

1

Principles of functional Magnetic Resonance Imaging

Martin A. Lindquist

Department of Biostatistics; Johns Hopkins University

Tor D. Wager

Department of Psychology & Neuroscience; University of Colorado at Boulder

CONTENTS

1.1	Introduction	3
1.2	The Basics of fMRI Data	5
1.2.1	Principles of Magnetic Resonance Signal Generation	6
1.2.1.1	The MRI Scanner	6
1.2.1.2	Basic MR Physics	7
1.2.1.3	Image Contrast	8
1.2.2	Image Formation	9
1.2.3	From MRI to fMRI	11
1.3	BOLD fMRI	14
1.3.1	Understanding BOLD fMRI	15
1.3.2	Spatial Limitations	17
1.3.3	Temporal Limitations	18
1.3.4	Acquisition artifacts	19
1.4	Modeling Signal and Noise in fMRI	20
1.4.1	BOLD signal	20
1.4.2	Noise and nuisance signal	23
1.5	Experimental Design	25
1.6	Preprocessing	27
1.7	Data Analysis	30
1.7.1	Localization	30
1.7.2	Connectivity	33
1.7.3	Prediction	35
1.8	Resting-state fMRI	36
1.9	Data Format, Databases & Software	38
1.10	Future Developments	40

1.1 Introduction

Functional Magnetic Resonance Imaging (fMRI) is a non-invasive technique for studying brain activation. It measures changes in blood oxygenation and

blood flow related to neuronal activity, providing researchers with the means to study human brain function *in vivo*, either in response to a certain task or when at rest. During the past two decades fMRI has provided researchers with an unprecedented access to the inner workings of the human brain, which in turn has led to new insights into how the brain processes information.

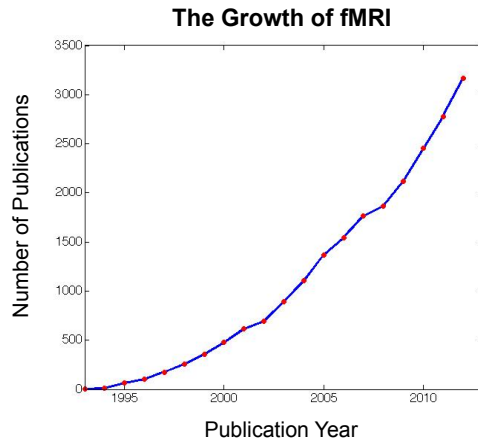
The data acquired in an fMRI study consists of a sequence of 3-D magnetic resonance images (MRIs), each made up of a number of uniformly spaced volume elements, or voxels. The voxels partition the brain into a large number of equally sized cubes. A typical image may consist of roughly 100,000 voxels, where the image intensity value corresponding to each voxel represents the spatial distribution of the nuclear spin density, which relates to blood oxygenation and flow, in the local area. During an fMRI experiment, 100 – 1,000 such 3-D images of the whole brain are acquired. In addition, a standard fMRI experiment consists of multiple subjects (e.g, 10 – 50), potentially brought in for multiple scanning sessions, each consisting of a number of replications of a certain experimental task.

Clearly, the amount of available data from a single experiment is extremely large, and the analysis of fMRI data is an example of the type of modern big-data problem that is fundamentally changing the quantitative sciences. In addition, the data exhibit a complicated temporal and spatial noise structure with a relatively weak signal (though, with appropriate methods, these signals across the brain can be highly predictive of psychological and clinical states). Hence, the available data is not only massive in scale but also complex making the statistical analysis of fMRI data a difficult task.

The field that has grown around the acquisition and analysis of fMRI data has experienced rapid growth in the past several years and found applications in a wide variety of fields, including neuroscience, psychology, medicine, economics and political science. The use of fMRI data is also central to a number of emerging fields, such as cognitive neuroscience, affective neuroscience, social cognitive neuroscience, and neuroeconomics. In these areas fMRI data is being combined with data on performance and psychophysiology to yield new exciting models of human thought, emotion, and behavior.

This explosive growth is illustrated by the exponential increase, shown in Fig. 1.1, of the number of yearly publications in PubMed that mention the term ‘fMRI’ in either its title or abstract. In the early 1990s only a handful of such papers were published yearly, while in more recent years this number has increased to over three thousand papers/year. In addition, more and more methodological papers appear each year, and the field has become fertile ground for the development and application of cutting-edge statistical methods. Researchers entering the field of MRI methods development come from diverse backgrounds: statistics, computer science, engineering, mathematical psychology, mathematics, and physics.

The rapid pace of development, as well as the interdisciplinary nature of the diverse fields that use fMRI data, presents an enormous challenge to researchers. The ability to move the field forward requires strong collaborative

**FIGURE 1.1**

The yearly number of publications in PubMed that mention the term ‘fMRI’ either in its title or abstract between 1993 to 2012.

teams with expertise in a number of disciplines, including psychology, neuroanatomy, neurophysiology, physics, biomedical engineering, signal processing, and statistics. Of course, true interdisciplinary collaboration is extremely challenging, as all members of the research team must know enough about the other disciplines to be able to talk intelligently with experts in each discipline. Hence, making an impact in this exciting new area requires some initial start-up costs.

The goal of this chapter is to review the basic principles involved in the acquisition and analysis of fMRI data in enough detail to highlight the most important issues and concerns. The hope is that this will provide quantitative researchers with a basic understanding about the relevant research questions and how to apply their knowledge to these questions in an appropriate manner. We will also attempt to provide an overall road map to what kinds of study design and analysis options that are available and highlight some of their limitations. This chapter will be more focused on breadth rather than depth, with more detailed descriptions of many of the topics found in the later chapters of the book.

1.2 The Basics of fMRI Data

Functional MRI uses a standard magnetic resonance imaging (MRI) scanner to acquire a series of brain volumes that can be used to study dynamic changes

in brain activation. In order to understand the manner in which fMRI data is acquired, one must first focus on the acquisition of a single static 3-D image. For this reason, while the focus of this chapter is on functional imaging, we must necessarily begin by reviewing data acquisition and reconstruction techniques used to obtain a static MRI of the brain. This closely follows the description in Chapter ** (cite MRI CHAPTER). After this review, we will transition our focus towards the particular issues involved with acquiring data meant for use in an fMRI study.

Proper understanding of the data acquisition and reconstruction procedures associated with fMRI is complex, requiring background in both MR physics and signal processing. Thus, the description that follows is abbreviated. For a more in-depth discussion see, for example, excellent references such as [21] or [23].

1.2.1 Principles of Magnetic Resonance Signal Generation

In this section we outline the physical bases of fMRI. We begin by providing some basic background on the MR scanner and continue by illustrating how it can be used to generate signal, and in turn how this signal can be used to construct an image. While these topics are common with the acquisition of MRI images, we conclude the section with discussing particular issues associated with fMRI.

1.2.1.1 The MRI Scanner

An MRI scanner is a large and versatile piece of hardware. Its main component is a superconducting electromagnet with an extremely strong static magnetic field, typically varying from 1.5 – 7.0 Tesla in human brain research. To place this into context, the Earth's magnetic field is only 0.00005 Tesla. Thus, the field strength is strong enough to pull magnetic objects into its core. Because, the static field is always active, it is critical to observe caution when bringing objects into the MR scanner room. However, it is important to note that there are no known long-term effects on biological tissue, making the technique attractive for scanning humans.

A second critical component of the scanner is the radio frequency coils, hardware coils close to the object being imaged (e.g., the head) that can be used to generate and receive energy at the resonance frequency of the volume being imaged. They are turned on and off during the course of data acquisition.

A third component is the gradient coils, which are electromagnetic coils that can be used to create spatial variation in the strength of the magnetic field in a controlled manner. As we will see this is critical for the ability to encode spatial information into the signal that is necessary for the creation of images.

MR scanners are extremely versatile, as they can be used to study both brain structure and function in multiple ways. Different types of images can

be generated to emphasize contrast related to different tissue characteristics. In addition, the scanner can be used to study the directional patterns of water diffusion – diffusion-weighted imaging (DWI) used to measure white-matter tracts – elastic properties of brain tissue, flow of cerebrospinal fluid, and other properties. Hence, the same scanner is used to acquire structural MRIs, functional MRIs, and perform diffusion tensor imaging (DTI) of white-matter tracts; see Chapters * and ** (cite MRI & DTI CHAPTERS). This is extremely beneficial as it allows for the acquisition of several different types of images on a specific subject during a given scanning session. In particular, structural images are always acquired as part of a standard fMRI scanning session, as they play an integral part in subsequent preprocessing of the data; see Section 6 and Chapter **** (cite PREPROCESSING CHAPTER).

1.2.1.2 Basic MR Physics

All magnetic resonance imaging techniques rely on a core set of physical principles. To properly understand these principles, one should begin by looking at a single atomic nucleus and illustrate its impact on the generated MR signal. In particular we focus on hydrogen atoms consisting of a single proton (^1H atoms), as they are the most commonly used nuclei in MRI due both to their magnetic properties and abundance in the human body.

Protons can be viewed as positively charged spheres that are always spinning about their axis. This gives rise to a net magnetic moment along the direction of the axis of the spins, which is the source of the signal we seek to measure. Unfortunately, it is not possible to measure the magnetization of a single proton using an MRI scanner. Instead, we must focus our attention on measuring the net magnetization of the ensemble of all nuclei within a chosen volume. The net magnetization, denoted M , can be represented as a vector with two components. The first is a longitudinal component, which is parallel to the magnetic field, and the second a transverse component perpendicular to the field.

In the absence of an external magnetic field, the individual nuclei are randomly oriented with respect to one another and therefore do not give rise to a net magnetization. However, when placed into a strong magnetic field, the nuclei align with the field, creating a net longitudinal magnetization in the direction of the field. While aligned the nuclei precess about the field with an angular frequency determined by the Larmor frequency, but at a random phase with respect to one another.

In order to measure the net magnetization of the nuclei within a certain volume, one must perturb the equilibrium and observe the reaction. A radiofrequency (RF) electromagnetic field pulse causes the nuclei to absorb the energy at a particular frequency band, and become “excited”. Conceptually, we can imagine this process as the RF pulse aligning the phase of the precessing nuclei and tipping them over into the transverse plane. This causes

the longitudinal magnetization to decrease, and establishes a new transversal magnetization.

After the RF pulse is removed, the system seeks to return to equilibrium. Now the nuclei emit the absorbed energy as they “relax”. This causes the transverse magnetization to disappear, in a process known as transversal relaxation, while the longitudinal magnetization grows back to its original size in a process referred to as longitudinal relaxation. During this time a signal is created that can be measured using a receiver coil.

Longitudinal relaxation represents the restoration of net magnetization along the longitudinal direction as the nuclei return to their original aligned state. It is seen as an exponential recovery in magnetization described by a time constant T_1 . Transverse relaxation is the loss of net magnetization in the transverse plane due to loss of phase coherence. Since the net magnetization depends upon the combined contribution of a large number of nuclei, its value is largest when all the nuclei are in phase. However, the removal of the RF pulse causes the nuclei to de-phase, causing an exponential decay in magnetization described by a time constant T_2 . Both the T_1 and T_2 values depend upon tissue type, and it is this property that allows for the creation of structural MR images that can be used to differentiate between different tissue types.

The term T_2^* is similar to T_2 , but also depends on local inhomogeneities in the magnetic field caused by changes in blood flow and oxygenation. These inhomogeneities cause the nuclei to de-phase quicker than they normally would. Certain pulse sequences are able to eliminate the effects of these inhomogeneities, while others seek to emphasize them. Thus it is possible to produce images sensitive primarily to T_1 , T_2 , or T_2^* . T_2^* signal provides the basis for functional MRI, as it is sensitive to neurovascular changes that accompany psychological and behavioral function.

