2 Introduction to fMRI: experimental design and data analysis

2.1 Introduction to fMRI

Functional Magnetic Resonance Imaging (functional MRI or fMRI) is a non-invasive neuroimaging technique that can be used for studying human brain function *in vivo*. Functional MRI extends the use of Magnetic Resonance Imaging to provide information about biological function in addition to the anatomical information. Seiji Ogawa first demonstrated that by measuring the blood-oxygenation-level-dependent (BOLD) signal, Functional MRI could be used to visualize brain function (Ogawa et al., 1990).

The BOLD fMRI technique is designed to measure primarily, changes in the inhomogeneity of the magnetic field that result from changes in blood oxygenation. The fact that haemoglobin and deoxyhaemoglobin are magnetically different is exploited in the BOLD technique. Magnetic susceptibility refers to the amount of magnetization that can be achieved when a material is placed in a magnetic field. Deoxyhaemoglobin is paramagnetic and introduces an inhomogeneity into the nearby magnetic field, while oxyhaemoglobin is weakly diamagnetic and has little effect. Thus, the paramagnetic deoxyhaemoglobin induces a susceptibility difference between the blood vessels and the surrounding tissue can be used as an endogenous contrast (i.e. depends on intrinsic property of the biological tissue).

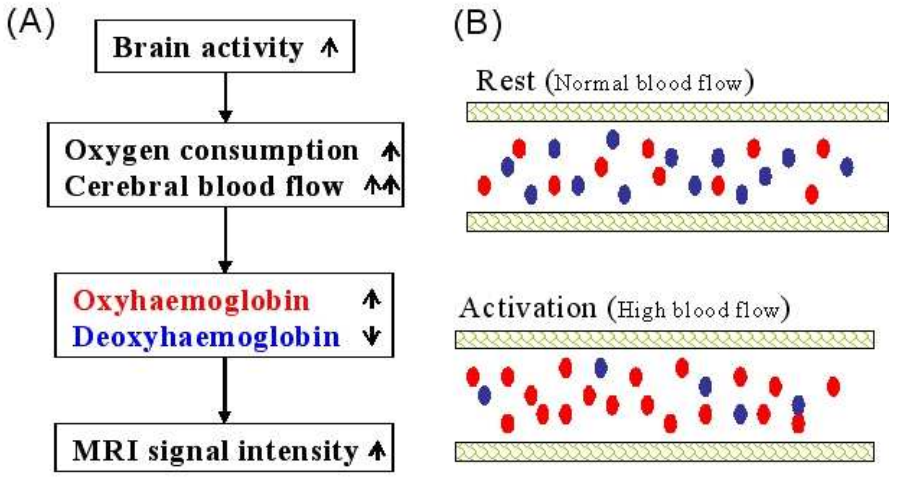

Figure 2.1 BOLD mechanism of functional MRI

(A) Blood-oxygen level-dependent signal mechanism in magnetic timbre imaging (B) oxyhaemoglobin and deoxyhaemoglobin blood flow during rest and activation

Hydrogen nuclei (protons) have magnetic properties, called nuclear spin. They behave like tiny rotating magnets. In presence of a magnetic field the hydrogen atoms, present in the water molecules of the brain, align themselves with this field and reach an equilibrium state. Exchange of energy between two systems at a specific frequency is called resonance. Magnetic resonance corresponds to the energetic interaction between spins and electromagnetic radio frequency (RF). When a brief radio frequency (RF) is applied, the hydrogen atoms absorb energy (excitation) and their equilibrium state is perturbed. These hydrogen atoms would emit energy (relaxation) at the same radio frequency until they gradually return to their equilibrium state. The magnetic vector of spinning protons can be broken down into two orthogonal components: a longitudinal or Z component, and a transverse component, lying on the XY plane. Relaxation gives rise to the magnetic resonance signal and is composed of two components. Longitudinal relaxation is due to energy exchange between the spins and surrounding lattice (spin-lattice relaxation, decay constant T1) and Transverse relaxation (spin-spin relaxation, decay constant T2) occurs due to the spins getting out of phase. T1 depends on the applied magnetic field strength with longer relaxation times for greater field

strengths. T2 is independent of the applied magnetic field strength and is always shorter than T1. The observed transverse relaxation time T2* is always shorter than T2 due to the combined effect of local field inhomogeneities and T2.

The fundamental concept underlying the formation of a magnetic resonance image is a magnetic gradient, i.e. a spatially varying magnetic field. Lauterbur (1973) demonstrated that by superimposing a magnetic field that varies linearly across space, hydrogen atoms would precess at different frequencies in a controlled fashion. Thus different points in space become identified by different resonance frequencies. The Fourier transform of the signal would show its strength at each frequency, and thus at each position. Mansfield (1977) proposed the technique of echo planar imaging (EPI) to obtain MRI images following a single excitation using a rapid gradient switching. A series of changing magnetic field gradients and oscillating magnetic fields is referred to as the pulse sequence. Presently, magnetic resonance imaging (MRI) instruments use three mutually orthogonal sets of electromagnetic 'gradient coils' to encode the three spatial co-ordinates of the MR signal (Cohen et al., 1994). The data acquisition is achieved in two steps: First, a particular slice is

