Invited Review

Design and Analysis of fMRI Studies With Neurologically Impaired Patients

Cathy J. Price, PhD,* Jenny Crinion, PhD, and Karl J. Friston, MD

Functional neuroimaging can be used to characterize two types of abnormality in patients with neurological deficits: abnormal functional segregation and abnormal functional integration. In this paper we consider the factors that influence the experimental design, analysis, and interpretation of such studies. With respect to experimental design, we emphasize that: 1) task selection is constrained to tasks the patient is able to perform correctly, and 2) the most sensitive designs entail presenting stimuli of the same type close together. In terms of data preprocessing, prior to statistical analysis, we note that structural pathology may call for constraints on nonlinear transformations, used by spatial normalization, to prevent distortion of intact tissue. This means that one may have to increase spatial smoothing to reduce the impact of inaccurate normalization. Important issues in statistical modeling concern the first level of analysis (estimation of activation within subject), which has to distinguish correct from incorrect responses. At the second level (between subjects), inference should be based on between-subjects variance. Provided that these and other constraints are met, deficits in functional segregation are indicated when activation in one or a set of regions is higher or lower in patients relative to control subjects. In contrast, deficits in functional integration are implied when the influence of one brain region on another is stronger or weaker in patients relative to control subjects.

Key Words: fMRI; patients; experimental design; statistical analysis
J. Magn. Reson. Imaging 2006;23:816–826.
© 2006 Wiley-Liss, Inc.

THIS PAPER IS CONCERNED with the experimental design and analysis techniques that are currently used in fMRI studies of patients with neurological deficits. Most fMRI studies of clinical cohorts aim to identify the brain areas in which changes in regional cerebral activation occur in response to pathology (i.e., a task by pathology interaction). The questions asked can be divided broadly into two classes involving either functional segregation or functional integration (Fig. 1). Functional segregation refers to the segregation of pro-

cesses in different brain regions. This is investigated by identifying brain regions that are activated by one task more than another. Functional integration refers to task-dependent processing that emerges from changes in the interactions among brain regions. The distinction between studies of functional segregation and integration is crucial for imaging patients because some patients suffer from abnormal functional segregation (i.e., the function of a discrete cortical area is abnormal) while others suffer from abnormal functional integration (i.e., abnormal interactions among different brain regions).

Studies of abnormal functional segregation attempt to correlate a syndrome, or particular symptoms, with changes in cortical responsiveness. For instance, studies of schizophrenia have shown that patients have abnormal responses in the frontal lobes (1–3), and that this hypofrontality correlates with the expression of psychomotor poverty (4). These results imply a regionally specific pathology that is consistent with abnormal functional segregation. However, this does not explain why, in schizophrenia, a region such as the frontal lobe might show normal activity in some contexts and abnormal activity in others. The alternative explanation is that the symptoms of schizophrenia reflect abnormal integration, i.e., abnormalities in how different brain regions interact with one another. Specifically, the frontal lobes may function normally when they interact with one set of regions but abnormally when they interact with another set. To demonstrate abnormal functional integration in schizophrenia, Friston and colleagues (5) showed that for word generation relative to word repetition, prefrontal and superior temporal activity was negatively correlated in control subjects but positively correlated in schizophrenia. This complete reversal of large-scale prefronto-temporal interactions in schizophrenics indicates abnormalities in regionally specific functional connectivity. The reversed correlations can be regarded as a task-specific failure of prefrontal cortex to suppress activity in the temporal lobes (or vice versa) (6–9).

In the following, we provide a fairly comprehensive and practical guide to some key issues that should be considered in the design and analysis of fMRI studies of patients. The paper is divided into five sections. The first deals with issues related to experimental design of patient studies, including constraints on task and stimulus selection, and issues related to presentation parameters such as the interstimulus interval and order of conditions. The second section is concerned with the

Figure 1. Functional segregation and integration. Functional segregation (left side) refers to the segregation of different brain regions according to their function. In this illustration, regions A1 and A2 are activated by task 1, while regions B1 and B2 are activated by task 2. Damage to one region (e.g., A1 indicated by dotted lines) impairs performance on task 1 but not on task 2. Functional integration refers to functions that depend on how regions interact with one another. In the illustration, region A1 is involved in two different tasks: one that depends on its interactions with region A2, and one that depends on its interaction with region A3. Damage to the connections between regions A1 and A3 will disrupt responses in A1 when the task requires regions A1 and A3, but not when the task requires regions A1 and A2.

Functional segregation

Task 1 activates A1 and A2
Task 2 activates B1 and B2

Functional integration

Task 1 activates A1 and A2
Task 2 activates A3 and A4
Task 3 activates A1 and A3
Task 4 activates A2 and A4
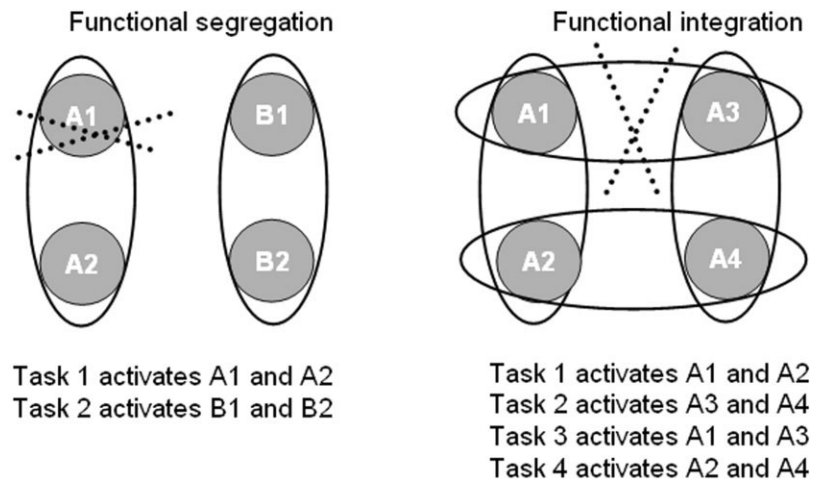
preprocessing steps that are required to correct for head motion and compare different patient populations. The third is concerned with how regional activation is estimated. The fourth and fifth sections deal with the analysis of abnormal segregation and functional integration, respectively.

EXPERIMENTAL DESIGN

Task and Stimulus Selection

In fMRI, brain regions are activated by changing stimuli or tasks. Sensory and perceptual processes are typically investigated with stimulus manipulations (e.g., view colored vs. black and white shapes), and motor or high-level cognitive processes are typically investigated with task manipulations (e.g., moving the left hand vs. the right hand, or verbal fluency vs. verbal repetition). The most informative designs manipulate two or more variables, including both the task and the stimuli, as this allows the effects of each variable to be estimated in different contexts, thereby enabling context-specific regional responses to be detected. These "factoral" designs, (10) are particularly appropriate for studying functional integration, which usually focuses on context-dependent changes in the coupling among different brain regions.

