



ELSEVIER

Speech Communication 41 (2003) 7–21

SPEECH
COMMUNICATION

www.elsevier.com/locate/specom

Functional imaging and language: A critical guide to methodology and analysis

Sophie K. Scott^{a,b,*}, Richard J.S. Wise^b

^a *Departments of Psychology and Phonetics, University College, Gower Street, WC1E 6BT London, UK*

^b *MRC Clinical Sciences Centre, Hammersmith Hospital, London, UK*

Abstract

This paper summarizes the methodology involved in functional neuroimaging, both experimental designs and data analyses. It is intended as a general introduction to the techniques and terminology involved, and aimed at speech scientists new to the area. The methods covered are positron emission tomography (PET) and functional magnetic resonance imaging (fMRI). Other imaging methods, reliant on the pattern of electrical discharges associated with neural activity, have also been used clinically and experimentally, and provide excellent temporal resolution but poor spatial resolution. It is not within the scope of this review to address these. The emphasis is on potential criticisms and problems concerning PET and fMRI, since much has already been published about the advantages, real or perceived. The strengths and weaknesses of PET and fMRI are addressed, with reference to language studies.

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

Keywords: PET; fMRI; Functional imaging terminology; Functional imaging methodology; Language

1. Introduction

1.1. Functional specificity

Early hypotheses about the function of the human brain did not address its possible role in cognition: for example, Aristotle felt that such a senseless and bloodless (post-mortem) organ could not have an important role. This position changed radically when more became known about the structure and function of the human brain through post-mortem studies.

Specifically, ‘lesion deficit’ analysis related the loss of a mental faculty in life to the distribu-

tion of a cerebral lesion found at post-mortem. Famously, Paul Broca demonstrated that the faculty of speech was lateralized to the left cerebral hemisphere, more precisely to the left inferior frontal gyrus. Building on this work, Karl Wernicke identified a role for the posterior left temporal lobe in the perception of speech, and linked the function of his area to that of Broca through a white matter tract (the arcuate fasciculus) connecting the two. This neat anatomical relation between areas of cortex involved in speech perception and production was not devoid of a certain scientific license, as it has become apparent that the arcuate fasciculus does not project to the left frontal operculum, considered to be the centre of classic Broca’s area.

Based on these observations, Lichtheim, using minimal behavioural data, often gathered indirectly

* Corresponding author. Tel.: +44-20-76795342.

E-mail address: sophie.scott@ucl.ac.uk (S.K. Scott).

and with only limited post-mortem data, constructed a model of single word perception, comprehension, repetition and production that is still used by modern neurologists to examine aphasic patients at the bedside. However, the limitations of inadequate information about the patient's behaviour in life and the delay until the patient died of natural causes to allow post-mortem examination made this technique for establishing structure–function relations in the human brain ponderous and inexact. Furthermore, the language ‘faculties’ considered were usually basic, and the models created by the early neurological aphasiologists at best only crudely captured the complex deficits of patients.

Since its introduction in the late 1960s, single-case cognitive neuropsychology has had a major impact on language research. The method relies on a detailed behavioural analysis of individual patients to identify novel, and often counterintuitive, dissociations of function. From the data collected on a limited number of intensively studied patients, sophisticated information processing models of language have been constructed: with a number of assumptions, it is argued that the data gives insight into normal language structure–function relationships based on the missing functions (‘subtracted’ by more or less focal brain pathology) observed across a variety of theoretically interesting patients. In general, however, the main emphasis of this approach has been to dissociate different cognitive stores and processes rather than to relate any one function to a particular region of the cerebral cortex.

The advent of anatomical imaging techniques in the 1970s has enabled the functional neuroanatomy of the human brain to be directly addressed by relating behaviour to the site of a lesion. The early X-ray CT scans had poor spatial resolution and were comprised of large pixels (picture elements) that made the image slices through the brain very ‘grainy’.

Modern CT scanners give resolution several orders of magnitude better. The images from magnetic resonance imaging (MRI) however, with excellent contrast between gray and white matter structures and between normal and abnormal brain, yield image slices, oriented at any desired




Fig. 1. Sagittal slice of a T1 weighted 1.5 T MRI of a normal subject. Individual sulci and gyri can be easily distinguished in the cortex. Note that the air in the sinuses shows black.

angle, that are almost indistinguishable from the photographed brain slices used to illustrate anatomical textbooks (Fig. 1). Combining anatomical neuroimaging with cognitive neuropsychological case studies would seem to have enormous potential in answering questions about the relationships between structure and function.

However, the overwhelming limitation has been the uncontrolled distribution of focal brain pathology. Cerebral infarcts, haemorrhages, tumours, infections and focal degenerations do not respect functional boundaries and are often extensive. All destructive lesions that involve the cortex also affect underlying white matter tracts to a greater or lesser extent. Although the functional deficit is usually attributed to the site of the lesion, the impact of disconnection of remote, intact cortical regions by the white matter lesion may be considerable and is unassessable except in a few instances. However, the use of anatomical imaging has had notable successes. A recent example was the localization of a very precise region, 1 cm across, in the left hemisphere that seems to be

important in articulation (Dronkers, 1996). At the other end of the spectrum, a very precise functional dissociation in the ability to process regular and irregular verbs was identified in patients, some of whom had, over time, infarcted large volumes of both cerebral hemispheres (Marslen-Wilson and Tyler, 1997): an illustration that theoretically important dissociations can be found in patients whose lesions on scans look unpromisingly ‘crude’ (i.e. large).

Therefore, functional neuroimaging techniques are the method of choice to investigate structure–function relationships: studies can be performed on normal subjects as well as patients; the techniques can explore activated regions across the whole brain; and the spatial resolution will allow functions to be attributed to individual gyri or sulci. The rest of this article discusses the principles of functional neuroimaging and the strengths and weaknesses of its application to language research.

