Functional Connectivity in the Cortical Circuits Subserving Eye Movements

Christopher R. Genovese* and John A. Sweeney+

* *Department of Statistics*

Carnegie Mellon University

\+ *Neurobehavioral Studies Program*

Western Psychiatric Institute and Clinic

Running Head: Cortical Circuits Subserving Eye Movements

Keywords: Oculomotor System, Functional Neuroimaging, Hierarchical Models

Acknowledgments: This work is supported by NSF Grants DMS 9505007 and DMS 9705034 and NIH Grants MH 422969, MH 45156, MH 01423, NS 35949, HD 35469, and NARSAD. The author would also like to thank Rob Kass, the discussants, and the referees for helpful comments.

Correspondence:

Christopher R. Genovese
Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213

Phone: (412) 268-7836
Fax: (412) 268-7828
E-mail: genovese@stat.cmu.edu

ABSTRACT

The eyes move continually during visual processing, usually without explicit control or awareness, to scan the key features of a scene. This is essential to our ability to attend to multiple features of our environment and to extract useful information from complex visual stimuli. Eye movement abnormalities are a significant and reliable neurobehavioral marker for a number of major neurological, developmental, and psychiatric disorders, so enhanced understanding of how the brain controls eye movements can yield insights into the brain abnormalities at the root of these conditions.

Eye movements have been studied extensively in both humans and monkeys. A general picture has emerged from these data regarding what areas of the brain subserve eye-movement processes. Although there is a strong homology between humans and monkeys, there are also notable differences, and a detailed delineation of the system in the human brain is still needed. To develop and test a complete theory of how eye movements are implemented requires study of the component sub-processes of the human system, of the interactions among these sub-processes, and of their functional connectivity. The advent of neuroimaging has opened the door to exciting new advances in this area.

We use functional Magnetic Resonance Imaging (fMRI) to study the eye-movement system in humans. Among neuroimaging techniques, fMRI offers a superior combination of spatial and temporal resolution. Each experiment yields as data the realization of a complicated spatio-temporal process that contains information about the dynamics of neural processing during eye movements. We apply a Bayesian hierarchical model for these data to make inferences about the system. In particular, we address an open question: What is the functional relationship between the neural systems subserving saccadic eye movements (rapid repositionings) and smooth visual pursuit of a target? We also illustrate several computational and statistical issues that arise in making inferences from these data.

1. Introduction

The human brain is roughly three pounds of protoplasm, blood, and salt water that could be held in one hand, but it is arguably the most complex object known. The brain contains approximately 100 billion neurons, each of which typically has thousands of synaptic connections to other neurons. The operation of an individual neuron is conceptually simple: it maintains a resting electrical potential across its membrane and, when suitably stimulated, "fires" to send chemical messages (neurotransmitters) to other neurons. It is the richness and plasticity of these connections that enable the brain to store information and build associations at many levels. The challenge of understanding the brain is to understand how complex behaviors, from sensory and motor functions to intelligence and problem solving, can emerge from a combination of such relatively simple units.

In recent years, this challenge has been addressed from several scientific directions. Neurophysiologists have used animal models to learn how neuronal activity changes in response to various tasks and stimuli, and from this they have uncovered the workings of many specialized brain regions and neural circuits. Neuropsychologists study behavioral processes in both humans and animals to discover how specific brain areas subserve different sensory, motor and cognitive functions. The field adopts a primarily "bottom up" methodology in which the behavioral effects of circumscribed brain lesions are measured in a controlled fashion. In contrast, cognitive psychologists take a "top down" approach, building abstract models of memory, cognition, and problem solving that characterize the structure of the mind and that make testable predictions for empirical studies. Psychiatrists and neurologists gain insight into the brain by studying the effects of various clinical disorders in patients. Neurosurgeons learn about brain-behavior relationships by assessing the impact of brain surgery on their patients' functioning and from electrically stimulating the brain during neurosurgical procedures. The combined knowledge derived from these fields forms the basis of our current understanding of brain organization and function.

Great strides in understanding brain function have come in recent years from the application of new technologies that allow non-invasive physiological monitoring of the human brain in action. This field—called functional neuroimaging—enables characterization of regional changes in brain activity in response to varying task demands. Although it began with Positron Emission Tomography (PET), the field now has shifted largely to functional Magnetic Resonance Imaging (fMRI) as the technique of choice because of its superior temporal and spatial resolution. Functional neuroimaging has advanced each neuroscientific discipline separately and has spurred multiple integrative efforts across disciplinary boundaries. Animal physiologists can learn about the similarities and differences between humans and non-human primates and can study physiologic changes across the entire brain simultaneously rather than neuron by neuron. Clinical investigators can directly localize brain pathology non-invasively and *in vivo*. Neuropsychologists and cognitive psychologists can directly test their models. This synergy has greatly advanced our understanding of one of the great frontiers of science—discovering how the brain works.

