Birte U. Forstmann
Eric-Jan Wagenmakers
Editors

An Introduction to Model-Based Cognitive Neuroscience

Springer

Chapter 5
An Introduction to fMRI

F. Gregory Ashby

Abstract Functional magnetic resonance imaging (fMRI) provides an opportunity to indirectly observe neural activity noninvasively in the human brain as it changes in near real time. Most fMRI experiments measure the blood oxygen-level dependent (BOLD) signal, which rises to a peak several seconds after a brain area becomes active. Several experimental designs are common in fMRI research. Block designs alternate periods in which subjects perform some task with periods of rest, whereas event-related designs present the subject with a set of discrete trials. After the fMRI experiment is complete, pre-processing analyses prepare the data for task-related analyses. The most popular task-related analysis uses the General Linear Model to correlate a predicted BOLD response with the observed activity in each brain region. Regions where this correlation is high are identified as task related. Connectivity analysis then tries to identify active regions that belong to the same functional network. In contrast, multivariate methods, such as independent component analysis and multi-voxel pattern analysis identify networks of event-related regions, rather than single regions, so they simultaneously address questions of functional connectivity.

5.1 Introduction

Functional magnetic resonance imaging (fMRI) provides researchers an opportunity to observe neural activity noninvasively in the human brain, albeit indirectly, as it changes in near real time. This exciting technology has revolutionized the scientific study of the mind. For example, largely because of fMRI, there are now emerging new fields of Social Neuroscience, Developmental Neuroscience, Neuroeconomics, and even Neuromarketing.

This chapter provides a brief overview of fMRI and fMRI data analysis. This is a complex topic that includes many difficult subtopics, such as (1) MR physics, (2) a description of the complex machinery and equipment one finds in a typical

F. G. Ashby (✉)
Department of Psychological & Brain Sciences, University of California,
Santa Barbara, CA 93106, USA
Tel.: 805-893-7909
e-mail: ashby@psych.ucsb.edu

B. U. Forstmann, E.-J. Wagenmakers (eds.), *An Introduction to Model-Based Cognitive Neuroscience,* DOI 10.1007/978-1-4939-2236-9_5

brain-imaging center, (3) how to run this equipment effectively (e.g., set the many parameters that control the scanner; spot and avoid artifacts that can corrupt the data), (4) experimental design, and (5) fMRI data analysis. Obviously, covering all this material in depth is far beyond the scope of any single chapter, or even any single book. The reader interested in learning more about these topics is urged to consult any of the books listed at the end of this chapter under "Further Reading."

5.2 What Can Be Learned from fMRI?

Currently, the typical fMRI experiment records a sluggish, indirect measure of neural activity with a temporal resolution of 1–3 s and a spatial resolution of 25–30 mm^3. Nevertheless, as the thousands of fMRI publications attest, this highly imperfect technology has dramatically influenced the study of mind and brain. Because of its poor temporal resolution, fMRI is not appropriate for resolving small timing differences between different cognitive stages or processes. And although the spatial resolution is typically good enough to localize brain activity at the level of major brain structures, it is not good enough to localize activity, for example, at the level of the cortical column. Partly for these reasons, the use of fMRI is not without controversy. Although the mind sciences have been generally enthusiastic about fMRI, a smaller group of scientists remain skeptical. For example, fMRI has been labeled by some as "the new phrenology" [1].

Because of this controversy, before examining fMRI in more detail, it is worth considering what this relatively new technology has to contribute to the mind sciences. Because of its limited temporal and spatial resolution, fMRI is most appropriate for answering questions about gross neural architecture, rather than about neural process. But it can be very effective at addressing such questions. For example, consider the fMRI results shown in Fig. 5.1. The four panels show areas of significant activation within prefrontal cortex on four different days of training as subjects practiced a difficult perceptual categorization task [2]. Note that as automaticity develops in this task, prefrontal activation reduces significantly. In fact, by the 20th practice session (i.e., after approximately 11,000 trials of practice) there is no evidence of any task-related activity in prefrontal cortex. Therefore, these fMRI data show clearly that the neural architecture mediating this categorization behavior changes qualitatively with practice. This same question could be addressed without fMRI, and it seems possible that a similar conclusion might be reached after many clever behavioral experiments. But Fig. 5.1 paints a clear and compelling picture that one rarely sees with purely behavioral approaches.

FMRI can be used in a similar way to test purely cognitive theories that make no neuroscience assumptions. For example, suppose some cognitive theory predicts that the same perceptual and cognitive processes mediate performance in two different tasks. Then this theory should predict similar patterns of activation in an fMRI study of the two tasks, even if the theory makes no predictions about what those activation patterns should look like. If qualitatively different activation patterns are found in

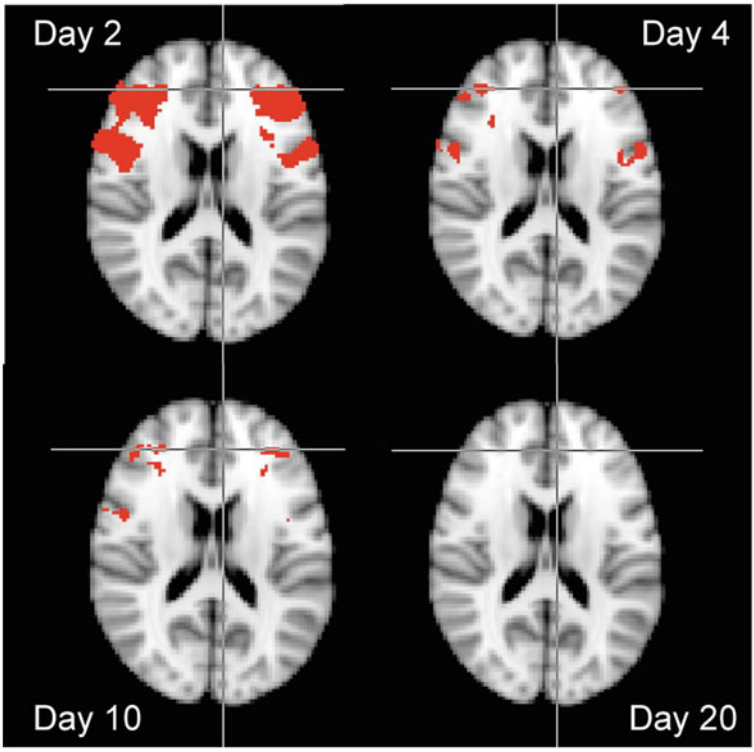

Fig. 5.1 Significant BOLD activation in prefrontal cortex on four different days of training on the same categorization task [8]

the tasks, then the theory probably needs some serious re-thinking. Applications like this are the primary reason that fMRI is popular even with cognitive scientists who have no fundamental interest in neuroscience.