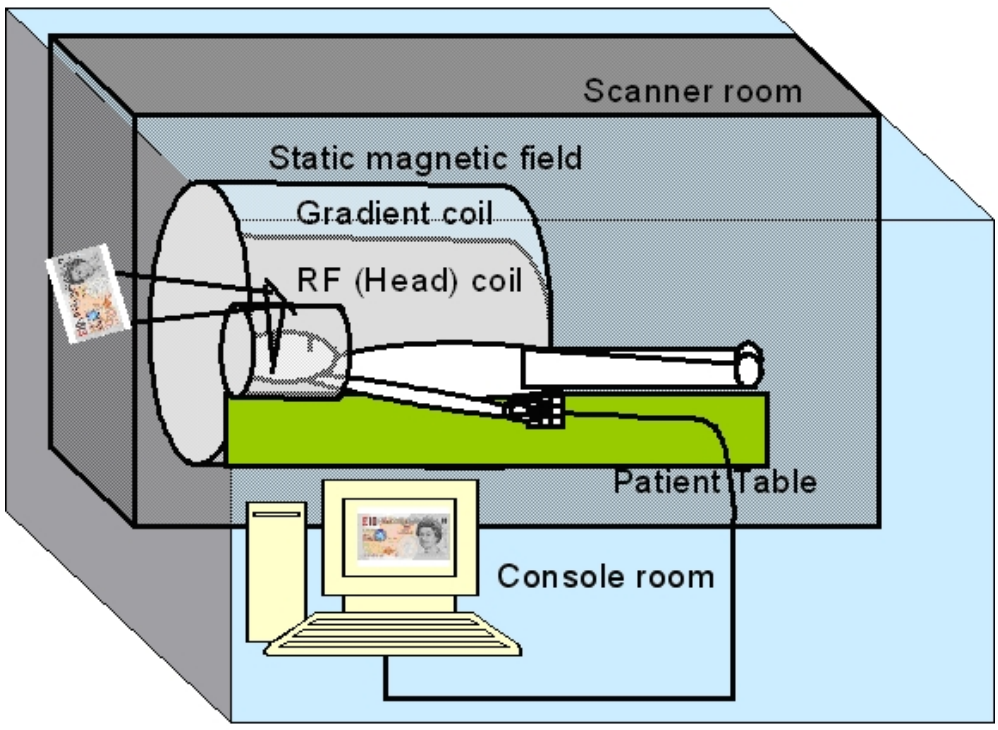

Figure 2.2 A schematic diagram of functional MRI scanning

selected within the total imaging volume using a one-dimensional excitation pulse. Then a two-dimensional encoding scheme (phase and frequency) is used to resolve the spatial distribution of the spin magnetizations. The field of view defines the spatial extent along different dimensions of the image space. Sequential excitation of adjacent slices may lead to off-resonance excitation (i.e. excitation of spins to intermediate state) that results in each slice being pre-excited by the previous excitation pulse. To overcome these effects, interleaved slice acquisition can be used.

There are two important factors that govern the time at which MR images are collected: (a) The time interval between successive excitation pulses, known as the repetition time, or TR and (b) The time interval between excitation and data acquisition, known as echo time or TE. The most commonly used contrast for structural anatomical images is T1-weighted. A number of methods exist for contrast generation in MRI images. In the following, a discussion of contrast mechanisms based on relaxation times is briefly described. To generate T1 contrast images, an intermediate TR and short TE is recommended. At short and long TRs, there is either little time for the longitudinal magnetization to recover or would recover completely. This would result in loss of contrast between tissues. Further, TE should be much less than T2 to have exclusively T1 contrast. Similarly to generate T2 contrast images, intermediate TE is recommended to observe differences in transverse magnetization between tissues and long TR to eliminate T1 effects. MR signal changes that are measured at data acquisition can be generated by using the gradients only (Gradient Echo sequence) or by a second 180º electromagnetic pulse, called a refocusing pulse (Spin Echo sequence). The refocusing pulse corrects for phase dispersion due to T2 effects, so that all spins are approximately in phase during the data acquisition period. T2 weighted images can only be generated using spin echo sequences, while T1 weighted images can be generated by any of the gradient or spin echo sequences. Spin echo sequences provide true spin-spin relaxation that does not depend on the field inhomogeneity (e.g. T2* effects) using the 180º refocusing pulse. Hence spin-echo sequences can be used to avoid the susceptibility artefacts, which are caused by magnetic field inhomogeneities near air-tissue interfaces, usually observed as signal losses or

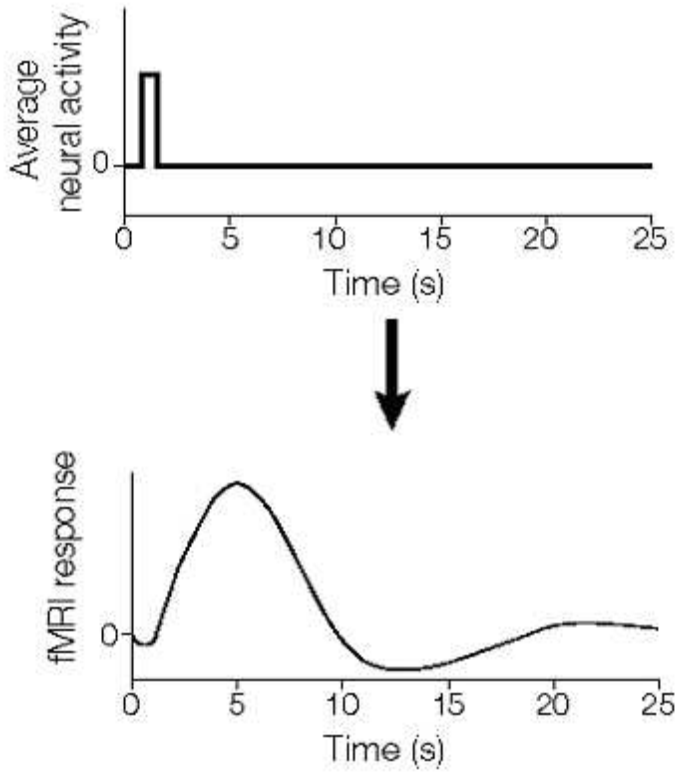

Figure 2.3 The haemodynamic response function for fMRI data

(reproduced from Heeger and Ress, 2002)

dropouts in the orbitofrontal region and temporal lobes of the brain. The T2* contrast forms the basis of BOLD fMRI. T2* contrast requires long TR and medium TE and the MR signal needs to be generated using the magnetic field gradients rather than using the refocusing pulse that would eliminate field inhomogeneity effects. Due to the reduced T2* sensitivity, spin-echo sequences are less frequently used for BOLD fMRI.

The measured RF signal decays over time depending on many factors including the presence of inhomogeneities in the magnetic field. Greater inhomogeneity results in decreased image intensity.