2. Positron emission tomography

The brain is a metabolically demanding organ and receives 20% of cardiac output in the resting individual, although it comprises <2% of body weight. The distribution of the blood supply within the brain is not uniform, the grey matter (where the cell bodies of neurons are located) receiving three times as much blood flow as the white matter (which comprises the axons of neu-

rons and their supporting (glial) cells. Positron emission tomography (PET) and single photon emission tomography (SPET) provide images of the brain by tracing radio-labelled molecules that have diffused into the brain tissue from the blood vessels. The scanners estimate the concentration of local radioactivity. Fig. 2 shows an example of a structural PET scan; the lighter regions are areas of increased signal. Broadly speaking, PET can provide quantitative data (i.e. concentration of radioactivity per ml of brain tissue) whereas SPET gives qualitative images of relative concentrations of radioactivity (although it is now possible to obtain at least semi-quantitative data from SPET scans). The requirement is that the tracer molecule crosses the blood–brain barrier, that its fate in tissue can be mathematically modelled, and that radio-labelled metabolites of the original molecule do not recirculate to be taken up again by brain tissue and add to the recorded local activity to any significant extent. Using PET quantitatively, the rate of delivery of the tracer to the brain is estimated, usually by sampling the radioactivity in arterial blood. However, to avoid arterial cannulation, models have been developed that rely on measuring activity with the scanner in ‘reference’ tissue in which the tracer molecule is not specifically bound. Commonly, PET scans are obtained for glucose metabolism (using a radiolabelled metabolite of glucose, deoxyglucose, tagged with positron-emitting fluorine-18: ^{18}F FDG) or for neurotransmitter function (e.g. using 18-fluoro-dopa, ^{18}F -dopa, which is taken up by presynaptic




Fig. 2. PET scan of resting blood flow of a normal subject. The white and gray matter can be distinguished; there is greater perfusion throughout the gray matter, leading to higher counts, and these regions are lighter in colour.

terminals that have the dopa metabolite, dopamine, as the neurotransmitter).

These 'neurological' uses of PET contrast with the much more widely used method of 'activation'. Based on good evidence, it has been shown that measurements of regional cerebral blood flow (rCBF) correlate with net synaptic activity within that region. This technique is sufficiently sensitive to record changes in regional net synaptic activity in response to behavioural tasks, such as speech perception or production. At the start of a scan, a short-lived radio-labelled tracer (usually oxygen-15-labelled water, $H_2^{15}O$) that indexes rCBF is introduced into the blood by intravenous infusion. The subject lies on a table with their head in the scanner, which comprises a ring of coincidence detectors. When the radioactive molecule, ^{15}O , decays to the more stable ^{12}O form, a positron is emitted. This travels a short distance (approximately 1–2 mm) before annihilating, creating two gamma rays and a neutron. The two gamma rays travel at very close to 180° to one another and strike a pair of coincidence detectors almost simultaneously, within a time window of about 10 ns. By sampling a slice of tissue from many angles with multiple pairs of coincidence detectors (or, more cheaply but much less sensitively, by a rotating pair of detectors), the distribution of radioactivity across multiple slices of the brain is reconstructed by back projection of the recorded coincidences (corrections have to be made for estimates of the number of random coincidences). Thus the resolution of PET is dependent on the margin of error introduced by the distance the positron initially travelled before annihilation and the field-of-view of the coincidence detectors. SPET records emissions from isotopes that emit only single particles or photons, which places constraints on resolution and quantitation that are avoided in positron-emitting coincidence detection.

During a PET scan of rCBF, typically coincidence events are recorded over 40–120 s, with the actual temporal resolution being of the order of 15–20 s, the time course of rapid build-up of brain tissue radioactivity following slow bolus infusion of $H_2^{15}O$. Since the counts are obtained from tissue water, the signal reflects nutrient blood flow,

distinct from the signal obtained within functional magnetic resonance imaging (fMRI) which arises from blood within capillaries and venules. Although there are billions of neurons there are trillions of connections between neurons, and it is these synaptic connections, with the release and reuptake of neurotransmitters, that are the ultimate source of the signal in PET. A change in net synaptic activity within a brain region that can be spatially resolved by PET (i.e. the sum of activity within many tens of millions of synapses) is indexed by a change in rCBF of the order of 2–15%, depending on the brain region. It is important to realize that inhibitory synaptic activity (i.e. release of a neurotransmitter that suppresses the rate of firing of the post-synaptic neuron) consumes energy and increases local rCBF, so that a local increase in rCBF may be due to either excitatory or inhibitory synaptic activation. Thus a deactivation, i.e. a reduction of rCBF in the activation scan relative to a baseline scan, is not synonymous with inhibition, but a reflection that the baseline scan was associated with greater neural activity within that region. As even a so-called 'rest' scan is performed on a conscious, thinking subject, and 'rest' is associated with its own cognitive activations (such as stimulus-independent thoughts), it is not surprising that the reverse contrast, of baseline versus activation condition, is associated with its own pattern of distributed activations.

A very simple example of a PET design is the contrast of a condition when subjects listened to spoken nouns contrasted with a baseline condition when the subjects listened to non-speech sounds. The difference in activity seen in the contrast of the speech with the non-speech scan revealed where there was greater neural activity in the speech perception condition (e.g. Fig. 3). To improve the signal, these two conditions should be repeated as often as the design allows (the amount of radioactivity delivered in a PET study is tightly monitored and normally no more than 12–16 scans are possible).

To further improve sensitivity a number of subjects were studied (normally six or more subjects comprise one study). To compare activation both within and across subjects, certain physical adjustments have to be made to the scan data.




Fig. 3. Functional PET scan of activation from six normal subjects, rendered on slices from the MRI of a subject who most closely matches the average MRI template (MNI template) used in SPM99. The contrast shows regions more activated by speech than by SCN (from Mummery et al., 1999). The data is thresholded at $p = 0.05$ (corrected for multiple comparisons).

Since each subject will move over the course of their scanning session, even if their position is checked between scans, each subject's scans need to be spatially realigned to their first scan. The next stage, to enable comparison across subjects, is to normalize their brains into the same stereotactic space, a necessary stage since brains, like faces, vary in size and shape across individuals.

It is important to know the template being employed: a variety of methods are used, including normalizing to the brain in the (Talairach and Tournoux, 1988) brain atlas (which was constructed from the brain of a small French woman), or to an 'average' brain template, such as that constructed by the Montreal Neurological Institute, constructed from anatomical MRI scans obtained on 305 normal subjects. Methods exist for translating between different templates, to avoid discrepancies in the interpretation and comparison of data.

The final step, prior to statistical analysis, is to smooth the data. This compensates, at least in part, for interindividual differences in the distribution of gyri and sulci that persists after normalization. The amount of smoothing influences the final spatial resolution of the PET data. As the analysis of the scan data is designed to detect changes in rCBF, the changes in global blood flow (gCBF) between scans have to be removed as they may be of the order of 25% or more.