1.2.1.3 Image Contrast

One of the reasons for the versatility of the MR scanner is its ability to create images based on a variety of different contrasts that are sensitive to both the number and properties of the nuclei being imaged. To illustrate, assume the initial value of the net magnetization prior to excitation is given by the value M_0 . By altering how often we excite the nuclei (TR) and how soon after excitation we begin data collection (TE) we can control which characteristic of the tissue is emphasized. This relationship can be seen by noting that the measured signal is approximately equal to:

$$M_0(1 - e^{-TR/T_1})e^{-TE/T_2}. \quad (1.1)$$

If one, for example, choose a long TR and short TE value the signal will be approximately equal to M_0 , which in turn is proportional to the number of nuclei (or protons) in the tissue. Hence, these settings can be used to produce so-called ‘proton-density’ images that provide maps over how hydrogen is distributed across the sample. When the TE is short ($\sim 20ms$), but the TR is of

intermediate length, we instead get ' T_1 -weighted' images, which are typically used to reveal anatomical structure. Finally, for ' T_2 -weighted' images, another type of structural image, a long TR and an intermediate TE should be chosen. Because T_1 and T_2 vary with tissue type, T_1 - and T_2 -weighted images can be used to provide detailed representations of the boundaries between gray matter, white matter, and cerebrospinal fluid (CSF).



FIGURE 1.2

Examples of proton density, T_1 , and T_2 -weighted images.

Because T_2^* is sensitive to flow and oxygenation, T_2^* -weighting can be used to create images of brain function. T_2^* -weighted images are obtained in a similar manner to T_2 -weighted images. The difference lies in manner in which the pulse sequence uses the magnetic gradients. This is beyond the scope of this chapter, but interested readers should be referred to [21] or [23]. See Fig. 1.2 for examples of the difference in proton density, T_1 , and T_2 -weighted images. In particular note that the images highlight different anatomical properties of the underlying sample, and their usage depends upon the goals of the study.

1.2.2 Image Formation

The goal of exciting nuclei in the MRI scanner is ultimately to obtain enough information to be able to construct an image of the underlying sample. Any image is represented by a matrix of numbers that correspond to spatial locations. The images generally depict the spatial distribution of some property of the nuclei within the sample. This could be the density of nuclei, their mobility, or the relaxation time of the tissues in which they reside. Pulse sequences define particular manipulations of RF pulses and the shape of the magnetic field that allow us to reconstruct the acquired data into a map of the underlying signal sources, i.e. the hydrogen atoms, and obtain images of the brain.

While most MRIs are 3-D representations of the brain, they are almost always constructed through the acquisition of a series of 2-D slices. This ac-

quisition can be performed in either a sequential or interleaved manner. To illustrate, consider that we are interested in acquiring a total of N_z slices of the brain. Using a sequential scheme the slices are acquired in order, in either an ascending or descending manner. Using an interleaved scheme one instead collects data in an alternating slice order. This can minimize the risk of signal bleeding from an adjacent slice that was previously excited. Both sequential and interleaved sequences have different pros and cons, though a detailed discussion is beyond the scope here.

In general, the process of exciting nuclei only provides information about the net magnetization within the slice. In order to construct a meaningful image of the brain we must find ways to extract information about the spatial contributions to the net magnetization, i.e. how much each individual voxel contributes to its value. This is done using a system of gradient coils that manipulate the magnetic field within the chosen slice and sequentially control its spatial inhomogeneity. In short, this process allows one to express each measurement of the signal as the Fourier transformation of the spin density at a single point in the frequency domain, or k-space as it is commonly called in the field.


Conceptually, we can think of k-space being sampled at a number of discrete locations $(k_x(t_j), k_y(t_j))$ at time t_j . Here $t_j = j\Delta_t$ is the time of the j^{th} measurement, where Δ_t depends on the sampling bandwidth of the scanner; typically it takes values in the range of 250 – 1000 μ s. See Chapter ** (cite MRI CHAPTER) or [35] for more detail. Mathematically, the measurement of the MR signal at the j^{th} time point of a readout period can be written

$$S(t_j) \approx \int_x \int_y M(x, y) e^{-2\pi i(k_x(t_j)x + k_y(t_j)y)} dx dy, \quad (1.2)$$

where $M(x, y)$ is the spin density at the point (x, y) . This is the entity that we seek to measure at each voxel of the brain.

To reconstruct a single MR image, one needs to sample a large number of individual k-space measurements. The exact number depends upon the desired image resolution. For example, to fully reconstruct a 64×64 image, a total of 4096 separate measurements are required, each sampled at a unique coordinate of k-space. There is a time cost involved in sampling each point, and therefore the time it takes to acquire an image is directly related to its spatial resolution.

There exist a variety of approaches towards sampling k-space. In echo-planar imaging (EPI) k-space is sampled uniformly around its origin ([42]). This allows for the quick and easy reconstruction of the image using the Fast Fourier Transform (FFT); see Fig. 1.3 for an illustration. Alternatively, one can use non-uniform trajectories, such as the Archimedean spiral ([17]). While these trajectories provide a number of benefits relating to speed and signal-to-noise ratio, the FFT algorithm cannot be directly applied to the non-uniformly sampled raw data. Instead the raw data are typically interpolated onto a Cartesian grid in k-space and thereafter the FFT is applied to reconstruct the

**FIGURE 1.3**

The raw data obtained from the MRI scanner is sampled in k-space, typically on a uniform grid. Using the inverse Fourier transform (IFT) the data can be transformed into image space where data analysis is performed.

image. Image reconstruction is described in more detail in Chapter ** (cite RECONSTRUCTION CHAPTER).

The k-space signal is measured over two channels, and therefore the raw k-space data will be complex valued. It is assumed that both the real and imaginary components are measured with independent normally distributed error. Since the Fourier transformation is a linear operation, the reconstructed data will also be complex valued in each voxel, with both parts following a normal distribution. In the final stage of the reconstruction, the complex valued measurements are separated into magnitude and phase components. In the vast majority of studies only the magnitude portion of the images are used in the data analysis, while the phase portion is discarded.

It is important to note that the magnitude images no longer follow a normal distribution, but rather a Rice distribution ([20]). The shape of this distribution depends on the signal-to-noise (SNR) ratio within the voxel. For the case when no signal is present, it behaves like a Rayleigh distribution. When the SNR is high, the distribution will be approximately Gaussian.

1.2.3 From MRI to fMRI

The data acquisition and reconstruction techniques discussed so far provide the means for obtaining a static 3-D image of the brain. However, changes in brain hemodynamics in response to neuronal activity impact the local intensity of the MR signal. Therefore, a sequence of properly acquired T_2^* -weighted brain images allow for the study of changes in brain function over time. This can be achieved by repeatedly measuring T_2^* -weighted images of the brain every few seconds. The time between successive 3-D brain volumes is referred

to as the repetition time, or TR . In most fMRI experiments the TR is on the order of 2 seconds. However, depending on the research question values can vary from anywhere between 0.1 – 6 seconds.




FIGURE 1.4

Example of a structural T_1 -weighted image and corresponding T_2^* -weighted functional images measured during two conditions.

During the course of the scanning session the subject is either asked to perform a certain task, experience an induced psychological or behavioral state, or simply rest. The types of tasks that are performed depend upon the research question and can vary from the relatively mundane (e.g. tapping your fingers) to the complex. Fig. 1.4 shows an example of a T_1 -weighted structural scan, as well as T_2^* -weighted functional images collected under two conditions. In recent years it has become increasingly common to perform so-called resting-state scans, where the subject lies still without actually performing an explicit task. These types of studies have been used to investigate synchronous activations in spatially distinct regions of the brain, which are thought to reflect functional systems supporting cognitive processes, and this will be discussed further in Section 1.8.

A standard fMRI experiment gives rise to massive amounts of data. It can consist of multiple subjects (usually 10 – 50, though some larger studies now include hundreds or thousands of participants) potentially brought in for multiple scanning sessions. Each session consists of a number of runs, or replications of a certain experimental task. Each run consists of a series of brain volumes, each volume is made up of multiple slices, and each slice contains many voxels that has an intensity values associated with it. On top of that, multiple high-resolution structural scans are collected for pre-processing and

presentation purposes, and often diffusion tensor imaging (DTI) is also collected to inform subsequent network analyses. Hence, the amount of available data from a single experiment is enormous.

There are a number of critical determinations that go into designing the manner in which fMRI data is acquired. Certain issues directly related to the type of magnet and pulse sequences that are used, may ultimately depend on the particular lab where the data is being collected and could be outside of the researchers control. However, there still remain many decisions that should be carefully discussed within the research team before beginning the process of acquiring the data.

One set of decisions concerns the desired spatial and temporal resolution of the study. The temporal resolution determines our ability to separate brain events in time. In fMRI its value depends upon how quickly each individual image is acquired, i.e. the TR . In contrast, the spatial resolution determines our ability to distinguish changes in an image across different spatial locations. The manner in which fMRI data is collected makes it impossible to simultaneously increase both, as increases in temporal resolution limit the number of k-space measurements that can be made in the allocated sampling window and thereby directly influence the spatial resolution of the image. Therefore there are inherent trade-offs required when determining the appropriate spatial and temporal resolutions used in an fMRI experiment.




FIGURE 1.5

Each brain volume consists of multiple axial slices measured over a certain spatial extent, denoted the field-of-view (FOV). The matrix size, or the number of voxels acquired in the xy -plane, determines the spatial resolution with higher values giving better resolution. Together the slice thickness, FOV and matrix size determine the size of the voxel.

As previously mentioned, 3-D brain volumes are typically acquired in a series of axial slices (over the xy -plane with a fixed z value) of a certain slice thickness. Each slice is measured over a certain spatial extent, referred to as the field-of-view (FOV). The matrix size, or the number of voxels acquired

in the xy -plane, determines the spatial resolution with higher values giving better resolution. Together the slice thickness, FOV and matrix size determine the size of the voxel. See Fig. 1.5 for an illustration. For example, consider that we choose to acquire a series of thirty $5mm$ slices with a FOV of $192mm$, as this is typically sufficient to cover the entire brain. In addition, suppose the matrix size is 64×64 , corresponding to 4096 voxels in the brain. Thus, the dimensions of each voxel will be $3 \times 3 \times 5mm$.

Structural images tend to have high spatial resolution, but as they are static images, they lack any temporal resolution and generally do not reflect function at all. These are typically T_1 -weighted images, as these are useful for distinguishing between different types of tissue. One of the benefits of MRI as an imaging technique is its ability to provide detailed anatomical scans of gray and white matter with a spatial resolution well below $1mm^3$. However, the time needed to acquire such an image is prohibitively high and currently not feasible for use in functional studies. Hence, by necessity functional images have lower spatial resolution, but higher temporal resolution. As such they can be used to relate changes in signal to an experimental manipulation. The spatial resolution is typically on the order of $3 \times 3 \times 5mm$, corresponding roughly to image dimensions on the order of $64 \times 64 \times 30$, which can readily be sampled with a resolution of approximately 2 seconds. However, with modern high-resolution imaging, combined with higher field strengths and new acquisition techniques, it is possible to achieve much higher resolution. For example, with simultaneous multi-slice acquisition at 7T, it is possible to acquire $2 \times 2 \times 2mm$ data across the brain in less than 1 second.

Regardless of these limitations, fMRI still provides relatively high spatial resolution compared with many other functional imaging techniques, including positron emission tomography (PET), electroencephalography (EEG), and magnetoencephalography (MEG). This balance of spatial and temporal resolution, together with its non-invasiveness, accounts for the popularity of the technique compared to other modalities. Advances in high-field imaging and parallel acquisition methods promise to further push the boundaries of both spatial and temporal resolution.

1.3 BOLD fMRI

The ability to connect measures of brain activation, obtained using fMRI, with the underlying neuronal activity that caused them, will greatly impact the choice of statistical procedure as well as the subsequent conclusions that can be made from the experiment. Therefore from a modeling perspective, it is critical to gain some basic understanding of brain physiology to better understand the data at hand. Again, the overview presented here is by necessity brief

and interested readers are referred to textbooks dealing specifically with the subject (e.g., [23]).