The increase in neuronal activity in a brain area results in an initial increase in oxygen consumption. After a delay of about 2 sec, a large increase in localized cerebral blood flow is triggered, which over-compensates the oxygen consumption. Therefore, localized increases in blood flow increase blood oxygenation and consequently reduce deoxyhaemoglobin. As a result, better

visibility in MRI images is thought to correlate with neuronal activity. Simultaneous fMRI and electrophysiological recordings by Logothetis and colleagues (Logothetis et. al. 2001) have confirmed that the BOLD contrast mechanism directly reflects the neural responses elicited by a stimulus. However, fMRI activation in an area is correlated with the local field potentials reflecting processing of the incoming input rather than the spiking activity. Hence, the absence of an FMRI signal does not necessarily mean that no information processing is taking place in a particular brain area. After fifteen years of fMRI studies, there is still much to learn about the source of these signals (see Heeger and Ress, 2002 for review).

The fMRI provides a non-invasive method to access indirectly neuronal activity in the brain with a relatively good spatial and temporal resolution. Before the emergence of functional MRI, radio isotope based techniques such as Positron Emission Tomography (PET) which measures regional cerebral blood flow (rCBF), were widely used for mapping the brain function. However, these techniques are invasive and have a low spatial and temporal resolution.

Although animal studies provide an unprecedented approach to study neural mechanisms at cellular level, the limited communication and cognitive capabilities restricts the investigation of brain function in animals. Electrophysiological methods due to their invasive nature (i.e. require insertion of electrodes directly into the brain) have limited use for studying brain function in humans. Electroencephalography (EEG) measures of the electrical activity of the brain by recording on a millisecond time scale from electrodes placed on the scalp. The magnetoencephalography (MEG) and EEG techniques signals derive from the net effect of ionic currents flowing in the dendrites of neurons during synaptic transmission. While EEG has poor spatial resolution, MEG technique promises good spatial and temporal resolution. The inverse problem of uniquely identifying the locations of neural sources giving rise to pattern of activity on the skull has by and large limited the value of EEG and MEG in mapping brain function.

Lesion studies provide clear evidence that a brain region is necessary for a particular behaviour but do not specify the time course of the region's activity. Lesion studies result in a permanent loss of a brain region, thus lending itself to

be an irreversible process. Hence, human lesion studies can only be done by finding patients with isolated damage to a particular brain area. The temporary interruption of function within a brain region is possible using transcranial magnetic stimulation (TMS). Due to several considerations as outlined above, the functional MRI technique offers a suitable method for investigating human brain function.

2.2 Issues related to Experimental Design

Developing successful fMRI experiments requires careful attention to experimental design, data acquisition techniques, and data analysis (Chein and Schneider, 2003). Experimental design is at the heart of any cognitive neuroscience investigation.

As fMRI does not measure absolute neural activity, neuroimaging studies must be designed to quantify relative changes of activity. Further, the brain is constantly engaged in several controlling tasks such as respiration, heart-beat etc. Hence, to measure specific task-related activity, we need to scan subjects while at rest or while performing a simple baseline task (Gusnard and Raichle, 2001). Assuming that brain activity scales in a linear fashion and that cognitive processes are additive, we can test for brain activations pertaining to certain cognitive processes (Berns, 1999). Although there is no inherent baseline associated with the blood oxygen-level-dependent (BOLD) signal (Gusnard and Raichle, 2001) that is measured in traditional functional MRI (fMRI) studies, researchers often have attempted to establish such a baseline by using periods of rest. Rest periods may be 10- to 30-s long blocks of rest or fixation (blocked fMRI), the final seconds of long intertribal intervals (ITIs; in the case of slow, or non-overlapping, event related fMRI), or 2- to 4-s null trials (in the case of rapid event-related fMRI). Because no task is being performed during rest, it has seemed reasonable to assume that this baseline represents something akin to a zero-activity condition that then can be compared with activity during cognitive tasks. Therefore, when activity in a particular region of the brain during a cognitive task is no greater than during rest, it often has been supposed that this particular region of the brain is not involved in the task. However, periods of rest have often been associated with significant cognitive activity (Stark and

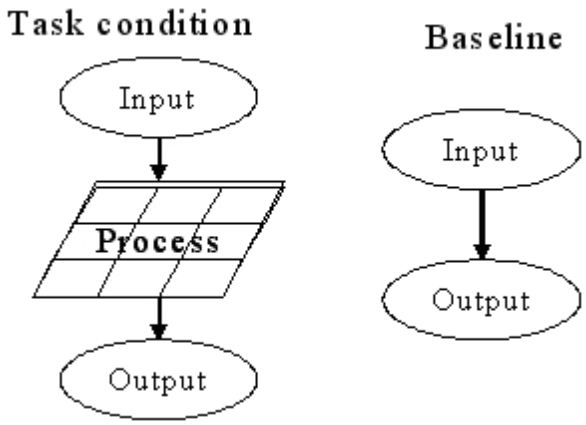

Figure 2.4 The principle of subtractive design

Squire, 2001), suggesting the crucial role of baseline tasks in design and interpretation of fMRI studies.

Overall, designs can be classified into three types i.e., categorical, factorial or parametric (Friston, 1997). The categorical designs assume that the cognitive processes can be dissected into sub-cognitive processes. That is one can remove and add different cognitive processes by the assumption of pure insertion. In other words, pure insertion requires that one cognitive component does not affect the effect of another cognitive component. The categorical designs are further divided into subtraction type or conjunction type. Cognitive subtraction designs are used to test the hypothesis pertaining to activation in one task as compared to that in another task considering the fact that the neural structures supporting cognitive and behavioural processes combine in a simple additive manner. Whereas in the cognitive conjunctions type designs, several hypotheses are tested, asking whether all the activations in a series of task pairs are jointly significant. Cognitive conjunctions can be thought as an extension of the subtraction technique in the sense that they combine a series of subtractions. While cognitive subtraction studies are designed such that a pair of tasks differ only by the processing components of interest, cognitive conjunction studies are designed such that two or more distinct task pairs each share a common processing difference. The problem of finding a baseline that activate all cognitive processes except the process of interest can be overcome by conjunction design (Price et al., 1997). The only constraint on selecting the

baseline is that the component of interest is the only process that differs in each task pair (Price and Friston, 1997).

Factorial designs involve combining two or more factors within a task and looking at the effect of one factor on the response to other factor. The problem of interactions (i.e., the effect that the added component in the activation task has on pre-existing components) can be overcome when the experimental design is factorial. Price et al. (1997) demonstrated that when the design is factorial, conjunction analysis reveals commonalities in activation, while the interactions reveal task-specific effects. In particular, the effect of a cognitive component (i.e., an effect that is independent of other components) is best captured by the main (activation) effect of that component and that the integration among components (i.e., the expression of one cognitive process in the context of another) can be assessed with the interaction terms (Friston et al., 1996).

