

Functional MRI experiments: acquisition, analysis and interpretation of data

N.F. Ramsey^{a,*}, H. Hoogduin^b, J.M. Jansma^c

^aNeuroscience Division, Department of Psychiatry, University Medical Center of Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands

^bBCN, University of Groningen, Oostersingel 59, 9713 EZ Groningen, The Netherlands

^cLFMI, NINDS, National Institutes of Health, 10 Center Drive, Building 10, Room BID118, MSC 1065, Bethesda, MD 20892, USA

Abstract

Functional MRI is widely used to address basic and clinical neuroscience questions. In the key domains of fMRI experiments, i.e. acquisition, processing and analysis, and interpretation of data, developments are ongoing. The main issues are sensitivity for changes in fMRI signal that are associated with brain function, and the design of tasks with which brain functions are invoked. In this paper we address these issues, in terms of strengths, weaknesses and future developments. Acquisition of data is commonly achieved with techniques that measure blood oxygen level-dependent (BOLD) signal changes. Although the mechanisms that affect BOLD signal are complex and not well understood, fMRI yields results that agree with known functional topography. Sensitivity for task-related brain activity is expected to benefit from technological advances in acquisition, i.e. SENSE or parallel imaging, and higher field scanners (3 T). Data analysis is geared towards modelling sources of signal variation, i.e. reducing noise in the data time-series, and the cerebrovascular response to task-related changes in neuronal activity. Analytical algorithms such as connectivity and component analysis contribute to the extraction of meaningful information from fMRI datasets. The choice of tasks, and consequently of the statistical evaluation procedures, is best guided by the specific questions that are formulated a priori. The future is expected to bring more sophisticated questions, and tasks that allow for accurate modelling of involved brain functions. An example of a hypothesis-driven experiment is presented, where we investigated whether practise of a working memory task caused a shift in the neuronal representation of working memory or not.

© 2002 Elsevier Science B.V./ECNP All rights reserved.

Keywords: Functional neuroimaging; fMRI; Methodology

1. Introduction

Functional MRI is currently one of the most widely used techniques for basic as well as clinical neuroscience. It has evolved in various domains, into a method for addressing questions with regard to functional topography, for testing hypotheses about functions of brain structures and networks, and for elucidating the underlying neurophysiological mechanisms. Some have started to explore the potential of fMRI measures as endophenotype markers for psychopathological disorders, exploiting the potential to use neuroanatomically localised neurophysiological measures as read-out variables for genetic mechanisms (Raemaekers et al., 2002). In this paper the various elements of fMRI are addressed in terms of strengths, weaknesses and future developments. Each element is associated with particular scientific disciplines, and with particular conceptual and technical developments: data acquisition (MR physics,

engineering, signal processing), data processing (physiology, signal processing, mathematical statistics), and interpretation (biomedical sciences, psychology, image visualisation). Particularly in technological areas, developments are expected to continue into the next decade, and will undoubtedly improve some of the weaknesses of fMRI. Detailed discussions of imaging and analysis techniques are beyond the scope of this paper. For excellent reading material on the various technical and methodological issues of (f)MRI the following books are recommended: Moonen and Bandettini (1999), Jezzard et al. (2001), Tofts (2003), and Vlaardingerbroek and Den Boer (1999). The paper is concluded with an example of a study on working memory, to show how fMRI can be used to test hypotheses about brain function.

2. The fMRI experiment from A to Z

An fMRI experiment consists of several components.

*Corresponding author.

Firstly, there is the acquisition of scans. This typically involves continuous series of scans, each lasting a few seconds and covering much or all of the brain. A scan consists of several thousand data points, each of which is derived from a cube of brain tissue. The size of the cube, also called a voxel (volume element) determines the resolution of the scan, which in turn determines the minimum distance between two adjacent points in the brain that can be distinguished from one another. The series of scans is stored as a time-series of 3D volumes, where each voxel is associated with a series of intensity values. In MRI, the intensity is determined by various factors, which depend on the type of scan that is executed. The basis of the signal is provided by protons. The scan type, called the pulse sequence, determines the particular set of factors that affect the basis of the signal. Some pulse sequences for instance are sensitive to the type of biological material such as grey matter, white matter, CSF or blood (e.g. 'T1' or 'T2' weighted imaging, angiography), while others are sensitive to motion (e.g. cardiac imaging) or diffusion of water along nerve fibres (diffusion-weighted imaging) (Vlaardingerbroek and Den Boer, 1999). The distinction is evident from the signal intensity contrasts that are created between the material of the mechanism of interest, and the rest of the scanned material. In fMRI, the pulse sequence is sensitive to blood dynamics, i.e. blood flow, blood volume and oxygenation state. The essence of fMRI is that it enables visualisation and measurement of transient changes in blood dynamics that are coupled to changes in neuronal activity.

The second component of an fMRI experiment is the task. The task is presented to the subject while lying in the scanner, and indirectly invokes a selective set of brain functions. It requires the subject to perform specific acts, and thereby induces series and complexes of neuronal events in a controlled manner. Acts can be of a receptive nature, e.g. perception of visual or auditory material such as looking at a flashing light or listening to music, or of a reactive nature, such as responding to a stimulus according to certain rules conveyed to the subject or memorising words for a certain amount of time. In practise both types are invoked with any task, and the challenge is to isolate a particular set of brain functions. To do this, a second task has to be designed which invokes all of the functions that one is not interested in, i.e. a reference (control) task. Many brain functions cannot be isolated that way, because they are coupled to other functions (Friston et al., 1996a). For example, memory cannot be invoked in a controlled manner without using external stimuli, and external stimuli invariably invoke memory processes to some degree. For those functions, a more complex set of tasks is required. The way tasks are constructed and are organised during the experiment is referred to as the paradigm or design. More complex designs involve multiple tasks (conjunction design) (Price and Friston, 1997; Ramsey et al., 2001), or tracking of events in time by spacing stimuli far enough

apart to assess the vascular response (event-related design) (d'Esposito et al., 1999; Liu et al., 2001). As the design of the tasks and the constellation in which they are presented during the experiment determine how the results should be interpreted, this component is discussed below as the interpretation of fMRI experiments.