1.3.1 Understanding BOLD fMRI

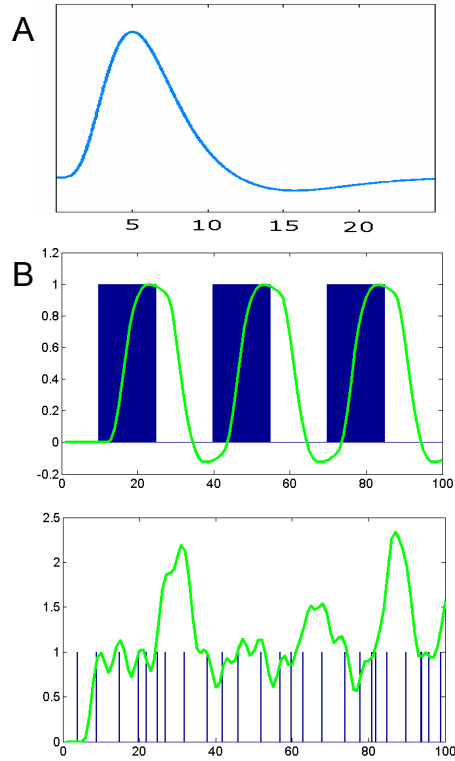
The most popular approach for performing fMRI uses the Blood Oxygenation Level Dependent (BOLD) contrast ([50, 28]), measured using the difference in signal between a series of T_2^* -weighted images. Other methods for performing fMRI are available, but less widely used. These include several varieties of Arterial Spin Labeling (ASL), which use pulse sequences sensitive to blood volume or cerebral perfusion. Because it is by far the most common method currently used, our focus in this Chapter is exclusively on BOLD physiology.

BOLD imaging takes advantage of differences in the magnetic properties of oxygenated and deoxygenated hemoglobin. As neural activity increases, so do the metabolic demands for oxygen and nutrients in affected regions of the brain. Neural firing signals the extraction of oxygen from hemoglobin in the blood. This extraction causes the hemoglobin to become paramagnetic as iron atoms are more exposed to the surrounding water. This creates small distortions in the magnetic field that cause a decrease in T_2^* , leading to a faster decay of the signal and a local decrease in BOLD signal. A subsequent over-compensation in blood flow increases the amount of oxygenated hemoglobin, leading to reduced signal loss and increased BOLD signal in the affected region.

Initially, fMRI was performed by injection of contrast agents with paramagnetic properties such as iron. However, the discovery that the T_2^* relaxation rate of oxygenated hemoglobin was longer than that of deoxygenated hemoglobin led to the advent of BOLD imaging, which has since come to dominate the field.

BOLD fMRI allows us to study the hemodynamic responses to neural firing. The change in the MR signal caused by a neural event is typically referred to as the hemodynamic response function (HRF); see Fig. 1.6A for an illustration. The increased metabolic demands due to neuronal activity lead to increases in the inflow of oxygenated blood to active regions of the brain. Since more oxygen is supplied than actually consumed, this leads to a decrease in the concentration of deoxygenated hemoglobin, which leads to an increase in signal. This positive rise in signal has an onset approximately 1 – 2 seconds after the onset of neural activity and peaks 5 – 8 seconds after peak neural activity. After reaching its peak level, the BOLD signal decreases to a below baseline level which is sustained for roughly 10 seconds. This effect, known as the post-stimulus undershoot, is due to the fact that blood flow decreases more rapidly than blood volume thereby allowing for a greater concentration of deoxygenated hemoglobin in previously active brain regions.

A number of studies have shown evidence of a decrease in oxygenation levels in the time immediately following neural activity, giving rise to decreased BOLD signal in the first half second following activation. This is believed to

**FIGURE 1.6**

(A) The standard canonical HRF model used in fMRI data analysis. (B) The BOLD response modeled as the convolution of the stimulus function and the HRF. Here we see examples using both a block and event-related design.

be due to oxygen extraction taking place prior to the inflow of oxygenated blood, and is usually referred to as the initial dip ([45, 41]). The ratio of the amplitude of the dip compared to the positive BOLD signal depends on the strength of the magnet and has been reported to be roughly 20% at 3 Tesla. There is also evidence that the dip may be more localized to areas of neural activity (e.g., [70, 34]) than the subsequent rise, which appears less spatially specific. Due in part to these reasons, the negative response has so far not been reliably observed and its existence remains controversial ([37]).

Clearly, the BOLD signal only provides an indirect measure of the quantity we actually seek to measure, which is the underlying neural activation. It is therefore important to understand how well the BOLD signal reflects actual increases in neuronal firing. The answer to this question is complex, and understanding the physiological basis of the BOLD response has long been a

topic of intense research interest (e.g. [7]). In short, it has been shown that the BOLD signal corresponds closely to the local electrical field potential surrounding a group of cells, which is likely to reflect changes in post-synaptic activity under many conditions ([38]). However, neural activity and BOLD signal may under other conditions become decoupled. Thus, BOLD signal is only likely to reflect a portion of the changes in neural activity in response to a task or psychological state. For this reason many regions may exhibit changes in neural activity that is missed because they do not change the net metabolic demand of the region.

1.3.2 Spatial Limitations

One of the primary benefits of fMRI is that it provides relatively high spatial resolution compared with many other functional imaging modalities, such as PET, MEG and EEG. The spatial resolution of fMRI can be made to be less than 1mm^3 in high-field imaging of animals, but is typically on the order of $27 - 36\text{mm}^3$ for human studies. For these reasons, features such as cortical columns and sub-nuclei cannot easily be identified, and it may be difficult to study certain small-scale features using fMRI.

The limiting factors in fMRI include signal strength and the point-spread function of BOLD imaging, which typically extends beyond the actual neural activation sites into draining veins. The fact that only a fraction of the oxygen that flows to an active region is actually extracted, leads to oxygen rich blood entering the venous system, thus increasing the BOLD signal in areas removed from the active neurons. This point-spread function decreases as the magnetic field strength increases, and interacts with head movement and physiological noise.

In addition, there are multiple analysis choices that ultimately limit the spatial resolution in most studies. First, it is common to spatially smooth fMRI data prior to analysis, which leads to a decrease in the effective resolution of the data. Second, making inferences about populations of subjects requires analyzing groups of individuals, each with inherent differences in brain structure. Usually, individual brains are aligned to one another through a registration or warping process, which introduces substantial blurring and noise in the group average.

One can potentially improve inferences in space by advances in both the manner data is acquired and preprocessed. For example, an important innovation in the area of acquisition has been the use of multiple coils with different spatial sensitivities to simultaneously measure k-space ([61, 54]). This approach, known as parallel imaging, allows for an increase in the amount of data that can be collected in a given time window. Hence, it can be used to either increase the spatial resolution of an image or decrease the amount of time required to sample an image with a certain specified spatial resolution. Parallel imaging techniques have already had a great influence on the way data is collected, and their role will only increase. The appropriate manner to deal

with the acquisition and reconstruction of multi-coil data is a key direction for future research.

Similarly, in the area of preprocessing, the introduction of enhanced spatial inter-subject normalization techniques and improved smoothing techniques could help researchers avoid the most dramatic effects of blurring the data. In addition, adaptive smoothing techniques can be used to more efficiently retain boundaries between different tissue types.

1.3.3 Temporal Limitations

The temporal resolution of an fMRI study depends on the TR , which in most fMRI studies ranges from 0.5 – 4.0 seconds. Clearly, these values indicate a fundamental disconnect between the underlying neuronal activity, which takes place on the order of tens of milliseconds, and the temporal resolution of the study. However, the statistical analysis of fMRI data is primarily focused on using the positive rise in the BOLD response to study the underlying neural activity. Hence, the limiting factor in determining the appropriate temporal resolution is generally not considered the speed of data acquisition, but instead the speed of the underlying evoked hemodynamic response to a neuronal event. Since inference is based on oxygenation patterns taking place 5 – 8 seconds after activation, TR values in the range of 2 seconds have generally been considered adequate.

However, the currently used resolutions are not conducive to efficient modeling of physiological artifacts present in the fMRI signal. For example, heart-rate and respiration give rise to periodic fluctuations that are difficult to model due to violations of the Nyquist criteria, which states that it is necessary to have a sampling rate at least twice as high as the frequency of the periodic function one seeks to model. At standard temporal resolutions this is clearly violated, and heart rate and respiration are often left un-modeled. Because of aliasing, these fluctuations tend to be distributed throughout the time course, giving rise to temporal autocorrelation in the signal. As fMRI signal generally suffers from low signal-to-noise ratio (SNR) and physiological artifacts potentially make up a large portion of the noise component, this may be a serious impediment. There has recently been active research in increasing the temporal resolution of fMRI studies, making TRs on the order of hundreds of milliseconds possible ([34, 29, 48]). These advances may ultimately allow us to circumvent many of these issues.

Because of the relatively low temporal resolution and the sluggish nature of the hemodynamic response, inference regarding when and where activation is taking place is based on oxygenation patterns outside of the immediate vicinity of the underlying neural activity we want to base our conclusions on. Since the time-to-peak positive BOLD response occurs in a larger time scale than the speed of brain operations, there is a risk of unknown confounding factors influencing the ordering of time-to-peak relative to the ordering of brain activation in different regions of interest. For these reasons it is difficult to

determine the absolute timing of brain activity using fMRI. However, studies have shown ([46, 47]) that the relative timing within a voxel in response to different stimuli can be accurately captured in a well-designed experiment. There are also indications that focusing inference on features related to the initial dip can help alleviate concerns ([33]) regarding possible confounders. However, these types of studies require significant increases in temporal resolution and the ability to rapidly acquire data becomes increasingly important.

Another way of improving inferences in time is through appropriate experimental design. In principal it is possible to measure the HRF at a finer temporal resolution than the TR as long as the onsets of repeated stimuli are jittered in time relative to when the data is collected ([12]). For example, if the onset of a repeated stimulus is shifted by half a TR in a fraction of the stimuli, it may be possible to estimate the HRF at a temporal resolution of $TR/2$, compared to a resolution of TR if jittering is not performed.

A series of recent technological developments referred to as “multiband” or “simultaneous multi-slice” MRI ([29, 48]), have sped up the temporal resolution of fMRI acquisitions by approximately an order of magnitude (i.e., from 2 s to 0.2 s, for whole-brain imaging), and appear likely to offer the possibility for even further acceleration. In contrast to standard acquisition techniques, multiband MRI excites multiple slices simultaneously, and the MR signals from these slices are then separated using multiple receiver coils and the aid of special encoding techniques. The introduction of multiband MRI promises to change the manner in which fMRI data is acquired and how it is analyzed.

1.3.4 Acquisition artifacts

As with almost all types of physical measurements, fMRI data can be corrupted by acquisition artifacts. These artifacts arise from a variety of sources, including head movement, brain movement and vascular effects related to periodic physiological fluctuations, and reconstruction and interpolation processes. In particular, fMRI data often contain transient spike artifacts and slow drift over time related to a variety of sources, including magnetic gradient instability, RF interference, and movement-induced and physiologically induced inhomogeneities in the magnetic field. These artifacts will likely lead to violations of the assumptions of normally and identically distributed errors that are commonly made in subsequent statistical analysis. Unless they are properly dealt with, the consequences will include reduced power in group-level analysis and potentially increased false positives in single-subject inference. As always it is critical to examine the data (preferably in as raw a form as possible), in order to diagnose problems. However, this can be challenging given the massive amount of fMRI data collected.

A significant source of signal variations includes the substantial slow drift of the signal across time. The presence of this low-frequency noise component in fMRI can obscure results related to a psychological process of interest and produce false positive results. Therefore, it is usually removed statistically

prior to or during analysis. A consequence of slow drift is that it is often impractical to use fMRI for designs in which a process of interest only unfolds slowly over time or only happens once, such as the experience of strong emotions, though alternative techniques such as Arterial Spin Labeling or BOLD functional connectivity may be suitable for this purpose. The vast majority of fMRI designs use discrete events that can be repeated many times over the course of the experiment. For example, the most common method for studying emotional responses in fMRI is the repeated presentation of pictures with emotional content.

Susceptibility artifacts in fMRI occur because magnetic gradients near air and fluid sinuses and at the edges of the brain cause local inhomogeneities in the magnetic field that affects the signal, causing distortion in echo-planar sequences and blurring and dropout in spiral sequences. These problems increase at higher field strengths and provide a significant barrier for performing effective high-field fMRI studies. Not all scanner/sequence combinations can reliably detect BOLD activity near these sinuses. The regions most affected include the orbitofrontal cortex, inferior temporal cortex, hypothalamus, and amygdala. Some signal may be recovered by using optimized sequences or improved reconstruction algorithms.


Figure 1.7 shows examples of several types of artifacts, including susceptibility artifacts that are endemic to fMRI and several other types that can usually be controlled.

1.4 Modeling Signal and Noise in fMRI

In order to appropriately model fMRI data, it is important to gain a better understanding of the components present in a BOLD fMRI time series. In general, it consists of the BOLD signal, which is the component of interest, a number of nuisance parameters and noise. Here we discuss each component in detail and discuss modeling strategies.