As a final illustration of the benefits of fMRI, consider the following example. Unstructured categories are categories in which the stimuli are assigned to each category randomly. As a result, there is no similarity-based or logical rule for determining category membership. An example might be the category of numbers that have personal meaning to you (e.g., social security number, phone number, etc.). I had hypothesized in print that unstructured categories must be learned by explicit memorization, even when feedback-based training is provided [3]. But then Seger and her colleagues published several fMRI papers showing that feedback-based unstructured category learning elicits elevated task-related activation in the striatum, not the hippocampus [4–6]. These results suggested that unstructured category learning might be mediated by procedural memory, not declarative memory. Because of these papers, my colleagues and I decided to look for behavioral evidence that unstructured category learning is mediated by procedural memory. In fact, we found that switching the locations of the response buttons interfered with the expression of

unstructured category learning, but not with a (rule-based) version of this task that used similar stimuli and was known to depend on declarative memory. This sensitivity to response location is a hallmark of procedural memory, so our behavioral results were in agreement with the fMRI results. The important point here is that this behavioral experiment would not even have been run if the fMRI experiments had not identified a neural network that previously had been associated with procedural memory.

In summary, fMRI is a powerful method for studying neural and cognitive architecture, but it is not as effective at addressing questions about process. For this reason, it is not some magic method that will supplant all others in the scientific study of the mind. But it does provide an important new tool for mind scientists. When used in a converging operations approach that includes more traditional behavioral methodologies, it has the potential to dramatically improve our understanding of the human mind.

5.3 MR Physics and BOLD Imaging

The MR scanner uses superconducting electromagnets to produce a static, uniform magnetic field of high strength. Ten years ago, the standard field strength used in fMRI research was 1.5 Tesla (T), whereas the standard today is 3 T. Even so, a number of research centers have scanners considerably stronger than this (e.g., above 10 T). Some of these are used with human subjects, but many are only used for non-human animal research.

The static field, by itself, does not produce an MR signal. An MR signal requires radiofrequency coils that generate magnetic pulses. Turning a pulse on changes the magnetization alignment of protons (typically within water molecules) within the magnetic field. When the pulse is turned off, the protons relax to their original equilibrium alignment, which releases energy detected by the coils as the raw MR signal. Spatial resolution is provided by additional magnetic fields known as gradients. The strength of each gradient changes linearly along a single spatial dimension. Thus, three mutually orthogonal gradients are used to localize a signal in three spatial dimensions. The software that controls all these magnetic fields is typically called the pulse sequence.

The pulse sequence is run on the main computer that controls the scanner. In most fMRI experiments, a second computer creates the stimuli that are presented to the subject and records the subject's behavioral responses. This second computer is synchronized with the first, so that the onset of each stimulus presentation occurs at a precisely controlled moment during image acquisition. Visual stimuli are most often presented by directing a computer-controlled projector at a mirror directly above the subject's face, and responses are collected on some device held in the subject's hands (e.g., that has buttons or a joystick).

Two general types of pulse sequences are common, depending on whether the goal is structural or functional imaging. The goal of structural MR is usually to measure

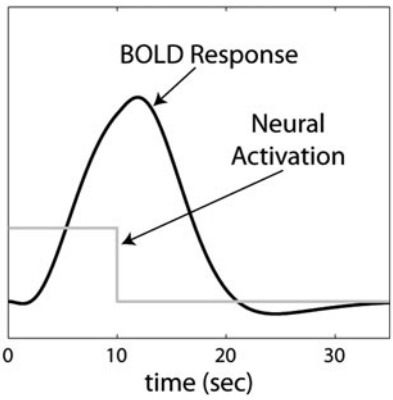

Fig. 5.2 A hypothetical BOLD response (*black curve*) to a constant 10 s neural activation (*gray curve*)

the density of water molecules, which differs, for example in bone, gray matter, cerebrospinal fluid, and tumors. The vast majority of functional MR (fMRI) experiments measure the blood oxygen-level dependent (BOLD) signal. The physics of this process is complex and far beyond the scope of this chapter. For our purposes, it suffices to know that the BOLD signal is a measure of the ratio of oxygenated to deoxygenated hemoglobin. Hemoglobin is a molecule in the blood that carries oxygen from the lungs to all parts of the body. It has sites to bind up to four oxygen molecules. A key discovery that led eventually to BOLD fMRI was that hemoglobin molecules fully loaded with oxygen have different magnetic properties than hemoglobin molecules with empty binding sites [7].

The theory, which is not yet fully worked out, is that active brain areas consume more oxygen than inactive areas. When neural activity increases in an area, metabolic demands rise and, as a result, the vascular system rushes oxygenated hemoglobin into the area. An idealized example of this process is shown in Fig. 5.2. The rush of oxygenated hemoglobin into the area causes the ratio of oxygenated to deoxygenated hemoglobin (i.e., the BOLD signal) to rise quickly. As it happens, the vascular system over compensates, in the sense that the BOLD signal actually rises well above baseline to a peak at around 6 s after the end of the neural activity that elicited these responses. Following this peak, the BOLD signal gradually decays back to baseline over a period of 20–25 s.

5.4 The Scanning Session

An experimental session that collects fMRI data also commonly includes a variety of other types of scans. At least four different types of scans are commonly acquired. Typically, the first scan completed in each session is the localizer. This is a quick structural scan (1–2 min) of low spatial resolution and is used only to locate the subject's brain in 3-dimensional space. This knowledge is needed to optimize the

location of the slices that will be taken through the brain in the high-resolution structural scan and in the functional scans that follow.

The ordering of the other scans that are commonly done is not critical. Frequently, however, the second type of scan completed is the high-resolution structural scan. Depending on the resolution of this scan and on the exact nature of the pulse sequence that is used to control the scanner during acquisition, it may take 8–10 min to collect these data. The structural scan plays a key role in the analysis of the functional data. Because speed is a high priority in fMRI (i.e., to maximize temporal resolution), spatial resolution is sacrificed when collecting functional data. The high-resolution structural scan can compensate somewhat for this loss of spatial information. This is done during preprocessing when the functional data are aligned with the structural image. After this mapping is complete, the spatial coordinates of activation observed during fMRI can be determined by examining the aligned coordinates in the structural image.

The third step is often to collect the functional data. This can be done in one long run that might take 20–30 min to complete, or it can be broken down into 2 or 3 shorter runs, with brief rests in between. There are many parameter choices to make here, but two are especially important for the subsequent fMRI data analysis. One choice is the time between successive whole brain scans, which is called the repetition time and abbreviated as the TR. If the whole brain is scanned, typical TRs range from 2–3 s, but TRs as low as 1 s are possible on many machines, especially if some parts of the brain are excluded from the scanning.

Another important choice is voxel size, which determines the spatial resolution of the functional data. When a subject lies in the scanner, his or her brain occupies a certain volume. If we assign a coordinate system to the bore of the magnet, then we could identify any point in the subject's brain by a set of three coordinate values (*x, y, z*). By convention, the *z* direction runs down the length of the bore (from the feet to the head), and the *x* and *y* directions reference the plane that is created by taking a cut perpendicular to the *z* axis. The brain, of course, is a continuous medium, in the sense that neurons exist at (almost) every set of coordinate values inside the brain. FMRI data, however, are discrete. The analog-to-digital conversion is performed by dividing the brain into a set of cubes (or more accurately, rectangular right prisms). These cubes are called voxels because they are three-dimensional analogues of pixels – that is, they could be considered as volume pixels.