The third component is processing and statistical analysis of the data. The primary objective of algorithms devised for this purpose is to detect and isolate changes in blood dynamics that reflect neuronal events. Several analysis programs are available that enable statistical evaluation of fMRI datasets (e.g. SPM at www.fil.ion.ucl.ac.uk/spm, Brain Voyager at www.brainvoyager.de, AFNI at <http://afni.nimh.nih.gov/afni>). They contain methods of searching for changes in fMRI signal that is associated with the task, using predominantly multiple regression, and assigning statistical significance to these changes. In addition, several programs allow one to test for correlations between regions (connectivity analysis) (Horwitz et al., 1999), or to extract the main factors in the data time-series (principal or independent component analysis) (e.g. Andersen et al., 1999; Friston et al., 2000). Each of these components is discussed in more detail below, and an example of an experiment is described to explain how fMRI can be used to test a hypothesis about a particular brain function.

3. Acquisition of fMRI data

Key to the use of fMRI for studying brain activity is that the scans are sensitive to blood dynamics. In order to associate changes in blood flow with specific neural events in the brain, scans have to be fast. Images covering the whole cortex can be acquired in as little as 1 s, but a price has to be paid. In MRI, speed, resolution (i.e. the size of the units of which images are composed) and image quality (expressed as signal to noise ratio SNR) affect each other significantly (e.g. Vlaardingerbroek and Den Boer, 1999). Fast scans are typically relatively coarse in that the number of points in space that is acquired is reduced to a minimum. For example, whereas anatomical scans, which last about 15 min, can be conducted at a resolution of less than 1 mm (i.e. 0.001 cc), functional scans typically provide images with a voxel size of 4 mm (i.e. 0.064 cc) or more. Also SNR is relatively low in fMRI, resulting not only in noisy images, but also in increased variability in image signal intensities over repeated scans (Ramsey, 1999). This variability, also called image stability, is important in that it determines the sensitivity for small changes in blood dynamics. Detection of changes in blood dynamics that accompany neural events in the brain, such as making fingers move, are the primary objective of statistical analysis of the fMRI time-series data. Each voxel in the scanned volume is tested for a correlation

between signal intensity in that voxel and the neural events that are evoked by a particular task (see next section).

The relationship between blood flow and MRI signal in fMRI is a complicated matter (e.g. Ogawa et al., 1998). In PET research the relationship between blood flow and signal is fairly straightforward, and the models that describe this relationship appear to be adequate for incorporation in data analysis. For functional MRI some models exist, but they apply to a specific family of scan techniques (perfusion, arterial spin labelling) (e.g. Jezzard and Ramsey, 2003). There is no complete model for the family that is most widely employed in neuroscience, i.e. the blood oxygen level-dependent signal change detection ('BOLD') techniques. These techniques make use of the fact that the effect of haemoglobin on the surrounding magnetic field changes with the occupation of oxygen: oxygenated haemoglobin has a negligible effect on the field, whereas deoxygenated haemoglobin causes inhomogeneity in the field. The latter causes the fMRI signal to decay faster, resulting in a drop in signal intensity, on the order of 1–5% in a 1.5 T scanner. Essentially the drop in signal that follows an increase in deoxyhaemoglobin is proportional to the local quantity of the (deoxygenated) molecule. This is, in turn, determined by multiple variables, notably haematocrit, blood flow, blood volume and oxygenation state (Ogawa et al., 1992). The last three, taken together in the term 'blood dynamics', change when neuronal metabolism changes with activation. The change in fMRI signal is the end result of interactions between these variables, which may act together in the same direction, or may interfere with one another with respect to their effects on signal intensity. The effect of deoxyhemoglobin on MR signal is also a complicated matter, as it affects the signal of the blood itself, and signal of the surrounding tissue (Ogawa et al., 1993).

Blood flow is the variable of primary interest in fMRI, because it is closely associated with neuronal activity. The relationship is however not as straightforward as one would like it to be, for several reasons. It is not related completely linearly to neuronal activity. This has been shown in various studies, and it is attributed to the complexity of metabolic processes in neurones and adjacent glia cells, and to either a dissociation (Fox and Raichle, 1986) or a non-linear association (Buxton and Frank, 1997) between neuronal glucose utilisation and vascular oxygen delivery. The latter is held responsible for the finding that oxygen delivery exceeds oxygen demand (calculated from glucose consumption). This principle is the basis for fMRI, in that neuronal activation causes an increase instead of a decrease in MR signal. The cerebrovascular system responds to an increase in oxygen demand by providing so much oxygenated blood that the local presence of deoxyhemoglobin is reduced, resulting in less local inhomogeneity of the magnetic field (Ogawa et al., 1992).

At high neuronal firing rates oxygen delivery does not

increase at the same pace as at lower firing rates, resulting in an apparent saturation of blood flow (e.g. Vafaee et al., 1999). An example is shown in Fig. 1, where the BOLD response, in primary motor cortex, to a single instance of finger movement is compared to the response to multiple movements. The BOLD response was modelled for several stimulus rates (inter-stimulus intervals of 3, 2 and 1 s), on the basis of the single stimulus, and compared to the measured response. The figure shows that the BOLD response increases linearly with stimulus rates below 0.5 Hz. At faster rates (1 Hz) the measured response fell short of the modelled response, indicating saturation of the BOLD effect. Another issue is that blood flow appears to be more prominent in neuronal terminal areas than at the cell bodies (Logothetis et al., 2001), potentially resulting in a spatial bias in sensitivity for blood flow changes. It is also not clear whether blood flow changes reflect excitatory or inhibitory neuronal activity, or both, although excitatory activity seems to dominate fMRI signal changes (Waldvogel et al., 2000).

What makes fMRI more sensitive to multiple sources of signal fluctuation than PET is probably that fact that fMRI deals with changes in the basic MR signal provided by protons. Effects of changes in blood dynamics manifest themselves as a percent reduction or increase in the basic signal. That signal is also affected by a host of other factors, including field homogeneity, motion of the head and of the body, oxygen flow, respiration, cerebrovascular

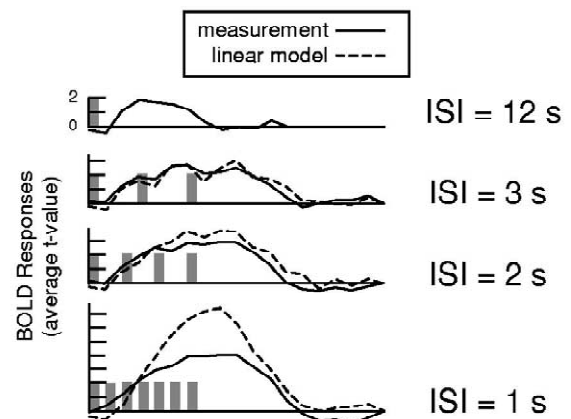


Fig. 1. Linearity of the BOLD response. BOLD response for a single subject, to pressing a button with the thumb, measured with PRESTO (Ramsey et al., 1998). For acquisition of the BOLD response, 50 movements were made with 12-s interval to allow the response to return to baseline (event-related design). This curve, obtained from the primary motor region, was used to model linear responses to trains of movements. The primary motor region was identified using a block-design session immediately before the event-related session. Three conditions are shown: with 3, 2 and 1 s interval (movements are indicated with grey bars). The measured BOLD response is shown as a black line. The modelled linear response, composed of the single movement curve, is shown as a dotted line. The BOLD response deviates from the linear model only at the 1 Hz condition. This indicates that saturation occurs in motor cortex, but only when inter-stimulus interval is shorter than 2 s.