In this paper, we concentrate on the oculomotor system, the brain system that controls eye movements [37, 33]. From neuronal recordings in non-human primates and from clinical studies of humans, this system is known to consist of multiple components that are widely distributed across the brain. It involves both low-level functions for motor control and high-level processes that integrate sensory input and cognitive functions in more sophisticated ways. Neuroimaging is especially informative regarding this system for several reasons: (i) a great deal is already known about the neurophysiology of the oculomotor system from extensive monkey studies, and neuroimaging allows for the study of this system in humans, with guidance from the monkey studies and the well established parallels between human and non-human primates; (ii) the oculomotor system is known to involve the interaction of multiple brain areas that are widely distributed, and neuroimaging allows for the simultaneous study of distributed areas where other techniques cannot; (iii) critical distinctions in oculomotor processing occur at a fine-scale across both space and time; neuroimaging offers a excellent spatial and temporal resolution with the capability of choosing a balance between them to meet specific needs; and (iv) previously established associations between eye-movement impairment and various neurophysiological abnormalities yield specific predictions that can be tested with a neuroimaging experiment.

We study the oculomotor system using fMRI. During an fMRI experiment, a participant performs a sequence of behavioral tasks while Magnetic Resonance (MR) images of the subject's brain are acquired at regular intervals (on the order of a second). The brain's response to task performance gives rise to small changes in the images over time. A central challenge in the statistical analysis of fMRI data is to identify and characterize these task-related signal changes. The challenge of interpretation is even greater: to determine from the task-related signal changes (if any) how the brain subserves the processes under study. We apply a modeling framework described in [26] to address both these challenges by answering specific scientific questions about the workings of the oculomotor system.

We begin in the next section with a detailed description of the oculomotor system and the potential scientific and clinical gains to be achieved by studying this system. Within this context, we discuss the questions and issues that motivate our experiments, which in turn will be discussed in later sections. In Section 3, we present an overview of fMRI, from the fundamentals of signal acquisition to the statistical methods for analysis of the data. Section 4 describes the model for fMRI data on which we base our analysis. This is a hierarchical model that attempts to capture the critical sources of variation in fMRI data, originally proposed in [26]. We give the details of our experimental methods in Section 5 and describe the data and results of our analyses in Section 6. These data and analyses are only the first step in a long term study of the oculomotor system, and they give rise to a variety of more detailed questions. Finally, in Section 7, we discuss the impact of these results from a wider perspective and highlight some key issues and some future directions for our research.

To set these ideas in a proper context and to guide the reader through the details of later sections, it will be helpful to describe first the basic anatomical organization of the brain. To put the strengths and weaknesses of fMRI in perspective, we also briefly summarize the range of techniques that can be used to study the brain. Before proceeding, the reader should be warned that the words "function" and "functional" are used quite frequently in this paper but in several different senses. In a neuroscience context, "function" refers to the workings of the brain as a control and information processing system as opposed to "structure" which relates to the anatomical and physical layout of the brain. Thus, *structural* MRI attempts to infer anatomical features of the brain from the images; this is what one would use after being bumped on the head. *Functional* MRI, on the other hand, focuses not on the brain images themselves but on fluctuations in the measurements that relate to the firing of neurons in the active brain. In a mathematical context, function and functional take on their more familiar meanings. We hope that the intended meanings are clear from the text.

1.1. The Gross Anatomy of the Brain

The brain is subdivided into three basic parts: the cerebral hemispheres, the cerebellum, and subcortical areas including the brain stem. Refer to Figures 1, 2, and 3 throughout this discussion. It will be helpful, here and below, to have some terminology for discussing locations in the brain. Suppose we choose an arbitrary reference point in the brain. Any structure closer to the front (back) of the brain than the reference point is *anterior* (*posterior*) to that point. Any structure that is higher (lower) than the reference point is *superior* (*inferior*) or, equivalently, *dorsal* (*ventral*) to it. The plane between the left and right hemispheres demarcates the midline of the brain. Any structure closer (farther) to the midline than the reference point is *medial* (*lateral*) to it. Left and right are used to indicate lateral directions taken from the back of the brain looking forward. All of these directions are indicated on Figure 1. MR images whose slices are oriented perpendicular to the body axis are called *axial* images. By radiological convention, these images are arranged as though looking at brain from the perspective of a doctor facing a patient; thus, left in the image is right in the body and vice versa.

