CHAPTER 5

Magnetic Resonance Imaging of Brain Function

Stuart Clare

Department of Clinical Neurology
Centre for Functional Magnetic Resonance Imaging of the Brain
John Radcliffe Hospital
University of Oxford
Headington, Oxford OX39DU
United Kingdom

I. Update

The scope of human behavior that has been studied using fMRI and the range of uses to which that information has been put continues to grow and grow. From investigating new drug therapies, to informing marketing executives, many people want to use functional imaging to gain more insight into the brain. However, good experimental design and practice still lies at the heart of producing useful insights rather than just pretty pictures.

MRI systems operating at 3 T, offering increased sensitivity to brain activation, are now increasingly commonplace in hospital radiology departments. There are

DOI: 10.1016/B978-0-12-375043-3.00005-6

also an ever-increasing number of 7 T systems now available. This push to higher field strength, offering the possibility of localizing activations to a very high spatial extent (<1 mm), is also complemented by a number of other hardware developments such as receive RF coils with 32 or more independent channels, boosting sensitivity by factors of 4 or more (Wiggins *et al.*, 2006).

While the basic methods outlined in this chapter remain the same, the desire for a more quantifiable and specific measure of brain activity has lead to an increase in methods for fMRI. The use of MRI-based quantitative cerebral blood flow (CBF) measurements (Liu *et al.*, 2007) and cerebral blood volume measurements (Lu *et al.*, 2003) has increased, giving the ability to tease apart the different components of the hemodynamic response to neural activation. Combining measurements of CBF with traditional BOLD detection of the hemodynamic response allows comparison in activation levels in subjects between scans on different days (Leontiev and Buxton, 2007).

One other area that has seen an increase over the recent years is the development of targeted "molecular" contrast agents. These compounds contain an MR image-enhancing molecule (such as gadolinium or ultra small particles of iron oxide) bound to other molecules that will in turn bind to specific targets in the body (Frank *et al.*, 2004). While use of these compounds in humans to detect brain activity is not yet possible, this area is sure to increase over the coming years.

A. Hints and Tips for Performing fMRI

- Design of the stimulus paradigm is crucial to the success of the experiment: choose baseline and task conditions that tease apart the effect you are interested in.
- Be sure your stimulus presentation equipment is MR compatible and does not introduce radio waves into the MR environment.
- Estimate your effect size, with either literature values or pilot experiments, to ensure that you have enough trials or blocks in your main experiment.
- Understand what the fMRI analysis packages are doing "under the hood": they may offer simple interfaces to get your data analyzed, but you will get much more from your data if you understand the statistics they are computing.
- Analyze your data as you go along, rather than waiting to the end of the study: many flaws in the design can be picked up at an early stage.

II. Introduction

The rapid development of methods for noninvasive brain mapping, particularly over the past decade, has led to exciting advances in our understanding of the human brain. Foremost in these methodologies is the technique of functional

magnetic resonance imaging (fMRI). Utilizing the intrinsic magnetic properties of the blood, it is possible to identify the brain region associated with a specific sensory, motor, or even cognitive task to a high spatial precision. Unlike positron emission tomography (PET), MRI does not use radioactively labeled compounds and is essentially noninvasive and safe for repeat studies. Although fMRI does not share the temporal resolution of electroencephalography (EEG) or magnetoencephalography (MEG), it does have a spatial resolution of millimeters, and the most recent experiments suggest that it may be able to detect activations at the level of the cortical layers (Silva and Koretsky, 2002).

While fMRI is a complex methodology, developments in recent years, particularly by the scanner manufacturers and other commercial and academic groups, have meant that the tools for fMRI, while expensive, are more commonly available. In particular, the available software for both fMRI stimulus presentation and image analysis are highly sophisticated and user-friendly. This means that human fMRI is now achievable by research groups without dedicated physics and image analysis support.

III. Experimental Procedures

A. Overview of Methods

Functional MRI relies on detecting the small changes in image brightness on MRI scans, associated with the hemodynamic changes in the brain, in response to a specific external stimulus or ''internal'' cognitive process (Belliveau *et al.*, 1991). Carrying out an fMRI experiment therefore consists of three primary components: presenting or otherwise cueing the stimulus, scanning the brain rapidly using MRI, and analyzing the MRI scans to detect changes in image intensity.

While the subject is being scanned repeatedly, ideally covering the whole brain every 3 s, a stimulus or cue is presented to them. This could be a simple visual stimulus, such as a flashing light, or a more complex stimulus, such as a list of numbers to remember and recall at some point. This stimulus is repeated a number of times to build up confidence in determining the brain regions that are truly responding to the stimulus, while averaging out other ''random'' brain processes. The resulting images are then analyzed using computer software to detect those regions in the images that show a significantly time-locked response to the stimulus. These regions are displayed as bright ''activations'' overlaid on a conventional brain scan or brain atlas, such as shown in Fig. 1.