In parametric designs, rather than assuming that the cognitive processes are composed of different cognitive components, they are considered as belonging to different psychological dimensions. The systematic changes in the brain responses according to some performance attributes of task can be investigated in parametric designs. In parametric designs one can also look at the linear and non-linear types of relations to be determined empirically.

The experimental design can be either a within-group or a between-group design. Due to the difficulty of matching all parameters (including age, IQ, gender etc) between groups, within-group designs are generally preferred, except when comparing special populations such as patients with a control group.

2.2.1 Block and Event-related designs

An fMRI experiment to test a given biological hypothesis must be designed within the constraints of the temporal characteristics of the BOLD fMRI signal and of the various confounding effects to which fMRI signal is susceptible. Typically, two designs are possible 1) Epoch-based design using Blocks of stimulation (boxcar designs with alternating activation and rest) and 2) Event-

related design, where data may be recorded to monitor the BOLD response following a marked (pre-determined) event such as a single stimulus or task.

Blocked design (Epoch-based) experiments (Bandettini, 1994) are used mainly to average across many trials to obtain sufficient signal-to-noise ratios to generate functional activation images. The block design experiments descended from the low temporal resolution imaging based on blood dynamics (such as PET). However, such blocked trial procedures do not allow separate trials within the task blocks to be distinguished. Blocked-designs cannot be used if we want to consider trials that depend on subject's performance (e.g. correct or wrong; chooses among different alternatives) or need to present trials in a non-blocked fashion (e.g. the oddball paradigm). Dale and Buckner (1997) demonstrated the feasibility of using fMRI for selective averaging of rapidly presented individual trials, a technique that was used in event related potential (ERP) studies such as the EEG/MEG. It is shown that the haemodynamic response is delayed and lasts for several seconds even for brief stimulation (less than couple of seconds) (see Figure 2.3). As the haemodynamic response to individual trials extends temporally, the responses to successive trials may overlap. Hence the inter-trial interval between successive trials needs to last for about 15 seconds. However this severely limits the number of trials, which can be averaged per unit time, thus limiting the achievable signal-to-noise ratio. Dale and Buckner (1997) demonstrated that the haemodynamic response to successive events adds in an approximately linear fashion even at relatively

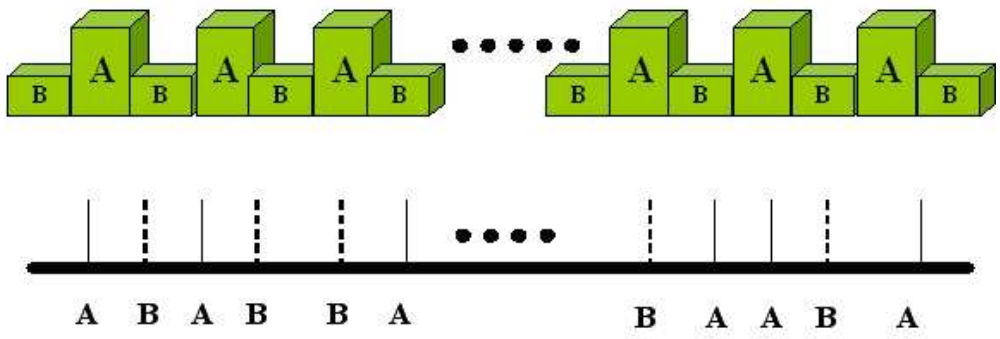

Figure 2.5 Block and Event-related designs

Block (top panel) and Event-related (bottom panel) fMRI design for two task conditions A and B. The x-axis is time from left to right.

short inter-trial intervals (2 sec and 5 sec) and hence selective averaging of rapidly presented individual trials is feasible. The findings of Dale and Buckner (1997) support the Linear Time Invariant model for the haemodynamic response function (Boynton et al., 1996). Dale (1999) has shown that the statistical efficiency of rapid event-related designs when the inter-trial interval is appropriately jittered can be up to 10 times greater than fixed inter-trial interval designs. Further, random intermixing of trial types eliminates strategy effects that might otherwise confound the results in blocked task paradigms.

In conducting a hypothesis-based experiment, we wish to be able to attribute any observed effects to experimentally manipulated conditions. This can be guaranteed only if conditions are randomly allocated to a presentation order for each subject in a sensible manner. Further, this randomisation should be appropriately balanced, both across and within subjects. With such random allocation of conditions, any unexpected effects are randomly scattered between the conditions, and therefore do not affect the designed effects.

2.3 Analysis of functional MRI Data

The main issue in analysing functional MRI images is comparing images, or groups of images, in a statistically meaningful way. In a typical fMRI experiment, a whole-brain functional image is acquired every 2-3 seconds resulting in a few hundred images to be analysed. Each image is acquired as a number of slices (e.g. 21 with thickness ~ 5 mm) with a typical in-plane resolution of 3x3 mm for a field of view of 192x192 mm. With these typical parameters, a single fMRI image would have dimensions of 64x64x21 mm. Statistical Parametric Mapping (SPM) is a form of data reduction, condensing information (in a statistically meaningful way) from a number of individual scans into a single image volume that can be more easily viewed and interpreted. Usually a univariate approach is followed in which the parametric map is computed by examining every voxel location across all images. In order to select a particular statistical distribution models (e.g. Poisson, normal, Gaussian), we need to know the underlying distribution of variance of the data being analysed, which is usually unknown in neuroimaging data. Further, univariate statistical models generally assume independent data points. Several

preprocessing steps are required before proceeding with statistical analysis in order to reduce artefacts and noise and to perform spatial transformations. The analysis of fMRI data within the framework of SPM2 software (http://www.fil.ion.ucl.ac.uk/spm Wellcome Department of Imaging Neuroscience, London) is presented here.