It is outside the scope of this review to address in detail the statistical methods used on functional

neuroimaging data sets, and the interested reader is directed towards far more specific articles (e.g. McColl et al., 1994; Friston et al., 1995, 1996a,b; Rabe-Hesketh et al., 1997). The unit of analysis is the voxel, a cube typically of the order of 2 mm^3 . As voxel size is smaller than the resolution of the scanner, activity within adjacent voxels is not independent. Activity at each voxel in standardized space is compared across scans (and, therefore, behavioural conditions). Although the various image analysis packages that are available use different statistical approaches, the general principle is to consider the peak difference in activity, the number of activated voxels that cluster together, and the number of these discrete clusters across the brain: all of these variables are amenable to statistical testing. Since the volume of the brain contains many tens of thousands of voxels, a correction for multiple comparisons is made, with allowance for the non-independence of activity within adjacent voxels.

The actual design of the experiment will affect the nature of the contrasts performed on the data. The results may be expressed as the specific coordinates of the peak voxels in stereotactic anatomical space, the cluster size of the individual activated regions and the number of activated regions across the volume of brain scanned (e.g. see Table 1). These results will reflect the choice of thresholds chosen by the experimenters, and Type I or Type II errors may be present. Lower thresholds are usually accepted when there is a

Table 1

Output from a statistical parametric mapping analysis (SPM99, Wellcome Department of Cognitive Neurology, London, UK) of a contrast of speech against SCN, where regions that increase in activity with increasing rates of spoken words, but not with increasing rates of SCN equivalents, have been identified (Mummery et al., 1999; see Fig. 3)

Set-level		Cluster level			Voxel-level			x (mm)	y (mm)	z (mm)
p	c	$p(c)$	K	$p(uc)$	$p(c)$	Z	$p(uc)$			
0.001	2	0.000	1889	0.000	0.000	7.15	0.000	64	-18	2
					0.000	6.81	0.000	62	-4	2
					0.000	6.08	0.000	52	-26	8
		0.000	1688	0.000	0.000	6.99	0.000	-60	-22	6
					0.002	5.49	0.000	-46	-22	6
					0.017	5.04	0.000	-58	4	-8

Voxel level significance is the probability that the difference in activity at that voxel is significant, with $p(c)$ reflecting significance corrected for whole brain volume multiple comparisons, and $p(uc)$ reflecting significance without correction for multiple comparisons. The numbers in bold are the peak voxels in that cluster, and the numbers beneath are subpeaks within that cluster. The co-ordinates (in Montreal Neurological Institute space) for each peak are given on the right hand side, expressed as x , y , z , where x expresses the right-to-left dimension (positive numbers = right), y refers to the anterior–posterior dimension (positive numbers = anterior to the anterior commissure), and z refers to the superior–inferior dimension (positive values = the anterior–posterior commissure line).

Cluster level significance reflects the probability of seeing this number of adjacent voxels being co-activated, where $p(c)$ reflects significance corrected for whole brain volume multiple comparisons, and $p(uc)$ reflects significance without correction for multiple comparisons.

Set level significance. At this level of peak and cluster level significance, this reflects the significance of the number of clusters seen, against the number of clusters that would be expected by chance.

prior hypothesis about a particular region but this is method is open to abuse with ‘hypotheses’ being constructed after the data set has been analyzed. Results are much more likely to be reproducible when an appropriate correction has been made for analysis across the whole brain volume.

2.1. Design

Design is again an issue that is covered in depth in other sources (Friston et al., 1997), and this section is a general summary. As described in the earlier section, the simplest PET design would compare two different behavioural conditions, one that it is predicted will activate the process of interest (the ‘activation’ task) and a baseline task which does not engage the process under investigation but does engage processes that are not under investigation and are also engaged by the activation task. This ‘subtraction’ methodology allows, all other things being equal, to relate an activated region to a specific brain function. Ideally the baseline task controls for all cognitive, motor and perceptual processes that are not of interest in the experimental condition, but to select

such a baseline is far from easy. The frequent result is one or more unpredicted activations, because of processing differences between tasks that were neither intended nor, even after the event, understood (although this does not prevent speculative, post-hoc rationalizations to appear in the discussion of many functional imaging publications). There also remains the problem of what to make, if anything, of the reverse contrast: brain regions more active in the baseline task are common, although with correct subtraction design they should be absent. They are often inexplicable, and further evidence that hypotheses about the processes engaged by both tasks are incomplete (Friston et al., 1995, 1996a,b).

For these reasons, simple subtraction designs are now used less frequently. There are two general approaches to improving on this design: factorial and parametric designs. In factorial designs, as in an ANOVA with behavioural data, main effects and interactions can be considered. It has also been proposed that conjunction analyses can be performed (Price and Friston, 1997; Friston et al., 1999): these permit the identification of brain regions that are common to two (or more) contrasts,

and are thus possibly implicated in one or more processes or representations common to all the activation tasks relative to their baseline tasks. This approach has been used to attempt to disentangle cognitive processes that operate at a fairly generic level. For example, brain regions have been identified that are activated when naming objects, versus viewing non-objects and saying 'OK', and also activated when reading aloud written words aloud versus viewing false font strings silently and saying 'OK' (Moore and Price, 1999). The use of conjunction analysis enabled the demonstration that there was common activation for both words and real objects (over their respective baselines) in left medial anterior fusiform gyrus and the right inferior frontal gyrus, as opposed to activation specifically to read words (on the superior temporal gyrus) and specific to real objects (bilateral occipitotemporal regions). It was assumed that the inclusion of speech responses (silently mouthed) in each condition meant that activations associated with articulation were not seen in the contrasts. Although this approach is intuitively attractive, it should be borne in mind that the theoretical basis for conjunction analyses (as opposed to revealing commonly activated regions as main effects) is not established in the statistical literature.

Parametric designs (Friston et al., 1997) are similar to correlation analysis: a parameter is varied across scans, and used as a co-variate in the analysis. A simple example is the use of varying the rates of presentation of acoustic stimuli. Thus, the contrast in Fig. 3, where speech is contrasted with non-speech (signal-correlated noise (SCN)), demonstrates the activation seen when increasing rates of single, bisyllabic concrete nouns are heard but not when increasing rates of matched SCN stimuli are heard (from Mummery et al., 1999). This effectively performs a correlation, identifying the neural activity that co-varies positively with the increasing number of speech events. Other stimuli can be used in rate paradigms, since the PET camera is sensitive to the number of events that occur during a scan (e.g. read words, Leff et al., 2000; repeated words, Wise et al., 1999; pitch, Griffiths et al., 1998) or a conceptual property of the stimuli (e.g. imageability, Wise et al., 2000), or of the subjects' responses (e.g. proportion of 'no'

responses, Scott et al., 2000). Parametric designs have the advantage that there is no baseline comparison, and regional correlations with experimenter-determined input variables or with patients' on-line behavioural responses are more compelling evidence when it comes to interpreting the results.