Human fMRI can be performed using most modern MRI scanners found in radiology departments of many hospitals, operating at a field strength of 1.5 T. However, in recent years, the desire to detect these small hemodynamic changes has led to the successful use of field strengths of 3–4 T in research MRI systems. A small number of research sites worldwide are experimenting with the use of even

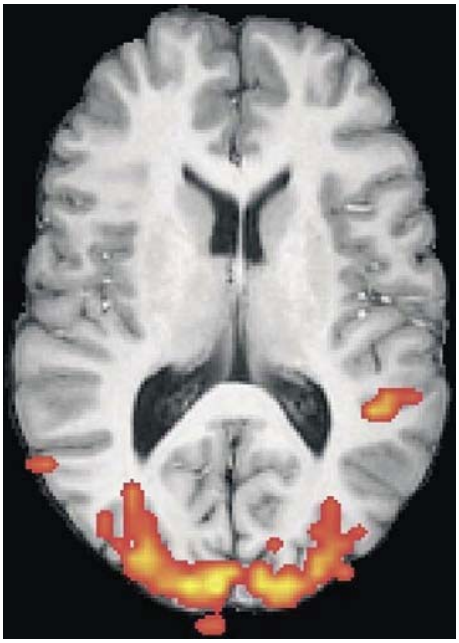

Fig. 1 Example of an fMRI result showing activation in the visual areas resulting from the subject looking at a contrast reversing checkerboard stimulus.

higher field strengths for human imaging, the highest currently being 9 T; however, the difficulties with producing high-quality human brain images diminish the benefits in signal detection offered at such high field strengths (Chan, 2002).

Functional MRI of nonhuman primates and rodents is covered in more detail in other chapters in this volume. For the rest of this chapter, the use of human subjects is assumed; however, much of the underlying methodology is the same for human or animal subjects.

B. Blood Oxygenation and MRI

The oxygen carrier in blood, hemoglobin, gets its red color from the iron molecule, which forms the binding site of oxygen. This presence of iron in the molecule makes it magnetically sensitive. In its deoxygenated state, hemoglobin displays paramagnetic properties, meaning that the local magnetic field is increased in the presence of an external magnetic field (Thulborn *et al.*, 1982). In contrast, oxygenated hemoglobin is diamagnetic and has little effect on the local magnetic field. The effect of these local changes in magnetic field can be detected in type of MR scan that is said to be T_2* weighted. On a T_2*-weighted MR image, a pixel that contains predominantly oxygenated hemoglobin will appear brighter

than a pixel that contains predominantly oxygenated hemoglobin. This form of image contrast in MRI is termed blood oxygenation level-dependent (BOLD) contrast (Ogawa *et al.*, 1990).

Although it is possible to quantify the change in hemoglobin oxygenation using MRI, it is typical in fMRI to just detect the relative signal changes. The exact link between neuronal firing and the BOLD signal change that is detected is complex and not entirely understood. Upon the metabolic demand that synaptic activity produces, oxygen is removed from the blood and the concentration of deoxyhemoglobin increases. This would result in a small dip in the MR image intensity, which is sometimes observed in fMRI experiments. However, the much stronger effect is the large increase in image intensity that follows, peaking at about 6 s after the neural activity. This represents a large increase in the concentration of oxyhemoglobin, far greater than its resting state level. This results from a large increase in the local blood flow rate and local blood volume due to capillary expansion. While this apparently excessive overcompensation in oxygen delivery was initially a puzzling result, recent physiological models have demonstrated the need for such increases to maintain the necessary oxygen delivery rate to the mitochondria (Buxton *et al.*, 1998).

Following this peak in local oxygenation level, the signal returns back toward its baseline state, but is often observed to decrease still further (known as the undershoot), as the relative contribution of oxygen extraction, blood flow, and blood volume return to their baseline state. Most fMRI "activations" are detected as regions that display the large increase in signal, peaking several seconds after the stimulus. In fact, the presence of either the initial dip in signal or the poststimulus undershoot is not detected in many experiments, as it seems to vary by brain location and can be obscured by image noise, particularly at lower field strengths (Buxton, 2001).

The complex physiological processes that give rise to the signal changes observed in fMRI mean that there are a number of reasons for caution in interpreting the experimental results. First, and most obviously, is that the signal arises from hemodynamic effects and not directly from neural activity. The location of peak BOLD effects could indeed be some distance from the site of the activating neurons. This is particularly the case when imaging using methods that are more sensitive to the signal from the large draining veins, which could be centimeters from the actual activation site (Kim *et al.*, 1994). To guard against this particular problem it is advisable to interpret the spatial location of fMRI activations with reference to a map of veins (as can be acquired easily using MRI or from standard atlases).

Second, thought must be given to the time characteristics of the fMRI response. The timing of peak activation relative to the signal needs to be taken into account when analyzing the images, as it may vary over brain regions. Care must also be exercised in interpreting any differences in signal timing as representing temporal differences in the onset of neural activation. While it is certainly possible to obtain an indication of neural timing from the BOLD response, a lag in signal in one region relative to another does not necessarily mean a difference in neural timing

and may just represent a difference in blood supply in those regions (Miezin *et al.*, 2000). The large delay after activation before the signal returns to baseline also has implications for experimental design and is discussed in more detail later.

It should also be noted here that BOLD contrast is not the only way to perform fMRI, although it is by far the easiest and most commonly used method. MR images can also be made sensitive to the blood flow rate alone. Such experiments suffer from low signal-to-noise ratio (SNR) and do not have the same temporal resolution as BOLD fMRI, but are very useful in interpreting the BOLD signal changes and may turn out to be more spatially specific than BOLD (Liu *et al.*, 2002).