2.3.1 Preprocessing

Spatial transformations are important in many aspects of functional image analysis and involve both within- and between-subject registration followed by spatial smoothing with a Gaussian kernel. Preprocessing includes several steps, all of which are aimed at massaging the data so that it is suitable to be statistically analysed. The first several steps put each image volume into a standardized spatial reference frame. The last preprocessing step applies a Gaussian spatial filter. Few scans at the beginning of each session are discarded to account for transients in magnetic field of scanner. The origin of the images is set to match the line joining anterior-commissure to the posterior-commissure (AC-PC line).

Within-subject registration

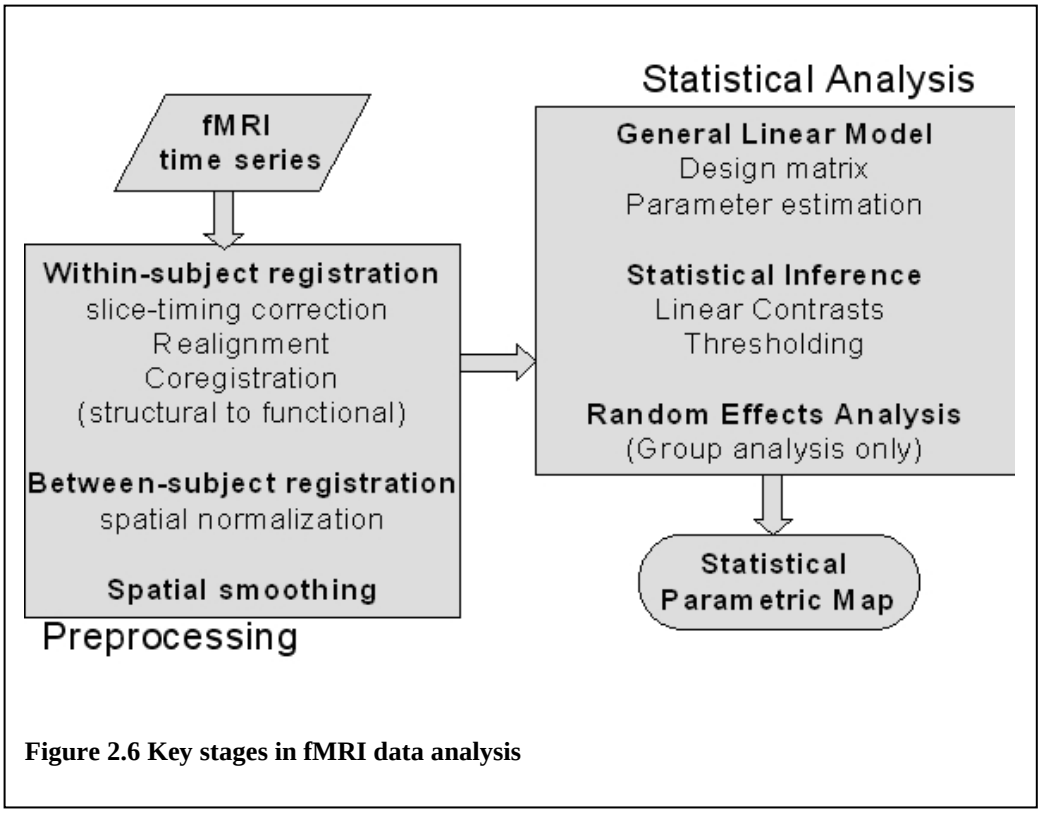

Figure 2.6 Key stages in fMRI data analysis

The 3-dimensional functional brain images are usually acquired as a number of slices in 2-dimensions. Hence, there will be a time difference approximately equal to the TR (repetition time or inter scan interval) between the first slice and the last slice acquired in a single whole-brain acquisition. One option to compensate for the time difference between bottom and top slices of the brain is to acquire the slices in an interleaved fashion. Hence all odd numbered slices are acquired first followed by even numbered slices. During preprocessing stage, it is desirable to temporally interpolate the slices so that it would be equivalent to acquiring the whole brain image at a single time point. This is usually done with respect to a reference slice (e.g. middle slice or bottom slice of the brain), which depends on the regions of particular interest for a given experiment. This procedure is referred to as slice timing correction.

In functional imaging, the signal changes due to any haemodynamic response can be small compared to signal changes that can result from subject motion. So, prior to performing the statistical tests, it is important that the images are as closely aligned as possible. Although the subjects are asked to keep their head's still, movement does occur. The realignment algorithm follows a rigid-body registration procedure (Friston et al., 1995a). A rigid body can have a linear translational movement or a rotational movement in each of the three directions (X, Y and Z). Correspondingly, there are six parameters that need to be estimated (X, Y, Z translations, pitch, roll and yaw). For multi-session data, realignment works in two steps. First, the first functional images from each session are realigned to each other taking the first session as reference. Second, the remaining images within each session are realigned to the first image. As a consequence, all images are realigned to the first image from the first session.

When applying slice-timing correction and realignment, the order of these two preprocessing steps needs special consideration. Applying realignment procedure first would account for large movements, but the images will no longer correspond to the specific time that the slice was supposed to have been acquired after being realigned. On the other hand, slice-timing correction essentially interpolates the data temporally and the realignment procedure would need to work on resliced images after the slice timing correction has been applied. The disadvantage of reslicing the data several times

during the preprocessing stage would incur loss in the image quality. The movement-related activation can be substantially large compared to the task-related BOLD changes. Hence, often the realignment parameters are included as covariates of no interest in the statistical analysis stage. An additional way to account for differences in timing of haemodynamic responses for different brain regions would be to include the temporal derivatives of the canonical HRF as part of the basis functions during statistical analysis.

Between-subject registration

Sometimes, it is desirable to warp images from a number of individuals into roughly the same standard space to allow signal averaging across subjects. A further advantage of using spatially normalized images is that activation sites can be reported according to their coordinates within a standard space such as the one described by Talairach and Tournoux (1988). SPM2 uses the average brain template created by the Montreal Neurological Institute, that is an average of 152 brain images and hence more representative of the population as compared to the Talairach and Tournoux atlas. The Normalization process (Friston et al., 1995a) not only considers the rigid-body transformations but also considers shears and zooms to match the individual subject's images to the template. For accurate normalization, it would also be required to use nonlinear transformations that would account for deformations that do not vary in a linear fashion. SPM2 uses cosine basis functions as part of nonlinear transformations for normalization procedure.