For patient studies, the desired task is usually directed toward the sensory, motor, or cognitive processes that are impaired relative to normal subjects. However, task selection is not easy to perform in patients with behavioral deficits, because if a process is impaired, how do you engage the impaired process to measure the neuronal correlates? Moreover, if a patient is presented with a task they can not perform normally, how do you control for abnormal processing demands in areas that are fully functional? fMRI studies of patients therefore need to employ tasks that the patients can perform. Since this contrasts with the design of neuropsychological studies that typically require tasks with impaired responses, we will start this section with two examples of how task difficulty affects neuronal responses in functionally preserved areas (see also Fig. 2). Our first example is concerned with how impaired perceptual analysis influences subsequent cognitive re-

sponses. The second example describes how impaired motor responses affect perceptual processing.

Our example of how perceptual difficulties affect subsequent cognitive responses comes from the developmental dyslexia literature. By definition, developmental dyslexics have poor reading accuracy. Therefore, if they are engaged in an fMRI study involving reading, they are likely to read fewer words accurately than the control subjects with whom they are compared. Fewer reading responses will result in less activation in reading areas (11–13), but how can these abnormal neuronal responses be interpreted? They could reflect brain regions that are physiologically compromised, or they could occur in fully functional brain regions that do not activate normally because of problems that originate in
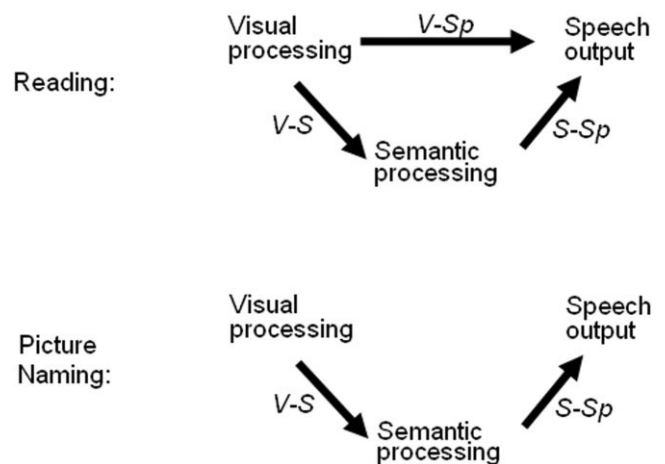
Reading:

Visual processing →(V-Sp)→ Speech output
Visual processing →(V-S)→ Semantic processing
Semantic processing →(S-Sp)→ Speech output

Picture Naming:

Visual processing →(V-S)→ Semantic processing
Semantic processing →(S-Sp)→ Speech output

Figure 2. Task selection. This figure summarizes two points. The first is that brain regions and their functions are interdependent. A deficit at any functional level will affect the functions in all other regions. It then becomes difficult to determine whether abnormal responses are due to a deficit in that region or in another region. The second point is that functional responses in any given region can be tested by selecting a range of tasks that tap similar processes to the impaired task. In the example given, one can assess visual, semantic, and speech output processing in dyslexics who cannot read by attempting to tap the same processes during picture-naming.

other brain regions. For instance, if dyslexic participants have a problem at the level of perceptual processing, normal responses in speech production areas will be prohibited. It then becomes impossible to determine whether abnormal neuronal responses are due to a primary physiological deficit or are the secondary consequence of impaired processing in an earlier region.

Our example of how motor difficulties influence perceptual responses is based on patients with speech production deficits. If such patients are engaged in an fMRI paradigm that involves speech production (e.g., auditory repetition), their neuronal activation is likely to be abnormal. However, without further investigation, it is not possible to distinguish abnormal neuronal responses that cause the speech production deficit from those that are a consequence of the speech production deficit. Specifically, if the patient is unable to repeat the words he or she hears, the level of auditory attention and perceptual analysis may be less than normal, which may lead to less activation in areas associated with auditory perception. In other words, the abnormal responses in these areas are a consequence of the task not being performed, rather than an indication that there is a deficit in the perceptual processing areas.

How can these task-performance confounds be avoided? The obvious approach is to engage subjects in tasks they can perform at the same level as the control subjects. This can be achieved by making the task relatively easy or selecting a task that does not directly engage the impaired behavior (see Fig. 2). For example, dyslexics (and their control group) could be scanned while they read familiar words ("cat," "dog," "house," "and," etc.) (14,15). Alternatively, if dyslexics are not able to read any words but otherwise have intact language skills, they could be scanned while they perform tasks that rely on processes that are similar to reading and engage the same neuronal systems (e.g., picture naming) (16). Another approach is to vary the task difficulty systematically (17). This allows the performance confounds to be modeled statistically and discounted as an explanation for regional effects and their interaction with pathology. Similarly, in group studies, when performance varies among patients, the impact of performance on activation can be tested directly in the statistical model. For example, if activation in some brain regions is higher in patients with better performance, there is a direct link between good behavior and activation (18–20). In the section entitled Second-Level Analyses of Abnormal Functional Segregation, we will return to ways of dealing with performance confounds at the level of statistical analysis.

One of the distinctions that arises in the discussion of task performance above is the difference between explicit paradigms that directly tap the process of interest (e.g., reading in dyslexics) and implicit paradigms in which the process of interest is engaged indirectly. In functional imaging paradigms, implicit paradigms need not involve any task change, but can be based solely on stimulus changes. Indeed, functional imaging can detect activation elicited by a change in stimulus even when the task and response times remain constant. For instance, face-naming may elicit emotional responses that do not affect naming-time but do change the dis-

tribution and composition of hemodynamic responses. Functional imaging can therefore detect activation related solely to changes in facial expression (21,22). In this context the effect of facial expression is not necessary for the task but is invoked by the stimulus change (i.e., it induces incidental processing). Detecting implicit and incidental processing can be particularly useful for studying patients with limited behavioral responses. For example, Morris et al (21,22) reported a functional imaging study of a patient who was blind in his right hemifield. Although the patient was not aware that fearful faces were being presented to his right hemifield, discriminatory amygdala responses were detected.

In summary, when the motivation for a functional imaging study of the patient is a better understanding of the patient's sensorimotor or cognitive deficit, it is difficult but necessary to select a paradigm the patient can perform. Task manipulations are appropriate when the patient can perform the task, and necessary if the question relates to whether the neuronal systems that are necessary to perform the task are abnormal. Stimulus manipulations are appropriate when the patient has limited abilities, and are necessary to make deductions about normal activation in the context of other cognitive deficits. For studies of abnormal integration, both stimulus and task manipulations are usually required to allow context-dependent changes in regional interactions to be assessed.