1.4.1 BOLD signal

The evoked BOLD response is a complex, nonlinear function of the results of neuronal and vascular changes ([66, 7]), which complicates our ability to appropriately model its behavior. The shape of the response depends both upon the applied stimulus and the hemodynamic response to neuronal events. A number of methods for modeling the BOLD response, as well as the underlying HRF, exist in the literature. A major difference between the methods lies in how the relationship between the stimulus and BOLD response is modeled. In particular, we differentiate between non-linear physiological-based models,

**FIGURE 1.7**

Examples of common fMRI artifacts: (A) k-space artifact; (B) ghosting in a phantom; (C) susceptibility artifact (dropout); and (D) Normalization artifact.

such as the Balloon model ([7, 56]), and models that assume a linear time invariant (LTI) system.

The Balloon model describes the dynamics of cerebral blood volume and de-oxygenation and their effects on the resulting BOLD signal. It consists of a set of ordinary differential equations that model changes in blood volume, blood inflow, deoxyhemoglobin and flow-inducing signal and describes how these changes impact the observed BOLD response. While models of this type tend to be more biophysically plausible than their linear counterparts, they have a number of drawbacks. First, they require the estimation of a large number of model parameters. Second, they do not always provide reliable estimates with noisy data, and third, they do not provide a direct framework for performing inference. In general, they are not yet considered feasible alternatives for performing whole-brain multi-subject analysis of fMRI data in cognitive neuroscience studies. However, they are a critical component in the study of brain connectivity using Dynamic Causal Modeling (DCM). This is discussed in Section and in more detail in Chapter ** (cite EC CHAPTER).

While the flexibility of nonlinear models is attractive, linear models allow for robust and interpretable characterizations in noisy systems. It is common to assume a linear relationship between neuronal activity and BOLD response, where linearity implies that the magnitude and shape of the evoked HRF does not depend on any of the preceding stimuli. Studies have shown that under certain conditions the BOLD response can be considered approximately linear with respect to the stimulus ([6]), particularly if events are spaced at least 5 seconds apart, though there are still some nonlinearities (10%) at 5 second spacing ([47]). However, other studies have found that nonlinear effects in rapid sequences (e.g., stimuli spaced less than 2 seconds apart) can be quite large ([3, 66]).

The ability to assume linearity is important, as it allows the relationship between stimuli and the BOLD response to be modeled using a linear time invariant system, in which assumed neuronal activity (based on task manipulations) constitutes the input, or impulse, and the HRF is the impulse response function. In a linear system framework the signal at time t , $x(t)$, is modeled as the convolution of a stimulus function $v(t)$ and the hemodynamic response $h(t)$, i.e. $x(t) = (v * h)(t)$. Here $h(t)$ is either assumed to take a canonical form, or alternatively modeled using a set of linear basis functions. See Fig. 6B for two examples.

An LTI system is characterized by its scaling, superposition and time-invariance properties. Scaling implies that if the input is scaled by a factor b then the BOLD response will be scaled by the same factor. This is important as it implies that the amplitude of the measured signal provides a measure of the amplitude of neuronal activity. Therefore the relative difference in amplitude between two conditions can be used to infer that the neuronal activity was similarly different. For this reason much of the activation analyses performed on fMRI data is focused on studying contrasts between the responses to stimuli at different levels. Superposition implies that the response to two different stimuli applied together is equal to the sum of the individual responses. Finally, time-invariance implies that if a stimulus is shifted by a time t , then the response is similarly shifted by t . These three properties allow us to differentiate between responses in various brain regions to multiple closely spaced stimuli.

When using an LTI system, as with any model, one makes a number of assumptions. First, it is assumed that the BOLD response is linear. Studies have shown that this is reasonable [6], though some departures from linearity have been observed. For example, there is some evidence of refractory effects, which are reductions in amplitude of a response as a function of inter-stimulus intervals. This may cause the amplitude of closely spaced stimuli to be overestimated. Second, it is assumed that the neural activity function is correctly modeled. As this is typically assumed to be equal to the experimental paradigm, one must assume this provides a reasonable proxy for the underlying neuronal activity. Third, it is assumed that the HRF is correctly modeled. Often researchers assume a canonical shape for the HRF. A popular

choice is the linear combination of two gamma functions, i.e.

$$h(t) = \left(\frac{t^{\alpha_1-1} \beta_1^{\alpha_1} e^{-\beta_1 t}}{\Gamma(\alpha_1)} c \frac{t^{\alpha_2-1} \beta_2^{\alpha_2} e^{-\beta_2 t}}{\Gamma(\alpha_2)} \right) \quad (1.3)$$

where $\alpha_1 = 6$, $\alpha_2 = 16$, $\beta_1 = \beta_2 = 1$, $c = 1/6$ and Γ represents the gamma function. This particular shape, seen in Fig. 1.6A, is based on empirical findings from data extracted from the visual cortex.

However, it is critically important to note that the timing and shape of the HRF are known to vary across the brain, within an individual and across individuals [1, 59]. Part of the variability is due to the underlying configuration of the vascular bed, which may cause differences in the HRF across brain regions in the same task for purely physiological reasons ([64]). Another source of variability is differences in the pattern of evoked neural activity in regions performing different functions related to the same task.

One of the major shortfalls when analyzing fMRI data is that users typically assume a canonical HRF ([19]), which may lead to mis-modeling of the signal in large portions of the brain ([39]). It is therefore important that these regional variations are accounted for when modeling the BOLD signal. This is often handled by modeling the HRF using multiple basis functions, and using a linear combination to better fit the evoked BOLD responses. To illustrate, suppose we model the HRF as a linear combination of temporal basis functions, $f_i(t)$, such that

$$h(t) = \sum \beta_i f_i(t). \quad (1.4)$$

Then the BOLD response can be rewritten:

$$x(t) = \sum \beta_i (s * f_i)(t) \quad (1.5)$$

where each corresponding β_i describes the weight of the i^{th} component.

The ability to use basis sets to capture variations in hemodynamic responses depends both on the number and shape of the reference waveforms that are used in the model. For example, the finite impulse response (FIR) basis set, consists of one free parameter for every time-point following stimulation in every cognitive event-type that is modeled [16, 18]. It can be used to estimate HRFs of arbitrary shape for each event type in every voxel of the brain. Another approach is to use the canonical HRF together with its temporal and dispersion derivatives to allow for small shifts in both the onset and width of the HRF. Other choices of basis sets include those composed of principal components [1, 68], cosine functions [71], radial basis functions [56], spectral basis sets [30] and inverse logit functions [32]. For a critical evaluation of a number of commonly used basis sets, see [32] and [31].

1.4.2 Noise and nuisance signal

The measured fMRI signal is corrupted by random noise and various nuisance components that arise from hardware limitations and the subjects themselves.

One source of variability are the fluctuations in the MR signal intensity caused by thermal motion of electrons within the subject and the scanner gives rise to noise that tends to be highly random and independent of the experimental task. The amount of thermal noise increases linearly as a function of the field strength of the scanner, with higher field strengths giving rise to more noise. However, it does not tend to exhibit spatial structure and averaging the signal over multiple voxels can minimize the effects.

Another source of variability in the signal is due to scanner drift, caused by scanner instabilities, which result in slow changes in voxel intensity over time, so-called low-frequency noise. The amount of drift varies across space, and it is important to include this source of variation in your models. Because of drift most of the power in the time course lies in the low-frequency portion of the signal. To remove the effects of drift it is common to remove fluctuations below a specified frequency cutoff using a high-pass filter. This can be performed either by applying a temporal filter as a preprocessing step, or by including covariates of no interest into the model. As an example of the latter, the drift can be modeled using a p^{th} order polynomial function or a series of low frequency cosine functions. The most important issue when using a high-pass filter is to ensure that the fluctuations induced by the task design are not in the range of frequencies removed by the filter, as this may cause us to throw out the signal of interest. Hence, the ultimate choice in how to model the drift needs to be made with the experimental design in mind.

When subjects move their heads in the MRI scanner, the sequence of measurements corresponding to a given voxel in the resulting images may actually be composed of values originating from different brain locations. This necessitates motion correction, in which, prior to analysis, researchers estimate the between scan movement using a rigid body transformation, and then realign the images. However, this procedure does not correct for so-called spin history artifacts, or changes in the magnetic field caused by head motion that lead to nonlinear, time-varying distortion of the resulting brain images, and there has been some debate on how to deal with these residual artifacts [25].

Some researchers ([25, 40]) suggest including motion regressors as nuisance covariates in the model for the BOLD response to adjust for this error, arguing that this yields estimates that seem more reasonable than those obtained when these covariates are omitted. But as head motion tends to be task related, there is concern that inclusion of these covariates can lead to underestimating the signal component due to genuine activation [53]. [10] argue that inclusion of motion regressors has a generally detrimental effect on activation for subjects with minimal head motion. In contrast, in recent work we have shown [60] that omitting signal components due to systematic error correlated with the task will generally lead to biased estimates of the effects of interest in a standard GLM analysis. Thus, we recommend that motion regressors be included in models of the BOLD response.

Physiological noise due to patient respiration and heartbeat can, as previously discussed, cause fluctuations in signal across both space and time.

Physiological noise can in certain situations be directly estimated from the data ([33]). Some of it can be removed using a properly designed band-pass filter. However, in most studies, with TR values ranging from 2 – 4s, one cannot hope to estimate and remove the effects of heart rate and respiration based solely on the observed fMRI time series. According to the Nyquist theorem, it is necessary to have a sampling rate at least twice as high as the frequency of the periodic function one seeks to model. If the TR is too slow, which is true in most fMRI studies, there will be problems with aliasing. In this situation the periodic fluctuations will be distributed throughout the time course, giving rise to temporal autocorrelation. Some groups have therefore begun directly measuring heartbeat and respiration during scanning and using this information to remove signal related to physiological fluctuations from the data [9]. This is done either as a pre-processing step, or by including these terms as covariates in the model.

In standard time series analysis, model identification techniques are used to determine the appropriate type and order of a noise process. In fMRI data analysis this approach is not feasible due to the large number of time series being analyzed, and noise models are typically specified *a priori*. Noise in fMRI is typically modeled using either an AR(p), with p set to either 1 or 2, or an ARMA(1,1) process [55]. Here the autocorrelation is generally thought to be due to unmodeled nuisance signal. If these terms are properly removed there is evidence that the resulting error term corresponds to white noise *citelund*.

In our own work, we typically use an auto-regressive process of order 2. The reason we choose an AR model over an ARMA model is that it allows us to use method of moments rather than maximum likelihood procedures to estimate the noise parameters. This significantly speeds-up computation time when repeatedly fitting the model to tens of thousands of time series. Choosing the order of the AR process to be 2 has been empirically determined to provide the most parsimonious model that is able to account for autocorrelation present in the signal due to aliased physiological artifacts.

1.5 Experimental Design

The experimental design of an fMRI study is complicated as it not only involves the standard issues relevant to all psychological experiments, but also issues related to data acquisition and stimulus presentation. Not all designs with the same number of trials of a given set of conditions are equal, and the spacing and ordering of events is critical. What constitutes an optimal experimental design depends on the psychological nature of the task, the ability of the fMRI signal to track changes introduced by the task over time, and the specific comparisons that one is interested in making. In addition, as the efficiency of the subsequent statistical analysis is directly related to

the experimental design, it needs to be carefully considered during the design process.

A good experimental design attempts to maximize both statistical power and psychological validity. The statistical performance can be characterized by its estimation efficiency (the ability to estimate the shape of the HRF), or its detection power (the ability to detect which voxels are active). The psychological validity is often measured by the randomness of the stimulus presentation (e.g., balanced transitional probabilities across trial types), as this helps control for issues related to anticipation, habituation and boredom; however, whether a predictable vs. unpredictable design is psychologically undesirable depends heavily on the particular paradigm and task. Thus, when designing an experiment there is inherent trade-offs between estimation efficiency, detection power and randomness. The optimal balance between the three ultimately depends on the goals of the experiment and the combination of conditions that are of primary interest. For example, a design used to localize areas of brain activation may stress high detection power at the expensive of estimation efficiency and randomness. Conversely, designs that attempt to link activity to particular events or time periods during processing may emphasize estimation efficiency at the expense of detection power.