pulsation and blood flowing into the imaged volume. The effect of some of these can exceed the BOLD effect significantly, causing dramatic fluctuations from one scan to the next. For instance, signal in tissue near cavities is affected so strongly that it either disappears completely, or causes excessive sensitivity to motion (the position of the head affects field homogeneity). The magnitude of artefacts depends in part on the particular pulse sequence used, be it EPI (Belliveau et al., 1991), FLASH (Frahm et al., 1993), Spiral (Yang et al., 1996) or PRESTO (Van Gelderen et al., 1995), but in any of these, artefacts remain an issue of concern. These issues are discussed in more detail in Moonen and Bandettini (1999) and Jezzard and Ramsey (2003).

To summarise, the relationship between neuronal activity and fMRI signal is quite complicated, involving multiple (neuro)physiological mechanisms. Nevertheless, many studies have shown that fMRI yields activation patterns that agree with what is known from human lesion studies and from non-human primate studies about functional topography. The fact that there is no satisfactory model available to describe the whole set of mechanisms, does by no means invalidate the use of fMRI. To fully understand what really happens during brain activation as measured with fMRI, and to understand what neuronal events cannot be measured, requires further development of a comprehensive model. Currently, this is an ongoing endeavour involving mathematical modelling as well as cross validation with imaging methods that measure specific physiological processes, such as arterial spin labelling to measure cerebral blood flow (e.g. Yang et al., 2000), perfusion imaging (using contrast agents) to measure blood volume and elaborate T2 measurement (van Zijl et al., 1998) techniques to measure blood oxygenation. None of these methods can currently compete with fMRI, as they are either significantly slower, or less sensitive to changes in neuronal activity.

In the near future, several improvements can be expected. Recent improvements on the receiver side of fMRI equipment, i.e. radiofrequency coils and amplifiers, have boosted the speed of imaging considerably. In particular the use of multiple receiver coils, named ‘SENSE’ or parallel imaging (Golay et al., 2000) can reduce image acquisition speed by a factor of 4 or even more. The strength of the magnet is also an important variable in the magnitude of activation-induced fMRI signal change. Where 1.5 T is now the norm for clinical scanners, 3 T is rapidly becoming the standard for neuroscience. High field scanners (the next upgrade is aimed at 7 T) make it possible to reduce scan speed, but they also magnify some of the problems inherent in MR imaging, such as field distortions near the nasal cavity and the ear channels, and motion artefacts. Another problem is that higher fields require faster ‘gradient switching’ (i.e. the additional magnets that are switched as part of the data acquisition scheme), resulting in higher levels of electromagnetic

power and subsequently in the risk of inducing nerve stimulation. The most effective use of the stronger scanners will be the combination with SENSE techniques, as the latter reduce acquisition time, allowing for reduction of the frequency of gradient switching.

4. Processing and analysis of fMRI data

Analysis of fMRI data has two objectives: (1) to identify brain regions that are involved in the function that the investigator is interested in (function of interest FOI) and/or to quantify relationships between task characteristics, MR signal and regions, and (2) to minimise residual noise in the data in order to optimise extraction of function-related signal changes. In both cases it is of critical importance to model the sources of signal changes in the time-series as accurately as possible. The more variation in signal can be modelled, the more sensitive the analysis is for signal changes that are coupled to the neuronal events one is interested in. Part of signal instability can be dealt with in the data acquisition. For instance, motion artefacts can be detected during scanning, and corrected during reconstruction by using indicators of movements acquired during scanning (Hu and Kim, 1994; Ramsey et al., 1998), or by replacing part of the raw data with parts of earlier scans (in so-called K-space).

Another effect of motion is displacement of scan volumes in space. This is typically the result of slow changes in the position of the head, or of brief displacements due to movement of the body. Before analysing image time-series, these displacements have to be corrected, and this can be done with registration algorithms (Thevenaz et al., 1998; Woods et al., 1992). Most of the commonly available programs will do an adequate job, both in terms of finding the best position and of realigning and reslicing the volumes. Failure to register the volumes will result in large signal fluctuation in voxels at the cortical surface and bordering CSF, where the contrast in MRI signal is largest, effectively reducing sensitivity for the BOLD effect (Ramsey, 1999). Fig. 2 shows an example of the effect of motion correction in image acquisition (by means of a navigator echo; Ramsey et al., 1998). The navigator echo was acquired during scanning, but image reconstruction was performed both with and without this information. The figure displays a schizophrenic patient in a fingertapping experiment, who moved throughout the scan session. Without correction, the statistical image was devoid of activity, due to large signal fluctuations throughout the brain, whereas with correction activity was clearly present in the contralateral primary motor cortex. Alternatively, motion may be correlated with the task, for instance when one moves the body (even slightly) during a motor task, and moves back during the rest condition, resulting in ‘activation’ at the cortical surface. There is no way to distinguish real activation from




Fig. 2. Motion correction during scanning. A schizophrenic patient participated in a fingertapping experiment with a simple block design. Images were acquired with PRESTO (Ramsey et al., 1998). Five slices, out of the 26 scanned, are shown from bottom to top, covering the primary motor cortex. Slices are in radiological orientation (left is right and vice versa). Due to motion, no activity was found in the statistical map (A). During scanning, navigator echoes were acquired, which measure phase changes in the raw MRI signal caused by motion. When these echoes were used to correct motion artefacts during image reconstruction, image stability was greatly improved, resulting in emergence of motor cortex activity (B). This shows the importance of motion correction, especially when studying psychiatric patients as they tend to move more than healthy controls.

stimulus-correlated motion artefacts in the statistical maps (Hajnal et al., 1994) if the volumes have not been registered.

