Introduction

The goal of this book is to provide the reader with a solid background in the techniques used for processing and analysis of functional magnetic resonance imaging (fMRI) data.

1.1 A brief overview of fMRI

Since its development in the early 1990s, fMRI has taken the scientific world by storm. This growth is easy to see from the plot of the number of papers that mention the technique in the PubMed database of biomedical literature, shown in Figure 1.1. Back in 1996 it was possible to sit down and read the entirety of the fMRI literature in a week, whereas now it is barely feasible to read all of the fMRI papers that were published in the previous week! The reason for this explosion in interest is that fMRI provides an unprecedented ability to safely and noninvasively image brain activity with very good spatial resolution and relatively good temporal resolution compared to previous methods such as positron emission tomography (PET).

1.1.1 Blood flow and neuronal activity

The most common method of fMRI takes advantage of the fact that when neurons in the brain become active, the amount of blood flowing through that area is increased. This phenomenon has been known for more than 100 years, though the mechanisms that cause it remain only partly understood. What is particularly interesting is that the amount of blood that is sent to the area is more than is needed to replenish the oxygen that is used by the activity of the cells. Thus, the activity-related increase in blood flow caused by neuronal activity leads to a relative surplus in local blood oxygen. The signal measured in fMRI depends on this change in oxygenation and is referred to as the blood oxygenation level dependent, or BOLD, signal.

Figure 1.1.   A plot of the number of citations in the PubMed database matching the query ["fMRI" OR "functional MRI" OR "functional magnetic resonance imaging"] for every year since 1992.

Figure 1.2 shows an example of what is known as the *hemodynamic response*, which is the increase in blood flow that follows a brief period of neuronal activity. There are two facts about the hemodynamic response that underlie the basic features of BOLD fMRI and determine how the data must be analyzed. First, the hemodynamic response is slow; whereas neuronal activity may only last milliseconds, the increase in blood flow that follows this activity takes about 5 seconds to reach its maximum. This peak is followed by a long undershoot that does not fully return to baseline for at least 15–20 seconds. Second, the hemodynamic response can, to a first approximation, be treated as a *linear time-invariant* system (Cohen, 1997; Boynton et al., 1996; Dale, 1999). This topic will be discussed in much greater detail in Chapter 5, but in essence the idea is that the response to a long train of neuronal activity can be determined by adding together shifted versions of the response to a shorter train of activity. This linearity makes it possible to create a straightforward statistical model that describes the timecourse of hemodynamic signals that would be expected given some particular timecourse of neuronal activity, using the mathematical operation of *convolution*.

1.1.2 Magnetic resonance imaging

The incredible capabilities of magnetic resonance imaging (MRI) can hardly be overstated. In less than 10 minutes, it is possible to obtain images of the human

Figure 1.2. An example of the hemodynamic responses evoked in area V1 by a contrast-reversing checker-board displayed for 500 ms. The four different lines are data from four different individuals, showing how variable these responses can be across people. The MRI signal was measured every 250 ms, which accounts for the noisiness of the plots. (Data courtesy of Stephen Engel, University of Minnesota)

brain that rival the quality of a postmortem examination, in a completely safe and noninvasive way. Before the development of MRI, imaging primarily relied upon the use of ionizing radiation (as used in X-rays, computed tomography, and positron emission tomography). In addition to the safety concerns about radiation, none of these techniques could provide the flexibility to image the broad range of tissue characteristics that can be measured with MRI. Thus, the establishment of MRI as a standard medical imaging tool in the 1980s led to a revolution in the ability to see inside the human body.

1.2 The emergence of cognitive neuroscience

Our fascination with how the brain and mind are related is about as old as humanity itself. Until the development of neuroimaging methods, the only way to understand how mental function is organized in the brain was to examine the brains of individuals who had suffered damage due to stroke, infection, or injury. It was through these kinds of studies that many early discoveries were made about the localization of mental functions in the brain (though many of these have come into question

subsequently). However, progress was limited by the many difficulties that arise in studying brain-damaged patients (Shallice, 1988).