A typical voxel size in functional imaging might be 3 mm $\times$ 3 mm $\times$ 3.5 mm. In this case, in a typical human brain, 33 separate slices might be acquired each containing a 64 $\times$ 64 array of voxels for a whole brain total of 135,168 voxels. In each fMRI run, a BOLD response is recorded every TR seconds in each voxel. Thus, for example, in a 30 min run with a TR of 2 s, 135,168 BOLD responses could be recorded 900 separate times (i.e., 30 times per minute $\times$ 30 min), for a total of 121,651,200 BOLD values. This is an immense amount of data, and its sheer volume greatly contributes to the difficulties in data analysis.

Many studies stop when the functional data acquisition is complete, but some other types of scans are also common. A fourth common type of scan is the field map. The ideal scanner has a completely uniform magnetic field across its entire

bore. Even if this were true, placing a human subject inside the bore distorts this field to some extent. After the subject is inside the scanner, all inhomogeneities in the magnetic field are corrected via a process known as shimming. If shimming is successful, the magnetic field will be uniform at the start of scanning. Sometimes, however, especially in less reliable machines, distortions in the magnetic field will appear in the middle of the session. The field map, which takes only a minute or two to collect, measures the homogeneity of the magnetic field at the moment when the map is created. Thus, the field map can be used during later data analysis to correct for possible nonlinear distortions in the strength of the magnetic field that develop during the course of the scanning session.

5.5 Experimental Design

Almost all fMRI experiments use a block design, an event-related design, or a free-running design. In a block design, the functional run consists of a series of blocks, each of which may last for somewhere between 30 s to a couple of minutes. Within each block, subjects are instructed to perform the same cognitive, perceptual, or motor task continuously from the beginning of the block until the end. In almost all block-design experiments, subjects will simply rest on some blocks. For example, a researcher interested in studying the neural network that mediates rhythmic finger tapping might use a block design in which blocks where the subject is resting alternate with blocks in which the subject taps his or her finger according to some certain rhythm.

Event-related designs are run more like standard psychological experiments, in the sense that the functional run is broken down into a set of discrete trials. Usually each trial is one of several types, and each type is repeated at least 20 times over the course of the experiment. As in a standard experiment, however, the presentation order of the trial types within each run is often random. When analyzing data from an event-related design, it is critical to know exactly when the presentation of each stimulus occurred, relative to TR onset. A common practice is to synchronize stimulus presentation with TR onset. The first event-related designs included long rests between each pair of successive trials. In these slow event-related designs, rests of 30 s are typical. These are included so that the BOLD response in brain regions that participate in event processing can decay back to baseline before the presentation of the next stimulus. This makes statistical sense, but it is expensive since it greatly reduces the number of trials a subject can complete in any given functional run. Another problem is that because subjects have so much time with nothing to do, they might think about something during these long rests, and any such uncontrolled cognition would generate an unwanted BOLD response that might contaminate the stimulus-induced BOLD response.

Most current event-related designs use much shorter delays. These rapid event-related designs became possible because statistical methods were developed for dealing with the overlapping BOLD responses that will occur anytime the BOLD

response in a brain region has not decayed to baseline by the time another stimulus is presented. It is important to realize however, that even in rapid event-related designs the delay between trials is still significantly longer than in standard laboratory experiments. For example, a typical rapid event-related design might use random delays between successive trials that cover a range between 2 and 16 s. There are several reasons for this. First, because of the need to synchronize stimulus presentation with the TR, it is often necessary to delay stimulus presentation until the onset of the next TR. Second, in order to get unique estimates of the parameters of the standard statistical models that are used to analyze fMRI data, delays of random duration must be used. The process of adding such random delays between events is called *jittering*.

Finally, in free-running designs, events are presented to the subject continuously in time and typically discrete events are impossible to define. For example, subjects might watch a movie in the scanner, or simply lay there passively. The activities that subjects perform in free-running designs are often more natural than is possible with more structured designs, but this increased freedom comes at a cost because the data that result are more challenging to analyze than the data collected from block or event-related designs.

5.6 Data Analyses

A number of features of fMRI data greatly complicate its analysis. First, as mentioned above, a typical scanning session generates a huge amount of data. Second, fMRI data are characterized by substantial spatial and temporal correlations. For example, the sluggish nature of the BOLD response means that if the BOLD response in some voxel is greater than average on one TR then it is also likely to be greater than average on the ensuing TR. Similarly, because brain tissue in neighboring voxels will be supplied by a similar vasculature, a large response in one voxel increases the likelihood that a large response will also be observed at neighboring voxels. A third significant challenge to fMRI data analysis is the noisy nature of fMRI data. Typically the signal that the data analysis techniques are trying to find is less than 2 or 3 % of the total BOLD response.

The analysis of fMRI BOLD data is broken down into two general stages—preprocessing and task-related analysis. Preprocessing includes a number of steps that are required to prepare the data for task-related analysis. This includes, for example, aligning the functional and structural scans, correcting for any possible head movements that might have occurred during the functional run, and various types of smoothing (to reduce noise). Typically, the same preprocessing steps are always completed, regardless of the particular research questions that the study was designed to address. In contrast, the task-related analyses include all analyses that are directed at these questions.

A wide variety of software packages are available for fMRI data analysis. Many of these are free, and they each have their own advantages and disadvantages. The most widely used package is SPM (Statistical Parametric Mapping), which is written

and maintained by the Wellcome Trust Centre for Neuroimaging at the University College London. SPM is freely available at http://www.fil.ion.ucl.ac.uk/spm/. SPM is a collection of Matlab functions and routines with some externally compiled C code that is included to increase processing speed. A thorough description of the statistical foundations of SPM was provided by Friston, Ashburner, Kiebel, Nichols, and Penny [8].

Another widely used fMRI data analysis software package is called FSL, which is an acronym for the FMRIB Software Library. FSL is produced and maintained by the FMRIB Analysis Group at the University of Oxford in England. FSL is also freely available and can be downloaded at http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSL. Descriptions of the statistical foundations of the FSL routines were provided by Smith et al. [9] and by Woolrich et al. [10].

5.6.1 Modeling the BOLD Response

The goal of almost all fMRI experiments is to learn something about neural activity. Unfortunately however, the BOLD response measured in most fMRI experiments provides only an indirect measure of neural activation [11, 12]. Although it is commonly assumed that the BOLD signal increases with neural activation, it is known that the BOLD response is much more sluggish than the neural activation that is presumed to drive it. As a result, for example, the peak of the BOLD signal lags considerably behind the peak neural activation (e.g., see Fig. 5.2).

Logothetis and colleagues have presented evidence that the BOLD response is more closely related to local field potentials than to the spiking output of individual neurons [13, 14]. Local field potentials integrate the field potentials produced by small populations of cells over a sub-millimeter range, and they vary continuously over time. Most applications of fMRI make no attempt to model neural activation at such a detailed biophysical level. Rather, neural activation is typically treated as a rather abstract latent (i.e., unobservable) variable. It is assumed to increase when a brain region is active and to decrease during periods of inactivity. As with any latent variable, however, to make inferences about neural activation from observable BOLD responses requires a model of how these two variables are related.