Presentation Parameters

Once the tasks and stimuli have been selected, the presentation parameters (i.e., the rate of stimulus presentation and the ordering of conditions) must be specified. As a general rule, the interstimulus interval should be kept as short as possible (given psychological constraints) and stimuli of each type should be presented together (e.g., AAAABBBB) rather than randomized or alternating (e.g., ABAABABB). This is because the hemodynamic response to each stimulus is prolonged, reaching its peak after approximately 4–5 seconds, and returning to baseline after 8–10 seconds. If subsequent stimuli of the same type arrive in close succession, the responses to each stimulus will summate and thus increase the amplitude of the signal. Indeed, the summation of responses from successive stimuli is thought to be approximately linear when the interstimulus interval (ISI) is more than a second or so (23). The summated response from a train of similar stimuli will not only be greater than that from a single stimulus, its time course will also be closer to the known shape of the hemodynamic response (see Fig. 3). This allows for greater sensitivity in the design because signal estimation is more efficient. Put simply, if you can elicit hemodynamic responses with the same time-constants as the hemodynamic response function (i.e., several seconds), you will be able to detect them more efficiently.

For single-case functional imaging studies of patients, sensitivity is particularly important. This will be maximized when the ISI is selected to be long enough to enable intact performance but short enough to capital-
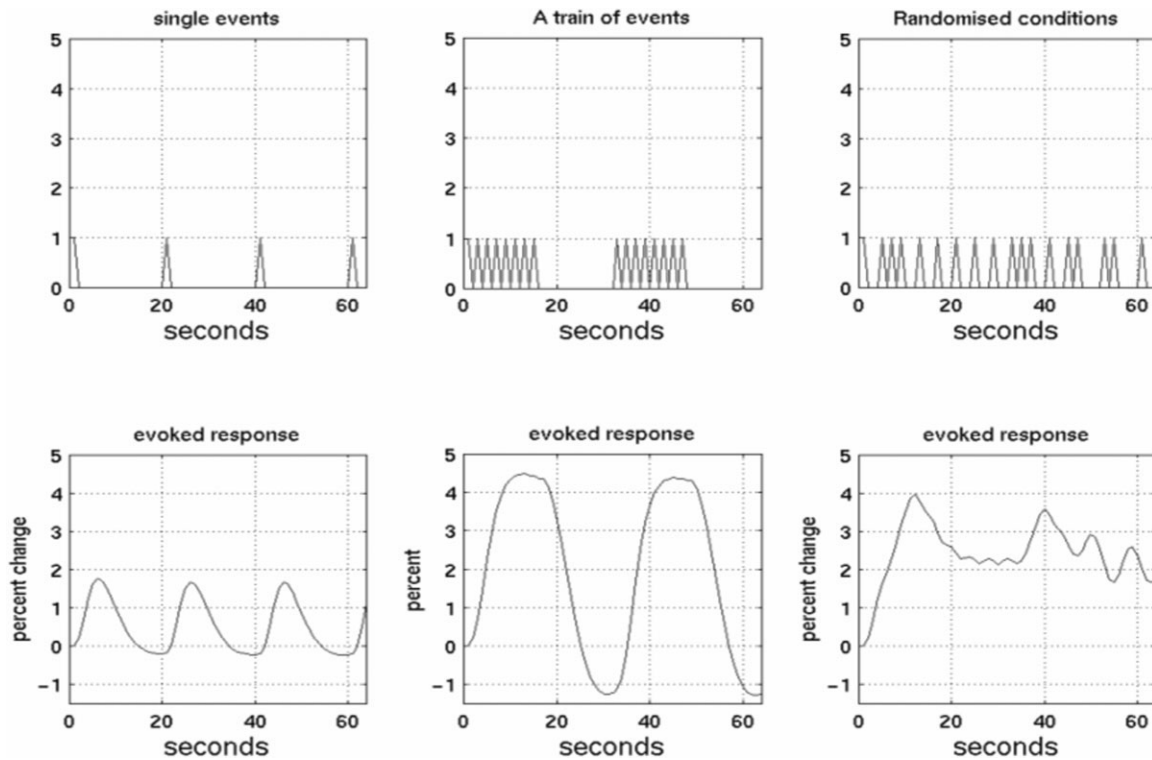
Figure 3. Stimulus presentation. The effect of stimulus timing (above) on the evoked hemodynamic response (below). When stimuli are presented in close succession, experimental variance is increased because the amplitude of the hemodynamic response increases within a block (i.e., the signal summates from trains of stimuli of the same type) and decreases between blocks (i.e., the signal has time to return to baseline). The increased experimental variance increases the likelihood of detecting small activations.

ize on the peak activation generated from the summation of responses to successive stimuli. In some paradigms, it may also be necessary to ensure that the trials are not predictable (e.g., to avoid expectations and attentional sets that may confound differences among blocks). In this case, the degree of bunching of the stimuli of one type can be varied over blocks (e.g., AABBBBBAAAAABBAAABBB).

In summary, the nature of the hemodynamic response means that the most sensitive designs for detecting activation generally involve blocked rather than randomized trials and an ISI that is long enough to optimize the number of correct responses but short enough to maximize signal.

PREPROCESSING FMRI DATA

Four preprocessing steps are used for fMRI data: realignment, coregistration, spatial normalization, and spatial smoothing. Each will be discussed in turn (see left side of Figs. 4 and 5).

Realignment

The purpose of realignment is to remove changes in signal intensity introduced by head motion during the scanning session. If the variance associated with head motion is not removed, it will be attributed to error. Therefore, realignment is important for increasing sensitivity to signals of interest. It is implemented by esti-

mating and minimizing linear transformations between the first scan and each subsequent scan. Realignment can be regarded as a simple form of spatial normalization (see below), in which subsequent scans are spatially normalized to match the first. It is simple because the transformation is a rigid-body linear transformation (i.e., the changes are the same across the whole image), whereas in spatial normalization transformations can also be nonlinear (i.e., changes in one area of the brain differ from those in other areas of the brain).

Coregistration (for Single-Subject Analyses)

For examination of single-subject responses, coregistration can be used to realign the functional images to a high-resolution structural image. This allows for a more accurate localization of the signal's origin and its anatomical designation.

Spatial Normalization (for Group Analyses)