C. Rapid MRI for fMRI

MR images are essentially maps of water content in the brain generated by the NMR phenomenon that certain atoms, when placed in a strong magnetic field, will absorb and emit radiofrequency energy at a specific frequency dictated by the strength of that applied magnetic field. However, it is straightforward in MRI to modulate this basic signal such that the intensity in a region of the image is not just dependent on water content but also on the local structural environment or other physiological parameters. Examples of this are the T_1- and T_2-weighted images often used in clinical diagnosis, where, for example, the region of cell damage produced by a stroke can be seen very clearly. Another example of this is the so-called T_2*-weighted image, which is highly sensitive to the local magnetic environment and is particularly sensitive to the BOLD effect (Haacke *et al.*, 1999).

A typical diagnostic MR image is optimized for spatial resolution and contrast to detect the particular pathology of interest. This typically means a scan time of several minutes. Although it is possible to do fMRI with a scan lasting minutes, the requirement of needing to keep a discrete set of brain regions active for such a time makes this impractical for anything other than the simplest experiment.

Speeding up the scanning process requires not only very high-performance scanner hardware, but comes at a cost to image quality. However, with the advent of fMRI, most modern scanners have the technological capability to run very fast imaging methods.

The most common fast imaging method used for fMRI is echo planar imaging (EPI). This method is able to collect data from a single "slice" through the brain in less than 100 ms, meaning that, at coarse resolution, it is possible to scan the whole brain in around 3 s (Stehling *et al.*, 1991). EPI also has the benefit of being inherently a T_2*-weighted sequence, so it is ideally suited to BOLD fMRI.

The largest drawback to using EPI is that the images often contain image distortion and signal loss. An example of this is shown in Fig. 2. The air-filled sinuses that sit below the frontal lobes of the brain cause a "hole" to appear in the EPI images, compared to the standard MRI, and the frontal lobes also appear

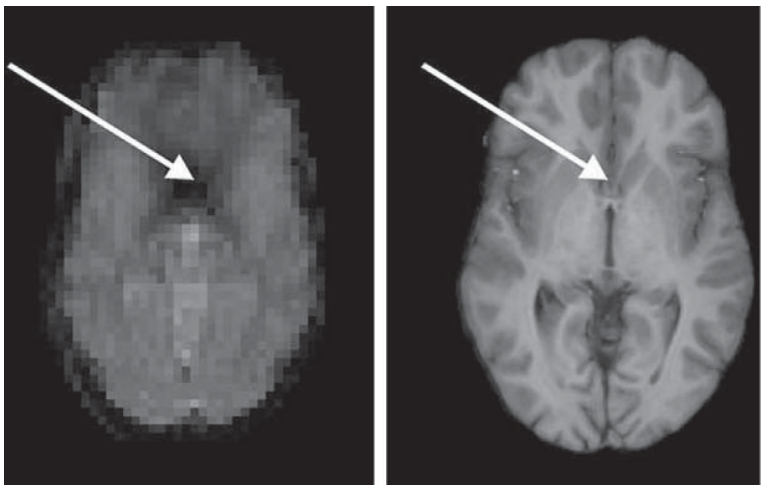

Fig. 2 An example of signal loss in the frontal lobes seen in echo planar imaging (EPI). (Left) A typical EPI scan used for fMRI and (right) the same slice as seen in a conventional MRI scan. Brain tissue that is clearly visible in the conventional MRI scan appears missing in the EPI scan, particularly in the region indicated by the arrow.

smeared out and distorted (Jezzard and Clare, 1999). Such effects mean that it is difficult to accurately detect activations in these regions of the brain. Several methods may be used to try and address this problem, such as correcting the distortions by using a "field map" (Jezzard and Balaban, 1995) by using specially designed mouthpieces graphite to compensate for the effect of air sinuses (Wilson *et al.*, 2002), or by using imaging methods similar to EPI that do not suffer from distortion (Glover and Law, 2001), although these often have their own disadvantages.

Unfortunately, the characteristics of the EPI method that make it susceptible to signal loss near sinuses are also those that make it sensitive to the BOLD effect. A typical fMRI experiment on a human subject using EPI would have image pixel sizes of 3×3 mm^2 and use a slice thickness of between 3 and 6 mm. Using thinner slices is one way of reducing signal loss near sinuses, but again this could reduce sensitivity to the BOLD effect (Merboldt *et al.*, 2000). Another parameter that affects signal loss is known as the "echo time" or TE. By reducing the echo time, signal loss is reduced, but again this comes at a cost to BOLD sensitivity. In practice, an echo time of 30–50 ms gives good results over the whole brain, but this is a parameter worth varying in initial pilot experiments to find the optimum balance in the brain regions of interest (Clare *et al.*, 2001).

One final note on rapid MRI is that these methods produce a high level of acoustic noise, often over 100 dB. This means that it is essential to provide the subject with adequate ear protection and warn them prior to the experiment.