Almost all current applications of fMRI assume that the transformation from neural activation to BOLD response can be modeled as a linear, time-invariant system. Although it is becoming increasingly clear that the transformation is, in fact, nonlinear (e.g., [15–17]), it also appears that these departures from linearity are not severe so long as events are well separated in time (e.g., at least a few seconds apart) and brief exposure durations are avoided [17]. These two conditions are commonly met in fMRI studies of high-level cognition.

In the linear systems approach, one can conceive of the vascular system that responds to a sudden oxygen debt as a black box. The input is neural activation and the output is the BOLD response. Suppose we present a stimulus event E_i to a subject at time 0. Let $N_i(t)$ denote the neural activation induced by this event at

time t and let $B_i(t)$ denote the corresponding BOLD response. Then from the systems theory perspective, the box represents the set of all mathematical transformations that convert the neural activation $N_i(t)$ into the BOLD response $B_i(t)$. For convenience, we will express this mathematical relationship as

$$B_i(t) = f[N_i(t)]$$

where the operator f symbolizes the workings of the black box.

A system of this type is said to be linear and time-invariant if and only if it satisfies the superposition principle, which is stated as follows:

If $f[N_1(t)] = B_1(t)$ and $f[N_2(t)] = B_2(t)$, then it must be true that

$f[a_1 N_1(t) + a_2 N_2(t)] = a_1 B_1(t) + a_2 B_2(t)$, for any constants a_1 and a_2.

In other words, if we know what the BOLD response is to neural activation $N_1(t)$ and to neural activation $N_2(t)$, then we can determine exactly what the BOLD response will be to any weighted sum of these two neural activations by computing the same weighted sum of the component BOLD responses.

If the superposition principle holds then there is a straightforward way to determine the BOLD response to *any* neural activation from the results of one simple experiment. All we need to do is to measure the BOLD response that occurs when the neural activation is an impulse—that is, when it instantly increases from zero to some large value then instantly drops back to zero. Denote the BOLD response in this idealized experiment by $h(t)$. In linear systems theory the function $h(t)$ is called the impulse response function because it describes the response of the system to an impulse. In the fMRI literature, however, $h(t)$ is known as the hemodynamic response function, often abbreviated as the hrf. Note that "hemodynamic response function" is not a synonym for "BOLD response". Rather the hrf is the hypothetical BOLD response to an idealized impulse of neural activation.

If the relationship between neural activation and the BOLD response satisfies superposition, then once we know the hrf, the BOLD response to any neural activation $N(t)$, no matter how complex, can be computed exactly from the so-called convolution integral:

$$B(t) = \int_0^t N(\tau)h(t-\tau)d\tau \tag{5.1}$$

The convolution integral massively simplifies the analysis of fMRI data, and as a result it forms the basis for the most popular methods of fMRI data analysis.

Given that the hrf plays such a critical role in analyzing fMRI data, the natural next question to ask is: how can we determine numerical values of the hrf? The most obvious method for determining the hrf, which is suggested by the name "impulse response function", is simply to input an impulse to the system and record the output. If the system is linear and time-invariant, then the output will exactly equal $h(t)$. With traditional fMRI experiments, of course, we cannot directly input a neural activation, so using this method to estimate the hrf is highly problematic. Even so, this method has been used to estimate the hrf in primary visual cortex (e.g., [18, 19]).

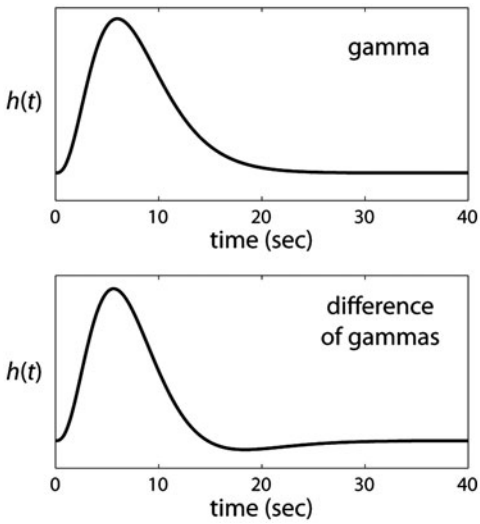

Fig. 5.3 Two popular models of the hrf

A much more popular method is to select a specific mathematical function for the hrf based on our knowledge of what we think this function should look like. For example, we know the hrf should peak at roughly 6 s and then slowly decay back to baseline. So we could select a mathematical function with these properties and then just assume that this is a good model of the hrf. In fact, this is, by far, the most popular method for determining the hrf in fMRI data analysis. The most popular choices are a gamma function or the difference of two gamma functions. Examples of both of these models are shown in Fig. 5.3.

5.6.1.1 Preprocessing

The most common goal of fMRI research is to identify brain areas activated by the task under study. The data that come directly out of the scanner, however, are poorly suited to this goal. The preprocessing of fMRI data includes all transformations that are needed to prepare the data for the more interesting task-related analyses. Preprocessing steps typically are the same for all experiments, so any analyses that do not depend on the specific hypotheses that the experiment was designed to test are typically called preprocessing.

The variability in raw fMRI data is so great that it easily can swamp out the small changes in the BOLD response induced by most cognitive tasks. Some of this variability is unavoidable in the sense that it is due to factors that we cannot control or even measure (e.g., thermal and system noise). But other sources of variability are systematic. For example, when a subject moves his or her head, the BOLD response sampled from each spatial position within the scanner suddenly changes in a predictable manner. The analyses done during preprocessing remove as many of these systematic non-task-related sources of variability as possible.

Typically the first preprocessing step is slice-time correction. Almost all fMRI data are collected in slices. If the TR is 2.5 s, then the time between the acquisition of the first and last slice will be almost this long. Slice-time correction corrects for these differences in the time when the slices are acquired.

The second step is to correct for variability due to head movement. Arguably, this is probably the most important preprocessing step. Even small, almost imperceptible head movements can badly corrupt fMRI data. Huettel et al. [20] give an example where a head movement of 5 mm increases activation values in a voxel by a factor of 5. When a subject moves his or her head, brain regions will move to new spatial locations within the scanner, and as a result, activation in those regions will be recorded in different voxels than they were before the movement occurred. Mathematical methods for correcting for head movements depend heavily on the assumption that when a subject moves his or her head, the brain does not change shape or size and therefore can be treated as a rigid body. Head movement correction then becomes a problem of rigid body registration (e.g., [21]). The BOLD responses from one TR are taken as the standard and then rigid body movements are performed separately on the data from every other TR until each of these data sets agrees as closely as possible with the data from the standard.

The third step, called coregistration, is to align the structural and functional data. This is critical because the spatial resolution of the functional data is poor. For example, with functional data a voxel size of $3 \times 3 \times 3.5$ mm is common. With structural images, however, the voxel size might be $.86 \times .86 \times .89$ mm, which is an improvement in resolution by a factor of almost 50.