There are two major classes of designs used in most fMRI experiments, namely blocked designs and event-related designs (though these can be intermixed or hybridized); see Fig. 6B for examples of each. In a blocked design, experimental conditions are separated into extended time intervals, or blocks. For example, one might repeat the process of interest during an experimental block (A) and have the subject rest during a control block (B). The A-B comparison can then be used to compare differences in signal between the two conditions. Increasing the length of each block will lead to a larger evoked response during the task. This increases the separation in signal between blocks, which in turn leads to higher detection power. However, it is still important to include multiple transitions between conditions as otherwise differences in signal due to low-frequency drift may be confused for differences in task conditions and to ensure that the same mental processes are evoked throughout each block. If block lengths are too long, this assumption may be violated due to the effects of fatigue and/or boredom.

Another benefit of blocked designs is that they are robust to uncertainties in the shape of the HRF. This holds because the predicted response depends upon the total activation caused by a series of stimuli, making it less sensitive to variations in the shape of responses to individual stimulus. However in contrast, block designs provide imprecise information about the particular processes that activated a brain region and cannot readably be used to directly estimate important features of the HRF, such as the onset, width or time-to-peak.

In an event-related design the stimulus consists of short discrete events, such as brief light flashes, whose timing can be randomized. These types of designs are flexible and allow for the estimation of key features of the HRF

(e.g., onset and width) that can be used to make inference about the relative timing of activation across conditions and about sustained activity. Event-related designs allow one to discriminate the effects of different conditions as long as one either inter-mix events of different types or vary the inter-stimulus interval between trials. Another advantage to event-related designs is that the effects of fatigue, boredom, and systematic patterns of thought unrelated to the task during long inter-trial intervals can be avoided. A drawback is that the power to detect activation is typically lower than for block designs, though the ability to obtain data over more trials per unit time can counter this loss of power.

What constitutes an optimal experimental design ultimately depends on the task, as well as on the ability of the fMRI signal to track changes introduced by the task over time. It also depends on what types of comparisons are of interest. The delay and shape of the BOLD response, scanner drift and physiological noise all conspire to complicate experimental design for fMRI. Several methods have been introduced that allow researchers to optimally select the design parameters, as well as the sequencing of events that should be used in an experiment ([65, 36, 26, 27]). These methods define fitness measures for the estimation efficiency, detection power and randomness of the experiment, and apply search algorithms (e.g., the genetic algorithm) to optimize the design according to certain specified criteria. When defining the fitness metrics it is typically assumed that the subsequent data analysis will be performed in the general linear model (GLM) framework described in Section 1.7.1.

The use of more complex nonlinear models requires different considerations when defining appropriate fitness metrics. An important consideration relates to assumptions made regarding the shape of the HRF and the noise structure. The inclusion of flexible basis functions and correlated noise into the model will modify the trade-offs between estimation efficiency and detection power, and potentially alter what constitutes an optimal design. Hence, even seemingly minor changes to the model formulation can potentially have a large impact on the efficiency of the design. Together these issues complicate the design of experiments and significant work remains to find the appropriate balance between them. As research hypotheses ultimately become more complicated, the need for more advanced experimental designs will only increase further.

1.6 Preprocessing

Prior to statistical analysis, fMRI data typically undergoes a series of preprocessing steps aimed at removing artifacts and validating model assumptions. The main goals are to minimize the influence of data acquisition and physiological artifacts, to validate statistical assumptions and to standardize the

locations of brain regions across subjects in order to achieve increased validity and sensitivity in group-level analysis.

When analyzing fMRI data it is typically assumed that all of the voxels in a particular 3-D brain volume were acquired simultaneously. Further, it is assumed that each data point in a specific voxel's time series only consists of signal from that particular voxel; an assumption that is invalid if the subject moves between scans. Finally, when performing group analysis and making population inference, all individual brains are assumed to be registered, so that each voxel is located in the same anatomical region for all subjects. Pre-processing is used to condition the data in ways that increase the plausibility of all of these assumptions. Without appropriate pre-processing of the data prior to analysis, none of these assumptions hold and the resulting statistical analysis would be invalid.

The major steps in the fMRI preprocessing pipeline are slice acquisition-timing correction ('slice-time' correction), realignment, co-registration of structural and functional images, normalization of brains to a group template, and smoothing. Below we briefly describe each step. For more detail see Chapter ** (cite PREPROCESSING CHAPTER).

Slice-time Correction: A typical assumption in the analysis of fMRI data is that all voxels within a 3-D image are acquired simultaneously. In reality, different slices from the same volume are acquired sequentially in time relative one another. Thus, many researchers seek to estimate the signal intensity in all voxels at a common standardized time point in the acquisition period. This can be done by interpolating the signal intensity at the chosen time point from the same voxel in previous and subsequent acquisitions. Some researchers do not use slice timing, as it adds interpolation error to the data, and instead use (a) more flexible hemodynamic models to account for variations in acquisition time across the brain, or (b) more rapid acquisition sequences in which multiple slices are acquired simultaneously.

Realignment: A major source of error is subject movement during the course of the experiment. Excessive motion may cause the intensity in a given voxel to be contaminated by signal from neighboring voxels. For these reasons it is critical to realign each individual image to compensate for movement. This is typically done by first choosing a reference image, either the first or mean image, and then applying a rigid body transformation to all other images in the time series to match it. This allows the images to be translated (shifted in the x , y , and z directions) and rotated (altered roll, pitch, and yaw) to match the reference image. An iterative algorithm is used to search for the parameter estimates that provide the best match between a target image and the reference image. Matching is typically performed by minimizing the sums of the squared differences between the two images. Once the best match is found, the data is interpolated into the new space.

Realignment is critically important when analyzing fMRI data. It is able

to correct for small movements of the head. However, it is important to note that it is unable to correct for the more complex spin-history artifacts created by the motion. For this reason the estimated motion parameters at each time point are saved for later inspection and are often included in subsequent analysis as covariates. It is not uncommon to have to exclude subjects that move their heads too much during the course of the scan. While there do not exist firm rules stating how much movement should be considered too much, more than 1.5mm displacement within a scanning session is typically considered to be problematic.

Co-registration: Functional MRI data is typically of low spatial resolution and therefore provides relatively little anatomical detail. It is therefore common to map the results obtained from the analysis of functional data onto a high-resolution structural MR image for presentation purposes. This process is referred to as co-registration, and is typically performed using either a rigid body (6 parameters) or an affine (12 parameters) transformation. Because functional and structural images are collected with different sequences and different tissue classes have different average intensities, using a least squares difference method to match images is often not appropriate. Instead, it is preferable to estimate the parameters of the transformation by maximizing the mutual information between the two images. Typically, a single structural image is co-registered to the first or mean functional image. Co-registration is also a necessary step for subsequent normalization procedures (described below). Here high-resolution structural images (T_1 and/or T_2) are used for warping to a standard template and localization. The same transformations are thereafter applied to the functional images, which produce the activation statistics, so the results can be mapped onto a standard space.

Normalization: For group analysis, it is necessary for each voxel to lie within the same brain structure for all subjects. However, individual brains clearly have differences in both shapes and features. That said there are certain regularities shared by every non-pathological brain, and normalization attempts to register each subjects anatomy with a standardized atlas space defined by a template brain, reported in the standard coordinate systems of the Montreal Neurologic Institute (MNI), or that of Talairach and Tournoux [63]. Normalization can be linear, involving simple registration of the gross shape of the brain, or nonlinear, involving warping to match local features. This warping consists of shifting the locations of voxels by different amounts depending on their original location.

Inter-subject registration is perhaps the largest source of error in group-level analysis of fMRI data. For these reasons it is important to inspect each normalized brain as a quality control step. This can be done in a number of ways, and researchers should develop a set of standardized procedures to assess the results.

Smoothing: Many researchers apply a spatial smoothing kernel to the functional data, blurring the image intensities in space. One reason is to improve inter-subject registration, by eliminating intra-subject differences. A second reason is that Gaussian Random Field Theory, a popular multiple-comparisons correction procedure, assumes that the variations across space are continuous and normally distributed. Smoothing typically involves convolution with a Gaussian kernel, which is a 3-D normal probability density function often described by the full width of the kernel at half its maximum height (FWHM) in mm. One estimate of the amount of smoothing required to meet the assumption is a FWHM of 3 times the voxel size (e.g., *9mm* for *3mm* voxels).

1.7 Data Analysis

There are several common objectives in the analysis of fMRI data. These include localizing regions of the brain activated by a certain task, determining distributed networks that correspond to brain function and making predictions about psychological or disease states. Many of these objectives are related to understanding how induced or measured psychological states leads to changes in brain activity (a combination of neural and glial function), and others are related to the analysis of ongoing spontaneous fluctuations. All these objectives are intrinsically statistical in nature, and this area is the primary domain of statisticians currently involved in the field. For this reason, the material covered in this section is the focus of many subsequent chapters of this book. Here we simply provide a brief overview of relevant topics, but refer to future chapters for a more in-depth treatment.

1.7.1 Localization

To date, the most common use of fMRI has been to localize areas of the brain that activate in response to a certain task. These types of human brain mapping studies are instrumental for increasing our understanding of brain function. The general linear model (GLM) approach is the most common statistical method for assessing relationships between tasks and brain activity [69]. Here the data is considered to be a linear combination of model functions plus noise. The model functions are assumed to have known shapes, but with unknown amplitudes that need to be estimated. The GLM can be used to estimate whether the brain responds to a single event type, to compare different types of events, and to assess correlations between brain activity and behavioral or other psychological variables.

In a typical fMRI experiment, the predictors are related to psychological events, and the outcome variable is the signal from a certain brain voxel

or a region of interest. Analysis is typically massively univariate, meaning that a separate GLM analysis is performed at every voxel in the brain, and summary statistics are saved in maps of statistic values across the brain. This approach assumes an improbable independence between voxel pairs. Typically dependencies between voxels are dealt with later using random field theory, which makes assumptions about the spatial dependencies between voxels.

Using the GLM, the data at each voxel can be expressed as

$$\mathbf{Y} = \mathbf{X}\beta + \epsilon, \quad (1.6)$$

where $\epsilon \sim N(\mathbf{0}, \mathbf{V})$. Here \mathbf{Y} represents the data (a vector of length T), \mathbf{X} is a design matrix containing information about various signal components, and \mathbf{V} is a covariance matrix that incorporates information about temporal autocorrelations in the data. The latter is typically modeled using either an autoregressive (AR) or auto-regressive moving-average (ARMA) model.

Though the model formulation is simple, difficulties arise when attempting to construct an appropriate design matrix \mathbf{X} . This process is complicated by a number of factors, including the fact that the BOLD response contains low-frequency drift and artifacts related to head movement and cardiopulmonary-induced brain movement, the neural response shape may not be known, and the hemodynamic response varies in shape across the brain. For these reasons the design matrix \mathbf{X} usually consists of both nuisance parameters (corresponding to drift components and the estimated motion parameters) and signal of interest.

The simplest version of the GLM assumes that both the stimulus function and the HRF are known. As discussed in Section 1.4.1 the stimulus is assumed to be equivalent to the experimental paradigm, while the HRF is typically modeled using a canonical HRF; see Fig. 1.6A. If unwilling to assume a fixed canonical HRF, it is possible to model the shape by expressing the HRF as the linear combination of a number of basis functions. Here each basis function is convolved with the stimulus and entered as a separate column of the design. See Chapter **** (cite GLM CHAPTER) for more information about the GLM.

Often we are interested in combining the results for individual subjects in order to perform group-level inference. Because of the hierarchical structure of the data, an appropriate approach towards analyzing multi-subject fMRI data is to use either a multi-level or mixed-effects GLM model. For example, we can express the first-level model for subject i as $\mathbf{Y}_i = \mathbf{X}\beta_i + \epsilon_i$, where $\epsilon_i \sim N(\mathbf{0}, \mathbf{V}_i)$. The second level model can in turn be written, $\beta_i = \beta + \eta_i$, where $\eta_i \sim N(0, \sigma^2)$. This can easily be fit using most statistical software packages. However, in the neuroimaging community it is often approximated by performing a GLM on each subject, and thereafter using the resulting activation parameter estimates in a ‘second level’ group analysis; see Chapter **** (cite GROUP ANALYSIS CHAPTER).