Spatial normalization is a technique that transforms each brain image into a standard space defined by a template image. It is necessary for intersubject averaging and statistical comparisons between different groups of subjects, and when standard coordinates are required to specify the location of activations. The algorithms work by minimizing the difference between the image to be normalized and the template image. This involves both linear and nonlinear transformations.
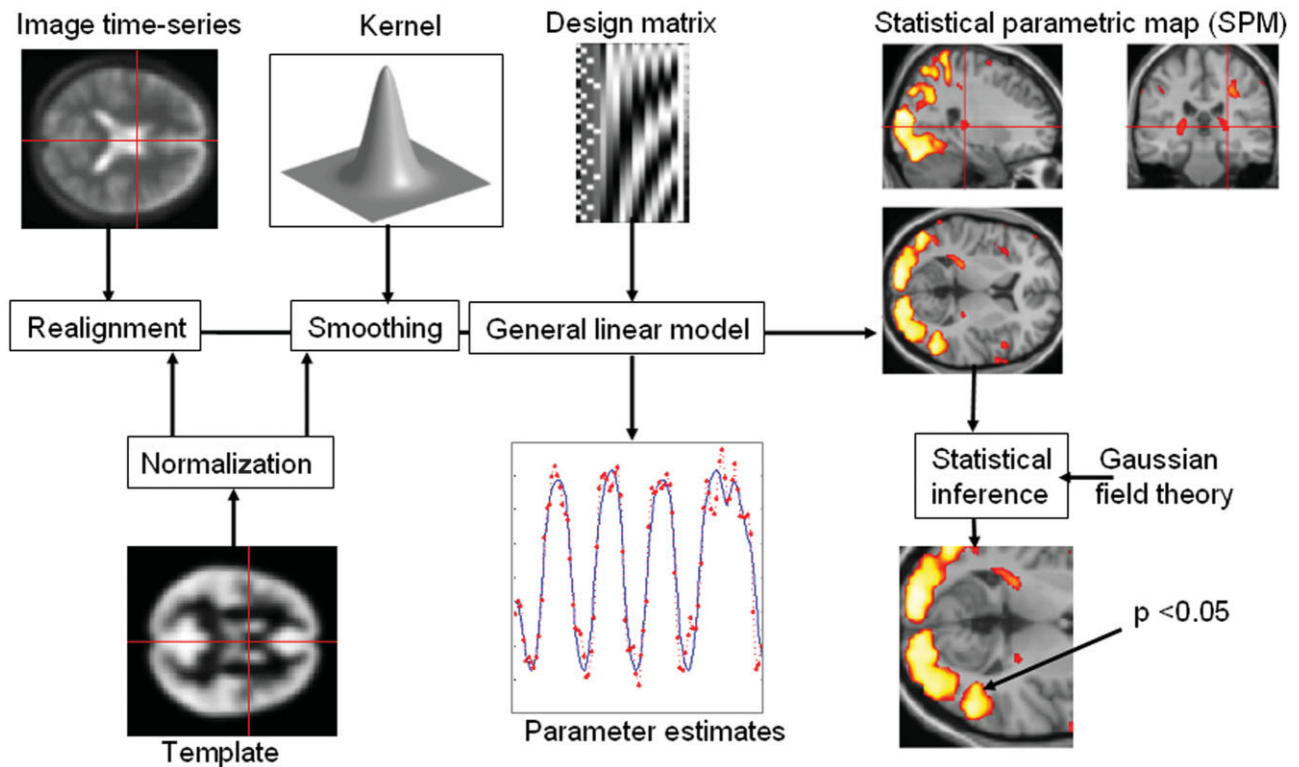
Figure 4. Data analysis stages. This schematic depicts the transformations that start with an imaging data sequence and end with an SPM. SPMs can be thought of as "X-rays" of the significance of an effect. Voxel-based analyses require the data to be in the same anatomical space. This is implemented by realigning the data (and removing movement-related signal components that persist after realignment). After realignment the images are subjected to spatial normalization so that they match a template that already conforms to a standard anatomical space. After smoothing, the GLM is employed to estimate the parameters of the model, and derive the appropriate univariate test statistic at every voxel. The test statistics that ensue (usually T or F statistics) constitute the SPM. The final stage is to make statistical inferences on the basis of the SPM and random field theory, and to characterize the responses observed using the fitted responses or parameter estimates.

The linear (affine) transformations scale the volume by the same amount (e.g., uniform shrinkage or stretching to meet the size of the template). The nonlinear transformations are concerned with local shape (e.g., shrinking one region within a plane but stretching another region), and their contribution is determined by the degree of regularization. High regularization limits nonlinear changes, whereas low regularization allows more nonlinear changes.

In the presence of a focal lesion, automated algorithms attempt to reduce image mismatch between the image to be normalized and the template to which it is being matched. As the template is typically derived from normal brains, nonlinear transformations can lead to inappropriate image distortion at the site of the lesion where the brains to be normalized have areas of signal intensity that are very different compared to the corresponding area of the template. In other words, the spatial normalization may distort the location of intact tissue to diminish the contribution of the lesion. For studies of patients with lesioned brains, one solution to aid the normalization process is to limit the nonlinear
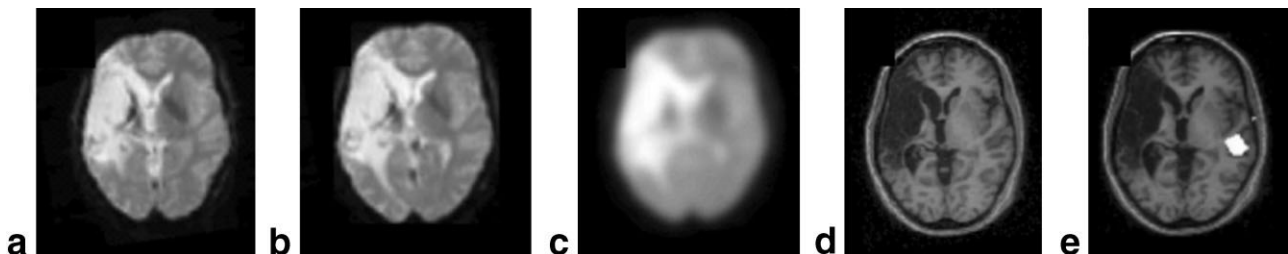

Figure 5. Spatial normalization and coregistration. Axial views of fMRI images at different stages of processing: **(a)** realigned fMRI scan from a patient with an extensive left hemisphere infarct, which appears as an area of low signal in the left frontal and temporal regions; **(b)** after spatial normalization; **(c)** after spatial normalization and smoothing; **(d)** a high-resolution structural T1 MR image from the same patient after coregistration to the functional scan; **(e)** significant activation differences (in white) displayed on the coregistered MRI scan for more accurate localization.

contributions (i.e., use high regularization). However, this negates the contribution of nonlinear deformations, which can greatly improve the quality of the normalization, both for normal brains (24), and for lesioned brains (25). Although the affine transforms usually provide an acceptable match of the brain outline, they cannot match local brain detail without disturbing alignment elsewhere. For example, nonlinear normalization can match the ventricular outline more accurately to the template compared to affine-only registration. This is particularly important for the study of stroke patients, who tend to have enlarged ventricles consequent to their age or atrophy associated with the lesion. The affine-only procedure may therefore result in a significant difference between each subject's mean image and the template to which the images are being matched.

Another solution that was introduced to reduce spatial normalization problems in patients is cost-function masking (25). Cost-function masking excludes the lesioned area from the calculation of image difference. Therefore there is no attempt to minimize image differences in the area of the lesion, and the lesion does not bias either the linear or nonlinear transformations elsewhere in the brain. Note that masking the abnormal region does not mean that areas under the mask remain untransformed; rather, a continuation of the solution for the unmasked portions of the image is applied to the area under the mask. This continuation will be constrained to be smooth by the use of the nonlinear regularization term in the normalization process. Brett et al (25) evaluated the cost-function masking technique, and their results suggest that it is superior to affine-only normalization.