The fourth step, normalization, warps the subject's structural image to a standard brain atlas. There are huge individual differences in the sizes and shapes of individual brains, and these differences extend to virtually every identifiable brain region. These differences make it difficult to assign a task-related activation observed in some cluster of voxels to a specific neuroanatomical brain structure. A researcher particularly skilled in neuroanatomy could coregister the functional activation onto the structural image and then look for landmarks in the structural scan that would allow the neuroanatomical locus of the cluster to be identified. An alternative is to register the structural scan of each subject separately to some standard brain where the coordinates of all major brain structures have already been identified and published in an atlas. Then we could determine the coordinates of a significant cluster within this standard brain, look these coordinates up in the atlas, and thereby determine which brain region the cluster is in. The process of registering a structural scan to the structural scan from some standard brain is called normalization [22].

Among the earliest and still most widely used brain atlases is the Talairach atlas [23], which is based entirely on the detailed dissection of one hemisphere of the brain of a 60-year old French woman. The atlas is essentially a look-up table containing major brain areas and their anatomical (x, y, z) coordinates. For many years, the Talairach atlas was almost universally used in neuroimaging, primarily because of the lack of any reasonable alternatives. But there has always been widespread dissatisfaction with this atlas because it is based on one hemisphere of a single, rather

unrepresentative brain. More recently, an atlas produced by the Montreal Neurological Institute (MNI) has become popular. The MNI atlas was created by averaging the results of high resolution structural scans that were taken from 152 different brains. The coordinate system was constructed to match the Talairach system, in the sense that it uses the same axes and origin. Whichever atlas is used, it is important to note that the registration problem in normalization is considerably more complex than in head motion correction or coregistration. This is because normalization requires more than rigid body registration. Not only will there be rigid body differences between the standard brain and the brain of typical subjects, but there will also be size and shape differences. Size differences can be accommodated via a linear transformation, but a nonlinear transformation is almost always required to alter the shape of a subject's brain to match either the Talairach or MNI standards.

Step five spatially smoothes the data with the goal of reducing nonsystematic high frequency spatial noise. In this step, the BOLD value in each voxel is replaced by a weighted average of the BOLD responses in neighboring voxels. The weight is greatest at the voxel being smoothed and decreases with distance. There are a number of advantages to spatially smoothing fMRI data. Most of these are due to the effects of the smoothing process on noise in the data. First, because smoothing is essentially an averaging operation, it makes the distribution of the BOLD responses more normal (i.e., because of the central limit theorem). Because the statistical models that dominate fMRI data analysis assume normally distributed noise, smoothing therefore transforms the data in a way that makes it more likely to satisfy the assumptions of our statistical models. A second benefit is that smoothing is required by a number of popular methods for solving the multiple comparisons problem (i.e., those that depend on Gaussian random field theory). A third benefit of smoothing, which is the most important of all, is that it can reduce noise and therefore increase signal-to-noise ratio.

Finally, in step six, temporal filtering is done primarily to reduce the effects of slow fluctuations in the local magnetic field properties of the scanner.

5.6.1.2 Task-Related Data Analyses

After pre-processing is complete, the next step is to try to identify brain regions that were activated by the task under study. The most popular approach to this problem is a correlation-based technique that is the foundation of most fMRI software packages [24, 25]. The idea is to first predict as accurately as possible what the BOLD response should look like in task-sensitive voxels. Next, the observed BOLD response in each voxel is correlated with this predicted signal. Voxels where this correlation is high are then identified as task related.

The first step in predicting the BOLD response to each stimulus event is to make an assumption about how long the neural activation will last in brain regions that process this event. A common assumption is that the neural activation induced by the event onset will persist for as long as the stimulus is visible to the subject. Another possibility is that the neural activation persists until the subject responds (so the

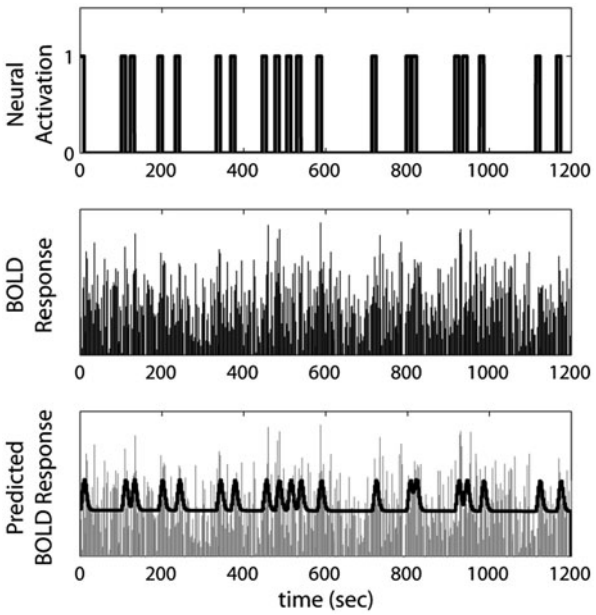

Fig. 5.4 A hypothetical example of the standard correlation-based analysis of fMRI data. The *top* panel shows the boxcar function that models the presumed neural activation elicited by the presentation of 20 separate stimuli. The *middle* panel depicts the hypothetical BOLD response in this experiment in a voxel with task-related activity. The *bottom* panel shows the best-fitting predicted BOLD response that is generated by convolving an hrf with the boxcar function shown in the *top* panel and then adding a constant baseline activation level

duration of neural activation equals the subject's response time). The second step is to model all presumed neural activations via a boxcar function. This is simply a function that persists for the duration of fMRI data acquisition and equals 1 when neural activation is assumed to be present and 0 when neural activation is absent. The top panel of Fig. 5.4 shows a hypothetical example of a boxcar function that describes the presumed neural response to the presentation of 20 separate stimuli. The stimulus presentations were spaced irregularly in time (i.e., jittered) to improve the statistical properties of the analysis. The middle panel of Fig. 5.4 shows the hypothetical BOLD response recorded in this experiment from one task-related voxel. Note that from visual inspection alone, it is not at all obvious that this is a task-related voxel.

The correlation method assumes linearity, so the third step in the analysis is to choose a model of the hrf. As mentioned earlier, the most popular choice is to select a specific mathematical function for the hrf that has no free parameters (e.g., either function shown in Fig. 5.3). Step four is to compute the predicted BOLD response by convolving the neural boxcar function with the hrf (using Eq. 5.1). Because there are no free parameters in either the boxcar function or the hrf, this integral can be evaluated numerically. In other words, for every TR in the experiment, a numerical value of the predicted BOLD response can be computed from Eq. 5.1. The bottom panel in Fig. 5.4 shows the predicted BOLD response that results when the boxcar function in the top panel is convolved with a gamma function hrf (and an estimated baseline activation is added).

The final step is to correlate these predicted BOLD values with the observed BOLD response in every voxel. Voxels where this correlation is high are presumed

to show task-related activity. The correlation is typically done within the context of the familiar General Linear Model (GLM) that is the basis of both multiple regression and analysis-of-variance. The outcome of this analysis in each voxel is the value of a statistic—most often a z or t value—that tests the null hypothesis that activation in that voxel is not correlated with the predicted BOLD response, or in other words, that activity in the voxel is not task related. Extreme values of the statistic are therefore evidence for task-related activity. In Fig. 5.4, the t statistic that results from this correlation has a numerical value of 7.78.

