# 1 Introduction

Functional magnetic resonance imaging (fMRI) provides researchers the 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. The effects have probably been greatest within cognitive psychology and perception, but the influence of fMRI has spread to almost every area of the mind sciences. For example, there are now emerging new fields of social neuroscience, developmental neuroscience, and neuroeconomics, all largely because of fMRI. There is even a new field of neuromarketing.

fMRI has provided exciting new opportunities to study topics that had long seemed out of reach of rigorous scientific investigation. For example, the past few years have seen studies published in reputable journals in which researchers used fMRI to study the nature of consciousness (e.g., Lloyd, 2002), the effects of meditation on brain function (e.g., Cahn & Polich, 2006), and the neural basis of moral judgments (e.g., Greene, Sommerville, Nystrom, Darley, & Cohen, 2001). Researchers interested in using this new technology in their own research, however, have some notable challenges ahead when the time comes to analyze the data they collect. An fMRI experiment produces massive amounts of highly complex data. The statistical methods that most mind science researchers were trained on in graduate school, such as analysis of variance (ANOVA) and regression, provide a useful background for the most basic methods of fMRI analysis, but even the most straightforward fMRI analysis is considerably more complex than most classic treatments of ANOVA and regression. Furthermore, many other statistical methods that are now routinely used to analyze fMRI data are almost never covered in traditional statistics courses. Among many other topics, this includes, for example, Gaussian random field theory, false discovery rate, coherence analysis, Granger causality, and independent component analysis. Even worse (or better, depending on your perspective), complex new fMRI data analysis techniques are being proposed all the time. In fact, the statistical analysis of fMRI data is now a popular research area in statistics departments around the world.

This text introduces and surveys the most widely used current statistical methods of analyzing fMRI data. Every step is covered—from preprocessing to advanced methods for

assessing functional connectivity. Because understanding the data analysis process is always a critical prerequisite to designing an efficient and powerful experiment, a naïve reader who works through this book will learn much about fMRI experimental design, though it should be understood by readers at the outset that this text focuses exclusively on data analysis. Just as a text on ANOVA and regression would not typically describe the computer equipment that may have been used to collect the data, this book does not include a description of the complex machinery and equipment one finds in a typical brain-imaging center or of 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). Neither is there any description of the physics of fMRI. The reader interested in learning more about these topics is urged to consult any of the several good books that concentrate on these issues (e.g., Buxton, 2002; Huettel, Song, & McCarthy, 2004).

This book assumes no previous background in fMRI. It does assume some statistical background, however, including basic univariate statistical inference (e.g., t-tests) and some exposure to ANOVA and regression. Anyone who has completed the basic first-year statistics sequence that is required, for example, in almost every doctoral psychology program in the United States should have more than enough statistical background to understand this book. Multivariate normal distributions will also be encountered in several chapters, but a brief overview of the necessary material on this topic is provided in appendix B. At the mathematics level, a few integrals will be encountered that require some calculus to understand. Nevertheless, motivated readers without calculus should be able to follow 95% of the material in the book. Basic matrix algebra is also used extensively, but a survey of everything a reader would need to know about this topic is included in appendix A.

### What Is fMRI?

Magnetic resonance imaging (MRI) provides a method to study the structure and function of the brain by measuring differences in the magnetic properties of certain molecules. The first human MRI scanner was built in 1977, and in 1985 the Food and Drug Administration approved MRI for clinical use. Within 10 years, thousands of MRI instruments were installed in hospitals throughout the United States, and today MRI is a routine medical procedure. Perhaps the most important reason for the dramatic rise in popularity in MRI for diagnostic and scientific purposes is that MRI is completely noninvasive and carries few health risks. In this sense, MRI was a significant improvement over other available neuroimaging techniques. For example, computed tomography (CT) scanning uses x-rays, and positron emission tomography (PET) scanning requires injecting the subject with a drug containing a radioactive label.

Currently, the most common clinical application of MRI is to assess brain structure (or the structure of other tissue) by measuring the density of water molecules (most typically).

Because the density of water is different in air, white matter, gray matter, blood vessels, and tumors, MRI becomes an effective method for visualizing brain structure. In most hospitals, structural MRI is a routine procedure, but few hospitals currently perform fMRI, though medical applications of fMRI are likely to rise in the future. For example, one useful clinical application of fMRI could be as a presurgical procedure to map out the functional architecture of a patient's brain. Such a map would be useful to a neurosurgeon who wants to avoid excising tissue associated with some critical skill (e.g., speech).