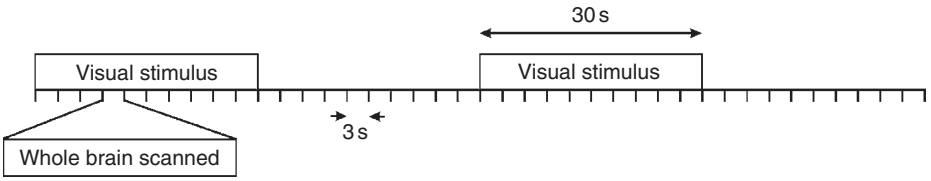

Fig. 3 A schematic diagram of a typical fMRI experiment where the whole brain is scanned every 3 s and alternating periods of visual stimulus and no stimulus are given every 30 s.

D. Experimental Design

The simplest form of fMRI experiment consists of blocks of stimulus presentation and "rest," interleaved such as illustrated in Fig. 3 for a visual stimulus. Typical timings of such a block paradigm are 30 s of stimulus and 30 s of rest, both repeated four times, with the whole brain being scanned once every 3 s. Critical to the success of the experiment is that there is a single clear difference between the task and the rest condition. It is not conceivable that the brain is completely inactive during the "rest" period, but it needs to be assumed that such activity is equally likely to be present during the task period as the rest period.

Because activations are detected by comparing serial MRI scans, any subject movement between each scan will reduce the ability to detect them. It is, therefore, important to minimize subject movement. For most short fMRI experiments, this is best done by using foam pads around the subject's head. Such pads or pillows not only provide support but also act as points of reference for the subjects as they try and keep their own head still. Many experimenters advocate the use of thermoplastic masks or bite bars to keep the subject still. While these do minimize head movement very effectively, they are often uncomfortable for the subject and usually require some time to get used to using them.

Whatever stimulus is used, it is likely that some form of visual stimulus or cue will need to be given to the subject either to stimulate the visual cortex directly, to instruct the subject in the timing and pacing of a movement task, or present cognitive stimuli such as patterns of letters to compare or remember. The most versatile way of presenting such stimuli is to use a high-quality video projector connected to a computer. This can project text or images onto a screen near the end of the scanner, at the subject's feet. Then an angled mirror above the subject's face, or a prism arrangement, enables the subject to view the screen. There are a number of more specialized systems that deliver the picture directly to glasses worn by the subject; however, the complexity and expense of these systems only make their use justified for particular applications that need higher control on what the subject sees.

Visual presentation is not the only way to present stimuli to subjects. As indicated earlier, the high acoustic noise environment of the scanner is not ideal for using auditory stimuli; however, it is possible to use MRI-compatible headphones (either pneumatic or electrostatic). Successful auditory studies have

been accomplished by scanning at a slower rate, with short gaps in the scanning during which the speech or sound is played (Hall *et al.*, 1999). Other devices, such as vibrotactile stimulators to stimulate the somato-sensory cortex, thermal devices to stimulate the pain network, or olfactometers to deliver smells, have all been used in the MRI environment. As with all equipment that comes in contact with a subject, it is essential that there is no chance of it causing any harm. The scanner environment adds additional constraints, both on safety and on the ability to get high-quality MR images, so it is important to check carefully before even taking a device into the magnet room.

As well as presenting a stimulus to a subject, it is often desirable to record some response from the subject. The most simple and versatile way to do this is with an MRI compatible, four button box. The subject can rest their four fingers of one hand on each of the buttons on the box, and the response can be fed back to the control computer for recording which button was pressed and the precise timing. Interface with the computer is usually best done via a dedicated analog and digital interface card (such as from National Instruments, Austin, TX), which should be able to handle not only the signals from button boxes but also from joysticks or other analog devices, and be able to control other stimulus presentation modalities that require a digital or analog signal. Getting the subject to verbalize a response is not advisable in general because the movement of the head caused by speaking can reduce the ability to detect small activations. Additionally, it can be hard to hear any response in the noisy environment of the scanner during the experiment.

Software for cueing the experiment and presenting the stimuli is available commercially (e.g., Presentation, Neurobehavioral Systems, Albany, CA) or from research groups (e.g., DMDX, www.u.arizona.edu/~kforster/dmdx/dmdx.htm; Cogent, www.vislab.ucl.ac.uk/Cogent). Such software enables a series of stimuli to be cued up then played out sequentially and will record the nature and timing of any responses by the subject. The pacing of an experiment relative to the acquisition of the scans is critically important, as it is the final images that contain the signal of interest. While this might seem trivial, it is often the case that neither the scanner nor the stimulus presentation computer can be relied upon to keep exact time. While this timing difference may only be one-hundredths of a second, over a long scan run this can make a significant difference. It is usually possible to arrange for the scanner to output a timing (TTL) trigger at the start of acquiring each scan. This trigger can then be detected by the stimulus presentation computer and be used to start the next stimulus.

