design depends on the specific hypothesis (contrast). With a randomised design involving two event types, (ABBBABAABAABB etc) the efficiency of the differential effect is maximal at minimal SOA_{min}. However, efficiency of the common effect (versus baseline) is then minimal. When stimulus ordering is constrained, an alternating design may result (ABABABABA etc), for which the optimal SOA_{min} for a differential effect is approximately 8s. When the design needs to be sensitive to both the differential and common effects, 'null events' can be introduced, when no stimulus is presented (ABB-B--A-ABBAA- etc). The most efficient estimation of both the differential and common effects in this case is with minimal SOA_{min}. All experiments presented in this thesis examine differential effects within randomised designs, hence a minimal SOA_{min} has been adopted in all of the event-related experiments described. The duration of the physiological haemodynamic response, and, therefore, of the fitted response, is approximately 30 s. Although the stimulus onset asynchrony in the three event-related experiments presented in this thesis ranged from 2.5 - 4 s, leading to overlap of successive haemodynamic responses, SPM accounts for this overlap using an implicit convolution regression model (Friston *et al*., 1998b).

2.9 Inferences about subjects and populations: Random vs Fixed effects analyses

The statistical inference drawn from fMRI time series may be of two types. Firstly, the results may be specific to the particular subjects at the time of scanning. This fixed effects inference is drawn from the effect size relative to the within

subject variability. Highly significant results typically obtain from a fixed effects analysis because the degrees of freedom are high as they pertain to the number of scans across all subjects. The effect size is averaged across subjects hence 6 subjects are normally included in these analyses to provide a representative mean across subjects. However, the limitation of this type of analysis is that an effect size may be primarily driven by a few subjects. To overcome this one could perform a conjunction analysis across the six subjects which tests for regions commonly activated by a particular condition in all subjects. In essence, the fixed effects analysis is an extension of a case report, commonly used in clinical studies and animal lesion experiments, where an effect is observed in a particular subject and then this effect is replicated in further subjects.

Random effects analyses allow inferences to be made about the population from which the sample of subjects was drawn. One observation per subject per condition is entered into a random effects analysis (usually a contrast of parameter estimates from a 1^{st} level analysis). Hence the effect size is compared against the between subject variability in these contrasts. This type of analysis is, therefore, not at risk of being biased by strong effects in a subset of subjects. It follows that more subjects are required to achieve a significant result with random effects analyses, as the degrees of freedom depend on the number of subjects scanned, a suitable minimum number of subjects being 12. A random effects analysis was the default analysis for data presented in this thesis.

Random effects analyses are typically a one-sample t-test testing whether the estimated effect size (i.e. contrast) is significantly greater than zero across all

subjects. In a fMRI experiment it is, however, unlikely that there is only one effect of interest. Take for example a 2x2 factorial design. Using one-sample t-tests to draw second level inferences would mean running separate analyses for each main effect and the interaction. An alternative is to conduct a repeated measures ANOVA, entering one observation for each of the four cells of the 2x2 design for each subject. This approach was adopted for the oddball experiment described in chapter 5. It should, however, be noted that if more than one condition is entered into the second level analysis, the assumption of sphericity (e.g. homogeneity of co-variance) must be made. This assumes that the between subject, within-condition variability is at the same level for each observation contrast (e.g. cell in the 2x2 factorial design).