In order to better understand how mental functions relate to brain processes in the normal state, researchers needed a way to image brain function while individuals performed mental tasks designed to manipulate specific mental processes. In the 1980s several groups of researchers (principally at Washington University in St. Louis and the Karolinska Institute in Sweden) began to use positron emission tomography (PET) to ask these questions. PET measures the breakdown of radioactive materials within the body. By using radioactive tracers that are attached to biologically important molecules (such as water or glucose), it can measure aspects of brain function such as blood flow or glucose metabolism. PET showed that it was possible to localize mental functions in the brain, providing the first glimpses into the neural organization of cognition in normal individuals (e.g., Posner et al., 1988). However, the use of PET was limited due to safety concerns about radiation exposure, and due to the scarce availability of PET systems.

fMRI provided exactly the tool that cognitive neuroscience was looking for. First, it was safe, which meant that it could be used in a broad range of individuals, who could be scanned repeatedly many times if necessary. It could also be used with children, who could not take part in PET studies unless the scan was medically necessary. Second, by the 1990s MRI systems had proliferated, such that nearly every medical center had at least one scanner and often several. Because fMRI could be performed on many standard MRI scanners (and today on nearly all of them), it was accessible to many more researchers than PET had been. Finally, fMRI had some important technical benefits over PET. In particular, its spatial resolution (i.e., its ability to resolve small structures) was vastly better than PET. In addition, whereas PET required scans lasting at least a minute, with fMRI it was possible to examine events happening much more quickly. Cognitive neuroscientists around the world quickly jumped on the bandwagon, and thus the growth spurt of fMRI began.

1.3 A brief history of fMRI analysis

When the first fMRI researchers collected their data in the early 1990s, they also had to create the tools to analyze the data, as there was no "off-the-shelf" software for analysis of fMRI data. The first experimental designs and analytic approaches were inspired by analysis of blood flow data using PET. In PET blood flow studies, acquisition of each image takes at least one minute, and a single task is repeated for the entire acquisition. The individual images are then compared using simple statistical procedures such as a t-test between task and resting images. Inspired by this approach, early studies created activation maps by simply subtracting the average activation during one task from activation during another. For example, in the study by Kwong et al. (1992), blocks of visual stimulation were alternated with blocks of no stimulation. As shown in Figure 1.3, the changes in signal in the visual cortex

Figure 1.3.   Early fMRI images from Kwong et al. (1992). The left panel shows a set of images starting with the baseline image (top left), and followed by subtraction images taken at different points during either visual stimulation or rest. The right panel shows the timecourse of a region of interest in visual cortex, showing signal increases that occur during periods of visual stimulation.

were evident even from inspection of single subtraction images. In order to obtain statistical evidence for this effect, the images acquired during the stimulation blocks were compared to the images from the no-stimulation blocks using a simple paired t-test. This approach provided an easy way to find activation, but its limitations quickly became evident. First, it required long blocks of stimulation (similar to PET scans) in order to allow the signal to reach a steady state. Although feasible, this approach in essence wasted the increased temporal resolution available from fMRI data. Second, the simple t-test approach did not take into account the complex temporal structure of fMRI data, which violated the assumptions of the statistics.

Researchers soon realized that the greater temporal resolution of fMRI relative to PET permitted the use of event-related (ER) designs, where the individual impact of relatively brief individual stimuli could be assessed. The first such studies used trials that were spaced very widely in time (in order to allow the hemodynamic response to return to baseline) and averaged the responses across a time window centered around each trial (Buckner et al., 1996). However, the limitations of such slow event-related designs were quickly evident; in particular, it required a great amount of scan time to collect relatively few trials. The modeling of trials that occurred more rapidly in time required a more fundamental understanding of the BOLD hemodynamic response (HRF). A set of foundational studies (Boynton et al., 1996; Vazquez & Noll, 1998; Dale & Buckner, 1997) established the range of event-related fMRI designs for which

the BOLD response behaved as a linear time invariant system, which was roughly for events separated by at least 2 seconds. The linearity of the BOLD is a crucial result, dramatically simplifying the analysis by allowing the use of the General Linear Model and also allowing the study of the statistical efficiency of various fMRI designs. For example, using linearity Dale (1999) and Josephs & Henson (1999) demonstrated that block designs were optimally sensitive to differences between conditions, but careful arrangement of the events could provide the best possible ER design.