The goal of fMRI is to observe the brain as it is functioning in as close to real time as possible. The ideal fMRI methodology would measure neural activity with high spatial resolution in real time. This goal has not yet been realized, and in fact, the best available current methods fall far short of this goal. For example, currently the typical fMRI experiment records a sluggish, indirect measure of neural activity with a temporal resolution of 1–3 seconds and a spatial resolution of 3–5 $mm^3$. Nevertheless, as the thousands of fMRI publications attest, this highly imperfect technology has dramatically influenced the study of mind and brain.

The vast majority of fMRI experiments measure the blood oxygen level–dependent (BOLD) signal. The physics of this process is complex and far beyond the scope of this text. Interested readers should consult Hashemi, Bradley, and Lisanti (2004) for a mostly nontechnical description; for those readers with a background in physics, Haacke, Brown, Thompson, and Venkatesan (1999) provide a much more rigorous treatment. 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 six oxygen molecules. A key discovery that eventually led to BOLD fMRI was the observation that hemoglobin molecules fully loaded with oxygen have different magnetic properties than those of hemoglobin molecules with empty binding sites (Pauling & Coryell, 1936).

The theory, which is not yet fully worked out, is that active brain areas consume more oxygen than do 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. Immediately after the neural activity, there is (typically) an oxygen debt, so the ratio of oxygenated to deoxygenated hemoglobin often falls below baseline levels. The rush of oxygenated hemoglobin into the area causes the ratio (i.e., the BOLD signal) to rise quickly. As it happens, the vascular system overcompensates, in the sense that the ratio of oxygenated to deoxygenated hemoglobin actually rises well above baseline to a peak at around 6 seconds after the neural activity that elicited these responses. After this peak, the BOLD signal gradually decays back to baseline over a period of 20–25 seconds.

Current evidence suggests that the neural activity most closely related to changes in the BOLD signal is the local field potential (Logothetis, 2003; Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001). This is the summed total electrical activity in a small region

around the recording site. Chapter 3 explores these issues in more detail and reviews mathematical models of the relationship between neural activation and the BOLD signal. These models are critical in fMRI data analysis because in most cases, we are interested in an unobservable, latent construct; namely, neural activity. For this reason, we can use fMRI data to make inferences about neural activity only if we have a rough understanding of how neural activity is related to the observable BOLD response.

### The Scanning Session

An experimental session that collects fMRI data also commonly includes a variety of other types of scans. Typically, the first scan completed in each session is the *localizer*. This is a very quick structural scan (1–2 minutes) of low spatial resolution that is used only to locate the subject's brain in three-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 performed is not critical. Frequently, however, the *high-resolution structural scan* would follow the localizer. Depending on the resolution of this scan and on the exact nature of the pulse sequences that are used to control the scanner during acquisition, it may take 8–10 minutes to complete this protocol. This 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 (see chapter 4 for details). 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 minutes to complete or can be broken down into two or three shorter runs, with brief rests in between. There are many parameter choices to make here, but two are especially important for the subsequent analysis of fMRI data. One choice is the time between successive whole-brain scans, which is called the *repetition time* and is abbreviated TR. If the whole brain is scanned, typical TRs range from 2 to 3 seconds, but TRs of 1 second or faster 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 might be 3 mm × 3 mm × 3.5 mm. In this case, in a typical human brain, 33 separate slices might be acquired, each containing a 64 × 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-minute run with a TR of 2 seconds, 135,168 BOLD responses could be recorded 900 separate times (i.e., 30 times per minute × 30 minutes), 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 two other types of scans are also common. One is to collect a *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 of the bore will distort 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 reappear 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.

A final common type of scan is *diffusion tensor imaging* (DTI). The goal here is to measure the major fiber tracts (e.g., bundles of axons) of the subject's brain. Although all human brains will theoretically contain the same major tracts, there is surprising variability across individuals in the robustness or thickness of these tracts. DTI can be useful for correlating performance in a task across subjects with the measured robustness of a theoretically relevant fiber tract.

DTI measures the distance and direction along which water molecules diffuse during a short but fixed amount of time. During any fixed time interval, water molecules outside of cells will tend to diffuse the same distance in every direction, but inside of a neuron, for example, water will diffuse farther up and down the length of an axon than it will diffuse in a direction perpendicular to the axon. Thus, the first step in analyzing DTI data is to find locations where diffusion in one direction is much greater than in any other direction. By linking together such directions from neighboring points, it may be possible to trace out fiber tracts. This linking process is called *tractography*. DTI is not covered in this book. Readers interested in learning about this important topic should consult any of several excellent reviews (e.g., Le Bihan et al., 2001; Mori, 2007).