3. Functional magnetic resonance imaging

MRI has been a major development in the *in vivo* study of the human brain. The subject is placed in a strong and homogeneous magnetic field. For anatomical imaging the signal is derived from the protons of tissue water. With the tissue of interest placed within the magnetic field, the protons align to the field. This is then disrupted with a brief radio-frequency pulse at an angle orthogonal to the main field, and the protons realign to their original position over time. A head coil detects this disruption in the 'steady' magnetic field. Water molecules in different tissues 'relax' back into their orientation at different speeds. This recovery time is called T1 and is longer in cerebrospinal fluid (3 s) than in brain white matter (0.50 s). Varying the time between the radio pulses allows discrimination between tissues with long and short T1 values. Thus T1 acquisition sequences are used for structural MRI scans (see Fig. 1).

MRI provides clinically useful images of the brain, with good contrast between white and gray matter. The development of MRI into a method for functional neuroimaging was made possible by another factor that affects the relaxation: small local field alterations, caused by a contrast agent or paramagnetic particles. To paraphrase a lot of physics, this T2* or T2 relaxation factor (depending on the sequence—gradient echo or spin echo respectively—this refers to the way the field is created, although T2 is not normally collected in a functional scan) was developed as a tracer of cerebral blood flow, first with a contrast agent (Villringer et al., 1988) and later using the oxygenation state of the blood as a marker (Ogawa et al., 1990a,b; Turner et al., 1991), in combination with a fast acquisition method, echo planar imaging (EPI) (Mansfield, 1977). This led to

functional scans of visual stimulation (Kwong et al., 1992; Ogawa et al., 1990a,b). Deoxygenation in brain regions, due to functional activation, could thus be detected. This is the blood oxygen level dependent contrast (BOLD). There is an initial dip in the BOLD response, corresponding to the direct deoxygenation of the blood, i.e. an increase in deoxyhaemoglobin. This dip can be difficult to distinguish, and requires high magnetic fields or intrinsic optical imaging. It is spatially very closely tied to the neural activity, since it is seen in very small vessels. The later rise in the BOLD response, due to the appearance of more oxygenated blood on the venous side of the circulation within a brain region, is measured in most studies. This change in oxygen saturation arises because the balance between oxygen delivery (blood flow) and demand (oxygen metabolism) changes when net local synaptic activity briefly increases. There is a decrease in deoxyhaemoglobin, associated with a change in the MR signal, and the location of this measured change is within venules and not brain parenchyma (which degrades, to some extent, the spatial resolution of the technique). fMRI offers non-invasive scans free of the use of ionizing radiation. A single study can comprise many scans, and repeated scanning of a subject across sessions is permitted. In addition, fMRI offers good spatial resolution. Although the temporal resolution does not equate with electrophysiological techniques, activity associated with single events (e.g. the presentation of a single word) can be recorded, whereas PET averages activity across a number of repeated events.

To optimize the signal across the brain with respect to the cognitive processes studied, the time to acquire an image of the whole head (repetition time) should be as fast as possible. This can be achieved by using the fast EPI method, which also gives a reasonable signal-to-noise ratio, at the cost of a lower but still good spatial resolution (2 mm). There is generally a trade-off in the design of any study between the thickness of the slices acquired and their number: it will take longer to acquire thinner than thicker slices of the whole brain, although the resolution will be better. If a smaller region of the brain is studied, less time is needed to acquire the image. These are parameters that are

decided at the start of the study, and frequently designs are employed that do not acquire whole head volumes.

A serious problem is that EPI is susceptible to distortion and signal loss (e.g. Devlin et al., 2000). This is particularly severe around brain/bone/air interfaces where local field inhomogeneities arise naturally, and thus signal can be lost around the anterior and medial temporal lobe regions, orbitofrontal cortex, and near the ears. This worsens at higher field strengths. The signal-to-noise ratio of the BOLD signal, however, improves at higher magnet strengths.

Some correction for signal attenuation can be introduced by shimming. This procedure is normally done before a scan, or at the start of the day (depending on magnet size), and equalizes the magnetic field. This corrects for inhomogeneities caused by the introduction of a body into the scanner. This can be done to emphasize certain brain regions e.g. superior temporal plane, to improve the signal from auditory cortex (e.g. Talavage et al., 2000). However artifacts caused by air–bone interfaces are harder to correct by shimming.

For multi-subject studies, the data needs to be normalised and smoothed. Smaller spatial filters are used to smooth the data than is common in PET. Since more data can be collected from any one subject (there is no dose limitation as in PET) single subject analyses are feasible: in such cases the subjects' data can be normalized to their own structural MRI to take even more advantage of the good spatial resolution.

3.1. Design

In fMRI, data can be acquired continuously; without the scan-pause-scan cycle required in PET to allow radioactivity to decay to background levels before the next scan, more data can be acquired. Because of this continuity, the activation measured via the BOLD response can be considered to be a time series analysis, locked to the condition. As in the analysis of PET data, there are several ways of doing this, but for blocked fMRI these tend to be based around identifying BOLD variations at the voxel level that correlate with the temporal change of the experimental conditions

(epoch-related designs). At its simplest, the subject might lie in the scanner, watching a cross-hair fixation point for 30 s, followed by 30 s of reading single words presented at fixed interstimulus interval (ISI), followed by cross hairs etc., for a total of 5 min. This would give five cycles of the experimental-baseline conditions (necessary, since it is relative activation that is being measured). The analysis program then attempts to fit haemodynamically smoothed curve functions to the imaging data (at each voxel point), with the same periodicity as the condition variation. Since the BOLD response can lag the neural activity by up to 10 s, this is taken into account. As in the PET analysis, the significance of this co-variation is determined, but in terms of the percentage of signal change, not in terms of the number of counts. The number of contiguous voxels and the number of overall blobs are also analysed in order to determine the significance of the activation. The continuous collection of fMRI data and the good temporal resolution allows an event-related design, in which transient responses are identified (as opposed to epoch-related responses). This means that BOLD responses to single events (e.g. stimulus presentations, motor responses) can be estimated. The faster the time to acquire an image of the whole head, or repetition time (TR), the better the event-related resolution. If one slice is measured, TR could be very short (100 ms). The duration of TR affects the design, since it is important to ensure that the ISI of experimental events is different to, and not a multiple or divisor of, TR. This sampling of the BOLD response over a number of points in the cycle of the experimental paradigm (jittering), enables the shape of the BOLD response to be characterised more accurately. Since the BOLD response lags the neural activity by up to 10 s, the curve fitted to the data at each voxel is generally allowed a range of delay functions in the analysis. At its simplest, this could be done by testing for an early or late response (Ni et al., 2000) or by fitting a family of response curves with different onsets (as in SPM99, Wellcome Department of Cognitive Neurology). Since the BOLD lag can vary across subjects and across a single brain, this is a stage of the analysis whereby the basic assumptions need to be carefully identified.