A block design paradigm, where there are relatively long blocks of stimulation and rest, is suitable for many applications, but can suffer from a number of problems. First, it may not be possible to reliably get a brain region to be active for such a long period of time. Often habituation will occur as the subject easily learns a task or loses interest in the stimulus. There are also some stimuli that need to be presented for short periods of time, such as painful ones. Second, in some cases, particularly cognitive paradigms, it is not desirable to use the same type of stimuli continually. For example, if the experiment requires the subject to

discriminate between stimuli, they may quickly learn that the response stays the same for 30 s and make the cognitive part of the paradigm invalid. Third, there are some cases where it is desirable to separate out brain regions that are involved in different aspects of a task. This is particularly the case in experiments on memory, where different brain networks may be involved in the storage of information to the retrieval. To deal with the second problem, it is possible to carry out a block design experiment where the majority of stimuli are of one type in the ''task'' period and of the other type in the ''rest'' period, but there are a few of the other type added to keep the novelty component of the experiment.

An alternative to block designs are ''event-related'' designs. This design type is similar to that used in evoked potential EEG recordings in the brain where a single stimulus is presented at some repetition rate. In EEG, this is particularly suitable, as the electrical activity associated with the stimulus lasts less than a few hundred milliseconds. In fMRI, however, where the BOLD stimulus response takes 15–20 s to return to baseline following the stimulus, waiting such a long time between stimuli can be less suitable. This method is useful for looking for brain networks responding to different parts of a complex task (such as the memory illustration given earlier), but a more efficient use of event-related fMRI is to present the single stimuli at shorter intervals, often randomized in time. The analysis of such data is more complex, as it requires a model of how the overlapping BOLD responses to individual, closely spaced stimuli combine. However, such experiments can be highly efficient, particularly when stimuli need to be presented in some random order (Friston *et al.*, 1999).

E. Analysis of Images

In the simplest case, analysis of data consists of subtracting the average of all the images acquired during the ''rest'' phase of the experiment from the average of those acquired during the ''task'' period. While this gives a general qualitative indication of activated regions, in order to assign statistical significance to the result it is necessary to carry out a more detailed analysis. Also, this simple subtraction does not take account of the fact that the peak BOLD signal is delayed by around 6 s from the start of stimulation. In practice, typical fMRI analysis consists of three parts: preparation of fMRI data, model-based detection of the BOLD signal, and statistical inference and thresholding of the activations. If the results of a number of subjects are to be combined and compared, such as between a patient group and a control group, then an additional step of ''group statistics'' is required. Each of these areas is looked at in turn here. Most of these methods are available in commercial or freely available software packages (e.g., FSL, www.fmrib.ox.ac.uk/fsl; SPM, www.fil.ion.ucl.ac.uk/spm; AFNI, afni.nimh.nih.gov/afni; Brain Voyager, Brain Innovation, The Netherlands).

1. Preparation of fMRI Data

As explained in the previous section, minimizing the head motion of the subject is vital to getting good results. However, even with the best restraint methods, there is often some small residual motion that occurs in the images. This can be reduced by performing motion correction, in software, on the data. The motion correction algorithm compares each individual image with the first in the series and applies a mathematical transform to rotate or move the image until they look as similar to each other as possible (Friston *et al.*, 1996). Because the BOLD signal changes in the image are very small, they generally do not bias the motion correction. Next, the images are often spatially smoothed (blurred). The optimal detection of activations occurs when the spatial smoothness of the images is the same as the size of the region of activation. By applying spatial smoothing to the images, the ability to detect activations is often increased (Shaw *et al.*, 2003). Finally, the time course of each pixel in the image is filtered to remove long-term drifts in the signal and is sometimes filtered to smooth the time course of the signal over time.

2. Model-Based Detection of the BOLD Signal

Figure 4 shows a representation of a number of scans from an fMRI time series. If a single pixel in an activated region is selected, and its intensity is plotted over time, it displays a clear delayed response with respect to the stimulus presentation. If a mathematical model for the amount of delay and smoothing that is seen in the theoretical BOLD response with respect to the stimulus pattern is assumed (such as shown in the bottom line of Fig. 4), then a statistical measure of how likely that pixel is truly activated in response to the stimulus can be obtained by calculating the correlation coefficient between this theoretical line and the actual time course. Critical to the success of this method is the choice of mathematical model that turns the stimulus time course into a theoretical BOLD response. All of the software packages have a range of choices for this function, but a commonly used one is a gamma function convolution of the stimulus time series. The mathematical framework for this correlation-based approach, known as the general linear model (GLM), is not the only one that can be taken, and it does indeed include assumptions that may not be fully appropriate for fMRI data; however, for the majority of fMRI experiments, it is sensitive and statistically reliable (Friston *et al.*, 1995). The GLM can also be used to analyze the "event-related" fMRI experiments described earlier, again producing a model for the theoretical BOLD response given the stimulus timing pattern used. Analyzing the response through time of each pixel in the image results in a statistical "map" of the strength of correlation throughout the brain.