### Experimental Design

In clinical applications where the only goal is to collect a high-resolution structural scan, the subject lies passively inside the scanner during the entire procedure. In fMRI, however, subjects are typically given some task to perform. In the standard setup, a mirror is attached to the top of the bore and directed at the subject's eyes. A computer-controlled projector directs visual information onto this mirror, and the subject responds to this material, often by pressing a button on a hand-held device.

fMRI experiments use either a *block design* or an *event-related design*. In a block design, the functional run consists of a series of blocks, each of which may last for somewhere between 30 seconds 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 (described more fully in chapter 5). 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. This is done in the following way. At the onset of each TR, the computer controlling the scanner sends a pulse to the computer that controls the experiment. The experiment is programmed in such a way that stimulus presentation is delayed until exactly the time when this pulse is received.

The first event-related designs included long rests between each pair of successive trials. In these *slow event-related designs*, rests of 30 seconds are typical. These are included so that the BOLD response in brain regions that participate in stimulus processing can decay back to baseline levels before the presentation of the next stimulus. This makes statistical sense, but it is expensive as 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. The most widely used of these methods are described in chapter 5. 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 might cover a range between, say, 2 and 16 seconds. 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, 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 (a process known as *jittering*). This topic is covered in detail in chapter 5.

### Data Analysis

A number of features of fMRI data make it especially challenging to analyze. 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 a voxel in which the BOLD response is greater than average on some particular TR 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. In other words, effect sizes are small. The noise has several sources. These can roughly be broken down into true noise and unaccounted-for signal.

One source of true noise is thermal motion of any electrons that are inside the bore of the magnet (i.e., including the brain) or in the equipment that is used to collect and process the raw data. A second true noise source is physiologic. For example, the same metabolic demand in the same brain region does not always elicit exactly the same BOLD response.

A number of other factors that contribute to observed noise in fMRI data might more accurately be described as unaccounted-for signal. These include head motion, scanner drift, and uncontrolled cognitive activity on the part of the subject. Relatively large head movements can be caused by the subject shifting his or her head position. Theoretically, these are corrected during preprocessing (as long as they are not too large; see chapter 4). Smaller movements occur as a result of heartbeat and respiration. One way to deal with these artifacts is to use biosensors to record heartbeat and respiration and then use these data as regressors during data analysis (see chapter 5). Scanner drift occurs when the strength of the magnetic field inside the bore slowly changes over the course of the scanning session. The possibility of such drift is often explicitly modeled during data analysis.

Finally, of course, anything the subject is thinking about that is unrelated to the task being studied will produce neural activation and changes in BOLD response. This is usually impossible to correct because such extraneous cognitive activity could presumably occur at any time and within almost any voxel. This seems especially likely in slow, event-related designs when the subject has long time periods without anything to do.

The analysis of fMRI BOLD data is broken down into two general stages: preprocessing and postprocessing. Preprocessing is covered in chapter 4, and the rest of this book is devoted to postprocessing. Preprocessing includes a number of steps that are required to prepare the data for statistical analysis. These include, 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, postprocessing includes all analyses that are directed at these questions. This is a complex and rapidly changing field of statistics and is the main focus of this book.

### Software Packages

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 available software is frequently updated, so no attempt will be made here to thoroughly review each package.

The most widely used package is Statistical Parametric Mapping (SPM), 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. At the time of this writing, the most current version is SPM8, which was released in April 2009. A thorough description of the statistical foundations of SPM was provided by Friston, Ashburner, Kiebel, Nichols, and Penny (2007).

Another widely used fMRI data analysis software package is called FSL, which is an acronym for 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://www.fmrib.ox.ac.uk/fsl/index.html>. Descriptions of the statistical foundations of the FSL routines were provided by Smith et al. (2004) and by Woolrich et al. (2009).

BrainVoyager is a commercially available software package that contains routines written in C++ to optimize speed and that uses a sophisticated three-dimensional graphics environment. BrainVoyager is a product of the company Brain Innovation B.V. located in The Netherlands. A single license costs upward of $8000. More information about the package

and its purchase are available at <http://www.brainvoyager.com/index.html>. A description of the underlying statistical foundations can be found in Goebel, Esposito, and Formisano (2006).

AFNI is a free software package created and maintained by neuroimaging researchers at the National Institute of Mental Health (NIMH) in Bethesda, Maryland. AFNI is an acronym for Analysis of Functional NeuroImages. AFNI is written in C and runs on Unix or Mac operating systems. It can be downloaded from <http://afni.nimh.nih.gov/afni>. The software is described by Cox (1996) and Cox and Hyde (1997).