Regardless of whether one performs single-subject or group-level analysis, the procedure follows the same general format. First one fits a statistical model




FIGURE 1.8

(Left) An example of a statistical image consisting of test statistic values at each voxel of the brain. (Right) The thresholded statistical map indicates active regions of the brain.

(e.g., the GLM) to data from a certain voxel in the brain. Next, the estimated model parameters are used to test for an effect of interest, for example $H_0 : \beta_1 - \beta_2 = 0$. This procedure is then repeated for each voxel across the brain, and the results are summarized in a statistical image; see the left panel of Fig. 1.8. The final step is to determine which voxels actually show a statistically significant effect. The results of neuroimaging studies are often summarized as a set of activated regions, such as those shown in the right panel of Fig. 1.8. These types of summaries describe brain activation by color-coding voxels whose test statistics exceed a certain statistical threshold for significance. The implication is that these voxels are activated by the experimental task.

Of course, a crucial decision is the choice of which threshold to use when deciding whether voxels should be deemed ‘active’ or not. In many fields, test statistics whose p-values are below 0.05 are considered sufficient evidence to reject the null hypothesis, with an acceptable false positive rate of 0.05. However, in brain imaging the order of 100,000 hypothesis tests (one for each voxel) are tested at a single time. Hence, using a voxel-wise alpha of 0.05 implies that 5% of the voxels on average will show false positive results. Hence, we would actually expect on the order of 5,000 false positive results. Thus, even if an experiment produces no true activation, there is a good chance that without a more conservative correction for multiple comparisons, the activation map will show a number of activated regions, leading to erroneous conclusions.

The traditional way to deal with this multiplicity problem is to adjust the threshold so that the probability of obtaining a false positive is simultaneously controlled for every voxel (i.e., statistical test) in the brain. In neuroimaging, a variety of different approaches towards controlling the false positive rate are commonly used; see [22] and Chapter ** for more detail (cite MC CHAPTER).

The fundamental difference between methods used is whether they control for the family-wise error rate (FWER) or the false discovery rate (FDR).

1.7.2 Connectivity

Previously, fMRI data was primarily used to construct maps representing regions of the brain activated by specific tasks. In recent years there has been increased interest in augmenting this type of analysis with connectivity studies that describe how various brain regions interact and how these interactions depend on experimental conditions. A number of methods have been suggested in the fMRI literature to quantify brain connectivity. Their appropriateness depend upon (i) what type of conclusions one is interested in making; (ii) what type of assumptions one is willing to make; (iii) the level of the analysis; (iv) the modality used to obtain the data; and (v) the number of brain regions that are included in the analysis.

The term connectivity is an umbrella term that has been used to refer to a number of related aspects of brain organization. In the neuroimaging literature it is common to distinguish between anatomical, functional and effective connectivity [14, 62]. Anatomical connectivity deals with the description of how different brain regions are physically connected, and can be approached using techniques such as diffusion tensor imaging (DTI; see Chapter ****). Functional connectivity (see Chapter ***) and effective connectivity (see Chapter ***) study the functional relationships between different brain regions.

Functional Connectivity is defined as the undirected association between two or more fMRI time series and/or performance and physiological variables. It makes statements about the structure of relationships among brain regions. However, methods used to access functional connectivity usually do not make any assumptions about the underlying biology and tend to be data-driven in nature.

A wide variety of methods have been proposed to study functional connectivity. The simplest approaches simply compare the bivariate correlations between regions of interest, or between a seed region and all other voxels across the brain. Recently, inverse covariance estimation methods have been applied, using the fact that for multivariate normal data, conditional independence between variables (regions) corresponds to zero entries in the inverse covariance matrix. This allows for the efficient investigation of a large number of regions simultaneously.

Other approaches towards investigating functional connectivity include using various multivariate decomposition methods, to identify task-related patterns of brain activation without making a priori assumptions about its form. These methods include principal components analysis (PCA) [2] and independent components analysis (ICA) ([8, 44]). These methods involve decomposing the time-by-voxel data matrix, \mathbf{Y} , into a set of spatial and temporal components according to some criteria (e.g., independence between components). They are discussed in detail in Chapter ** (cite MVDECOMP CHAPTER)

and have been especially fruitful for analyzing so-called resting-state data, where the subjects do not perform an explicit task. This is an area of intense research where standard methods for localizing brain activation are not applicable due to the lack of stimulus.

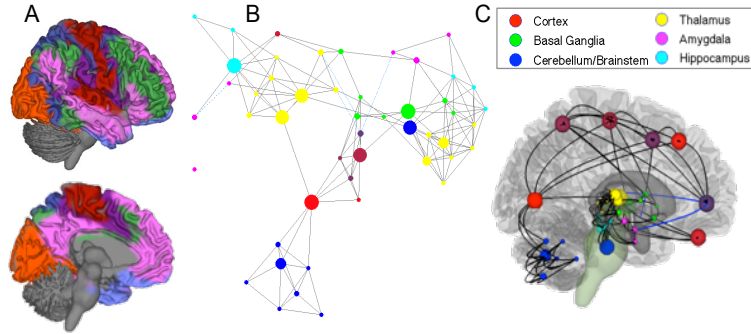
Effective Connectivity is defined as the directed influence of one brain region on the physiological activity recorded in other brain regions. It claims to make statements about causal effects among tasks and regions. Usually methods that assess effective connectivity make anatomically motivated assumptions and restrict inference to networks comprising of a number of pre-selected regions of interest.

Effective connectivity analyses are inherently model-dependent. Typically, a small set of regions and a proposed set of connections are specified *a priori*, and tests of fit are used to compare a small number of alternative models and assess the statistical significance of individual connections. Because connections may be specified directionally (with hypothesized causal influences of one area on another), the models are typically thought to imply causal relationships. Because there are many possible models, the choice of regions and connections must be anatomically motivated. Thus, most effective connectivity depends upon a neuroanatomical model that describes which areas are connected, and a mathematical model that describes how areas are connected. Popular methods for assessing effective connectivity include structural equation modeling (SEM) [43], Granger causality [57], and dynamic causal modeling (DCM) [13].

Effective connectivity is popular because it is thought to provide powerful conclusions. However, the validity of these conclusions depends on certain assumptions being correct. They are often poorly specified and difficult to check, which is a major shortcoming of the field. The main problem is that results are discussed in terms of the applied estimation algorithm rather than carefully defined estimands of interest. This is discussed in more detail in Chapter ** (EC CHAPTER).

In many situations we seek to create networks consisting of a large number of non-overlapping brain regions. Networks can be represented using graphs, which are mathematical structures that can be used to model pair-wise relationships between variables. They consist of sets of nodes (or vertices) V and corresponding links (or edges) E that connect pairs of vertices. A graph $G = (V, E)$ may be defined as being either undirected or directed with respect to how the edges connect one vertex to another; see Fig. 1.9 for an example.

As the number of regions included in the analysis approaches the hundreds, it can often be difficult to make sense of these vast amounts of data. Network analysis attempts to characterize these networks using a small number of meaningful summary measures. The hope is that comparisons of network topologies between groups of subjects may reveal connectivity abnormalities related to neurological or psychiatric disorders. As an example, there is a growing interest in using graph theoretical approaches to investigate the organizational principles of large-scale brain networks. These studies have shown clear

**FIGURE 1.9**

(A) A functional connectivity-based parcellation of the brain. (B) Co-activations in studies of disgust, from a meta-analysis of 148 neuroimaging studies. The nodes (circles) are regions or networks, color-coded by anatomical system. The edges (lines) reflect co-activation between pairs of regions or networks, assessed based on the joint distribution of activation intensity. The size of each circle reflects how strongly it connects disparate networks. (C) The same connections in the anatomical space of the brain.


topological organization of the brain, including modularity, small-worldness, and the existence of highly connected network hubs. These network properties are thought to provide important implications for health and disease. See [58] or Chapter ** for more detail (cite NETWORK CHAPTER).

1.7.3 Prediction

There is a growing interest in using fMRI data for the classification of mental disorders and predicting the early onset of disease. In addition, there is interest in developing methods for predicting stimuli directly from functional data. The ability to do so opens the possibility of inferring information about subjective human experience directly from brain activation patterns.

Predicting brain states is challenging and requires the application of novel statistical and machine learning techniques [52]. Various techniques have successfully been applied to fMRI data in which a classifier is trained to discriminate between different brain states and then used to predict the states in a new set of fMRI data. The application of machine learning methods to fMRI data is often referred to as multi-voxel pattern analysis (MVPA). As the name indicates, instead of focusing on single voxels, MVPA uses pattern-classification algorithms applied to multiple voxels to decode patterns of activity. In MVPA the goal is to determine the model parameters that allow for the most accurate prediction of new observations.

When applied to fMRI data the result is often a pattern of weights across

**FIGURE 1.10**

An MVPA example from [67]. The input features are measurements over all V voxels in the brain contained in the vector \mathbf{x} . These voxels are weighted by the vector \mathbf{w} of length V , so that if $xw > 0$ we categorize a subject as belong to group A and if $xw < 0$ we categorize them as belonging to group B .

brain regions that can be applied to new brain activation maps in order to quantify the degree to which the pattern responds to a particular type of event. For example, consider the situation where we want to categorize subjects into two groups (A or B) based on their brain activation. Here the input features are measurements over all V voxels in the brain contained in the vector x . We now seek to find a weighting of these voxels, represented by the vector w of length V , so that if $xw > 0$ we categorize a subject as belong to group A and if $xw < 0$ we categorize them as belonging to group B . Here the elements of the vector w consists of voxel-specific weights that can be mapped back onto 3D space and interpreted for scientific reasons; see Fig. 1.10.

When applying MVPA methods to fMRI data it is particularly important to make analysis choices that balance interpretability with predictive power. Certain methods may give good predictions but the resulting voxel weights may be difficult to interpret and may not generalize well to new subjects. In these situations the scientific value of the results may be unclear. MVPA is discussed in more detail in Chapter *** (cite MachineLearning).

Currently, one of the most exciting areas in neuroimaging is the work being done in decoding our thoughts based on brain activity. Here the idea is to study a persons activation and thereafter try to recreate what the subject is seeing or doing. In order to do this properly researchers need a good mathematical model of how the brain functions and high-speed computing. Although there are a number of companies that are starting to pursue brain decoding for purposes ranging from market research to lie detection, most researchers are more excited about what this process can teach us about the way the brain functions.

1.8 Resting-state fMRI

Clearly the brain is always at work, even in the absence of an explicit task. In fact, according to certain estimates, task-related changes in neuronal metabolism only account for about 5% of the brain's total energy consumption. Resting state fMRI (rs-fMRI) is a relatively new approach towards functional imaging that is used to identify synchronous BOLD changes in multiple brain regions while subjects lie in the scanner but do not perform a task [4].

Using rs-fMRI it has been shown that fluctuations in the low-frequency portion of the BOLD signal show strong correlations in spatially distant regions of the brain. While the exact mechanisms driving these correlations remain unclear, it is hypothesized that it may be due to fluctuations in spontaneous neural activity. Neuroscientists are increasingly interested in studying the correlation between spontaneous BOLD signals from different brain regions in order to learn more about the intrinsic functional connectivity of the brain.

Already rs-fMRI has revealed several large-scale spatial patterns of coherent signal in the brain during rest, corresponding to functionally relevant resting-state networks (RSNs). These networks are thought to reflect the baseline neuronal activity of the brain. A number of RSNs have been consistently observed both across groups of subjects and in repeated scanning sessions on the same subject. RSNs are localized to grey matter, and are thought to reflect functional systems supporting core perceptual and cognitive processes. Many regions that are co-activated during active tasks also show resting state connectivity, as brain regions with similar functionality tend to express similar patterns of spontaneous BOLD activation. Sometimes subsets of RSNs appear to be either up or down-regulated during specific cognitive tasks.

Resting-state fMRI is based on studying low-frequency BOLD fluctuations. Functionally relevant, spontaneous BOLD oscillations have been found in the lower frequency ranges (0.01 – 0.08 Hz). This is separable from frequencies corresponding to respiratory (0.1 – 0.5 Hz) and cardiovascular (0.6 – 1.2 Hz) signal. Typical resting experiments are of the order of 5 – 10 min, though the identification of an optimal duration of an rs-fMRI session and the possible need for multiple sessions remains an open issue. In addition, there is no consensus as to whether data should be collected while subjects are asleep or awake, and with eyes open or closed.