The normalization procedure can be performed in two different ways. First, the mean functional image from the output of realignment procedure can be used to match with the EPI template image in the MNI space and then the resultant parameters can be applied to all the functional images to be normalized. The spatial resolution of functional images is poor compared to the high-resolution structural images that give detailed anatomy of the subjects' brain. Hence, it is desirable to use the information provided by the structural images for better match with template brain. However, the functional and structural images are usually acquired using different imaging parameters and slice orientations. So prior to using a structural image to compute normalization parameters to a T1 template image, the structural images need to be co-registered with the functional images of the subject. This step forms part of the within-subject registration. Apart from a more precise spatial normalization, a further use of this registration is that the activations of the subject in the functional space can be overlaid onto the structural image of the subject for better visualization and localization of the activation.

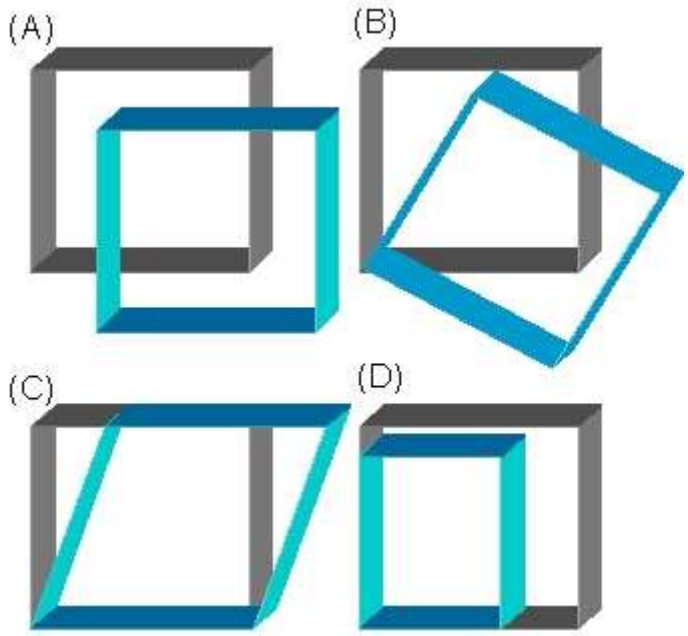

Figure 2.7 Rigid body and affine transformations

(A) Translation (B) Rotation (C) Shear (D) Zoom (adapted from Rik Henson, SPM Mini course, 2006, MRC-CBU)

As the normalization procedure invariably tries to warp the subject's brain into a template space, it can be problematic when there is a dropout in some regions or when the subjects have a lesion. Brett and Rorden (Brett et al., 2001) have suggested that such regions be masked prior to applying the normalization procedure to avoid the algorithm trying to fill in the lesion / dropout region with the surrounding tissue. This approach can however compromise the computation of parameter estimates during statistical analysis, particularly at group level, as data would not be available for regions excluded only in some subjects.

Spatial Smoothing

The matching of the brains in the Normalization step is only possible on a coarse scale, since there is not necessarily a one-to-one mapping of the cortical structures between different brains. Because of this, images are smoothed prior to the statistical analysis in a multi-subject study, so that corresponding sites of activation from the different brains are superimposed. Smoothing generally increases the signal relative to noise. From the matched filter theorem, to get optimum resolution of signal from noise, we need a filter that is matched to the signal. Since, haemodynamic responses are modelled to have a Gaussian shape; we need to use a Gaussian kernel of size at least twice the voxel size (FWHM of about 6 or 8 mm) for smoothing the functional images. The idea of smoothing is to replace the intensity value within each voxel with a weighted average (as determined by a Gaussian kernel centred on that particular voxel) that incorporates the intensity values of the neighbouring voxels. Smoothing is performed to compensate for residual between-subject variability after normalization. Smoothing also permits the application of Gaussian random field theory at the statistics inference stage.

2.3.2 Statistical analysis of fMRI images

Model setup and parameter estimation

After preprocessing, the images are ready for statistical analysis. FMRI data are high-pass filtered to remove physiological effects such as heartbeat, respiration, scanner-drift etc. Statistical analysis corresponds to Statistical Parametric

Mapping (Friston et al., 1995b) using the General Linear Model and theory of Gaussian fields. The GLM is used to specify the conditions in the form of a design matrix, which defines the experimental design and the nature of hypothesis testing to be implemented. The hypothesis is framed as a design matrix model. The design matrix has one row for each scan and one column for each effect one has built into the experiment or explanatory variables that may confound the results. The columns of the design matrix correspond to experimental conditions of interest (the hypothesis under test) and a set of columns that model effects of no interest. This is the stage where the groups designated for the images (e.g. reward/no reward) are specified. This stage corresponds to modelling the data in order to partition observed neurophysiological responses into components of interest, confounds, or components of no interest and an error term.

The general linear model (GLM) is an equation, which expresses the observed response variable in terms of a linear combination of explanatory variables plus a well-behaved error term. Commonly used parametric models, such as linear regression, t-tests and analysis of variance (ANOVA) are special cases of the general linear model. The GLM relates what one observes, to what one expected to see, by expressing the observations (response variable Y) as a linear combination of expected components (or explanatory variables x) and some residual error (ε), thereby equivalent to linear regression.

$$\begin{bmatrix} Y_1 \\ \mathrm{M} \\ Y_j \\ \mathrm{M} \\ Y_J \end{bmatrix} = \begin{bmatrix} x_{11} & \Lambda & x_{1l} & \Lambda & x_{1L} \\ \mathrm{M} & \mathrm{O} & \mathrm{M} & \mathrm{O} & \mathrm{M} \\ x_{j1} & \Lambda & x_{jl} & \Lambda & x_{jL} \\ \mathrm{M} & \mathrm{O} & \mathrm{M} & \mathrm{O} & \mathrm{M} \\ x_{J1} & \Lambda & x_{Jl} & \Lambda & x_{JL} \end{bmatrix} \begin{bmatrix} \beta_1 \\ \mathrm{M} \\ \beta_l \\ \mathrm{M} \\ \beta_L \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \mathrm{M} \\ \varepsilon_l \\ \mathrm{M} \\ \varepsilon_L \end{bmatrix}$$

This can be expressed in the matrix form as

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

Here, X is called the design matrix that contains the explanatory variables and β is the unknown parameter to be estimated. The ordinary least squares approach to calculate parameter estimates β would be

$$\beta^* = (X^T X)^{-1} X^T \mathbf{y}$$

The fitted response would be $Y = X \beta^*$ and the residual is y-Y. The assumption underlying least squares approximation is that the residuals are drawn from independent and identically distributed normal (Gaussian) distribution (white noise). This assumption is violated by the fMRI data, which are typically correlated from one scan to the next. Hence the effective degrees of freedom (df) cannot be assumed to be number of scans minus the dfs used in the model. SPM2 uses the restricted maximum likelihood (ReML) approach to estimate the non-sphericity (of which autocorrelation is one type) in fMRI data. Additional approach to deal with autocorrelation in fMRI data is to explicitly model using, for example, a first order auto regression model AR(1).