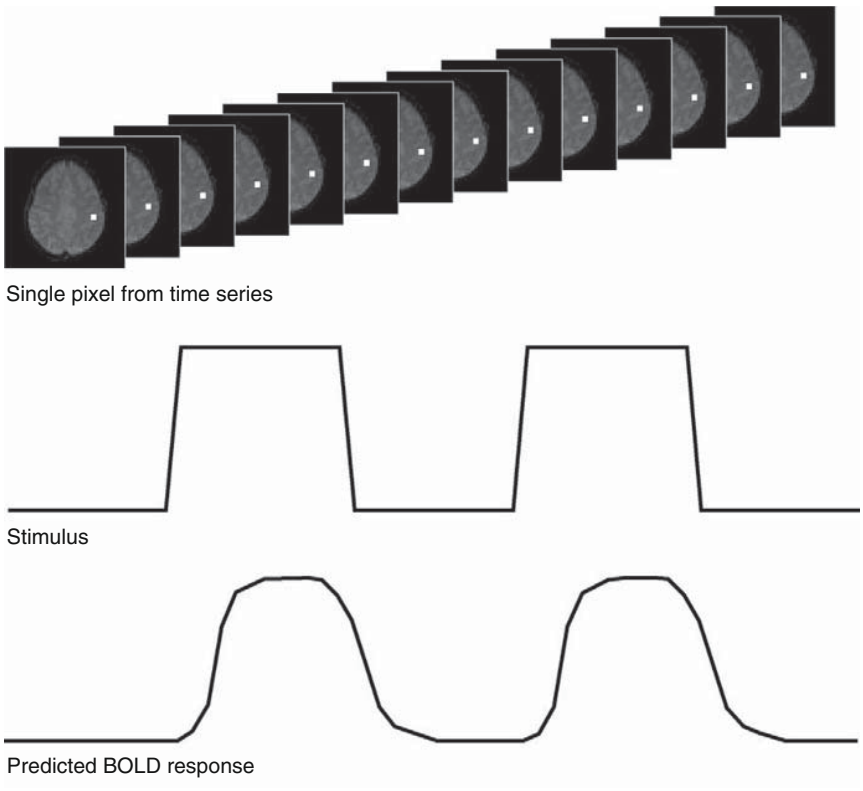

Fig. 4 Representation of the predicted BOLD response pattern (bottom line) to the stimulus (middle line) in single pixel that "activates" to that stimulus.

3. Statistical Inference and Thresholding of Activations

The output image produced by the GLM looks something like Fig. 5A. It is clear that there is a high correlation at the bottom of the image, but it is not clear yet if this is significant. It is possible to threshold the image so that only those individual pixels that have a correlation coefficient that is significant to better than $p < 0.5\%$, but in an image of approximately 10,000 pixels, we know that 50 pixels would be labeled as "activated" purely by chance. An alternative is to threshold at a much lower significant threshold (such as 0.5% divided by 10,000), but this stringent threshold risks missing genuinely activated pixels. As an alternative it is common to use information on the number of pixels near each other to increase our confidence in the result. For example, if we see 10 pixels in a block together all showing high correlation coefficients, then we can be more sure that this represents a genuine activation than if we saw one pixel on its own with a similarly high correlation coefficient. The fMRI analysis software packages all contain the

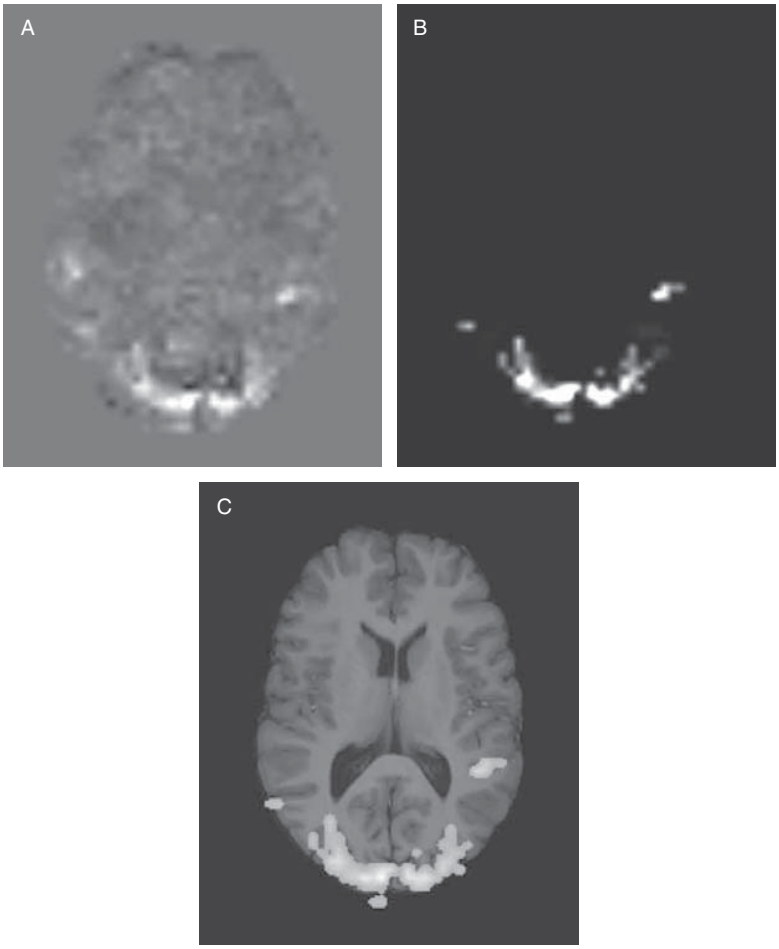

Fig. 5 (A) Statistical map obtained from the general linear output. This map is then thresholded at an appropriate *p* score (B) and is then overlaid on a high-quality MRI scan of the same subject's brain (C).

appropriate theory, such as that of Gaussian random field (Friston *et al.*, 1994) to threshold the activation images on the basis of both correlation coefficient (e.g., as reported as a "*Z* score") and pixel cluster size. What the software will output therefore is a list of pixel clusters that have a *Z* score of greater than say 2.3 (a good, if arbitrary, value to start with) and that have a pixel cluster size making the statistical significance (*p*-value) less than say 0.5%. These thresholded activation images can then be overlaid on the subject's high-quality MRI scan (to which the fMRI images have been aligned) or on a standard brain atlas, such as shown in Fig. 5C.