The final noise reduction is achieved in the statistical analysis. The time-series data for each voxel are commonly analysed by means of multiple regression, where both the input functions (derived from the paradigm) and known sources of noise are entered as independent factors. These noise factors typically include a global trend (modelling slow changes in hardware properties associated with continuous scanning at a high duty cycle), registration parameters (i.e. indicator of motion obtained from the registration procedure) to remove residual motion effects, and a series of sine and/or cosine functions with a range of frequencies (to eliminate rhythmic physiological fluctuations). For each input function a t - (or F -) value is computed from the regression coefficient and the residual error (Friston et al., 1995; Worsley and Friston, 1995). A statistical criterion (threshold) is finally computed on the basis of the number of voxels within the scanned and analysed volume, the spatial resolution and the autocorrelation in the time-series (e.g. Worsley, 1994; Zarahn et al., 1997).

After processing and analysis, statistical maps are composed of the t -value for each voxel, i.e. one map for each input function. From here on, there are several options. One can superimpose voxels where the t -value exceeds the preset threshold onto an anatomical volume for that subject, to identify and visualise the pattern of activities that are associated with the task. One may also co-register the map (without first thresholding voxels) to a standard brain (e.g. MNI; Collins et al., 1994), in order to perform a second statistical map analysis, this time on one or more groups of subjects. This generates a group-map which demonstrates the regions that are associated with the task across the group of subjects. Alternatively, the group-map may depict areas where two groups differ from one another. In these cases, statistical voxel-wise analysis is

based on the number of subjects rather than the number of scans (e.g. Aguirre et al., 1997).

An alternative to constructing a factor matrix based on known and hypothesised sources of signal variation, is principal or independent component analysis (Andersen et al., 1999). This approach extracts a set of patterns from the time-series of all voxels, and attributes weights to each pattern (i.e. factor) for each voxel. Component analyses primarily serve to identify activity patterns without a priori definition of the exact input function(s), as well as unknown (main) sources of noise. They do not allow for tests of activation like multiple regression, but attempts are ongoing to combine the advantages of both component and regression analysis. Component analysis is most useful for identifying input functions that can be used in subsequent experiments for factor matrix construction in multiple regression.

A different approach is indicated when one wants to test a hypothesis about activity in brain regions that are targeted a priori. One can first create a group-wise map based on one particular input function, and subsequently analyse the effect of other tasks on signal in the regions that are significant in the group-map. For this purpose, the statistical maps associated with the other tasks (generated in the individual regression analysis) are masked with the results of the group-map, and a mean t -value is computed for each region, for each task and for each subject. These data can then be analysed with a standard general linear model repeated measures analysis, allowing one to test for differences between tasks, and/or groups of subjects, in specific brain regions. This approach is described in more detail in Section 5.

In addition to identifying regions involved in a particular function, and testing hypotheses about their particular contributions, it is also possible to examine interactions between regions. In its simpler form correlations, across subjects, between identified regions can be determined, and mapped out in a 'functional' (with no directional spe-

cificity) or ‘effective’ (with specific directions) connectivity diagram (e.g. Horwitz et al., 1999). Although informative, connectivity analyses are quite sensitive to selection of brain areas to include in the equations, subject selection and quality of fMRI datasets. Comparison of connectivity maps between tasks or groups is a complex matter, but some studies have shown interesting results. For instance, Buchel et al. showed that activity between regions of the working memory system becomes more correlated after practise of a working memory task (Buchel et al., 1999), indicating that efficiency of cross-communication improves.

To summarise, analysis of fMRI data generally involves multiple regression, because it allows one to build a model (i.e. the factor matrix) for the sources of signal fluctuation. As the sources of fMRI signal change are better understood, the models will improve sensitivity for signal change associated with the FOI (function of interest). For instance, once scans can be acquired fast enough, factors can be included that represent the pulsation of blood vessels. Connectivity analyses are mathematically possible, but the interpretation of connectivity maps is limited as long as there is no firm, neuroanatomical, hypothesis about which brain structures are the real players of the network that is examined, or how they affect each other. Significant contribution to basic and clinical neuroscience requires that fMRI is used to test clear hypotheses, and it is to be expected that, with the accumulation of information about human brain function, this will be done at a larger scale, at the expense of descriptive ‘brain mapping’ studies.

5. Evaluation and interpretation of statistical maps

Experimental design is the critical feature of fMRI experiments for interpretation of results. The choice of tasks and the spacing between stimuli in time determine the meaning of the statistical maps. The statistical significance assigned to voxels essentially reflects the strength of the association between fMRI signal (over the time-series) and the input function or set of functions. In a simple design the association reflects the difference between two conditions (‘on’ versus ‘off’) (Bandettini et al., 1993). Although this is a powerful design (Liu et al., 2001), the interpretation does not go beyond indicating that a particular voxel responds differently to the two conditions. It does not indicate whether that voxel was important for the FOI, or whether it was involved but not essential. In a simple fingertapping experiment for instance, voxels may represent motor neurones (in primary motor cortex), but alternatively may represent attentional control (e.g. anterior cingulate cortex) or planning of the motor act (e.g. supplementary motor or premotor areas). Prior knowledge about neuroanatomy makes it possible to differentiate, but this undermines the hypothesis-testing aspect of an fMRI experiment. One cannot pose absence of brain activity as a

null-hypothesis, if it is known that functions other than the FOI introduce activity in the statistical maps. This touches on one of the conceptual problems in fMRI, namely the fact that statistical analysis is based on testing against the null-hypothesis of absence of activity, whereas the investigator does not want to know *whether* there is any proof of activity, but *where* activity is located. If the objective of an experiment is to search for regions that mediate the FOI, then there is no other choice than to adopt this null hypothesis, but there is a price to pay. The threshold for significance has to be adjusted for the number of voxels that is scanned and searched. As this typically amounts to tens of thousands of voxels in scans covering most or all of the brain, then the most straightforward way to determine the threshold is by means of a Bonferroni correction (e.g. Ramsey et al., 1996). The threshold for *t*-values will end up in the range of 4.5 or higher, corresponding to a *P*-value of 0.00001 or less, provided that a large number of scans are acquired (providing enough degrees of freedom). On the one hand this is rather stringent, because it does not take into account the possibility that time-series of voxels may be correlated in space (notably adjacent voxels). On the other hand, the noise in time-series is often correlated in time, resulting in an overestimation of the degrees of freedom, hence an overly liberal threshold. Various methods to correct for correlations in space and time are provided in data analysis programs, utilising properties of Gaussian fields (Worsley, 1994) or reverting to non-parametric statistics (e.g. Brammer et al. (1997) but see Aguirre et al. (1998)). The threshold may be brought down by imposing the additional requirement that only clusters of voxels (i.e. neighbouring in space) of a particular minimum size can qualify for significance (Friston et al., 1996b). Unfortunately, due to the fact that there are no clear neuroanatomical boundaries to regions that mediate an FOI, particularly when regions are searched for, choosing any particular cluster size is rather arbitrary (a cluster size of the order of 5–10 voxels appears to satisfy many investigators). From neurosurgery it is known that regions involved in language may be as small as 5 mm (Rutten et al., 1999), indicating that application of a cluster criterion may eliminate functionally relevant regions.