One principled way of doing this is to determine characteristics of individual subjects' BOLD response curves using motor cortex transient responses to thumb movements, and applying this to subsequent event-related cognitive studies (Aguirre et al., 1998; Zarahn et al., 1997; see Postle et al., 2000 for an example).

Other issues, particular to fMRI, need to be considered when designing the study and analysing fMRI data. The scan data needs to be filtered to remove periodic noise, such as that due to metabolic breathing (as opposed to breathing for speech) and to blood flow, which cause both movement and blood oxygen level variation, and also to remove low frequency signal drift (due to physiological changes and scanner changes). Since the measure in fMRI is dependent on the movement of the protons within the field, the movement of the subject must be kept to an absolute minimum. A movement of 0.5 mm could lead to an apparent signal change of 40%. Unlike PET, where the subject is lying with just their head in the scanner, in an fMRI experiment the subject's whole body is inserted into the scanner and their position cannot be monitored throughout the experiment. Movement associated with the subjects' responses can cause repetitive signal change, which will resemble real activations time-locked to the responses: this can be identified and controlled for to a degree (Bullmore et al., 1999). The issue of movement artifact is most dramatic when subjects speak in the scanner, since the movement of the musculature of the face and the air in the nasal cavities mean that only the top of the brain is visible to the scanner. Many fMRI studies have made a virtue of their use of silent verbal responding.

Obviously, due to the presence of the magnet, metal objects (including keyboards, headphones) cannot be used. Subjects with pacemakers, certain cochlear implants, permanent eyelining, glitter make up or tattoos cannot be scanned, and subjects who need to wear glasses must wear contact lenses or special glasses that have no metal in them (e.g. corrective swimming goggles). Subjects' movements are reduced due to the small bore of the scanner. Visual stimuli can be presented via a mirror, or prism spectacles.

fMRI presents particular problems for studies using auditory stimuli. When the radio pulse occurs, there is a very loud sound: a 1.5 T magnet in an EPI sequence generates 110 dB SPL complex tones with low frequency energy peaks from 50 to 1000 Hz (depending on the precise sequence used) (Counter et al., 1997). This sound increases with the Tesla value of the magnet: a 4 T magnet generates a level of approximately 130 dB SPL (Counter et al., 2000). Use of ear defenders is essential at these high noise levels, and can attenuate the level by 20–30 dB. Stimuli can be presented at levels high enough for the subjects to hear against the background, but a loud, repetitive background sound is suboptimal for any hearing study. Several different ways have been developed to counter this, including sound cancellation and wrapping the inside of the scanner and the subject in foam, in an attempt to reduce the transmission of sound by bone conduction (Ravicz and Melcher, 2001). An elegant development, known as sparse scanning, capitalizes on the delay inherent in the BOLD response (Hall et al., 1999). The auditory stimuli are presented in a silent period, followed by the onset of a cycle of scanning (and associated noise). This ensures that the BOLD response measured during scanning is associated with the stimuli, not the scanner noise. Likewise, during the silent period, when the stimuli are presented, the BOLD response due to the scanner noise can decay. This has the disadvantage of fewer scans per session, reducing statistical power. This same basic principle can be applied to the generation of overt verbal responses, i.e. when the subject speaks. The subjects' verbal responses are produced, and scanning immediately follows this, when any movement artifact due to speaking has ceased. Alternatively, some groups have used lower Tesla magnets and overt speech with head control, and have seen activation during speech (Kircher et al., 2000). Direct comparison with PET studies of the same task could reveal the effects of the movement of muscles and air on the BOLD signal. This study did not demonstrate effects in the anterior insula, a region shown to be active during speech production in a repetition PET study (Wise et al., 1999) or other regions shown in PET to be activated in propositional speech (Blank et al., 2001).

4. General issues in functional imaging

4.1. *Control/baseline tasks*

The selection of a contrasting baseline against which to examine the relative blood flow changes has immense implications for the pattern of observed activation. If brain regions are activated in both conditions, then they will be subtracted out of the resulting pattern of blood flow variation. On occasions this is desirable, on others not. A 'rest' condition was frequently used in early studies, where subjects were asked to 'empty their minds'. Any introspection on this sort of task reveals that this is an extremely activating condition, and an analysis of studies where this has been used as a baseline (Binder et al., 1999) has revealed that the cognitive processing during rest (related to stimulus-independent thoughts) has a distinct neural pattern. This, in turn, means that these brain regions are subtracted from the experimental condition, and are thus thought to play no part in that task. These issues can also arise with a non-rest baseline: the use of non-words in a passive reading study tends to reveal no difference between these and real words. Possibly, the reading system processes the non-words as far as possible, even to a lexical level, so that they are processed to the same degree as real words (Price et al., 1996). In contrast, baseline conditions that do not control for a sufficient number of features that are not of interest in the experimental condition can mean that much of the activation seen is a result of these properties of the task. For example, in an elegant study of autobiographical memories, subjects had previously been given questionnaires about their lives, which were then used to form true or false questions for the PET scans (Maguire and Mummery, 1999). The baseline task used was a decision about the number of syllables in a string of function words. This baseline task thus controls for the production of a motor response, the presentation of auditory stimuli, and a decision being performed on the heard input that cannot be made before the whole string has been processed. It does not control for the fact that the autobiographical questions and memories involve the processing of lexical and semantic structures, as well as general

memory retrieval processes. The authors in this case have very precise hypotheses about the role of particular brain regions in autobiographical memory, which this study is well able to address, but their activation patterns reveal additional cortical regions, which *may* be due to lexical aspects of the task. Similarly, response variation across scans can lead to activation differences which are due to an aspect of the task that is of less interest: if two tasks differ in their difficulty, then this can lead to activation differences. Thus, reaction time variation across conditions might generate specific activation differences (which could of course, be used as a co-variate of specific interest).

This criticism does not mean that all functional imaging studies that involve some form of baseline comparison are invalid. Careful analysis of the tasks used and pilot testing can help address many of these points, providing the caveats regarding the assumption of ‘cognitive insertion’ (Friston et al., 1996a,b) are followed.