In summary, detecting average activation across subjects is dependent on successful spatial registration of homologous brain areas. If scans from patients have been normalized less successfully than scans for controls, or for other patient groups, then average activation may be reduced or inaccurately located. Spatial normalization in patients with abnormal brains can be improved by reducing the impact of nonlinear transformations across the whole brain or specifically in the region of interest (ROI) using cost-function masking.

Spatial Smoothing

To increase the signal-to-noise ratio (SNR), validate the use of parametric statistical tests, and fulfill the lattice assumption of random field theory, fMRI data must be smoothed spatially. This is achieved by applying a point spread function (PSF, Gaussian kernel) with a width that is determined by the size of the voxels used in data acquisition and the size of the expected signal change. For patient studies, spatial smoothing at the second (between-subjects) level may have to be higher than normal because of increased between-subjects dispersion in the location of local structures.

FIRST-LEVEL STATISTICAL ANALYSES

Statistical analysis of patient studies usually involves estimation of activation within a patient, and then between-subjects comparisons of activation in patients and normal controls to assess the degree of abnormality (a test for task by pathology interaction). Thus there are two distinct stages to the analysis, which we will refer to as first- and second-level analyses. In this section we will discuss the first-level estimation of significant activation within a patient (see right side of Fig. 4). In the following section we will discuss the second-level comparison of patients with normal controls.

General Linear Model (GLM)

Estimation of regional activation in response to an experimental condition is usually based on the GLM, which partitions the observed neurophysiological response into a linear combination of 1) components of interest (e.g., variance of interest caused by a change in condition or performance), 2) confounds (e.g., motion artifacts or low-frequency variations in signal due to aliased biorhythms), and 3) error. Collectively these sources of variance constitute the design matrix.

Nonsphericity Correction

In parametric tests, error variance is usually assumed to be spherical (i.e., distributed identically and independently over observations). However, error variance is not always spherical. For example, in single-subject analyses, nonsphericity can arise due to serial autocorrelations (the error at any given time depends on that which occurred during a previous time step). Departures from sphericity require a nonsphericity procedure. This involves estimating the nonsphericity and 1) decorrelating the error terms so that they are independent, or 2) adjusting the degrees of freedom post hoc (c.f., the Greenhouse-Geisser correction).

Classic Inference

Deductions concerning the size of the activation relative to error variance are normally based on one of two types of statistical test. The F-statistic tests the null hypothesis that all of the effects are zero, while the T-statistic is used to test whether a particular contrast or difference between two conditions is zero. To find regionally specific effects, F and T statistics from functional imaging data are calculated in multiple voxels across the whole brain. These statistics constitute a statistical parametric map (SPM). The likelihood that voxels will be activated by chance is therefore very high. For instance, if T-statistics are generated in 10,000 voxels, then 500 voxels would be expected to be activated by chance using a conventional cutoff of $P < 0.05$. To correct for the number of false positives, the P-value has to be adjusted. However, it is not simply the case that the number of independent measurements corresponds to the number of voxels, because neighboring voxels will show similar responses (particularly after spatial smoothing). Correction for multiple comparisons therefore requires an estimation of the number of resolution elements that correspond to the number of independent measurements (RESELS). For example, in the SPM analysis package, the theory of random fields is introduced as a means of correcting for multiple com-

parisons in the context of continuous, spatially extended statistical fields. The ensuing adjustment to each voxel's P-value plays the same role as the Bonferroni correction in a family of discrete statistical tests.

Brain Inclusion Mask

In addition to correcting for multiple comparisons, the analysis also has to exclude data from areas that are not brain. This can be implemented in many ways, but requires some special consideration in the context of brain-damaged patients. For example, if a patient has extensive brain damage, the signal intensity may not match what is expected from brain tissue, and these voxels might therefore be excluded from the statistical analysis. Consequently, for statistical analysis of damaged brains, the brain inclusion mask must be specified very carefully.

SECOND-LEVEL ANALYSES OF ABNORMAL FUNCTIONAL SEGREGATION

In a single-case study, the inference regarding the significance of an effect is based on the error variance within the same subject (within-subject variance). The first-level analysis in a single subject therefore involves a fixed-effect analysis because the subject effect is treated as fixed. However, to infer abnormal regional brain responses, the patient's activation has to be compared with that in a group of normal controls. A fixed-effect analysis is not appropriate here because it does not account for how the activation varies from one subject to another (between-subjects variance). A comparison of patients with normal controls or other groups of patients therefore requires the activation from each subject (normal and patient) to be modeled independently, and the patient to be compared with normals using between-subjects variance. This is called a random effects analysis because subjects are assumed to be drawn randomly from a large population. The between-subjects variance is assumed to correspond to the variance of the general population. In other words, group studies need to assess the significance of an effect relative to between-subjects variability. In most neuroimaging analyses, random-effects models are implemented in a multistage procedure. This simply means that the effect (i.e., activation) is computed for each subject using a first-level analysis, and then images of the effects enter the second level in exactly the same way as the original images entered the first.

Patient Issues

There are five main factors that generate spurious or confounding differences between groups. First, confounds can be introduced when there are differences in the number of correct responses per subject/patient. This relates to the problem of equating task performance between the patient and normal controls (see Experimental Design section). One option might be to limit the statistical analyses to correct trials only (using an event-related analysis in fMRI). This is effected by modeling correct and incorrect trials separately and

comparing only like with like. However, to fully equate efficiency, the distribution of correct trials in the controls and the patients would have to be the same (c.f., arguments about "bunching" in Experimental Design). Therefore, it may not be possible to equate task performance post hoc, and it may be best to select stimulus paradigms that maximize the correct responses in the patient.

Second, artifactual differences between patient and control groups can be introduced by using fixed-effect analyses rather than random-effect analyses. As noted above, fixed-effect analyses are based on within-subject variance, and random-effect analyses are based on between-subjects variance. The reason why fixed-effect analyses are sometimes used for between-group comparisons stems from conventional analyses of PET data. In PET it is usually assumed that between-subjects variance is approximately equivalent to between-scans/within-subject variance. This is certainly not the case in fMRI, where within-subject variance is substantially lower than between-sessions or between-subjects variance. Random-effect analyses are now standard for between-groups comparisons in both PET and fMRI. Although fixed-effect analyses are appropriate when the inference is about some aspect of evoked responses in the group studied, they are not appropriate for making inferences about group differences. These require random-effects analyses.

Third, activation may not be detected when patients have been grouped together on the basis of similar pathology (e.g., damage to temporal lobe areas). Such analyses assume that patients do not vary in either the shape or size of their lesions, or in the degree of neuronal reorganization. Perilesional activation, for example, is unlikely to be detected when patients are grouped together. Inconsistent activations in patients may then generate positive average differences between normal and patient groups, and interpretation will be blind to patient-specific responses. This limitation can be overcome by modeling each patient as a single case study within the same statistical model. Fourth, when data from only one patient are available, there may be insufficient statistical power to detect activation in the patient. One way to improve sensitivity on the single-subject level is to increase the number of scanning sessions per subject and look for patient abnormalities that are consistent across sessions (for instance, by using conjunction analyses to look for abnormalities that replicate across sessions).