Pre-processing of rs-fMRI data typically follows the same pipeline applied to standard task-related BOLD fMRI. However, there are a few important differences. High pass temporal filtering applied to task fMRI data may be overly aggressive with respect to removing some of the relevant frequency information. Often the data is band-pass filtered at (0.01 – 0.08 Hz). As it has been shown that non-neuronal physiological signals may interfere with resting state BOLD data, the removal of confounding signals, such as respiratory or

cardiovascular noise has been shown to considerably improve the quality of data attributable to neural activity. It has therefore become common practice in rs-fMRI research to explicitly monitor these signals, and retrospectively correct for their confounding effects post-acquisition.

In addition the global mean signal, at least six motion parameters estimated in the pre-processing stage, the cerebrospinal fluid (CSF), and the white matter signals are also commonly removed prior to analysis in order to reduce the effects of head motion and non-neuronal BOLD fluctuations. However, the removal of the global signal has proven to be particularly controversial. In the past few years, there has been increased attention given to observed anti-correlations between RSNs. Anti-correlations between the components of the default-mode and attention networks have been consistently observed. However, recently there has been a lot of debate about these findings (e.g., [49]), as it is thought that global signal regression will induce a bias towards finding anti-correlations between RSNs.

Because of the lack of task, rs-fMRI is attractive as it removes the burden of experimental design, subject compliance, and training demands. It is particularly attractive for studies of development and clinical populations. In addition, it is easy to tack on a resting state scan even when performing task-based experiments. For these reasons, the amount of available resting state data has exploded and there is a growing subfield around the acquisition and analysis of rs-fMRI data.

As a final note, one of the primary benefits with R-fMRI is the ability to compare data across research labs, as experiments do not need to be synchronized. This has led to a number of large-scale data sharing initiatives (e.g. 1000 Functional Connectomes Project).

1.9 Data Format, Databases & Software

It is critical that new researchers interested in getting involved in fMRI research gain some basic understanding of the format of the data and have access to the tools needed to read, analyze and visualize them. In this section we discuss the formats that are typically used to store fMRI data, a variety of databases that will allow researchers to access fMRI data, and a number of freely available software packages that can be used to begin analyzing the data.

MRI data is usually stored in binary data files as either 8- or 16-bit integers. Additional information about the data, called meta-data, which includes image dimensions and type, are also stored. Structural MRI images are generally stored as 3-D data files, while functional MRI data can either be stored as a series of 3-D files, or as a single 4-D file.

Most MRI scanners save their reconstructed data to a file format called

DICOM (Digital Imaging and Communications in Medicine). The data is generally stored slice-wise and contains large amount of meta-data about image acquisition settings and the subject. Although DICOM is the standard format for outputting data from the MRI scanner, it is necessary to convert to other formats before performing the actual data analysis.

There are two main file formats typically used for data analysis, namely Analyze and NIFTI. The Analyze file format originates from the software package with the same name. It stores each data set in two separate files. The first is the data file, which contains the binary data and has the extension .img. The second is the header file, which contains the header file and has the extension .hdr. NIFTI (Neuroimaging Informatics Technology Initiative) is another file format designed to promote compatibility among programs using fMRI data sets. It extends Analyze by storing additional meta-data, such as affine matrices, data arrangement, and slice order information. It comes in two different formats. The first format combines the header and binary data into a single file with the extension .nii. The second format uses the same extensions as Analyze, i.e. .hdr and .img, keeping the meta- and binary data separate from one another.

In recent years there has been an increased movement towards openly sharing fMRI data between researchers in the field. The goal is to emulate similar data sharing initiatives in other disciplines, such as genetics, that have led to important advances. The first such effort was the fMRI Data Center (fMRIDC) which has come to consist of 107 fMRI datasets. More recently, there has been a particular focus on sharing rs-fMRI data, as this type of data is particularly easy to compare data across research labs, as experiments do not need to be synchronized. Perhaps, the most well known repository of this kind is the 1000 Functional Connectomes Project (FCP), which to date consists of resting-state data on roughly 5,000 subjects from sites all around the world. The OpenfMRI Project (openfmri.org) is a relatively recent project that is particularly dedicated to the free and open sharing of data from task-based fMRI studies. Finally, on an institutional level the Human Connectome Project (HCP), is a project supported by the National Institute of Health with the stated goal of mapping the human connectome. It will make available task-based and rs-fMRI on over 1,000 subjects, as well as data obtained using other modalities.

In addition to the above mentioned efforts, there has been longstanding interest in sharing the activation coordinates reported in papers for use in meta-analysis; see Chapter **** (cite meta-analysis chapter). This data consists of the spatial locations of peaks of activation (peak coordinates), reported in the standard coordinate systems of the Montreal Neurologic Institute (MNI). This information can be used to pool multiple separate but similar studies, which can be used to evaluate the consistency of findings across labs, scanning procedures, and task variants, as well as the specificity of findings in different brain regions or networks to particular task types. There are several such databases

containing peak coordinates that can be used for meta-analysis, including Brainmap (<http://brainmap.org>) and Neurosynth (<http://neurosynth.org>).

Finally, there are a number of free open source software packages used in the neuroimaging community that can be downloaded freely from the web, that are relatively easy to use. The three most popular are SPM, FSL and AFNI. The most commonly used software is SPM (Statistical Parametric Mapping) [51], which consists of a set of MATLAB functions for preprocessing, analyzing, and displaying fMRI and PET data. It has a very active development community and there exist a large variety of add-ons to the core tools. FSL (FMRIB Software Library) [24] is written in C++, and has a number of unique tools to its disposal. It provides a comprehensive library of image analysis and statistical tools for fMRI, MRI and DTI brain imaging data. Finally, AFNI (Analysis of Functional NeuroImages) [11] is written in C. It consists of a series of programs for processing, analyzing, and displaying fMRI data. Of particular interest to statisticians might be its use of functions from the statistical software package R.

1.10 Future Developments

It is truly an exciting time to be involved in neuroimaging research. More and more increasingly ambitious experiments are being performed each day. This is creating a significant new demand, and an unmatched opportunity, for quantitative researchers working in the neurosciences. Understanding the human brain is arguably among the most complex, important and challenging issues in science today. For this endeavor to be successful, a legion of neuro-quants is needed to make sense of the massive amounts of data being generated.

With this rapid development, new research questions are opening up every day. Below we have made a decidedly non-inclusive list of exciting research areas that we feel will be increasingly important in coming years. However, by the time this book is published this list could no doubt be expanded even further!

Longitudinal Imaging Studies: Longitudinal neuroimaging studies have become increasingly common in recent years. Here the same subjects are repeatedly scanned over a prolonged period of time and changes in brain structure and function is assessed. These types of studies promise to play an important role as they have the potential to answer critical questions regarding brain development, aging, neuro-degeneration, and recovery from traumatic brain injuries or stroke.

However, in the hand with the increased amount of information, the resulting data sets will be even larger and more complex than their single session counterparts. In addition to the standard problems of modeling within-

session relationships, there is here the additional problem of correctly modeling the between-session variation and the appropriate inclusion of potential time-varying covariates. The analysis of longitudinal and repeated measures data has a long history in statistics, and it is time these ideas are moved into the high-dimensional analysis of longitudinal fMRI data.

High Temporal Resolution Multiband Data: Recently, a series of technological developments referred to as multiband MRI, have sped up the temporal resolution of fMRI acquisitions by approximately an order of magnitude (i.e., from 2 s to 0.2 s, for whole-brain imaging), and appear likely to offer the possibility for even further acceleration. As previously mentioned, most BOLD fMRI data are acquired by sequentially acquiring a series of two-dimensional slices. In contrast, multiband MRI excites multiple slices simultaneously, and the MR signals from these slices are then separated using multiple receiver coils and the aid of special encoding techniques. Thus, multiband MRI combines hardware and software innovations to significantly speed-up fMRI acquisitions.

This unprecedented temporal resolution provides new challenges with regards to statistical analysis, but also offers new opportunities. Most analytic approaches to fMRI data are based on assumptions that may be inappropriate or suboptimal for the rapidly sampled data obtained using multiband approaches. Therefore, there is an opportunity to create new methods for the analysis of multiband data that will enhance the specificity and sensitivity of BOLD fMRI outcome measures and harness the increased information content. This is an exciting area where quantitative scientists promise to play an important role.

Harnessing Large-scale Data Bases: As previously mentioned, there have been numerous efforts to construct large-scale imaging databases. These endeavors have been performed on both a grassroots (e.g., the 1000 Functional Connectomes Project) and institutional level (e.g., the Human Connectome Project), and databases consisting of more than 1,000 subjects are becoming increasingly available.

Many times, the largest imaging data sets tend to be collected for reasons other than a specific targeted scientific hypothesis. In these cases, data quality, sampling bias, missing data, data quality and the availability of covariates or outcomes tend to be problematic. This is in contrast to experiments consisting of a small number of subjects, where tight experimental control leads to direct measures for testing hypotheses of interest at the expense of lower power. Broadly speaking, large data sets used in isolation are useful for exploratory or predictive exercises, while small data sets, useful for knowledge creation and confirmatory analyses are critically hampered by low power.

As the ultimate goal is to create new scientific knowledge, it would be ideal to use information from large data sets to inform the analysis of small sample data. Large data sets can be explored with the goal of establishing norms

and priors for the use in small, more targeted data sets. To achieve this goal we believe that there is a need for more quantitative researchers to become involved in the field in order to make sense of the massive amounts of data being generated.

Multi-modal Analysis: All methods used in the human neuro-behavioral sciences have limitations, and fMRI is of course no exception. Therefore the current trend is towards increasingly interdisciplinary approaches that use multiple methodologies to overcome some of the limitations of each method used in isolation. For example, fMRI data are increasingly combined with EEG and MEG data to improve temporal precision, among other benefits. Advances in engineering and signal processing, for example, allow EEG and fMRI data to be collected simultaneously. Neuroimaging data is also being combined with transcranial magnetic stimulation to integrate the powerful ability of neuroimaging to observe brain activity with the ability afforded by TMS to manipulate brain function and examine causal effects [5]. Finally, integrating genetics with brain imaging is seen as a way to study how a particular subset of polymorphisms may affect functional brain activity. In addition, quantitative indicators of brain function could facilitate the identification of the genetic determinants of complex brain-related disorders such as autism, dementia and schizophrenia [15].

Each of these multi-modal approaches promise to be important topics of future research, and to fully realize their promise, novel statistical techniques will be needed. Ultimately, combining information from different modalities will be challenging to data analysts, if for no other reason than that the amount of data will significantly increase. In addition, since different modalities are measuring fundamentally different quantities, it is not immediately clear how to best combine the information. However, clearly, this is an extremely important problem that has already started to become a major area of research.

Bibliography

- [1] G. K. Aguirre, E. Zarahn, and M. D’Esposito. The variability of human, BOLD hemodynamic responses. *NeuroImage*, 8(4):360–369, 1998.
- [2] A.H. Andersen, D.M. Gash, and Avison M.J. Principal component analysis of the dynamic response measured by fmri: a generalized linear systems framework. *Magnetic Resonance in Medicine*, 17:785–815, 1999.
- [3] R. M. Birn, Z. S. Saad, and P. A. Bandettini. Spatial heterogeneity of the nonlinear dynamics in the fmri bold response. *NeuroImage*, 14:817–826, 2001.
- [4] Bharat Biswal, F Zerrin Yetkin, Victor M Haughton, and James S Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic resonance in medicine*, 34(4):537–541, 1995.
- [5] D. E. Bohning, A. P. Pecheny, C. M. Epstein, A. M. Speer, D. J. Vincent, W. Dannels, and M.S. George. Mapping transcranial magnetic stimulation (tms) fields in vivo with mri. *Neuroreport*, 8:2535–2538, 1997.
- [6] G. M. Boynton, S. A. Engel, G. H. Glover, and D. J. Heeger. Linear systems analysis of functional magnetic resonance imaging in human v1. *J. Neurosci*, 16:4207–4221, 1996.
- [7] R. B. Buxton, E. C. Wong, and L. R. Frank. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magnetic Resonance in Medicine*, 39:855–864, 1998.
- [8] V. D. Calhoun, T. Adali, G.D. Pearlson, and J.J. Pekar. Spatial and temporal independent component analysis of functional mri data containing a pair of task-related waveforms. *Human Brain Mapping*, 13:43–53, 2001.
- [9] Catie Chang, John P Cunningham, and Gary H Glover. Influence of heart rate on the bold signal: the cardiac response function. *Neuroimage*, 44(3):857–869, 2009.
- [10] Nathan W Churchill, Anita Oder, Herve Abdi, Fred Tam, Wayne Lee, Christopher Thomas, Jon E Ween, Simon J Graham, and Stephen C Strother. Optimizing preprocessing and analysis pipelines for single-subject fmri. i. standard temporal motion and physiological noise correction methods. *Human brain mapping*, 33(3):609–627, 2012.