4.2. Use of overt responses. Are automatic or controlled processes being studied?

The other important issue in addition to the nature of the control conditions, is the nature of the experimental task. There is a distinct difference in the types of approaches used. One principal difference is between so-called ‘active’ and ‘passive’ tasks. Due to an apparent concern over whether or not the subject is doing what they have been asked to do, or whether they are sleeping or imagining abstract art, some studies have used explicit tasks (e.g. phoneme monitoring) when studying ‘automatic’ processes, (e.g. speech perception). Is the activation seen related to normal speech perception, or to the metalinguistic task? The issue of controlled versus automatic processing (Schneider and Shiffrin, 1977; Shiffrin and Schneider, 1977) is key. If the cognitive process(es) under consideration are themselves controlled, executive processes (e.g. working memory), or utilize them to some degree (e.g. memory retrieval), then the performance of the behavioural task associated with this is appropriate in scanning (e.g. Maguire and Mummery, 1999). If the process is an over-learned, automatic one or the representations

involved are obligatorily activated, then the use of an explicit task may be more problematic. It could be theoretically difficult to dissociate between activation due to the obligatory processing or due to more controlled processing, or some complex interaction between the two. Thus in semantic processing tasks medial prefrontal cortex is frequently activated (e.g. Mummery et al., 1998; Binder et al., 1997), but only when the subjects are required to make explicit decisions about the stimuli. These regions are not seen when the subjects ‘passively’ view or listen to meaningful stimuli (e.g. Mummery et al., 1999). The medial prefrontal activation is more likely to be due to the more generic aspects of assessing and responding to meaningful stimuli than to stored semantic memory/semantic representations. Indeed, patients with medial prefrontal lesions do not lose semantic knowledge. In contrast, several studies have used tasks that actively disguise from the subject what the point of the study is, to facilitate interpretation of any results seen as being due to automatic, obligatory processing. Studies of basic processing of emotional stimuli (e.g. Morris et al., 1996) have successfully used a non-emotional sex decision task, to focus the subjects’ attention upon different featural information in the faces. These studies are usefully complemented by functional imaging studies that contrast an explicit task (emotional decision) with an implicit task (sex decision) and demonstrate common areas of activation, thus emphasising the obligatory nature of the activation (Gorno-Tempini et al., 2001). Other papers have expressed in more detail the way that the task used in activation studies can influence the results seen (Poeppel, 1996), and this paper will not repeat these criticisms. In terms of useful, cognitively relevant functional imaging studies, we would argue for a parsimonious approach to experimental design: the data collected is going to be affected by the task administered, so we would press for a task that emphasizes the cognitive construct of interest.

4.3. Global blood flow changes

Most statistical packages for analysing PET data explicitly identify rCBF. Thus large, global

blood flow variations (which can contribute to low frequency drift in fMRI) are controlled out in the analysis. This procedure, however, makes the implicit assumption that the changes correlated to the neural activity will be small. If there are large regional changes (e.g. one hemisphere greater than the other), then this will not be detected and indeed will be controlled out of the data.

4.4. Big signal in one condition can distort other, smaller activations

Large signal changes, such as are seen when subjects make motor responses, can elevate the signal variation in PET sufficiently for activation from smaller signal changes to become harder to detect. This can, unfortunately, affect changes associated with purely cognitive processing. In fMRI a similar effect can be found, whereby a small BOLD response can be affected adversely by the presence of large activations elsewhere in the brain. This is again worth considering when designing a study, especially the use of a motor response when, possibly, a passive task might be predicted to activate the processes of interest.

4.5. Brain activation does not distinguish between excitation or inhibition

This is an important issue when trying to model functional imaging data or even ascribe basic properties to the activations seen. Since we are always measuring such data at the level of populations of millions of neurones rather than individual neurones, and since even then the valency of the activity cannot be labelled, this issue is not yet resolvable. The one way to deal with it is to use careful reference to ‘activations’ rather than excitations, and to be particularly careful when addressing ‘deactivations’.

5. PET versus fMRI

fMRI is growing in popularity. Compared to PET, it has many advantages. It is faster in collecting data, since there are no gaps between scans. There is no injection, which subjects prefer and

which means that a medical doctor need not be present. There is no radioactivity; while the amounts involved in a PET scan are trivially small, it remains a limit on any PET scan since the total dose given is strictly monitored. PET involves, in addition to a clinician, a radiographer and a chemist to run the camera and the tracer synthesis respectively, as well as a cyclotron to generate the tracer. As it is non-invasive, fMRI attracts less ethical concern. fMRI can be administered by a radiographer or trained individual, although maintenance of the machine is important, as is the role of a physicist. fMRI therefore has lower general overhead costs than PET.

Since theoretically limitless scans can be performed, fMRI is more powerful than PET: if a signal is not seen in a PET study it might be due to a lack of sensitivity, since the number of overall scans is limited. Conversely, since the PET signal relies on the number of counts measured, if the counts are high enough then a signal will be measured. The signal in fMRI is a physiological measure and thus generally smaller. Indeed in some individuals it can be very hard to determine (subjects are then labelled as ‘non-responders’).

fMRI has the spatial resolution to image structures in the brain stem, a very attractive prospect in studies of hearing. Unlike cortical regions, however, many brain stem structures move about due to cardiovascular pulsing. To enable the imaging of these structures, cardiac gating is used to image these structures at the same point in the cardiac cycle (Guimaraes et al., 1998). PET does not permit cardiac gating and the spatial resolution is too poor for subcortical regions to be frequently or reliably seen.

As mentioned earlier, event-related fMRI can be used to model transient neural responses, albeit over a much longer time scale than event-related scalp potentials or MEG. In contrast, PET sums neural activity over 40–50 s. While most of this activity is associated with the onset of the scan (the head curve), this still places a wide temporal window on the data, giving poor time resolution. However if the features under investigation are kept constant over the scan, and there is no overt fatiguing or boring of the subject, then a well-designed PET study (or blocked fMRI study) will

still deliver an effective picture of the processes and representations involved in the task. Furthermore, parametric variables associated with the subjects' behaviour (e.g. mean reaction time or percentage response class) can be used to see how activations are modified across conditions (Scott et al., 2000). Parametric uses of fMRI can be problematic, since the BOLD response may not be linear with respect to blood flow/neural activity, and may saturate (Dhankhar et al., 1997).