The correlation method applies to data from a single voxel at a time. Thus, if an experiment collects data from the whole brain, this analysis could easily be repeated more than 100,000 times to analyze all of the data collected in the experiment. The result of all these analyses is a value of the test statistic in every voxel that was analyzed. The resulting collection of statistics is often called a statistical parametric map, which motivated the name of the well-known fMRI data analysis software package, SPM.

A more recent variant of this correlation-based approach, called model-based fMRI, uses an independent computational model of the behavior under study to improve and refine the predicted BOLD signal [26]. In typical applications, the model is first fit to the behavioral data collected during the functional run separately for each subject. Next, parameter estimates from the model fits are used to build a model of the neural activation that is unique for every subject. From here, the analysis proceeds exactly as in the standard correlation-based method—that is, the predicted neural activation is convolved with an hrf to generate a predicted BOLD signal and then the GLM is used to generate a statistical parametric map. Model-based fMRI can be used to account for individual differences in fMRI data, but if the computational model is good, it can also be used to identify brain regions that respond selectively to components or sub-processes of the task. In particular, if the model has different parameters that describe different perceptual or cognitive processes that are presumed to mediate the behavior under study, then different regressors can be created that make specific predictions about each of these processes. For example, O'Doherty et al. [27] used this approach to identify separate brain regions associated with the actor versus the critic in actor-critic models of reinforcement learning.

Once a statistical map is constructed, the next problem is to determine which voxels show task-related activity. Of course if we only ask this question about a single voxel, then the answer is taught in every introductory statistics course. We simply decide what type 1 error rate we are willing to accept, find the threshold value of the statistic (e.g., the z or t value) that yields this error rate, and then decide that the voxel shows task-related activity if the value of this statistic exceeds this threshold. However, if the type 1 error rate equals 0.05 for each test, then with 100,000 independent tests, we would expect 5000 false positives if none of these voxels were task sensitive. This is clearly unacceptable. As a result, the criterion for significance must somehow be adjusted on each test to reduce the total number of false positives to some acceptable value. In statistics, this is called the multiple comparisons problem.

If the tests are all statistically independent then it is well known that the exact solution to this problem is to apply the Sidak or Bonferroni corrections. For example, if we want to limit the probability of a type 1 error to 0.05 in an overall collection of 100,000 z-tests (i.e., so that 95 % of the time there are no false positives in the 100,000 tests), then the Bonferroni correction sets the critical value on each test to approximately 0.0000005, which translates to a z-threshold for determining significance of 4.89. In the Fig. 5.4 example, the t value (i.e., 7.78) is so large that it would (correctly) be judged as significant, even if the Bonferroni correction is used, but in general the Bonferroni correction is so conservative that its use will typically cause us to miss true task-related activity in many voxels. The good news is that for fMRI data, the Bonferroni correction is much too conservative. The Bonferroni correction is exact when all the tests are statistically independent. With fMRI data, however, spatial correlations guarantee a positive correlation between test statistics in neighboring voxels. Thus, if significance is found in one voxel then the probability of obtaining significance in neighboring voxels is above chance. As a result, the critical value specified by the Bonferroni correction is too extreme. The bad news is that no exact solution to this problem is known. Even so, many different solutions to this problem have been proposed. Different methods are popular because they make different assumptions and have different goals. Most of the parametric methods rely on the theory of Gaussian random fields [28, 29]. Included in this list are all of the most popular cluster-based methods. These methods all require that spatial smoothing is performed during preprocessing. The most popular nonparametric methods include permutation methods [30] and methods that attempt to control the false discovery rate (FDR), rather than the false positive rate. The idea behind the FDR approach is that with many tests, a few false positives should not be feared [31]. So instead of trying to control the experiment-wise probability of a false positive, the FDR approach argues that a more important goal should be to limit the proportion of significant results that are false positives. In other words, consider the set of all voxels for which the null hypothesis of no signal is rejected. The goal of the FDR approach is to limit the proportion of these voxels for which the null hypothesis was incorrectly rejected.

In the standard correlation analysis so far considered, the GLM is applied separately to every voxel in the whole brain or region of interest. After the multiple comparisons problem is solved, a considerable challenge still remains to interpret the results of all these analyses. For example, suppose the analysis reveals strong task-related activation in the dorsolateral prefrontal cortex and in the dorsal striatum. Because these two significance decisions were based on independent applications of the GLM, we have no basis to conclude that these areas are functionally connected in the task we are studying. It could be that they are both part of independent neural networks that just happened to both be activated at similar times. So an important next step in the data analysis process is to identify functionally connected neural networks that are mediating performance in the task under study. This phase of the data analysis is known as connectivity analysis.

The idea underlying connectivity analysis is that a standard GLM analysis identifies clusters (or voxels) that show task-related activation, but it does not specify whether any pair of these clusters is part of the same or different neural networks.

If two clusters are part of the same network then they should be functionally connected in the sense that activation in one might cause activation in the other, or at least the separate activations in the two clusters should be correlated. If instead, the clusters are in separate neural networks then we would not expect either of these two conditions to hold. Connectivity analysis then, is done after a GLM analysis, with the goal of determining which brain regions in the task-related activation map are functionally connected.

An obvious method for testing whether two brain regions are functionally connected in a particular cognitive task is to measure the correlation between the BOLD responses in the two regions across TRs in an experiment where the task is performed. Regions that work together to mediate performance in the task should have correlated neural activations – that is, they should both be active while the task is being performed and they should both be inactive during rest periods. So one approach to connectivity analysis is to look for voxels or groups of voxels in different brain regions that show correlated BOLD responses (i.e., correlated across TRs). A simple solution to this problem is to compute the standard Pearson correlation between voxels or regions. A more sophisticated, yet similar approach uses Granger causality, which is a conceptually simple method that originated in the economics literature [32]. The idea is that if activation in region X causes activation in region Y, then knowledge of earlier activations in region X should improve our prediction of the current activation in region Y. Granger causality tests for such prediction using autoregressive models that are applied via the GLM [33].

One weakness of both Pearson correlation and Granger causality is that they both can fail because of high-frequency noise and/or because the hrf in the two brain regions might be different. A popular alternative is to compute correlations in the frequency domain using coherence analysis, rather than in the time domain, as is done with the Pearson correlation and with Granger causality [34]. Coherence analysis has an advantage over methods that compute correlations in the time domain because the coherence between BOLD responses in two brain regions is unaffected by hrf differences across the regions or by high-frequency noise.

The GLM-based methods of data analysis that compute correlations between predicted and observed BOLD responses require that we specify exactly when neural activation turns on and off. With free-running designs, this is often impossible. For this reason, data analysis choices are severely limited with free-running designs. Perhaps the most popular approach is to compute inter-subject correlations (ISCs; [35]). This method assumes that every subject experiences the same stimulation during the functional imaging. For example, if subjects are watching a movie, then every subject must see the same movie and the onset of the movie must begin at the same time for every subject. Under these conditions, the idea is to correlate the BOLD responses across TRs for every pair of subjects. If a brain region is responding to the movie then its activity should modulate up and down as the movie unfolds in similar ways in different subjects. So the ISC method identifies as task relevant, those voxels where the mean correlation across subjects is high.