Fifth, when neurovascular coupling is disrupted, abnormalities in the blood oxygen level-dependent (BOLD) signal can arise even when neural activity is effectively normal. As shown in Fig. 6, neurovascular coupling depends on the cerebral blood flow (CBF), cerebral blood volume (CBV), and cerebral blood oxygen consumption ($CMRO^2$). An extreme example of how neurovascular coupling can be disrupted was provided by Rother and colleagues (26), who reported a patient with extra-cranial carotid stenosis that resulted in a negative BOLD signal in the affected hemisphere during performance of a simple motor task. The carotid stenosis resulted in chronic vasodilation, which provided adequate tissue perfusion but reduced the reactivity of the
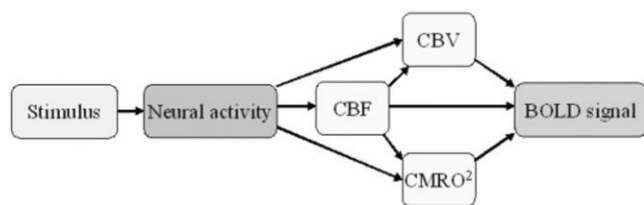
Figure 6. Neurovascular coupling. Schematic of the transformation of neural activity elicited by a stimulus to a hemodynamic response resulting in a BOLD signal. The BOLD signal reflects the ratio of nonparamagnetic oxygenated hemoglobin to paramagnetic deoxygenated hemoglobin. Neural activity alters this ratio by influencing several factors, including the CBF, CBV, and CMRO^2 (46).

cerebrovascular circulation and thereby reduced the dynamic range of the BOLD response. Although neuronal activity in this patient was considered to be normal, the difference in blood flow and oxygen concentration between the resting and active states (which is the dependent variable in fMRI studies) was limited and could potentially have led to erroneous conclusions regarding the underlying neural activity in the examined task. A less extreme example was provided by Huettel and colleagues (27), who demonstrated that lower SNRs in the visual cortex in normal aging resulted in statistically significant differences among younger subjects, even though the underlying distributions of neural activity were similar. Some of the conditions that can cause abnormal neurovascular coupling are listed in Fig. 7, and include cerebral ischemia (28), carotid stenosis (26), atherosclerosis in hypertension and diabetes (29,30), hypercholesterolemia (29), medications (31–34), and gliosis secondary to neural injury (see Ref. 35 for a comprehensive review of the topic). Simply screening patients and controls for these neurovascular factors pre-fMRI will greatly reduce the between subject variance and allow more meaningful group comparisons. Other solutions include not directly comparing between groups on a single task, but using group-by-task interactions to look at relative differences between tasks, and correlating the BOLD signal with behavioral measures during scanning.

In summary, a comparison of patient and control data is implemented in a two-stage procedure. The first-level (fixed-effect) analysis estimates the effect size for each patient or normal control. The second-level (random-effects) analysis compares the activation for different conditions or different subjects. In other words, estimates from the first level become the data for the second level, with inferences at the second level based on between-subjects variance.

STATISTICAL ANALYSES OF ABNORMAL FUNCTIONAL INTEGRATION

Studies of functional integration usually involve estimates of functional and effective connectivity among different regions. Essentially these are based on temporal correlations between activity in distant cortical regions (e.g., Refs. 36–39). In electrophysiological studies, which record spike trains of neural activity, the temporal scale is on the order of milliseconds. In functional neuroimaging, which measures hemodynamic changes, the temporal scale is on the order of seconds and a significant correlation simply implies that activity (pooled over the time scale) goes up and down together in distant regions. Such temporal correlations imply functional connectivity, which can be mediated in a number of different ways. One way rests on direct connections between the correlated regions (i.e., activity changes in one region cause activity changes in another region). The second way does not imply direct connections between correlated regions, but different regions may share connections from a region that is the source of correlated activity. The important point of this distinction is that functional connectivity does not necessarily imply direct anatomical connections between correlated regions. To the contrary, regional coupling among activities in different brain areas can be observed even when the underlying anatomical connections are indirect and polysynaptic. Functional connectivity refers simply to the presence of correlations, and is not concerned with how they were caused. Inferences about coupling among brain areas and changes in coupling involve the concept of effective connectivity. Effective connectivity is defined as the influence one neuronal system exerts over another, according to some model of this influence. We discuss some commonly used models below.

In this section we briefly discuss three different approaches for measuring functional integration: psychophysiological interactions, structural equation modeling (SEM), and dynamic causal modeling (DCM). Critically, all three are concerned with how interactions among brain regions are modulated by experimental manipulations. In a psychophysiological interaction analysis (40), the physiological response in one area of the brain is regressed on activity in all other voxels. The psychophysiological interaction is the change in the

Clinical factors that influence cerebral neurovascular dynamics and coupling

- Cerebral ischaemia *(Pinero et al.1999)*
- Carotid stenosis *(Rother et al. 2002)*
- Atheroschlerosis in hypertension and diabetes *(Claus et al. 1998, Shantaram et al. 1999)*
- Hypercholesterolaemia *(Claus et al. 1998)*
- Medications such as antipsychotics *(Miller et al 1997)*, antidepressants *(Nobler et al 2002)*, fluoxetine *(Pariente et al, 2001)*, or raloxifene *(Neele et al 2001)*
- Gliosis secondary to neural injury

Figure 7. Clinical factors that influence cerebral neurovascular dynamics and coupling.
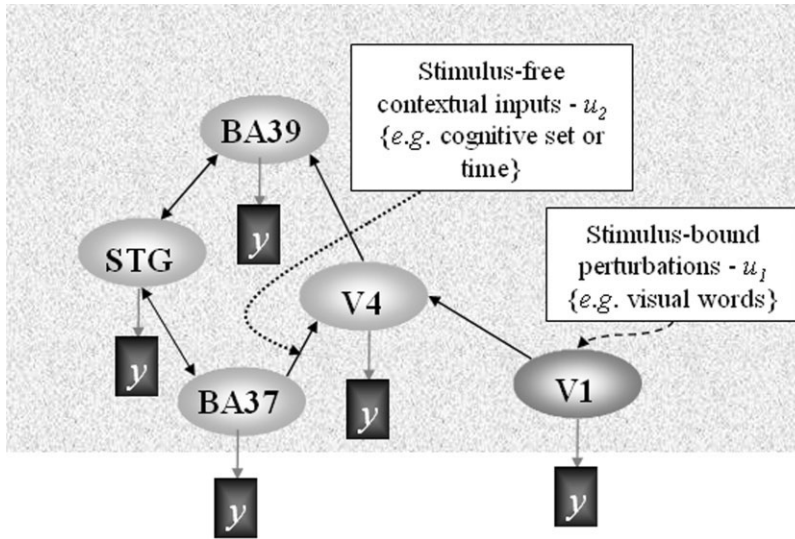
Figure 8. Functional integration and the modulation of specific pathways. This schematic illustrates the concepts behind DCM. In particular it highlights the two distinct ways in which inputs or perturbations can illicit responses in the regions or nodes that comprise the model. In this example there are five nodes, including visual areas, (V1, V4, BA 39, BA 37, and STG). Stimulus-bound perturbations designated as u_1 act as extrinsic inputs to the primary visual area (V1). Stimulus-free or contextual inputs, u_2, mediate their effects by modulating the coupling between V4 and BA39, and between BA37 and V4. For example, the responses in BA39 are caused by inputs to V1 that are transformed by V4, where the influences exerted by V4 are sensitive to the second input. The dark square boxes represent the components of the DCM that transform the state variables, z_1, in each region (neuronal activity) into a measured (hemodynamic) response, y_1.