The noise in BOLD data also was a challenge, particularly with regard to the extreme low frequency variation referred to as "drift." Early work systematically examined the sources and nature of this noise and characterized it as a combination of physiological effects and scanner instabilities (Smith et al., 1999; Zarahn et al., 1997; Aguirre et al., 1997), though the sources of drift remain somewhat poorly understood. The drift was modeled by a combination of filters or nuisance regressors, or using temporal autocorrelation models (Woolrich et al., 2001). Similar to PET, global variation in the BOLD signal was observed that was unrelated to the task, and there were debates as to whether global fMRI signal intensity should be regressed out, scaled-away, or ignored (Aguirre et al., 1997).

In PET, little distinction was made between intrasubject and group analyses, and the repeated measures correlation that arises from multiple (at most 12) scans from a subject was ignored. With fMRI, there are hundreds of scans for each individual. An early approach was to simply concatenate the time series for all individuals in a study and perform the analysis across all timepoints, ignoring the fact that these are repeated measures obtained across different individuals. This produced "fixed effects" inferences in which a single subject could drive significant results in a group analysis. The SPM group (Holmes & Friston, 1999) proposed a simple approach to "mixed effects" modeling, whose inferences would generalize to the sampled population. Their approach involved obtaining a separate effect estimate per subject at each voxel and then combining these at a second level to test for effects across subjects. Though still widely in use today, this approach did not account for differences in intrasubject variability. An improved approach was proposed by the FMRIB Software Library (FSL) group (Woolrich et al., 2004b; Beckmann & Smith, 2004) that used both the individual subject effect images and the corresponding standard error images. Although the latter approach provides greater sensitivity when there are dramatic differences in variability between subjects, recent work has shown that these approaches do not differ much in typical single-group analyses (Mumford & Nichols, 2009).

Since 2000, a new approach to fMRI analysis has become increasingly common, which attempts to analyze the information present in patterns of activity rather than the response at individual voxels. Known variously as multi-voxel pattern analysis (MVPA), pattern information analysis, or machine learning, these methods attempt to determine the degree to which different conditions (such as different stimulus classes) can be distinguished on the basis of fMRI activation patterns, and also

to understand what kind of information is present in those patterns. A particular innovation of this set of methods is that they focus on making predictions about new data, rather than simply describing the patterns that exist in a particular data set.

1.4 Major components of fMRI analysis

The analysis of fMRI data is made complex by a number of factors. First, the data are liable to a number of artifacts, such as those caused by head movement. Second, there are a number of sources of variability in the data, including variability between individuals and variability across time within individuals. Third, the dimensionality of the data is very large, which causes a number of challenges in comparison to the small datasets that many scientists are accustomed to working with. The major components of fMRI analysis are meant to deal with each of these problems. They include

- **Quality control**: Ensuring that the data are not corrupted by artifacts.
- **Distortion correction**: The correction of spatial distortions that often occur in fMRI images.
- **Motion correction**: The realignment of scans across time to correct for head motion.
- **Slice timing correction**: The correction of differences in timing across different slices in the image.
- **Spatial normalization**: The alignment of data from different individuals into a common spatial framework so that their data can be combined for a group analysis.
- **Spatial smoothing**: The intentional blurring of the data in order to reduce noise.
- **Temporal filtering**: The filtering of the data in time to remove low-frequency noise.
- **Statistical modeling**: The fitting of a statistical model to the data in order to estimate the response to a task or stimulus.
- **Statistical inference**: The estimation of statistical significance of the results, correcting for the large number of statistical tests performed across the brain.
- **Visualization**: Visualization of the results and estimation of effect sizes.
  The goal of this book is to outline the procedures involved in each of these steps.