Correlation-based analyses that use the GLM are univariate. This means that they analyze the data one voxel at a time. Univariate methods assign completely separate

parameters to every voxel, which means they assume that the data in neighboring voxels have no relationship to each other. *Post hoc* methods are then applied to try to overcome this obviously incorrect assumption. Included in this list are methods for correcting for the multiple comparisons problem that arises from this univariate approach, and the various connectivity analyses that attempt to recover information about spatial correlations from the many independent tests. An alternative approach is to perform multivariate data analyses that attempt to answer the significance and functional connectivity questions at the same time while also completely avoiding the multiple comparisons problem. The trick is that multivariate approaches identify task-related networks, rather than task-related clusters.

Two multivariate methods for analyzing fMRI data are especially popular. One is independent components analysis (ICA; [36, 37]). ICA, like its relative, principal components analysis, decomposes the observed BOLD data into a set of independent components. It operates on data from the whole brain at once, so the components it identifies are functionally independent neural networks that are simultaneously active during some fMRI experiment. For example, one network might mediate the processing of task-relevant stimulus information, one might mediate processing of any feedback provided during the course of the experiment, one might monitor the auditory stimulation provided by the scanning environment, and one could be the so-called default mode network, which can be seen when subjects lie quietly in the scanner with no task to perform (e.g., [38]). Each network defines a spatial pattern of activation across all voxels in the brain. When the network is active, voxels in brain regions that are part of the network will be active, and voxels that are not part of the network will be inactive. On any TR, ICA assumes that the observable BOLD response is a mixture of the activation patterns associated with each of these networks plus (perhaps) some noise. As the TRs change, the amount that each network is activated could change. For example, the network that processes the stimulus should become more active on TRs during and immediately after stimulus presentation and become less active during rest periods. So, on every TR, ICA estimates a weight for each neural network that measures how active that network is on that TR.

ICA has some significant advantages over standard, univariate GLM-based methods of data analysis. First, because it operates on all data simultaneously, ICA largely avoids the intractable multiple comparisons problem that plagues univariate analyses. Second, ICA identifies networks of event-related voxels, rather than single voxels, so it simultaneously addresses questions of functional connectivity. Third, GLM approaches assume that every time an event occurs during an experimental session, it elicits exactly the same BOLD response. In contrast, ICA allows the weights to differ on every TR so it allows the gradual ramping up or down of a network across TRs that might be seen during learning or habituation. Fourth, ICA makes no assumptions about the hrf or the nature of the BOLD response. In particular, it does not assume linearity between neural activation and the BOLD response. On the other hand, ICA does have weaknesses. First, it is time and resource consuming to run. Second, it is a purely exploratory data-analytic technique in the sense that it provides no straightforward method of testing specific *a priori* hypotheses about any of the components. Finally, it provides no foolproof method of identifying task-related components. In

many cases, an ICA analysis might identify several hundred components, only a few of which are likely related to the task under study. Finding these few components of interest among the hundreds identified can be a difficult challenge.

Recently, another multivariate method for analyzing fMRI data has become popular. This method, called multi-voxel pattern analysis (MVPA), applies machine learning classification methods to BOLD response data [39–41]. The idea is that if some brain region is responding differently to two different event types then it should be possible to find a classification scheme that can look at the responses of all voxels in that region and correctly identify which event triggered the response. This is the technique that is used in all of the well publicized claims that fMRI can be used to read one's mind. The first step in MVPA is to create a vector that represents the BOLD response to a specific event in a region of interest. In the simplest case, for every voxel in the region, the vector might have one entry that measures the BOLD response in that voxel to a single specific event [42]. For example, suppose we are interested in whether a region in ventral temporal cortex that includes 2000 voxels responds differently to pictures of shoes versus chairs [43]. For each picture presented to the subject, we estimate the BOLD response in every voxel in this region. Now imagine a 2000 dimensional space with a coordinate axis for each of the 2000 voxels. Each vector specifies a numerical value on each of these dimensions, and therefore we could plot the entries in each vector as a single point in this high dimensional space. The idea is that vectors from trials when a shoe is presented should cluster in a different part of the space compared to vectors from trials when a chair is presented if this brain region responds differently to these two stimulus classes. On the other hand, if this region does not discriminate between shoes and chairs then the points should all fall in the same region. Of course, it would be impossible to decide whether the points fall in the same or different regions by visual inspection. Instead, machine learning classification techniques are used (e.g., the support vector machine or the naïve Bayes classifier).

5.7 The Future

Despite its limitations, the future for fMRI is bright. It is likely to play an enduring role in psychology, cognitive neuroscience, and the mind sciences in general—at least until it is replaced by some similar, but more powerful technology (e.g., just as fMRI has now largely replaced PET scanning). There are several reasons for this optimism. Perhaps the most obvious is that fMRI allows scientists to investigate questions that before seemed unapproachable. But another reason that fMRI is likely to maintain its popularity is that it is rapidly improving on almost all fronts. New scanners and head coils are more reliable and produce cleaner data with higher signal-to-noise ratio. New pulse sequences allow for innovative types of scanning (e.g., as when diffusion spectrum imaging was developed as a superior alternative to diffusion tensor imaging). New methods of data analysis allow researchers to draw unique conclusions even from data collected using common pulse sequences on older

established machines. All these improvements also increase the flexibility of fMRI. Researchers continue to have more choices than ever when designing and running fMRI experiments and when analyzing the resulting data. As a result, potential applications of fMRI are largely limited by the creativity of fMRI researchers.

Exercises

1. Think of an experiment that is best addressed using fMRI instead of EEG or MEG. What are the key advantages of fMRI?
2. In an fMRI experiment with a TR of 2 s, the temporal resolution is considerably better than 2 s. How can the temporal resolution in fMRI be better than the TR?
3. Suppose we measure the height of 15 subjects, and then run each of them in an fMRI experiment. During data analysis we compute a whole-brain t-map for each subject. Next, for every voxel, suppose we correlate (across subjects) subject height with the value of the t statistic in that voxel. What would you conclude, if after correcting for multiple comparisons, we find a set of voxels where the correlation is significant? Does this outcome seem likely? Did the experiment identify a network that thinks about the subject's height?

Further Reading

Several books provide overviews of the whole fMRI field [20, 44], while others provide more depth on certain subtopics. For example, Hashemi, Bradley, and Lisanti [45] give a mostly nontechnical description of MR physics, whereas Haacke, Brown, Thompson, and Venkatesan [46] provide a much more rigorous treatment. In contrast, Ashby [33] and Poldrack, Mumford, and Nichols [47] focus exclusively on experimental design and fMRI data analysis.

Acknowledgments This research was supported in part by AFOSR grant FA9550-12-1-0355 and by the U.S. Army Research Office through the Institute for Collaborative Biotechnologies under grant W911NF-07-1-0072.