MRI gives us the blood flow signal, but we are interested in the neural activity. It is possible that the neural response is quicker and the changes in blood flow take place a little later. To account for these and to find the neural activity from the MRI signal, the columns of the design matrix are convolved with the canonical haemodynamic response function (HRF). The temporal and dispersion derivatives of the HRF are used additionally to account for variation in onset and width, respectively, of the HRF across different brain regions. An alternative approach is to use a basis functions that do not make any assumption about the shape of the haemodynamic response (e.g. using a finite impulse response model). Henson et al. (2001) have demonstrated that using the canonical HRF and its temporal and dispersion derivatives is sufficient for reliable detection of activation in event-related fMRI.

Statistical Inference

Brain activity specific to task is obtained by specifying linear contrasts. A contrast can be used to compare different conditions. The subtractive approach assumes that brain activity scales in a linear fashion. The conditions of interest are given a positive value, such as 1, and conditions that are to be subtracted from these conditions of interest take on a negative value, such as -1. The end result is a statistical parametric map. The activations thus obtained can be overlaid or rendered onto the high-resolution anatomical image of the subject in order to accurately locate the neural activity.

Statistical parametric mapping approach is a univariate approach. That is each voxel is analysed separately. Hence for a statistical threshold of $p<0.05$, 5% of the voxels would show activation by chance alone (false activation – type I error). This means a correction for multiple comparisons is needed. The traditional way of doing this is to use some version of a Bonferroni correction. However, due to large number of voxels involved, a straightforward implementation would severely reduce the estimated number of degrees of freedom. The individual voxels in most neuroimaging modalities (PET, fMRI, EEG, MEG etc.) are heavily correlated with neighbouring voxels. Hence, to the extent that the image data approximate a random Gaussian field (Worsley e al., 1996), correction for multiple comparison need to be only made for number of voxels that can be resolved independently (resolution elements or resels). The correction for multiple comparisons is controlled for family-wise error (FWE) rate. This assumption of random Gaussian field is assured by applying a Gaussian smoothing filter in the pre-processing stages.

A serious limitation of correcting for multiple comparisons is that the number of false negatives (type II error) is increased. Another approach is to determine the false discovery rate (FDR) that controls for 5% at ($p<0.05$) of observed activations can be false positives (Genovese et al., 2002). The FWE approach controls for a 5% chance of a single false positive. As a trade-off to correction for multiple comparison, alternative approaches have been described such as (i) using a strict uncorrected threshold (e.g. $p<0.001$), (ii) using an inference over the cluster size, so that it is unlikely to find activations in a cluster of size, say 30 voxels. (iii) small volume corrections in regions where a prior hypothesis exists (iv) a region of interest (ROI) analysis in which the average signal for all voxels in an anatomical or functional ROI is used, hence reducing the number of multiple comparisons voxel space to the number of ROIs.

Random Effects Analysis

In order to make an inference about brain activity in a task, the contrast images from a group of subjects are analysed using a random effects model (Holmes and Friston, 1998) using student's t-test or ANOVA like methods. The contrast

images represent spatially distributed images of the weighted sum of the parameter estimates for a particular contrast. In essence, it's like a difference image for (activation-rest) or (reward-no reward). When using a one-sample t-test, one contrast image for each subject is required. By doing that, the images are being collapsed over intra-subject variability (to only one image per contrast per subject) and the image-to-image residual variability is now between subject variance alone. When using ANOVA, a number of contrast images are entered from each subject. These need to be corrected for non-sphericity. If the contrast images being entered into ANOVA are main effects of a condition, a within-subjects model should be used. On the other hand, if the contrast images have already accounted for within-subject variability, then an ANOVA without constant term can be used.

The purpose of the Random Effects analysis is to find the areas that are activated in much the same way in all subjects, as opposed to a fixed effects model, which gives areas that are activated on the average across the subjects. This is really a crucial difference since a fixed effects analysis may yield significant results when one or a couple of subjects activate a lot even though the other subjects do not activate at all. The Random Effects analysis incorporates both within-subject variance, as well as between-subject variance. This allows generalization of the results to the population from which the subjects were drawn.

Bibliography

Bandettini, P.A. (1994). Magnetic resonance imaging of human brain activation using endogenous susceptibility contrast, PhD Thesis, Medical College of Wisconsin.

Bandettini, P., Wong, E., Jesmanowicz, A., Hinks, R., and Hyde, J. (1994). Spin-echo and gradient-echo EPI of human brain activation using BOLD contrast: a comparative study at 1.5 T. NMR Biomed, 7:12–20.

Berns, G. (1999). Functional neuroimaging. Life Sci., 65:2531– 2540.

Boynton, G., Engel, S., Glover, G., and Heeger, D. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. J. Neurosci., 16:4207–4221.

Brett, M., Leff, A., Rorden, C., and Ashburner, J. (2001). Spatial normalization of brain images with focal lesions using cost function masking. Neuroimage, 14:486–500.

Buxton, R., Wong, E., and Frank, L. (1998). Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magn Reson Med, 39:855–864.

Chein, J. M. and Schneider, W. (2003). Designing effective fMRI experiments. In Grafman, J. and Robertson, I., editors, Handbook of Neuropsychology. Elsevier Science B.V., Amsterdam.

Cohen, M. and Bookheimer, S. (1994). Localization of brain function using magnetic resonance imaging. Trends Neurosci., 17:268–277.