1.5 Software packages for fMRI analysis

In the early days of fMRI, nearly every lab had its own home-grown software package for data analysis, and there was little consistency between the procedures across different labs. As fMRI matured, several of these in-house software packages began to be distributed to other laboatories, and over time several of them came to be distributed as full-fledged analysis suites, able to perform all aspects of analysis of an fMRI study.

Table 1.1.   An overview of major fMRI software packages

Package	Developer	Platforms[a]	Licensing
SPM	University College London	MATLAB	Open-source
FSL	Oxford University	UNIX	Open source
AFNI	NIMH	UNIX	Open source
Brain Voyager	Brain Innovation	Mac OS X, Windows, Linux	Commercial (closed-source)

[a]Those platform listed as UNIX are available for Linux, Mac OS X, and other UNIX flavors.

Today, there are a number of comprehensive software packages for fMRI data analysis, each of which has a loyal following. (See Table 1.1) The Web sites for all of these packages are linked from the book Web site.

1.5.1 SPM

SPM (which stands for Statistical Parametric Mapping) was the first widely used and openly distributed software package for fMRI analysis. Developed by Karl Friston and colleagues in the lab then known as the Functional Imaging Lab (or FIL) at University College London, it started in the early 1990s as a program for analysis of PET data and was then adapted in the mid-1990s for analysis of fMRI data. It remains the most popular software package for fMRI analysis. SPM is built in MATLAB, which makes it accessible on a very broad range of computer platforms. In addition, MATLAB code is relatively readable, which makes it easy to look at the code and see exactly what is being done by the programs. Even if one does not use SPM as a primary analysis package, many of the MATLAB functions in the SPM package are useful for processing data, reading and writing data files, and other functions. SPM is also extensible through its toolbox functionality, and a large number of extensions are available via the SPM Web site. One unique feature of SPM is its connectivity modeling tools, including psychophysiological interaction (Section 8.2.4) and dynamic causal modeling (Section 8.3.4). The visualization tools available with SPM are relatively limited, and many users take advantage of other packages for visualization.

1.5.2 FSL

FSL (which stands for FMRIB Software Library) was created by Stephen Smith and colleagues at Oxford University, and first released in 2000. FSL has gained substantial popularity in recent years, due to its implementation of a number of

cutting-edge techniques. First, FSL has been at the forefront of statistical modeling for fMRI data, developing and implementing a number of novel modeling, estimation, and inference techniques that are implemented in their FEAT, FLAME, and RANDOMISE modules. Second, FSL includes a robust toolbox for independent components analysis (ICA; see Section 8.2.5.2), which has become very popular both for artifact detection and for modeling of resting-state fMRI data. Third, FSL includes a sophisticated set of tools for analysis of diffusion tensor imaging data, which is used to analyze the structure of white matter. FSL includes an increasingly powerful visualization tool called FSLView, which includes the ability to overlay a number of probabilistic atlases and to view time series as a movie. Another major advantage of FSL is its integration with grid computing, which allows for the use of computing clusters to greatly speed the analysis of very large datasets.

1.5.3 AFNI

AFNI (which stands for Analysis of Functional NeuroImages) was created by Robert Cox and his colleagues, first at the Medical College of Wisconsin and then at the National Institutes of Mental Health. AFNI was developed during the very early days of fMRI and has retained a loyal following. Its primary strength is in its very powerful and flexible visualization abilities, including the ability to integrate visualization of volumes and cortical surfaces using the SUMA toolbox. AFNI's statistical modeling and inference tools have historically been less sophisticated than those available in SPM and FSL. However, recent work has integrated AFNI with the R statistical package, which allows use of more sophisticated modeling techniques available within R.

1.5.4 Other important software packages

BrainVoyager. Brain Voyager, produced by Rainer Goebel and colleagues at Brain Innovation, is the major commercial software package for fMRI analysis. It is available for all major computing platforms and is particularly known for its ease of use and refined user interface.