regression slope in one psychological context relative to another (e.g., view objects or fixation). This involves a two-stage analysis. The first analysis extracts the activity in an ROI (see First-Level Statistical Analyses section) and the second analysis uses the activity in the ROI as an explanatory variable in the second-level statistical model to identify other regions in the brain where activity is significantly coupled with the ROI, and to determine how this coupling changes with condition. A significant psychophysiological interaction means that the contribution of one area to another changes significantly with the psychological context. With respect to patient studies, the psychophysiological interactions can be used to assess whether there is an abnormal coupling between brain regions that depends on the diagnosis (9).

In contrast to psychophysiological interactions, SEM and DCM look at the interactions between a limited number of regions. Therefore, both of these approaches are based on some prespecified anatomical model, which is usually constructed on the basis of known anatomical connections or functional systems. SEM of functional imaging data identifies connection strengths that best predict the variance-covariance structure of the empirical data (i.e., fMRI time series) under the constraints of the prespecified anatomical model. With respect to patient studies, when connectivity in one subject or group of subjects is compared with that in another, between-subjects differences in connectivity measures must be estimated (37,41). For example, in Ref. 41 the structural equation model encompassed responses for different subjects, so that regions within subjects were connected but regions from different subjects were not. Differences in functional integration across subjects could then be tested within this multi-subject model. Changes in effective connectivity are usually evaluated in SEM by comparing the goodness-of-fit of two models. In one model the connection of interest is held constant, and in the other it is allowed to change. If the second model is significantly better than the first, inferences can be made concerning changes in functional connections with the task or subject.

DCM (40,42–45) is based on a model of interacting neural regions supplemented with a description of how synaptic activity is transformed into a hemodynamic response (40,42–45). Unlike SEM, the characterization of coupling between brain regions is based on the rate of change of neuronal activity. As such, it does not depend on the units of activity per se, but rather the speed or rate of interregional coupling and how this is modulated by experimental manipulations (one condition relative to another). The baseline connectivity (i.e., coupling that does not depend on experimental manipulation) can be regarded as a baseline network established by the experimental context (i.e., task-set). In contrast, the modulation of intrinsic connectivity, by experimental factors, is modeled by the bilinear effects. The strength of the intrinsic connections indicates the speed with which a change in activity in one region produces a change of activity in another region. These intrinsic connections can be either positive or negative. A positive connection means that an increase in activity in one region results in an increase in activity in another region. Conversely, a negative connection means that an increase in activity in one region results in a decrease in activity in another region. Patients may show either a change in strength of instrinsic connectivity (stronger or weaker than controls) or a change in polarity (positive vs. negative).

Like SEM and psychophysiological interactions, a DCM analysis is conducted in several stages. At the first level, regional activations are extracted from each region in a subject-specific fashion. DCM is then used to estimate the strength of the effective connections and assess how these change from one experimental condition to another using a Bayesian approach. The intrinsic connections (which characterize the coupling between regions irrespective of experimental condition) and the bilinear terms (which capture how the intrinsic connections vary as a function of experimental condition) are estimated for each subject independently. These subject-specific parameters are then taken to a second level to allow T-tests on the bilinear terms. See Fig. 8 for an example.

CONCLUSIONS

To recap, fMRI can be used to address two types of neuronal abnormality in patients with neurological deficits: abnormal functional segregation and abnormal functional integration. In this paper we have considered the factors that influence the experimental design, analysis, and interpretation of such studies. With respect to experimental design, we emphasized that 1) task selection should be constrained to tasks the patient is able to perform correctly, and 2) the most sensitive designs present stimuli of the same type in a blocked or bunched fashion. With respect to preprocessing, abnormal brain structures may call for constrained spatial normalization to prevent distortion of intact tissue. We described the statistical fundamentals of first-level (within-subject) and second-level (between-subjects) analyses, and how they relate to inference using fixed- and random-effects models. We concluded with a brief introduction to the analysis of functional integration in the brain using the two most common models of effective connectivity: SEM and DCM.

We have covered a lot of issues in this review, with a special focus on those issues that are relevant to the experimental design and analysis of fMRI studies involving neurologically impaired patients. Clearly, we cannot discuss all of the details. However, we hope we have framed the key considerations and provided references for those who who are entering the field or wish to evaluate studies of this sort.