Constable, R. (1995). Functional MR imaging using gradientecho echo-planar imaging in the presence of large static field inhomogeneities. J Magn Reson Imaging, 5:746–752.

Cusack, R., Russell, B., Cox, S., De Panfilis, C., Schwarzbauer, C., and Ansorge, R. (2005). An evaluation of the use of passive shimming to improve frontal sensitivity in fMRI. Neuroimage, 24:82–91.

Dale, A. (1999). Optimal experimental design for event-related fMRI. Hum Brain Mapp, 8:109–114.

Deichmann, R., Gottfried, J., Hutton, C., and Turner, R. (2003). Optimized EPI for fMRI studies of the orbitofrontal cortex. Neuroimage, 19:430–441.

Friston, K. (1997). Testing for anatomically specified regional effects. Human Brain Mapping, 5:133–136.

Friston, K. J., Ashburner, J., Frith, C. D., Poline, J. B., Heather, J. D., and Frackowiak, R. S. J. (1995a). Spatial registration and normalisation of images. Human Brain Mapping, 2:165–189.

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. B., Frith, C. D., and Frackowiak, R. S. J. (1995b). Statistical parametric maps in functional imaging: A general linear approach. Human Brain Mapping, 2:189–210.

Friston, K., Price, C., Fletcher, P., Moore, C., Frackowiak, R., and Dolan, R. (1996). The trouble with cognitive subtraction. Neuroimage, 4:97–104.

Gusnard, D., Raichle, M., and Raichle, M. (2001). Searching for a baseline: functional imaging and the resting human brain. Nat. Rev. Neurosci., 2:685–694.

Heeger, D. and Ress, D. (2002). What does fMRI tell us about neuronal activity? Nat. Rev. Neurosci., 3:142–151.

Henson, R., Rugg, M., and Friston, K. (2001). The choice of basis functions in event-related fMRI. NeuroImage, 13(6):127. Supplement 1.

Holmes, A. and Friston, K. (1998). Generalisability, random effects and population inference. In NeuroImage, volume 7, page S754.

Hulvershorn, J., Bloy, L., Gualtieri, E., Leigh, J., and Elliott, M. (2005). Spatial sensitivity and temporal response of spin echo and gradient echo bold contrast at 3 T using peak hemodynamic activation time. Neuroimage, 24:216–223.

Lancaster, J., Woldorff, M., Parsons, L., Liotti, M., Freitas, C., Rainey, L., Kochunov, P., Nickerson, D., Mikiten, S., and Fox, P. (2000). Automated Talairach atlas labels for functional brain mapping. Hum Brain Mapp, 10:120–131.

Lauterbur, P.C. (1973). Image formation by induced local interactions. Examples employing nuclear magnetic resonance. Nature, 242:190–191.

Logothetis, N. (2002). The neural basis of the blood-oxygenlevel- dependent functional magnetic resonance imaging signal. Philos. Trans. R. Soc. Lond., B, Biol. Sci., 357:1003–1037.

Logothetis, N., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. Nature, 412:150–157.

Maldjian, J., Laurienti, P., Kraft, R., and Burdette, J. (2003). An automated method for neuroanatomic and cytoarchitectonic atlasbased interrogation of fMRI data sets. Neuroimage, 19:1233–1239.

Mansfield, P. (1977). Multi-planar image formation using NMR spin echoes. J. Phys. C 10:L55–L58.

Nichols, T., Brett, M., Andersson, J., Wager, T., and Poline, J. (2005). Valid conjunction inference with the minimum statistic. Neuroimage, 25:653–660.

Norris, D., Zysset, S., Mildner, T., and Wiggins, C. (2002). An investigation of the value of spin-echo-based fMRI using a Stroop colorword matching task and EPI at 3 T. Neuroimage, 15:719–726.

Ogawa, S., Lee, T., Kay, A., and Tank, D. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc. Natl. Acad. Sci. U.S.A., 87:9868–9872.

Ogawa, S., Menon, R., Tank, D., Kim, S., Merkle, H., Ellermann, J., and Ugurbil, K. (1993). Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. Biophys. J., 64:803– 812.

Ogawa, S., Tank, D., Menon, R., Ellermann, J., Kim, S., Merkle, H., and Ugurbil, K. (1992). Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. Proc. Natl. Acad. Sci. U.S.A., 89:5951–5955.

Ojemann, J., Akbudak, E., Snyder, A., McKinstry, R., Raichle, M., and Conturo, T. (1997). Anatomic localization and quantitative analysis of gradient refocused echo-planar fMRI susceptibility artifacts. Neuroimage, 6:156–167.

Osterbauer, R., Wilson, J., Calvert, G., and Jezzard, P. (2006). Physical and physiological consequences of passive intra-oral shimming. Neuroimage, 29:245–253.

Parkes, L., Schwarzbach, J., Bouts, A., Deckers, R., Pullens, P., Kerskens, C., and Norris, D. (2005). Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. Magn Reson Med, 54:1465–1472.

Price, C. and Friston, K. (1997). Cognitive conjunction: a new approach to brain activation experiments. Neuroimage, 5:261– 270.

Price, C.J., Moore, C.J., and Friston, K.J. (1997). Subtractions, conjunctions, and interactions in experimental design of activation studies. Human Brain Mapping, 5(4):264– 272.

Schmidt, C., Boesiger, P., and Ishai, A. (2005). Comparison of fMRI activation as measured with gradient- and spin-echo EPI during visual perception. Neuroimage, 26:852–859.

Schwarzbauer C., Mildner T., Heinke W., Zysset S., Deichmann R., Brett M., Davis M.H. (2006). Spin-echp EPI – The method of choice for fMRI of brain regions

affected by magnetic field inhomogeneities? Human Brain Mapping, Abstract No: 1049.

Stark, C.E., Squire, L.R. (2001). When zero is not zero: the problem of ambiguous baseline conditions in fMRI. Proc. Natl. Acad. Sci. U.S.A., 98, 22:12760-6.

Talairach, J. and Tournoux, P. (1988). Co-planar Stereotaxic Atlas of the Human Brain. Thieme, New York.

Worsley, K., Marrett, S., Neelin, P., Vandal, A. C., Friston, K., and Evans, A. C. (1996). A unified statistical approach for determining significant voxels in images of cerebral activation. Human Brain Mapping, 4:58–73.