4. Group Comparisons

If a typical result from a group of subjects is requires or if a comparison between two groups needs to be made, then there are two additional steps that need to be performed. First, the MRI scans must be morphed to align to some standard brain template. Such brain templates have been made up of MRI scans from hundreds of individuals and are generally supplied with the software packages. Although not an ideal template, it is very common to report results with referenced to the atlas of Talairach based on a dissection of a single brain. Software for morphing to template brain and alignment to the Talairach atlas is included in most fMRI analysis packages. Once all the subjects' scans from a single study are aligned to the same template, further statistical analysis can be performed on data to show regions that are significantly activated across all subjects or ones that are differentially activated in one group relative to another.

IV. An Application of Functional MRI

The ability of the brain to reorganize functionally after injury is a fascinating but hard-to-study phenomenon. Functional MRI has been used to demonstrate the cortical changes that occur upon rehabilitation after stoke (Johansen-Berg *et al.*, 2002). Here, the experiment is described as an example of the integration of all the components described earlier to investigate brain function.

Patients with mild to moderate injury following stroke were scanned on four separate occasions, both before and after a movement therapy aimed at regaining motor control over their damaged side. In the scanner, the subjects were cued visually to tap their hand, which was resting on a wooden board, at a rate that was either 25% or 75% of their maximum tapping rate. A 6-min series of EPI scans, with an echo time of 30 ms and a repetition rate of 3 s, was recorded during tapping of first the unaffected hand and then the affected hand. A simple block design was used, similar to that illustrated in Fig. 3, with the subject resting their hand for 30 s between 30 s of hand tapping. For all scanning sessions the hand-tapping rate was kept constant, even if the subject was able on a later occasion to tap at a faster rate. This was essential for determining whether the brain activation patterns associated with performing the movements changed over time. A recovery score for each subject was also determined based on their motor performance before and after therapy.

Scans from the four sessions of each subject were analyzed together as one GLM analysis. This analysis was set up to detect not only correlations between the pixel time course and a predicted BOLD response to the stimulus but also in the same analysis, differences between the two sessions before therapy and the two sessions after therapy. This illustrates the strength of a well-designed GLM analysis; it is not necessary to individually analyze the results from each fMRI session and then look for differences in activation patterns, but it is possible to set up the analysis in

such a way as to generate one statistical image with the particular result of interest, in this case the changes that have taken place after therapy.

The researchers then went on to perform a second level of analysis to obtain a summary result representing all the subjects. This was done by subtracting the pretherapy from the posttherapy activation image (unthresholded) for each subject and weighting this difference image by the recovery score for that subject. Combining these images across the group gave an image of areas where activation increases correlate with recovery across the group. The results of both the analysis of the individual subjects and the group analyses indicated that upon recovery, movement of the affected hand produced increased activity in the motor networks of the unaffected hand were being recruited to compensate for the damage on the affected side.

V. Discussion and Conclusion

The advent of fMRI has made a huge impact in the way that the brain is studied, both in the pure neuroscience setting and in a clinical context. The increasing availability of MRI scanners in hospitals throughout the world means that although the technology is expensive, it is increasingly available to researchers. This has been coupled with a solid development of tools for stimulus presentation and, importantly, for data analysis that has the sophistication and statistical rigor required for solid inference, coupled with an ease of use suitable for general laboratory use.

The next big challenge in the development of fMRI is to more fully understand, quantitatively, the relationship between the signals observed in the MRI scans and the underlying neural activity of which it is a marker. This will require not only more sophisticated imaging methodology and more complex physiological models but also way to get a closer measure of the working of the neuron *in vivo*. Here, the experiments performed on animals, as described in other chapters in this volume, will play a vital role.

References

Belliveau, J. W., Kennedy, D. N., Jr., McKinstry, R. C., Buchbinder, B. R., Weisskoff, R. M., Cohen, M. S., Vevea, J. M., Brady, T. J., and Rosen, B. R. (1991). Functional mapping of the human visual cortex by magnetic resonance imaging. *Science* **254,** 716.

Buxton, R. B. (2001). The elusive initial dip. *Neuroimage* **13,** 953.

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

Chan, S. (2002). The clinical relevance and scientific potential of ultra high-field-strength MR imaging. *AJNR Am. J. Neuroradiol.* **23,** 1441.

Clare, S., Francis, S., Morris, P. G., and Bowtell, R. (2001). Single-shot T2* measurement to establish optimum echo time for fMRI: studies of the visual, motor, and auditory cortices at 3.0 T. *Magn. Reson. Med.* **45,** 930.