ACKNOWLEDGMENTS

REFERENCES

1. Weinberger DR, Berman KF, Illowsky BP. Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia. III. A new cohort and evidence for a monoaminergic mechanism. Arch Gen Psychiatry 1988;45:609–615.
2. Berman KF, Illowsky BP, Weinberger DR. Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia. IV. Further evidence for regional and behavioral specificity. Arch Gen Psychiatry 1988;45:616–622.
3. Frith CD, Friston KJ, Herold S, et al. Regional brain activity in chronic schizophrenic patients during the performance of a verbal fluency task. Br J Psychiatry 1995;167:343–349.
4. Liddle PF, Friston KJ, Frith CD, Hirsch SR, Jones T, Frackowiak RS. Patterns of cerebral blood flow in schizophrenia. Br J Psychiatry 1992;160:179–186.
5. Friston KJ, Liddle PF, Frith CD, Hirsch SR, Frackowiak RS. The left medial temporal region and schizophrenia. A PET study. Brain 1992;115(Pt 2):367–382.
6. Fletcher P, McKenna PJ, Friston KJ, Frith CD, Dolan RJ. Abnormal cingulate modulation of fronto-temporal connectivity in schizophrenia. Neuroimage 1999;9:337–342.
7. Spence SA, Liddle PF, Stefan MD, et al. Functional anatomy of verbal fluency in people with schizophrenia and those at genetic risk. Focal dysfunction and distributed disconnectivity reappraised. Br J Psychiatry 2000;176:52–60.
8. Lawrie SM, Buechel C, Whalley HC, Frith CD, Friston KJ, Johnstone EC. Reduced frontotemporal functional connectivity in schizophrenia associated with auditory hallucinations. Biol Psychiatry 2002;51:1008–1011.
9. Noppeney U, Friston KJ, Price CJ. Effects of visual deprivation on the organization of the semantic system. Brain 2003;126(Pt 7):1620–1627.
10. Friston KJ, Price CJ, Fletcher P, Moore C, Frackowiak RSJ, Dolan RJ. The trouble with cognitive subtraction. Neuroimage 1996;4:97–104.
11. Shaywitz BA, Shaywitz SE, Pugh KR, et al. Disruption of posterior brain systems for reading in children with developmental dyslexia. Biol Psychiatry 2002;52:101–110.
12. Shaywitz SE, Shaywitz BA, Fulbright RK, et al. Neural systems for compensation and persistence: young adult outcome of childhood reading disability. Biol Psychiatry 2003;54:25–33.
13. Shaywitz BA, Shaywitz SE, Blachman BA, et al. Development of left occipitotemporal systems for skilled reading in children after a phonologically-based intervention. Biol Psychiatry 2004;55:926–933.
14. Brunswick N, McCrory E, Price CJ, Frith CD, Frith U. Explicit and implicit processing of words and pseudowords by adult developmental dyslexics—a search for Wernicke's Wortschatz? Brain 1999;122(Pt10):1901–1917.
15. Paulesu E, McCrory E, Fazio F, et al. A cultural effect on brain function. Nat Neurosci 2000;3:91–96.
16. McCrory EJ, Mechelli A, Frith U, Price CJ. More than words: a common neural basis for reading and naming deficits in developmental dyslexia? Brain 2005;128(Pt 2):261–267.
17. Ward NS, Brown MM, Thompson AJ, Frackowiak RS. Neural correlates of outcome after stroke: a cross-sectional fMRI study. Brain 2003;126(Pt 6):1430–1448.
18. Blank SC, Bird H, Turkheimer F, Wise RJ. Speech production after stroke: the role of the right pars opercularis. Ann Neurol 2003;54:310–320.
19. Rosen HJ, Petersen SE, Linenweber MR, et al. Neural correlates of recovery from aphasia after damage to left inferior frontal cortex. Neurology 2000;55:1883–1894.
20. Naeser MA, Martin PI, Baker EH, et al. Overt propositional speech in chronic nonfluent aphasia studied with the dynamic susceptibility contrast fMRI method. Neuroimage 2004;22:29–41.
21. Morris JS, DeGelder B, Weiskrantz L, Dolan RJ. Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. Brain 2001;124(Pt 6):1241–1252.
22. Morris JS, Ohman A, Dolan RJ. Conscious and unconscious emotional learning in the human amygdala. Nature 1998;393:467–470.
23. Friston KJ, Fletcher P, Josephs O, Holmes A, Rugg MD, Turner R. Event-related fMRI: characterizing differential responses. Neuroimage 1998;7:30–40.
24. Ashburner J, Friston KJ. Nonlinear spatial normalization using basis functions. Hum Brain Mapp 1999;7:254–266.
25. Brett M, Leff AP, Rorden C, Ashburner J. Spatial normalization of brain images with focal lesions using cost function masking. Neuroimage 2001;14:486–500.
26. Rother J, Knab R, Hamzei F, et al. Negative dip in BOLD fMRI is caused by blood flow—oxygen consumption uncoupling in humans. Neuroimage 2002;15:98–102.
27. Huettel SA, Singerman JD, McCarthy G. The effects of aging upon the hemodynamic response measured by functional MRI. Neuroimage 2001;13:161–175.
28. Pineiro R, Pendlebury S, Johansen-Berg H, Matthews PM. Altered hemodynamic responses in patients after subcortical stroke measured by functional MRI. Stroke 2002;33:103–109.
29. Claus JJ, Breteler MM, Hasan D, et al. Regional cerebral blood flow and cerebrovascular risk factors in the elderly population. Neurobiol Aging 1998;19:57–64.
30. Shantaram V. Pathogenesis of atherosclerosis in diabetes and hypertension. Clin Exp Hypertens 1999;21:69–77.
31. Miller DD, Andreasen NC, O'Leary DS, et al. Effect of antipsychotics on regional cerebral blood flow measured with positron emission tomography. Neuropsychopharmacology 1997;17:230–240.
32. Nobler MS, Olvet KR, Sackeim HA. Effects of medications on cerebral blood flow in late-life depression. Curr Psychiatry Rep 2002;4:51–58.
33. Neele SJ, Rombouts SA, Bierlaagh MA, Barkhof F, Scheltens P, Netelenbos JC. Raloxifene affects brain activation patterns in postmenopausal women during visual encoding. J Clin Endocrinol Metab 2001;86:1422–1424.
34. Pariente J, Loubinoux I, Carel C, et al. Fluoxetine modulates motor performance and cerebral activation of patients recovering from stroke. Ann Neurol 2001;50:718–729.
35. D'Esposito M, Deouell LY, Gazzaley A. Alterations in the BOLD fMRI signal with ageing and disease: a challenge for neuroimaging. Nat Rev Neurosci 2003;4:863–872.

36. McIntosh AR, Gonzalez-Lima F. Network interactions among limbic cortices, basal forebrain, and cerebellum differentiate a tone conditioned as a Pavlovian excitor or inhibitor: fluorodeoxyglucose mapping and covariance structural modeling. J Neurophysiol 1994;72:1717–1733.

37. Horwitz B, Rumsey JM, Donohue BC. Functional connectivity of the angular gyrus in normal reading and dyslexia. Proc Natl Acad Sci USA 1998;95:8939–8944.

38. Friston KJ, Frith CD, Liddle PF, Frackowiak RS. Functional connectivity: the principal-component analysis of large (PET) data sets. J Cereb Blood Flow Metab 1993;13:5–14.

39. Buchel C, Friston KJ. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modeling and fMRI. Cereb Cortex 1997;7:768–778.

40. Friston KJ, Buechel C, Fink GR, Morris J, Rolls E, Dolan RJ. Psychophysiological and modulatory interactions in neuroimaging. Neuroimage 1997;6:218–229.

41. Mechelli A, Penny WD, Price CJ, Gitelman DR, Friston KJ. Effective connectivity and intersubject variability: using a multisubject network to test differences and commonalities. Neuroimage 2002;17: 1459–1469.

42. Friston KJ, Harrison L, Penny W. Dynamic causal modelling. Neuroimage 2003;19:1273–1302.

43. Mechelli A, Price CJ, Friston KJ, Ishai A. Where bottom-up meets top-down: neuronal interactions during perception and imagery. Cereb Cortex 2004;14:1256–1265.

44. Stephan KE, Harrison LM, Penny WD, Friston KJ. Biophysical models of fMRI responses. Curr Opin Neurobiol 2004;14:629–635.

45. Penny WD, Stephan KE, Mechelli A, Friston KJ. Modelling functional integration: a comparison of structural equation and dynamic causal models. Neuroimage 2004;23(Suppl 1):S264–274.

46. Buxton RB, Uludag K, Dubowitz DJ, Liu TT. Modeling the hemodynamic response to brain activation. Neuroimage 2004;23(Suppl 1):S220–S233.