As mentioned earlier, fMRI can be used to test hypotheses about function characteristics of specific regions. If the design of the experiment is chosen carefully, then one can search for, and identify, the regions that are involved in the FOI, and subsequently test the effect of experimental manipulations on those regions, within the same experiment. This approach requires that the design include multiple conditions (i.e. versions of the same task that differ in one particular aspect). After analysis of the individual subjects’ dataset, as described earlier, each input function, coding for one of the conditions, is represented by a statistical map. All maps are transformed spatially into a common space (like the MNI brain; Collins et al., 1994), allowing for group-wise analyses. For identification

of the brain regions that are involved in the FOI, two of the conditions are contrasted. These should be carefully matched, to avoid contamination with activity associated with non-specific processes, such as for instance visual processing and motor response generation in cognitive experiments. A group-wise test for significance yields a statistical map with several regions. Next, each of those regions is labelled as a volume of interest (VOI), and the co-ordinates of all voxels within each VOI are stored. Then these co-ordinates are used to calculate the mean values of voxels within each VOI, in the statistical maps of other conditions. This results in one statistical measure for each VOI, for each condition. These data can be entered in any statistical program for analysis of effects of condition on the set of VOIs. The key advantage of this approach is that only several regions are analysed instead of all voxels in the scanned volume, resulting in greater statistical power. This approach also enables testing of a specific hypothesis about the effect of task manipulations on brain function, or about effects of pharmacological agents. To clarify this approach, an example is given of an experiment in healthy subjects, on the effect of practise on brain activity associated with a working memory task (Jansma et al., 2001).

In summary, there are many ways to design an fMRI experiment. The choice of tasks and conditions determines the interpretation of the data, and as such deserve considerable attention. In contrast with neuropsychological approaches where tasks are relatively complex in terms of the brain functions involved, fMRI tasks are preferably very simple. This is necessary if one is interested in a particular brain function, as it makes it easier to isolate it from other functions. The future may bring standardisation of a set of tasks that are representative of specific domains of human brain function, such as memory, language, working memory, attention, visual processing etc. This will make it possible to investigate the underlying neuronal mechanisms, including firing rates, neurotransmission, gene expression and such. An issue not dealt with in this paper is performance on tasks. It will be a challenge to incorporate performance measures in fMRI design, particularly when comparing patients to controls (e.g. Ramsey et al., 2002). The relationship between brain activity and functional output is by no means the least important aspect of brain function.

6. Testing a hypothesis: an example

A well-documented behavioural phenomenon is that performance of certain tasks improves with practise in an orderly fashion (Shiffrin and Schneider, 1977; Sternberg, 1966). When a task is new, performance requires conscious and effortful processing, but after practise it is less demanding yet more accurate. This improvement is attributed to ‘automatization’, i.e. a mechanism by which cognitive processes that are initially executed in a controlled

manner, become processed more efficiently and less dependent on conscious attention. The question is whether the same brain system is involved in both, or whether there are two different systems. We addressed this issue with an fMRI experiment, where we examined changes in the working memory system following practise of a simple task (Jansma et al., 2001).

Fifteen healthy subjects participated in an fMRI experiment where we investigated how the brain adjusts to automatization of working memory function. Subjects performed three versions (conditions) of one task. This task has the following basic format (Jansma et al., 2001). A set of five consonants is shown for 5340 ms (the target-set). After this a series of 10 consonants is displayed in sequence. Stimuli and their presentation are shown in Fig. 3. A new set of five consonants is then shown, followed by 10 new trials presented with an interval of 2670 ms. Subjects were instructed to memorise the target-set and subsequently press a button as quickly as possible when a consonant belonged to the target-set (50% were targets). Two experimental tasks were administered, which differed only with regard to the target-set(s): a novel task (NT) and a practised task (PT). In the PT one the same set was used repeatedly. In the NT the composition of the target-set was changed after every run of 10 trials. The target-set and set of non-targets for the NT were chosen from the 10 remaining consonants that were not used for the PT. During a training-session before fMRI scans, which lasted 21 min, only the PT was presented, in five series of 100 stimuli. During scanning, both tasks were presented, in 8 epochs (i.e. the length of a particular condition in seconds) of 10 stimuli each. An additional reaction time control task (CT, same numbers of epochs and stimuli) was included, during which subjects had to press the button as quickly as

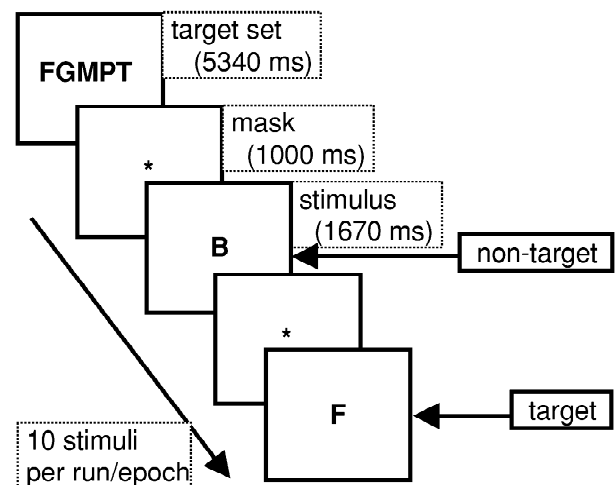


Fig. 3. Automatization task. The task is shown on the left. Stimuli were presented with a fixed interval. During the practised task (PT) the letter set at the start of each run of 10 stimuli was the same as that used for practise before scanning. During the novel task (NT), the letter set was different for every 10-trial run. See text for further details.

possible when the symbol $<>$ appeared, as well as rest periods of equal epoch duration. The sequence of the three tasks and rest periods was randomized. Reaction time for all correctly identified targets and accuracy for all stimuli were measured. The tasks contained equal numbers of targets. The critical feature of this experimental design is that it measures automatization, i.e. the difference between NT and PT, without the complication of learning effects within the scan session (Jansma et al., 2001). All subjects improve their performance in the PT after training, as evidenced by fewer errors and shorter reaction times as compared to the NT.