PET gives a signal that is constant over the whole brain (with the exception of small brainstem nuclei). The number of scans is limited, but if there are enough counts, then the activation is measurable (this does not mean it is necessarily meaningful!). The signal loss that is seen in EPI sequences in fMRI means this assumption cannot currently be made about fMRI. The signal loss in the anterior temporal lobes is of particular relevance to speech and language studies: while several studies have identified a role for anterior temporal lobe regions in semantic tasks, signals are rarely seen here in fMRI studies. Scientifically it is hard to determine whether this is due to signal loss or a lack of involvement of this region in the task. Comparative PET and fMRI studies (Devlin et al., 2000) have indicated that it is due to signal loss, and that no matter how much the threshold is dropped, the anterior temporal lobe regions show no activation, unlike the activity seen in the same task in PET. Different acquisition sequences and optimised shimming are techniques that may ameliorate things, but this proviso should be borne in mind when assessing fMRI studies that conflict with PET studies.

Finally, although PET has a poorer potential spatial resolution than fMRI, the signal from PET is detected within neural tissue, whereas in fMRI it is detected in the venous compartment. There is the danger, then, that the good spatial resolution of fMRI may be directed at vascular, not neural, tissue. This is principally a consequence of the use of the 'late' BOLD increase as a measure: the small 'initial dip' is derived from much smaller vessels and thus is nearer the initial neural activity. More powerful magnets (up to 9 T) have the potential to make reliable measurements of the initial dip.

6. Conclusions

The aim of this paper was to describe the methodology and terminology of functional neuroimaging with PET and fMRI, consider advantages and disadvantages of each, and particular issues that should be considered in the design and interpretation of any functional imaging study. The conclusions are coloured by the authors' perspectives, since one's topic of interest prejudices one's choice of technique. Thus visual scientists are unhampered in their research by the loud noise and dramatic loss of signal in the temporal lobes in fMRI, while for those who wish to study hearing and speech, the constraints of sparse sampling and loss of signal in relevant areas means PET has many advantages. For those who work in a very temporally defined topic, e.g. priming, PET studies are possible but event-related fMRI would seem appropriate. Alternatively, if a parametric design is used, removing the need for a baseline state, then PET has some advantages over fMRI due to the linearity of the response with neural activity (although the non-linear characteristics of fMRI responses can be characterised, Binder et al., 1994). In more general terms, it is important that the non-scanning scientific community understands some of the problems (and solutions) inherent in these approaches, the assumptions that are made in the analyses, and most of all the importance of the study design. Only then can functional imaging studies provide a useful source of converging information in the literature.

Acknowledgements

SKS is a Wellcome Career Development fellow, and RJSW is a Wellcome Senior Clinical fellow. We would like to thank Toby Scott and Steve Williams for the structural MRI shown in Fig. 1, and Lindsay Wickham for help with the structural PET image in Fig. 2.

References

- Aguirre, G.K., Zarahn, E., D'Esposito, M., 1998. The variability of human, BOLD hemodynamic responses. *Neuroimage* 8, 360–369.

- Binder, J.R., Frost, J.A., Hammeke, T.A., Bellgowan, P.S., Rao, S.M., Cox, R.W., 1999. Conceptual processing during the conscious resting state. A functional MRI study. *J. Cogn. Neurosci.* 11 (1), 80–95.
- Binder, J.R., Frost, J.A., Hammeke, T.A., Cox, R.W., Rao, S.M., Prieto, T., 1997. Human brain language areas identified by functional magnetic resonance imaging. *J. Neurosci.* 17, 353–362.
- Binder, J.R., Rao, S.M., Hammeke, T.A., Frost, J.A., Bandettini, P.A., Hyde, J.S., 1994. Effects of stimulus rate on signal response during functional magnetic resonance imaging of auditory cortex. *Brain Res. Cogn. Brain Res.* 2 (1), 31–38.
- Blank, S.C., Scott, S., Wise, R., 2001. Neural systems involved in propositional and non-propositional speech. *Neuroimage* 13 (6), S509–S509, Part 2, Suppl. S.
- Bullmore, E.T., Brammer, M.J., Rabe-Hesketh, S., Curtis, V.A., Morris, R.G., Williams, S.C., Sharma, T., McGuire, P.K., 1999. Methods for diagnosis and treatment of stimulus-correlated motion in generic brain activation studies using fMRI. *Hum. Brain Mapp.* 7 (1), 38–48.
- Counter, S.A., Olofsson, A., Borg, E., Bjelke, B., Haggstrom, A., Grahn, H.F., 2000. Analysis of magnetic resonance imaging acoustic noise generated by a 4.7 T experimental system. *Acta Otolaryngol.* 120 (6), 739–743.
- Counter, S.A., Olofsson, A., Grahn, H.F., Borg, E., 1997. MRI acoustic noise: sound pressure and frequency analysis. *J. Magn. Reson. Imaging* 7 (3), 606–611.
- Devlin, J.T., Russell, R.P., Davis, M.H., Price, C.J., Wilson, J., Moss, H.E., Matthews, P.M., Tyler, L.K., 2000. Susceptibility-induced loss of signal: comparing PET and fMRI on a semantic task. *Neuroimage* 11 (6, Part 1), 589–600.
- Dhankhar, A., Wexler, B.E., Fulbright, R.K., Halwes, T., Blamire, A.M., Shulman, R.G., 1997. Functional magnetic resonance imaging assessment of the human brain auditory cortex response to increasing word presentation rates. *J. Neurophysiol.* 77 (1), 476–483.
- Dronkers, N.F., 1996. A new brain region for coordinating speech articulation. *Nature* 384 (6605), 159–161.
- Friston, K.J., Holmes, A., Poline, J.B., Price, C.J., Frith, C.D., 1996a. Detecting activations in PET and fMRI: levels of inference and power. *Neuroimage* 4 (3, Part 1), 223–235.
- Friston, K.J., Holmes, A.P., Price, C.J., Buchel, C., Worsley, K.J., 1999. Multisubject fMRI studies and conjunction analyses. *Neuroimage* 10 (4), 385–396.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.P., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2, 189–210.
- Friston, K.J., Price, C.J., Buchel, C., Frackowiak, R.S.J., 1997. A taxonomy of study design. In: Frackowiak, R.S.J., Friston, K.J., Frith, C.D., Dolan, R.J., Mazziotta, J.C. (Eds.), *Human Brain Function*. Academic Press, San Diego, pp. 141–159.
- Friston, K.J., Price, C.J., Fletcher, P., Moore, C., Frackowiak, R.S.J., Dolan, R.J., 1996b. The trouble with cognitive subtraction. *Neuroimage* 4 (2), 97–104.
- Gorno-Tempini, M.L., Pradelli, S., Serafini, M., Pagnoni, G., Baraldi, P., Porro, C., Nicoletti, R., Umita, C., Nichelli, P., 2001. Explicit and incidental facial expression processing: An fMRI study. *Neuroimage* 14 (2), 465–473.
- Griffiths, T.D., Buchel, C., Frackowiak, R.S., Patterson, R.D., 1998. Analysis of temporal structure in sound by the human brain. *Nat. Neurosci.* Sep. 15, 422–427.
- Guimaraes, A.R., Melcher, J.R., Talavage, T.M., Baker, J.R., Ledden, P., Rosen, B.R., Kiang, N.Y., Fullerton, B.C., Weisskoff, R.M., 1998. Imaging subcortical auditory activity in humans. *Hum. Brain Mapp.* 6 (1), 33–41.
- Hall, D.A., Haggard, M.P., Akeroyd, M.A., Palmer, A.R., Summerfield, A.Q., Elliott, M.R., Gurney, E.M., Bowtell, R.W., 1999. “Sparse” temporal sampling in auditory fMRI. *Hum. Brain Mapp.* 7 (3), 213–223.
- Kircher, T.T., Brammer, M.J., Williams, S.C., McGuire, P.K., 2000. Lexical retrieval during fluent speech production: an fMRI study. *Neuroreport* 11 (18), 4093–4096.
- Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., Cheng, H.M., Brady, T.J., Rosen, B.R., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl. Acad. Sci. USA* 89 (12), 5675–5679.
- Leff, A.P., Scott, S.K., Crewes, H., Howard, D., Cowie, A., Wise, R.J.S., 2000. Impaired reading in patients with right hemianopia. *Ann. Neurol.* 47, 171–178.
- Maguire, E.A., Mummery, C.J., 1999. Differential modulation of a common memory retrieval network revealed by positron emission tomography. *Hippocampus* 9 (1), 54–61.
- Marslen-Wilson, W.D., Tyler, L.K., 1997. Dissociating types of mental computation. *Nature* 387 (6633), 592–594.
- McColl, J.H., Holmes, A.P., Ford, I., 1994. Statistical methods in neuroimaging with particular application to emission tomography. *Statist. Meth. Med. Res.* 3 (1), 63–86.
- Mansfield, P., 1977. Multiplanar image formation using NMR spin echoes. *J. Phys. C* 10, 55–58.
- Moore, C.J., Price, C.J., 1999. Three distinct ventral occipito-temporal regions for reading and object naming. *Neuroimage* 10 (2), 181–192.
- Morris, J.S., Frith, C.D., Perrett, D.I., Rowland, D., Young, A.W., Calder, A.J., Dolan, R.J., 1996. A differential neural response in the human amygdala to fearful and happy facial expressions. *Nature* 383 (6603), 812–815.
- Mummery, C.J., Ashburner, J., Scott, S.K., Wise, R.J.S., 1999. Functional neuroimaging of speech perception in six normal and two aphasic patients. *J. Acoust. Soc. Amer.* 106 (1), 449–457.
- Mummery, C.J., Patterson, K., Hodges, J.R., Price, C.J., 1998. Functional neuroanatomy of the semantic system: divisible by what? *J. Cogn. Neurosci.* 10 (6), 766–777.
- Ni, W., Constable, R.T., Mencl, W.E., Pugh, K.R., Fulbright, R.K., Shaywitz, S.E., Shaywitz, B.A., Gore, J.C., Shankweiler, D., 2000. An event-related neuroimaging study distinguishing form and content in sentence processing. *J. Cogn. Neurosci.* 12 (1), 120–133.

- Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W., 1990a. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl. Acad. Sci. USA* 87 (24), 9868–9872.
- Ogawa, S., Lee, T.M., Nayak, A.S., Glynn, P., 1990b. Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn. Reson. Med.* 14 (1), 68–78.
- Poeppel, D., 1996. A critical review of PET studies of phonological processing. *Brain Lang.* 55 (3), 317–351.
- Postle, B.R., Zarahn, E., D’Esposito, M., 2000. Using event-related fMRI to assess delay-period activity during performance of spatial and nonspatial working memory tasks. *Brain Res. Protocols* 5, 57–66.
- Price, C.J., Friston, K.J., 1997. Cognitive conjunction: a new approach to brain activation experiments. *Neuroimage* 5 (4, Part 1), 261–270.
- Price, C.J., Wise, R.J., Frackowiak, R.S., 1996. Demonstrating the implicit processing of visually presented words and pseudowords. *Cereb. Cortex* 6 (1), 62–70.
- Rabe-Hesketh, S., Bullmore, E.T., Brammer, M.J., 1997. The analysis of functional magnetic resonance images. *Statist. Meth. Med. Res.* 6 (3), 15–37.
- Ravicz, M.E., Melcher, J.R., 2001. Isolating the auditory system from acoustic noise during functional magnetic resonance imaging: examination of noise conduction through the ear canal, head, and body. *J. Acoust. Soc. Amer.* 109 (1), 216–231.
- Schneider, W., Shiffrin, R.M., 1977. Controlled and automatic human information processing II: Detection, search and attention. *Psychol. Rev.* 84, 1–66.
- Scott, S.K., Holmes, A., Friston, K.J., Wise, R.J.S., 2000. A fronto-thalamic system for response choice. *Neuroreport* 11 (7), 1523–1527.
- 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.
- Talairach, J., Tournoux, P., 1988. *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme-Verlag, Stuttgart.
- Talavage, T.M., Ledden, P.J., Benson, R.R., Rosen, B.R., Melcher, J.R., 2000. Frequency-dependent responses exhibited by multiple regions in human auditory cortex. *Hear. Res.* 150 (1–2), 225–244.
- Turner, R., Le Bihan, D., Moonen, C.T., Despres, D., Frank, J., 1991. Echo-planar time course MRI of cat brain oxygenation changes. *Magn. Reson. Med.* 22 (1), 159–166.
- Villringer, A., Rosen, B.R., Belliveau, J.W., Ackerman, J.L., Lauffer, R.B., Buxton, R.B., Chao, Y.S., Wedeen, V.J., Brady, T.J., 1988. Dynamic imaging with lanthanide chelates in normal brain: contrast due to magnetic susceptibility effects. *Magn. Reson. Med.* 6 (2), 164–174.
- Wise, R.J.S., Greene, J., Büchel, C., Scott, S.K., 1999. Brain systems for word perception and articulation. *The Lancet* 353 (9158), 1057–1061.
- Wise, R.J.S., Howard, D., Mummery, C.J., Fletcher, P., Leff, A., Büchel, C., Scott, S.K., 2000. Non-imageability and the temporal lobes. *Neuropsychologia* 38 (7), 985–994.
- Zarahn, E., Aguirre, G.K., D’Esposito, M., 1997. A trial-based experimental design for fMRI. *Neuroimage* 6, 122–138.