References

1. Dobbs D (2005) Fact or phrenology? Scientific American Mind
2. Waldschmidt JG, Ashby FG (2011) Cortical and striatal contributions to automaticity in information-integration categorization. *Neuroimage 56*:1791 -1802
3. Ashby FG, O'Brien JB (2005) Category learning and multiple memory systems. Trends Cogn Sci 2:83–89
4. Lopez-Paniagua D, Seger CA (2011) Interactions within and between corticostriatal loops during component processes of category learning. J Cogn Neurosci 23:3068 -3083

5. Seger CA, Cincotta CM (2005) The roles of the caudate nucleus in human classification learning. J Neurosci 25:2941–2951
6. Seger CA, Peterson EJ, Cincotta CM, Lopez-Paniagua D, Anderson CW (2010) Dissociating the contributions of independent corticostriatal systems to visual categorization learning through the use of reinforcement learning modeling and Granger causality modeling. NeuroImage 50:644–656
7. Pauling L, Coryell CD (1936) The magnetic properties and structure of hemoglobin, oxygenated hemoglobin, and carbonmonoxygenated hemoglobin. Proc Nat Acad Sci U S A 22:210–236
8. Friston KJ, Ashburner JT, Kiebel SJ, Nichols TE, Penny WD (eds) (2007) Statistical parametric mapping: the analysis of functional brain images. Academic, London
9. Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TEJ, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, Niazy R, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady JM, Matthews PM (2004) Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage 23:208–219
10. Woolrich MW, Jbabdi S, Patenaude B, Chappell M, Makni S, Behrens T, Beckmann C, Jenkinson M, Smith SM (2009) Bayesian analysis of neuroimaging data in FSL. NeuroImage 45:173–186
11. Ogawa S, Lee TM, Kay AR, Tank DW (1990) Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc Nat Acad Sci 87:9868–9872
12. Ogawa S, Lee TM, Nayak AS, Glynn P (1990) Oxygenation-sensitive contrast in magnetic resonance imaging of rodent brain at high magnetic fields. Magn Reson Med 16:9–18
13. Logothetis NK (2003) The underpinnings of the BOLD functional magnetic resonance imaging signal. J Neurosci 23:3963–3971
14. Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. Nature 412:150–157
15. Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. J Neurosci 16:4207–4221
16. Buxton RB, Frank LR (1998) A model for coupling between cerebral blood flow and oxygen metabolism during neural stimulation. J Cerebral Blood Flow Metab 17:64–72
17. Vazquez AL, Noll DC (1998) Non-linear aspects of the blood oxygenation response in functional MRI. NeuroImage 8:108-118
18. Huettel SA, Singerman JD, McCarthy G (2001) The effects of aging upon the hemodynamic response measured by functional MRI. NeuroImage 13:161–175
19. Richter W, Richter M (2003) The shape of the fMRI BOLD response in children and adults changes systematically with age. NeuroImage 20:1122–1131
20. Huettel SA, Song AW, McCarthy G (2004) Functional magnetic resonance imaging. Sinauer, Sunderland
21. Ashburner J, Friston K (2007) Rigid body registration. In: Friston KJ, Ashburner JT, Kiebel SJ, Nichols TE, Penny WD (eds) Statistical parametric mapping: the analysis of functional brain images. Academic, London, pp 49–62
22. Klein A, Andersson J, Ardekani BA, Ashburner J, Avants B, Chiang M-C, Christensen GE, Collins DL, Gee J, Hellier P, Song JH, Jenkinson M, Lepage C, Rueckert D, Thompson P, Vercauteren T, Woods RP, Mann JJ, Parsey RV (2009) Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. NeuroImage 46:786 -802
23. Talairach J, Tournoux P (1988) Co-planar stereotaxic atlas of the human brain: 3-Dimensional proportional system—an approach to cerebral imaging. Thieme Medical Publishers, New York
24. Friston KJ, Frith CD, Liddle PF, Frackowiak RS (1991) Comparing functional (PET) images: the assessment of significant change. J Cerebral Blood Flow Metab 11:690–699
25. Friston K, Holmes A, Worsley K, Poline J, Frith C, Frackowiak R (1995) Statistical parametric maps in functional imaging: a general linear approach. Hum Brain Mapp 2:189–210
26. O'Doherty JP, Hampton A, Kim H (2007) Model-based fMRI and its application to reward learning and decision making. Ann NY Acad Sci 1104:35–53
27. O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science 304:452-454

28. Worsley KJ (1995) Estimating the number of peaks in a random field using the Hadwiger characteristic of excursion sets with applications to medical images. Ann Stat 23:640–669
29. Worsley KJ, Marrett S, Neelin P, Vandal AC, Friston KJ, Evans AC (1996) A unified statistical approach for determining significant signals in images of cerebral activation. Hum Brain Mapp 4:58–73
30. Nichols TE, Holmes AP (2001) Nonparametric permutation tests for functional neuroimaging: A primer with examples. Hum Brain Mapp 15:1–25
31. Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. J R Stat Soc Series B Methodol 57:289–300
32. Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37:424–438
33. Ashby FG (2011) Statistical analysis of fMRI data. MIT Press, Boston
34. Sun FT, Miller LM, D'Esposito M (2004) Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. NeuroImage 21:647-658
35. Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R (2004) Intersubject synchronization of cortical activity during natural vision. Science 303:1634–1640
36. Calhoun VD, Adali T, Pearlson GD, Pekar JJ (2001) Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. Hum Brain Mapp 13:43–53
37. McKeown MJ, Makeig S, Brown GG, Jung T-P, Kindermann SS, Bell AJ, Sejnowski TJ (1998) Analysis of fMRI data by blind separation into independent spatial components. Hum Brain Mapp 6:160–188
38. Buckner RL, Andrews-Hanna JR, Schacter DL (2008) The brain's default network: Anatomy, function, and relevance to disease. Ann NY Acad Sci 1124:1–38
39. Haynes J, Rees G (2006) Decoding mental states from brain activity in humans. Nat Rev Neurosci 7:523–534
40. Norman K, Polyn SM, Detre G, Haxby JV (2006) Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. Trends Cogn Sci 10:424–430
41. Pereira F, Mitchell T, Botvinick M (2009) Machine learning classifiers and fMRI: a tutorial overview. NeuroImage 45(1 Suppl):S199–209
42. Mumford JA, Turner BO, Ashby FG, Poldrack RA (2012) Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. NeuroImage 59:2636–2643
43. Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science 293:2425–2430
44. Buxton RB (2002) Introduction to functional magnetic resonance imaging: principles and techniques. Cambridge University Press, New York
45. Hashemi RH, Bradley WG Jr, Lisanti CJ (2004) MRI: the basics, 2nd Ed. Lippincott Williams & Wilkins, Philadelphia
46. Haacke EM, Brown RW, Thompson MR, Venkatesan R (1999) Magnetic resonance imaging: physical principles and sequence design. Wiley, New York
47. Poldrack RA, Mumford JA, Nichols TE (2011) Handbook of fMRI data analysis. Cambridge University Press, New York