For fMRI a 3-dimensional technique was used (navigated PRESTO) that measures BOLD-sensitive signal changes (Ramsey et al., 1998). A single run of 384 scans was acquired over a period of 17 min (for scan parameters see Jansma et al., 2001). Data analysis consisted of three stages. Firstly, after motion correction, statistical activity maps were generated for each subject, for each of the three tasks (NT, PT and CT each compared to the rest condition, excluding scans obtained during presentation of the letter sets) by means of multiple regression (Worsley and Friston, 1995). Next, to identify regions that mediate working memory, these maps were smoothed (FWHM 8 mm) and normalized into standard space (Collins et al., 1994), and were analyzed for the whole group, contrasting the NT to the reaction time control task (CT) and using z -statistics (Worsley, 1994). This way, visual processing and motor responding was eliminated, as they were equally present in both tasks. The NT was chosen as working memory was invoked most strongly in this condition. The

CT was chosen, as it did not require working memory except for holding the task instruction online. This selection procedure yielded seven regions where activity differed significantly between NT and CT ($P < 0.0001$ per voxel, with at least 10 voxels in a cluster). The four largest regions, i.e. left and right dorsolateral prefrontal cortex, anterior cingulate cortex (extending into the supplementary motor area), right superior parietal cortex have been implicated in various studies as being important regions for working memory (e.g. Casey et al., 1998; Cohen et al., 1997). These are shown in Fig. 4, together with graphs displaying the magnitude of activity during the three conditions. Additional regions were right frontopolar cortex (the function of which is not well known), cuneus and precuneus. As the primary objective of the study was to assess the effect of practise on the working memory system, NT was compared to PT by entering the mean values of voxels within each VOI of each of the two statistical maps (for each subject) in a General Linear Model analysis with repeated measurements. This analysis yielded a significant overall effect of practise on activity in all the VOIs, except for the cuneus and precuneus which were not affected by practise. The frontal and parietal regions exhibited the same effect of practise, indicating that the network as a whole became less active. Fig. 4 shows the effect of practise in the frontal and parietal working memory regions.

Although being a straightforward approach, this VOI analysis is biased as it is unlikely that areas are found where activity emerges after practise. To address the question whether another brain system takes over after

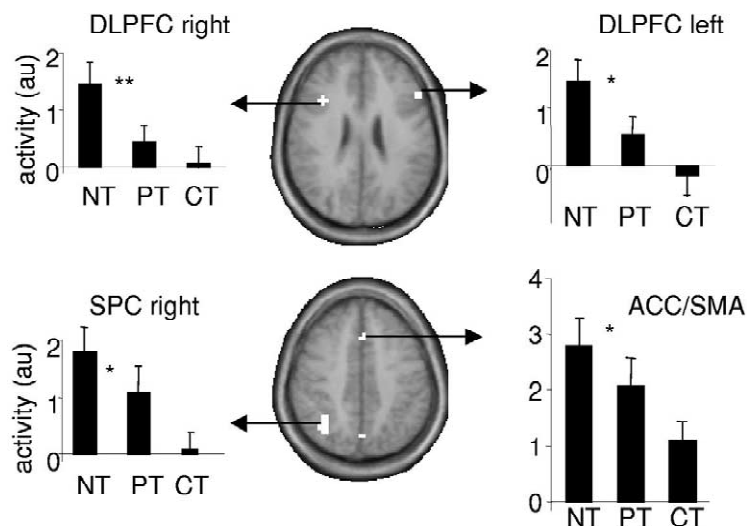


Fig. 4. Brain regions involved in working memory. In white the areas involved in working memory are superimposed on the averaged anatomical image. Two of the 30 slices imaged are shown, for display of four regions that emerged from the group-contrast between NT and CT, i.e. left and right dorsolateral prefrontal cortex (DLPFC), right superior parietal cortex (SPC) and anterior cingulate cortex extending into supplementary motor areas (ACC/SMA). Slices are in radiological orientation (left side is right hemisphere and vice versa). In these four regions brain activity declined with practise. The bars in the graphs represent levels of activity during the novel (NT), the practised (PT) and the control (CT) task, relative to the rest state. Activity is given in arbitrary units (a.u.), for the mean t -values of individual maps. Analysis of variance, effect of practise $*P < 0.05$, $**P < 0.01$. Figure is adapted from Jansma et al. (2001).

practise, the analyses were conducted again, but now we selected the regions that were active in the PT instead of the NT. This revealed that the same regions were active in the PT as in the first analysis. More importantly, there was no indication that there was a different cognitive network involved in the PT, even when applying a very liberal statistical threshold. Taken together, these data indicate that the same network is involved in processing of a working memory task before and after practise. Improvement of performance probably involves streamlining of neuronal communication within the network, rather than delegation to another system.

It should be noted that the bias resulting from VOI selection is not necessarily a problem as it depends on the hypothesis. In this study the question was whether automated cognitive processes were mediated by the working memory system, or by a different system dedicated to automated acts. The two VOI analyses were clearly relevant for this hypothesis. Had, however, the hypothesis pertained only to the working memory system, addressing for instance effects of medication on working memory function, then the second VOI analysis (on regions selected for PT activity) would not be as relevant (if regions are not active during NT, they are in principle not significant for working memory).

Another issue in VOI-driven analysis concerns the use of group averaging of maps. It is clear that functional anatomy, i.e. the precise location of functions on the cortex, varies greatly across subjects. To overcome this spatial variance, images have to be smoothed. If the VOIs derived from group analysis are projected back onto the individual subjects' maps, quite often only a small proportion of that region actually contains active voxels. Averaging across all voxels within that region will not only reduce the magnitude of activity, but also introduces noise from non-active voxels. One would preferably select the active regions for each subject separately, but the problem then arises how to decide how many VOIs there are for the group of subjects, and to which VOI active voxels should be assigned (if at all).

Apart from improving statistical power for detection of task-related changes in brain activity, the multi-stage approach carries another advantage particularly for pharmacological studies. Agents may affect blood dynamics non-selectively, for instance by altering vascular responsiveness to metabolic changes. In the current example, one can compare the difference between the control task (CT) and the resting state (reflecting visual and motor processing as well as attention), between compound and placebo. Assuming that visual and motor function are not affected by the compound, any difference between the two datasets will be indicative of a non-selective effect of the compound on fMRI signal change. This can be quantified in a VOI analysis, and can subsequently be used as a correcting factor or covariate in analysis of the FOI data.