Frank, J. A., Anderson, S. A., Kalsih, H., Jordan, E. K., Lewis, B. K., Yocum, G. T., and Arbab, A. S. (2004). Methods for magnetically labeling stem and other cells for detection by *in vivo* magnetic resonance imaging. *Cytotherapy* **6**(6), 621–625.

Friston, K. J., Worsley, K. J., Frackowiak, R. S., Mazziotta, J. C., and Evans, A. C. (1994). Assessing the significance of focal activations using their spatial extent. *Hum. Brain Mapp.* **1,** 214.

Friston, K. J., Holmes, A. P., Poline, J. B., Grasby, P. J., Williams, S. C., Frackowiak, R. S., and Turner, R. (1995). Analysis of fMRI time-series revisited. *Neuroimage* **2,** 45.

Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S., and Turner, R. (1996). Movement-related effects in fMRI time-series. *Magn. Reson. Med.* **35,** 346.

Friston, K. J., Zarahn, E., Josephs, O., Henson, R. N., and Dale, A. M. (1999). *Neuroimage* **10,** 607.

Glover, G. H., and Law, C. S. (2001). Stochastic designs in event-related fMRI. *Magn. Reson. Med.* **46,** 515.

Haacke, M. E., Brown, R. W., Thompson, M. R., and Venkatesan, R. (1999). "Magnetic Resonance Imaging." Wiley, New York.

Hall, D. A., Haggard, M. P., Akeroyd, M. A., Palmer, A. R., Summerfield, A. Q., Elliott, M. R., Gurney, E. M., and Bowtell, R. W. (1999). "Sparse" temporal sampling in auditory fMRI. *Hum. Brain Mapp.* **7,** 213.

Jezzard, P., and Balaban, R. S. (1995). Correction for geometric distortion in echo planar images from B0 field variations. *Magn. Reson. Med.* **34,** 65.

Jezzard, P., and Clare, S. (1999). Sources of distortion in functional MRI data. *Hum. Brain Mapp.* **8,** 80.

Johansen-Berg, H., Dawes, H., Guy, C., Smith, S. M., Wade, D. T., and Matthews, P. M. (2002). Correlation between motor improvements and altered fMRI activity after rehabilitative therapy. *Brain* **125,** 2731.

Kim, S. G., Hendrich, K., Hu, X., Merkle, K., and Ugurbil, K. (1994). Potential pitfalls of functional MRI using conventional gradient-recalled echo techniques. *NMR Biomed.* **7,** 69.

Leontiev, O., and Buxton, R. B. (2007). Reproducibility of BOLD, perfusion, and CMRO2 measurements with calibrated-BOLD fMRI. *Neuroimage* **35**(1), 175–184.

Liu, T. T., Wong, E. C., Frank, L. R., and Buxton, R. B. (2002). Analysis and design of perfusion-based event-related fMRI experiments. *Neuroimage* **16,** 269.

Liu, T. T., and Brown, G. G. (2007). Measurement of cerebral perfusion with arterial spin labeling: Part 1. Methods. *J. Int. Neuropsychol. Soc.* **13**(3), 517–525.

Lu, H., Golay, X., Pekar, J. J., and Van Zijl, P. C. (2003). Functional magnetic resonance imaging based on changes in vascular space occupancy. *Magn. Reson. Med.* **50**(2), 263–274.

Merboldt, K. D., Finsterbusch, J., and Frahm, J. (2000). Reducing inhomogeneity artifacts in functional MRI of human brain activation-thin sections vs gradient compensation. *J. Magn. Reson.* **145,** 184.

Miezin, F. M., Maccotta, L., Ollinger, J. M., Petersen, S. E., and Buckner, R. L. (2000). Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *Neuroimage* **11,** 735.

Ogawa, S., Lee, T. M., Kay, A. R., and Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl. Acad. Sci. USA* **87,** 9868.

Shaw, M. E., Strother, S. C., Gavrilescu, M., Podzebenko, K., Waites, A., Watson, J., Anderson, J., Jackson, G., and Egan, G. (2003). Evaluating subject specific preprocessing choices in multisubject fMRI data sets using data-driven performance metrics. *Neuroimage* **19,** 988.

Silva, A. C., and Koretsky, A. P. (2002). Laminar specificity of functional MRI onset times during somatosensory stimulation in rat. *Proc. Natl. Acad. Sci. USA* **99,** 15182.

Stehling, M. K., Turner, R., and Mansfield, P. (1991). Echo-planar imaging: magnetic resonance imaging in a fraction of a second. *Science* **254,** 43.

Thulborn, K. R., Waterton, J. C., Matthews, P. M., and Radda, G. K. (1982). Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochim. Biophys. Acta* **714,** 265.

Wiggins, G. C., Triantafyllou, C., Potthast, A., Reykowski, A., Nittka, M., and Wald, L. L. (2006). 32-channel 3 Tesla receive-only phased-array head coil with soccer-ball element geometry. *Magn. Reson. Med.* **56**(1), 216–223.

Wilson, J. L., Jenkinson, M., and Jezzard, P. (2002). Optimization of static field homogeneity in human brain using diamagnetic passive shims. *Magn. Reson. Med.* **48,** 906.