- [11] Robert W Cox. Afni: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research*, 29(3):162–173, 1996.
- [12] A. M. Dale. Optimal experimental design for event-related fmri. *Human Brain Mapping*, 8:109–114, 1999.
- [13] K. J. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *NeuroImage*, 19:1273–1302, 2003.
- [14] K.J. Friston. Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, 2:56–78, 1994.
- [15] D. C. Glahn, P. M. Thompson, and J. Blangero. Neuroimaging endophenotypes: strategies for finding genes influencing brain structure and function. *Human Brain Mapping*, 28:488–501, 2007.
- [16] G. H. Glover. Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage*, 9(4):416–429, 1999.
- [17] G. H. Glover. Simple analytic spiral k-space algorithm. *Magnetic Resonance in Medicine*, 42(2):412–415, 1999.
- [18] C. Goutte, F. A. Nielsen, and L. K. Hansen. Modeling the haemodynamic response in fmri using smooth fir filters. *IEEE Trans Med Imaging*, 19:1188–1201, 2000.
- [19] J. Grinband, T. D. Wager, M. Lindquist, V. P. Ferrera, and J. Hirsch. Detection of time-varying signals in event-related fmri designs. *NeuroImage*, 43(3):509–520, 2008.
- [20] H. Gudbjartsson and S. Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34:910–914, 1995.
- [21] Mark E. Haacke, Robert W. Brown, Michael R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Wiley-Liss, June 1999.
- [22] S. Hayasaka and T.E. Nichols. Combining voxel intensity and cluster extent with permutation test framework. *NeuroImage*, 23:54–63, 2004.
- [23] Scott A. Huettel, Allen W. Song, and Gregory Mccarthy. *Functional Magnetic Resonance Imaging*. Sinauer Associates, April 2004.
- [24] Mark Jenkinson, Christian F Beckmann, Timothy EJ Behrens, Mark W Woolrich, and Stephen M Smith. Fsl. *Neuroimage*, 62(2):782–790, 2012.
- [25] Tom Johnstone, Kathleen S Ores Walsh, Larry L Greischar, Andrew L Alexander, Andrew S Fox, Richard J Davidson, and Terrence R Oakes. Motion correction and the use of motion covariates in multiple-subject fmri analysis. *Human brain mapping*, 27(10):779–788, 2006.

- [26] Ming-Hung Kao, Abhyuday Mandal, Nicole Lazar, and John Stufken. Multi-objective optimal experimental designs for event-related fmri studies. *NeuroImage*, 44(3):849–856, 2009.
- [27] Ming-Hung Kao, Abhyuday Mandal, and John Stufken. Constrained multiobjective designs for functional magnetic resonance imaging experiments via a modified non-dominated sorting genetic algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 61(4):515–534, 2012.
- [28] Kenneth K Kwong, John W Belliveau, David A Chesler, Inna E Goldberg, Robert M Weisskoff, Brigitte P Poncelet, David N Kennedy, Bernice E Hoppel, Mark S Cohen, and Robert Turner. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences*, 89(12):5675–5679, 1992.
- [29] David J Larkman, Joseph V Hajnal, Amy H Herlihy, Glyn A Coutts, Ian R Young, and Gösta Ehnholm. Use of multicoil arrays for separation of signal from multiple slices simultaneously excited. *Journal of Magnetic Resonance Imaging*, 13(2):313–317, 2001.
- [30] C. Liao, K. J. Worsley, J-B. Poline, G. H. Duncan, and A. C. Evans. Estimating the delay of the response in fMRI data. *NeuroImage*, 16:593–606, 2002.
- [31] M. A. Lindquist, J.M. Loh, L. Atlas, and T. D. Wager. Modeling the hemodynamic response function in fmri: Efficiency, bias and mis-modeling. *NeuroImage*, 2008.
- [32] M. A. Lindquist and T. D. Wager. Validity and power in hemodynamic response modeling: A comparison study and a new approach. *Human Brain Mapping*, 28:764–784, 2007.
- [33] M.A. Lindquist, C.H. Zhang, G. Glover, and L.A. Shepp. Acquisition and statistical analysis of rapid 3d fmri data. *Statistica Sinica*, 2008.
- [34] M.A. Lindquist, C.H. Zhang, G. Glover, and L.A. Shepp. Rapid three-dimensional functional magnetic resonance imaging of the negative bold response. *Journal of Magnetic Resonance*, 2008.
- [35] Martin A Lindquist et al. The statistical analysis of fmri data. *Statistical Science*, 23(4):439–464, 2008.
- [36] T.T. Liu and L.R. Frank. Efficiency, power, and entropy in event-related fmri with multiple trial types: Part i: theory. *NeuroImage*, 21:387–400, 2004.
- [37] N. K. Logothetis. Can current fmri techniques reveal the micro-architecture of cortex? *Nature Neuroscience*, 3:413, 2000.

- [38] Nikos K Logothetis, Jon Pauls, Mark Augath, Torsten Trinath, and Axel Oeltermann. Neurophysiological investigation of the basis of the fmri signal. *Nature*, 412(6843):150–157, 2001.
- [39] J.M. Loh, M. A. Lindquist, and T. D. Wager. Residual analysis for detecting mis-modeling in fmri. *Statistica Sinica*, 2008.
- [40] T. E. Lund, K. H. Madsen, K. Sidaros, W. L. Luo, and T. E. Nichols. Non-white noise in fmri: Does modelling have an impact? *NeuroImage*, 29:54–66, 2006.
- [41] D. Malonek and A. Grinvald. The imaging spectroscopy reveals the interaction between electrical activity and cortical microcirculation: implication for optical, pet and mr functional brain imaging. *Science*, 272:551–554, 1996.
- [42] P. Mansfield. Multi-planar image formation using nmr spin echoes. *Journal of Physics*, C10:L55–L58, 1977.
- [43] A. McIntosh and F. Gonzalez-Lima. Structural equation modeling and its application to network analysis in functional brain imaging. *Human Brain Mapping*, 2:2–22, 1994.
- [44] M. J. McKeown and S. Makeig. Analysis of fmri data by blind separation into independent spatial components. *Human Brain Mapping*, 6:160–188, 1998.
- [45] Andrea Mechelli, Will D Penny, Cathy J Price, Darren R Gitelman, and Karl J Friston. Effective connectivity and intersubject variability: using a multisubject network to test differences and commonalities. *Neuroimage*, 17(3):1459–1469, 2002.
- [46] R.S. Menon, D. C. Luknowsky, and J. S Gati. Mental chronometry using latency resolved functional mri. *Proc. Natl. Acad Sci. USA*, 95, 1998.
- [47] F.M. Miezin, L. Maccotta, J.M. Ollinger, S.E. Petersen, and R.L. Buckner. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *NeuroImage*, 11, 2000.
- [48] Steen Moeller, Essa Yacoub, Cheryl A Olman, Edward Auerbach, John Strupp, Noam Harel, and Kâmil Uğurbil. Multiband multislice ge-epi at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fmri. *Magnetic Resonance in Medicine*, 63(5):1144–1153, 2010.
- [49] Kevin Murphy, Rasmus M Birn, Daniel A Handwerker, Tyler B Jones, and Peter A Bandettini. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *Neuroimage*, 44(3):893–905, 2009.

- [50] S. Ogawa, D.W. Tank, R. Menon, J.M. Ellerman, S.G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping and magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 89:5951–5955, 1992.
- [51] William D Penny, Karl J Friston, John T Ashburner, Stefan J Kiebel, and Thomas E Nichols. *Statistical Parametric Mapping: The Analysis of Functional Brain Images: The Analysis of Functional Brain Images*. Academic Press, 2011.
- [52] Francisco Pereira, Tom Mitchell, and Matthew Botvinick. Machine learning classifiers and fmri: a tutorial overview. *Neuroimage*, 45(1):S199–S209, 2009.
- [53] Russell A Poldrack, Jeanette A Mumford, and Thomas E Nichols. *Handbook of functional MRI data analysis*. Cambridge University Press, 2011.
- [54] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, and P. Boesiger. Sense: sensitivity encoding for fast mri. *Magnetic Resonance in Medicine*, 42:952–956, 1999.
- [55] P. L. Purdon, V. Solo, R. M. Weissko, and E. Brown. Locally regularized spatiotemporal modeling and model comparison for functional MRI. *NeuroImage*, 14:912–923, 2001.
- [56] J. J. Riera, J. Watanabe, I. Kazuki, M. Naoki, E. Aubert, T. Ozaki, and R. Kawashima. A state-space model of the hemodynamic approach: nonlinear filtering of bold signals. *NeuroImage*, 21:547–567, 2004.
- [57] A. Roebroeck, E. Formisano, and R. Goebel. Mapping directed influence over the brain using granger causality and fmri. *NeuroImage*, 25:230–242, 2005.
- [58] Mikail Rubinov and Olaf Sporns. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52(3):1059–1069, 2010.
- [59] D. L. Schacter, R. L. Buckner, W. Koutstaal, A. M. Dale, and B. R. Rosen. Rectangular confidence regions for the means of multivariate normal distributions. *Late onset of anterior prefrontal activity during true and false recognition: an event-related fMRI study*, 6:259–269, 1997.
- [60] Michael E Sobel and Martin A Lindquist. Causal inference for fmri time series data with systematic errors of measurement in a balanced on/off study of social evaluative threat. *Journal of the American Statistical Association*, (just-accepted):00–00, 2014.
- [61] D.K. Sodickson and W.J. Manning. Simultaneous acquisition of spatial harmonics (smash): fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38:591–603, 1997.

- [62] Olaf Sporns. *Networks of the Brain*. MIT press, 2011.
- [63] Jean Talairach and Pierre Tournoux. Co-planar stereotaxic atlas of the human brain. 3-dimensional proportional system: an approach to cerebral imaging. 1988.
- [64] A. L. Vazquez, E. R. Cohen, V. Gulani, L. Hernandez-Garcia, Y. Zheng, G. R. Lee, S. G. Kim, J. B. Grotberg, and D. C. Noll. Vascular dynamics and bold fmri: Cbf level effects and analysis considerations. *NeuroImage*, 32:1642–1655, 2006.
- [65] T. D. Wager and T. E. Nichols. Optimization of experimental design in fmri: a general framework using a genetic algorithm. *NeuroImage*, 18:293–309, 2003.
- [66] T. D. Wager, A. Vazquez, L. Hernandez, and D. C. Noll. Accounting for nonlinear BOLD effects in fMRI: parameter estimates and a model for prediction in rapid event-related studies. *NeuroImage*, 25(1):206–218, 2005.
- [67] Tor D Wager, Lauren Y Atlas, Martin A Lindquist, Mathieu Roy, Choong-Wan Woo, and Ethan Kross. An fmri-based neurologic signature of physical pain. *New England Journal of Medicine*, 368(15):1388–1397, 2013.
- [68] M. W. Woolrich, T. E. Behrens, and S. M. Smith. Constrained linear basis sets for HRF modelling using variational Bayes. *NeuroImage*, 21(4):1748–1761, 2004.
- [69] K. J. Worsley and K. J. Friston. Analysis of fMRI time-series revisited—again. *NeuroImage*, 2:173–181, 1995.
- [70] E. Yacoub, T.H. Le, and X. Hu. Detecting the early response at 1.5 tesla. *NeuroImage*, 7:S266, 1998.
- [71] E. Zarahn. Using larger dimensional signal subspaces to increase sensitivity in fmri time series analyses. *Hum Brain Mapp*, 17:13–16, 2002.