FreeSurfer. FreeSurfer is a package for anatomical MRI analysis developed by Bruce Fischl and colleagues at the Massachusetts General Hospital. Even though it is not an fMRI analysis package per se, it has become increasingly useful for fMRI analysis because it provides the means to automatically generate both cortical surface models and anatomical parcellations with a minimum of human input. These models can then be used to align data across subjects using surface-based approaches, which may in some cases be more accurate than the more standard volume-based methods for intersubject alignment (see Chapter 4). It is possible to import statistical results obtained using FSL or SPM and project them onto the reconstructed cortical surface, allowing surface-based group statistical analysis.

1.6 Choosing a software package

Given the variety of software packages available for fMRI analysis, how can one choose among them? One way is to listen to the authors of this book, who have each used a number of packages and eventually have chosen FSL as their primary analysis package, although we each use other packages regularly as well. However, there are other reasons that one might want to choose one package over the others. First, what package do other experienced researchers at your institution use? Although mailing lists can be helpful, there is no substitute for local expertise when one is learning a new analysis package. Second, what particular aspects of analysis are most important to you? For example, if you are intent on using dynamic causal modeling, then SPM is the logical choice. If you are interested in using ICA, then FSL is a more appropriate choice. Finally, it depends upon your computing platform. If you are a dedicated Microsoft Windows user, then SPM is a good choice (though it is always possible to install Linux on the same machine, which opens up many more possibilities). If you have access to a large cluster, then you should consider FSL, given its built-in support for grid computing.

It is certainly possible to mix and match analysis tools for different portions of the processing stream. This has been made increasingly easy by the broad adoption of the NIfTI file format by most of the major software packages (see Appendix C for more on this). However, in general it makes sense to stick largely with a single package, if only because it reduces the amount of emails one has to read from the different software mailing lists!

1.7 Overview of processing streams

We refer to the sequence of operations performed in course of fMRI analysis as a *processing stream*. Figure 1.4 provides a flowchart depicting some common processing streams. The canonical processing streams differ somewhat between different software packages; for example, in SPM spatial normalization is usually applied prior to statistical analysis, whereas in FSL it is applied to the results from the statistical analysis. However, the major pieces are the same across most packages.

1.8 Prerequisites for fMRI analysis

Research into the development of expertise suggests that it takes about ten years to become expert in any field (Ericsson et al., 1993), and fMRI analysis is no different, particularly because it requires a very broad range of knowledge and skills. However, the new researcher has to start somewhere. Here, we outline the basic areas of knowledge that we think are essential to becoming an expert at fMRI analysis, roughly in order of importance.

Figure 1.4.   A depiction of common processing streams for fMRI data analysis.

1. *Probability and statistics.* There is probably no more important foundation for fMRI analysis than a solid background in basic probability and statistics. Without this, nearly all of the concepts that are central to fMRI analysis will be foreign.
2. *Computer programming.* It is our opinion that one simply cannot become an effective user of fMRI analysis without strong computer programming skills. There are many languages that are useful for fMRI analysis, including MATLAB, Python, and UNIX shell scripting. The particular language is less important than an underlying understanding of the methods of programming, and this is a place where practice really does make perfect, particularly with regard to debugging programs when things go wrong.
3. *Linear algebra.* The importance of linear algebra extends across many different aspects of fMRI analysis, from statistics (where the general linear model is most profitably defined in terms of linear algebra) to image processing (where many operations on images are performed using linear algebra). A deep understanding of fMRI analysis requires basic knowledge of linear algebra.
4. *Magnetic resonance imaging.* One can certainly analyze fMRI data without knowing the details of MRI acquisition, but a full understanding of fMRI data requires that one understand where the data come from and what they are actually measuring. This is particularly true when it comes to understanding the various ways in which MRI artifacts may affect data analysis.
5. *Neurophysiology and biophysics.* fMRI signals are interesting because they are an indirect measure of the activity of individual neurons. Understanding how

neurons code information, and how these signals are reflected in blood flow, is crucial to interpreting the results that are obtained from fMRI analysis.

6. *Signal and image processing.* A basic understanding of signal and image processing methods is important for many of the techniques discussed in this book. In particular, an understanding of Fourier analysis (Section 2.4) is very useful for nearly every aspect of fMRI analysis.