In conclusion, sensitivity and modelling of the fMRI

time-series are the primary issues in fMRI. With the advent of hardware that allows for faster image acquisition, and paradigms that allow for modelling of neuronal mechanisms underlying brain functions, fMRI will continue to develop as a valuable tool for testing hypotheses in basic and clinical neuroscience. Modelling of brain functions will undoubtedly benefit from the combination of fMRI with other imaging techniques such as electroencephalography (EEG), magnetoencephalography (MEG) and near-infrared spectroscopy (NIRS), as well as methods to manipulate neuronal activity directly, such as transcranial magnetic stimulation (TMS).

Acknowledgements

We thank Arjen van der Schaaf and Matthijs Vink for their contribution to the BOLD response experiment.

References

- Aguirre, G.K., Zarahn, E., d'Esposito, M., 1997. Empirical analyses of BOLD fMRI statistics. II. Spatially smoothed data collected under null-hypothesis and experimental conditions. *NeuroImage* 5, 199–212.
- Aguirre, G.K., Zarahn, E., d'Esposito, M., 1998. A critique of the use of the Kolmogorov–Smirnov (KS) statistic for the analysis of BOLD fMRI data. *Magn. Reson. Med.* 39, 500–505.
- Andersen, A.H., Gash, D.M., Avison, M.J., 1999. Principal component analysis of the dynamic response measured by fMRI: a generalized linear systems framework. *Magn. Reson. Imag.* 17, 795–815.
- Bandettini, P.A., Jesmanowicz, A., Wong, E.C., Hyde, J.S., 1993. Processing strategies for time-course datasets in functional MRI of the human brain. *Magn. Reson. Med.* 30, 161–173.
- Belliveau, J.W., Kennedy, D.N., McKinstry, R.C., Buchbinder, B.R., Weisskoff, R.M., Cohen, M.S., Vevea, J.M., Brady, T.J., Rosen, B.R., 1991. Functional mapping of the visual cortex by magnetic resonance imaging. *Science* 254, 716–719.
- Brammer, M.J., Bullmore, E.T., Simmons, A., Williams, S.C., Grasby, P.M., Howard, R.J., Woodruff, P.W., Rabe-Hesketh, S., 1997. Generic brain activation mapping in functional magnetic resonance imaging: a nonparametric approach. *Magn. Reson. Imag.* 15, 763–770.
- Buchel, C., Coull, J.T., Friston, K.J., 1999. The predictive value of changes in effective connectivity for human learning. *Science* 283, 1538–1541.
- Buxton, R.B., Frank, L.R., 1997. A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. *J. Cereb. Blood Flow Metab.* 17, 64–72.
- Casey, B.J., Cohen, J.D., Ocraven, K., Davidson, R.J., Irwin, W., Nelson, C.A., Noll, D.C., Hu, X.P., Lowe, M.J., Rosen, B.R., Truwitt, C.L., Turski, P., 1998. Reproducibility of fMRI results across four institutions using a spatial working memory task. *Neuroimage* 8, 249–261.
- Cohen, J.D., Perlstein, W.M., Braver, T.S., Nystrom, L.E., Noll, D.C., Jonides, J., Smith, E.E., 1997. Temporal dynamics of brain activation during a working memory task [see comments]. *Nature* 386, 604–608.
- Collins, D.L., Neelin, P., Peters, T.M., Evans, A.C., 1994. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *J. Comput. Assist. Tomogr.* 18, 192–205.
- d'Esposito, M., Zarahn, E., Aguirre, G.K., 1999. Event-related functional MRI: implications for cognitive psychology. *Psychol. Bull.* 125, 155–164.
- Fox, P.T., Raichle, M.E., 1986. Focal physiological uncoupling of

- cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects. *Proc. Natl. Acad. Sci. USA* 83, 1140–1144.
- Frahm, J., Merboldt, K.D., Hanicke, W., 1993. Functional MRI of human brain activation at high spatial resolution. *Magn. Reson. Med.* 29, 139–144.
- Friston, K., Phillips, J., Chawla, D., Buchel, C., 2000. Nonlinear PCA: characterizing interactions between modes of brain activity. *Phil. Trans. R. Soc. Lond. B Biol. Sci.* 355, 135–146.
- Friston, K.J., Holmes, A., Poline, J.B., Price, C.J., Frith, C.D., 1996b. Detecting activations in PET and fMRI: levels of inference and power. *NeuroImage* 4, 223–235.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.B., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 3, 165–189.
- Friston, K.J., Price, C.J., Fletcher, P., Moore, C., Frackowiak, R.S., Dolan, R.J., 1996a. The trouble with cognitive subtraction. *NeuroImage* 4, 97–104.
- Golay, X., Pruessmann, K.P., Weiger, M., Crelier, G.R., Folkers, P.J.M., Kollias, S.S., Boesiger, P., 2000. PRESTO-SENSE: an ultrafast whole-brain fMRI technique. *Magn. Reson. Med.* 43, 779–786.
- Hajnal, J.V., Myers, R., Oatridge, A., Schwieso, J.E., Young, I.R., Bydder, G.M., 1994. Artifacts due to stimulus correlated motion in functional imaging of the brain. *Magn. Reson. Med.* 31, 283–291.
- Horwitz, B., Tagamets, M.A., McIntosh, A.R., 1999. Neural modeling, functional brain imaging, and cognition. *Trends Cogn. Sci.* 3, 91–98.
- Hu, X., Kim, S.G., 1994. Reduction of signal fluctuation in functional MRI using navigator echoes. *Magn. Reson. Med.* 31, 495–503.
- Jansma, J.M., Ramsey, N.F., Slagter, H.A., Kahn, R.S., 2001. Functional anatomical correlates of controlled and automatic processing. *J. Cogn. Neurosci.* 13, 730–743.
- Jezzard, P., Matthews, P.M., Smith, S.M. (Eds.), 2001. *Functional Magnetic Resonance Imaging: an Introductory Guide*. Oxford University Press, Oxford.
- Jezzard, P., Ramsey, N.F., 2003. Quantitative functional MRI. In: Tofts, P. (Ed.), *Quantitative Magnetic Resonance of the Brain: Measuring Change Produced by Disease*. Wiley, Indianapolis, in press.
- Liu, T.T., Frank, L.R., Wong, E.C., Buxton, R.B., 2001. Detection power, estimation efficiency, and predictability in event-related fMRI. *NeuroImage* 13, 759–773.
- Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412, 150–157.
- Moonen, C.T.W., Bandettini, P.A., 1999. *Functional MRI*. Springer Verlag, Berlin.
- Ogawa, S., Menon, R.S., Kim, S.G., Ugurbil, K., 1998. On the characteristics of functional magnetic resonance imaging of the brain. *Annu. Rev. Biophys. Biomol. Struct.* 27, 447.
- Ogawa, S., Menon, R.S., Tank, G.W., Kim, S., Merkle, H., Ellermann, J., Ugurbil, K., 1993. Functional brain mapping by blood level-dependent contrast magnetic resonance imaging. *Biophys. J.* 64, 803–812.
- Ogawa, S., Tank, G.W., Menon, R.S., Ellermann, J., Kim, S., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci. USA* 89, 5951–5955.
- Price, C.J., Friston, K.J., 1997. Cognitive conjunction: a new approach to brain activation experiments. *NeuroImage* 5, 261–270.
- Raemackers, M., Jansma, J.M., Cahn, W., Van Der Geest, J.N., Der Linden, J.A., Kahn, R.S., Ramsey, N.F., 2002. Neuronal substrate of the saccadic inhibition deficit in schizophrenia investigated with 3-dimensional event-related functional magnetic resonance imaging. *Arch. Gen. Psychiatry* 59, 313–320.
- Ramsey, N.F., 1999. Direct comparison of functional MRI and PET. In: Moonen, C.T.W., Bandettini, P.A. (Eds.), *Functional MRI*. Springer, Berlin, pp. 421–432.
- Ramsey, N.F., Kirkby, B.S., Vangelder, T., Berman, K.F., Duyn, J.H., Frank, J.A., Mattay, V.S., Vanhorn, J.D., Esposito, G., Moonen, C.T.W., Weinberger, D.R., 1996. Functional mapping of human sensorimotor cortex with 3D bold fMRI correlates highly with (H₂O) O 15 Pet rCBF. *J. Cereb. Blood Flow Metab.* 16, 755–764.
- Ramsey, N.F., Koning, H.A.M., Welles, P., Cahn, W., van der Linden, J.A., Kahn, R.S., 2002. Excessive recruitment of neural systems subserving logical reasoning in schizophrenia. *Brain*, in press.
- Ramsey, N.F., Sommer, I.E.C., Rutten, G.J., Kahn, R.S., 2001. Combined analysis of language tasks in fMRI improves assessment of hemispheric dominance for language functions in individual subjects. *NeuroImage* 13, 719–733.
- Ramsey, N.F., Van den Brink, J.S., van Muiswinkel, A.M., Folkers, P.J., Moonen, C.T., Jansma, J.M., Kahn, R.S., 1998. Phase navigator correction in 3D fMRI improves detection of brain activation: quantitative assessment with a graded motor activation procedure. *NeuroImage* 8, 240–248.
- Rutten, G.J.M., van Rijen, P.C., van Veelen, C.W.M., Ramsey, N.F., 1999. Language area localization with three-dimensional functional magnetic resonance imaging matches intrasulcal electrostimulation in Broca's area. *Ann. Neurol.* 46, 405–408.
- Shiffrin, R.M., Schneider, W., 1977. Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychol. Rev.* 84, 127–190.
- Sternberg, S., 1966. High speed scanning in human memory. *Science* 153, 652–654.
- Thevenaz, P., Ruttimann, U.E., Unser, M., 1998. A pyramid approach to sub-pixel registration based on intensity. *IEEE Trans. Med. Imag.* 7, 27–41.
- Tofts, P. (Ed.), 2003. *Quantitative Magnetic Resonance of the Brain: Measuring Change Produced by Disease*. Wiley, Indianapolis, in press.
- Vafae, M.S., Meyer, E., Marrett, S., Paus, T., Evans, A.C., Gjedde, A., 1999. Frequency-dependent changes in cerebral metabolic rate of oxygen during activation of human visual cortex. *J. Cereb. Blood Flow Metab.* 19, 272–277.
- Van Gelderen, P., Ramsey, N.F., Liu, G., Duyn, J.H., Frank, J.A., Weinberger, D.R., Moonen, C.T.W., 1995. Three dimensional functional MRI of human brain on a clinical 1.5 T scanner. *Proc. Natl. Acad. Sci. USA* 92, 6906–6910.
- van Zijl, P.C., Eleff, S.M., Ulatowski, J.A., Oja, J.M., Ulug, A.M., Traystman, R.J., Kauppinen, R.A., 1998. Quantitative assessment of blood flow, blood volume and blood oxygenation effects in functional magnetic resonance imaging. *Nat. Med.* 4, 159–167.
- Vlaardingerbroek, M.T., Den Boer, J.A., 1999. *Magnetic Resonance Imaging: Theory and Practice*. Springer Verlag, Berlin.
- Waldvogel, D., Van Gelderen, P., Muellbacher, W., Ziemann, U., Immisch, I., Hallett, M., 2000. The relative metabolic demand of inhibition and excitation. *Nature* 406, 995–998.
- Woods, R.P., Cherry, S.R., Mazziotta, J.C., 1992. Rapid automated algorithm for aligning and reslicing PET images. *J. Comput. Assist. Tomogr.* 16, 620–633.
- Worsley, K.J., 1994. Local maxima and the expected Euler characteristic of excursion sets of Chi square, *F* and *t* fields. *Adv. Appl. Prob.* 26, 13–42.
- Worsley, K.J., Friston, K.J., 1995. Analysis of fMRI time-series revisited—again. *NeuroImage* 2, 173–181.
- Yang, Y.H., Engelen, W., Xu, S., Gu, H., Silbersweig, D.A., Stern, E., 2000. Transit time, trailing time, and cerebral blood flow during brain activation: measurement using multislice, pulsed spin-labeling perfusion imaging. *Magn. Reson. Med.* 44, 680–685.
- Yang, Y.H., Glover, G.H., Vangelder, P., Mattay, V.S., Santha, A.K.S., Sexton, R.H., Ramsey, N.F., Moonen, C.T.W., Weinberger, D.R., Frank, J.A., Duyn, J.H., 1996. Fast 3D functional magnetic resonance imaging at 1.5 T with spiral acquisition. *Magn. Reson. Med.* 36, 620–626.
- Zarahn, E., Aguirre, G.K., d'Esposito, M., 1997. Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *NeuroImage* 5, 179–197.