The two cerebral hemispheres are approximate mirror images of each other, separated by the deep longitudinal or inter-hemispheric fissure and connected by a mass of fibers called the corpus callosum. The surface of the hemispheres consists of convoluted folds of "grey matter", a sheet of neurons varying in thickness from two to five millimeters; this is the *cerebral cortex* and contains the neurons responsible for most higher brain functions. The neurons in the cerebral cortex are arranged in six distinct and identifiable layers containing different types of neurons and connections. The thickness of these layers varies across the brain. The cortex can be thought of as a continuous sheet that is folded in a complicated but apparently systematic way. The crest of the convolutions are called *gyri* (singular gyrus), and the canyons between the gyri are called *sulci* (singular sulcus). Beneath the cortex is a mass of fibers ("white matter") that carry information between nearby cells in the grey matter and other parts of the brain. Deep in the brain, above the brain stem and under the cortex, is a collection of large, discrete clusters of cells called the *basal ganglia* that manage a variety of motor functions. Another large collection of nuclei (neuronal clusters) comprise the thalamus, which is a major relay station that transmits sensory information to the cerebral cortex and connects cerebellum, basal ganglia, cortex, and brain stem.

The cerebral cortex is divided into four main lobes: frontal, parietal, temporal, and occipital. (See Figure 1.) These lobes appear somewhat arbitrary anatomically but are distinct both histologically and functionally. The *occipital lobe*, in the most posterior section of cortex, is exclusively devoted to processing visual input. The *parietal lobe* controls higher visual and somato-sensory processing, navigation, and management of internal representations of the environment. The *temporal lobe* is primarily responsible for auditory processing, the integration of diverse sensory information, language comprehension, and object recognition. The *frontal lobe* controls motor function (including eye movements), speech, and many higher cognitive functions. The most anterior part of the frontal lobe, called the *prefrontal cortex*, appears to be quite special. Relative to other mammals, and even other primates, the prefrontal cortex is substantially enlarged in humans. The role of the prefrontal cortex in high-level cognition is still being investigated, but it currently appears to be involved in the context-sensitive modulation of attention and the voluntary control of behavior. The prefrontal cortex thus interacts with most other systems and is, in a sense, the primary player in high level functions such as planning and problem solving.

Knowledge of a few key neuro-anatomical landmarks is helpful for navigating around the brain. The *central sulcus* separates the parietal and frontal lobes. On the anterior side of this sulcus and coming up onto the gyrus there (called the precentral gyrus) is the primary motor area. On the posterior side of this sulcus and coming up onto the gyrus there (called the postcentral gyrus) is the primary somatosensory area governing perception of such sensations as touch and vibration. The lateral sulcus separates the temporal lobe from the frontal and parietal lobes. On the lateral surface of the temporal lobe itself, there are three identifiable gyri, one above another: the superior, middle, and inferior temporal gyri. Within the superior temporal gyrus near the base of the primary somatosensory area is the area involved in basic auditory perception. At the posterior aspect of this same gyrus in the left hemisphere is Wernicke's area which plays an important role in language comprehension. The lateral surface of the frontal lobe is similarly divided into gyri: superior, middle, and inferior. Roughly on the posterior aspect of the inferior frontal gyrus in the left hemisphere, anterior to the precentral gyrus, is Broca's area which plays the primary role in language expression.

Such a manifest of function-location pairs may give the mistaken impression that the brain is just a collection of specialized units. While true in terms of elementary processes, higher cognitive processing emerges from a complex interaction among these different systems. These functions tend to be organized into integrated pathways that allow hierarchically elaborative processing of stimuli. As we will see below, the oculomotor system is an exemplar of the hierarchical organization of the brain. Other examples include the so-called "what" and "where" pathways for object recognition [30]. The recognition of a visual pattern as a particular object starts in the occipital cortex with a decomposition of the input into various visual features, such as angles, edges, colors, shading, perspective, motion. This information is integrated along a pathway—the "what" pathway—that proceeds (ventrally) from primary visual cortex into the temporal lobe. The "where" pathway, which selectively localizes objects in space rather than identifying them, also begins with a collection of visual features but its outputs from visual cortex are different. They proceed (dorsally) up to the parietal lobe, which processes information to place an object within an internal representation of the environment.

The *cerebellum* lies beneath the occipital and posterior temporal lobes and directly posterior to the brain stem; it is clearly identifiable by its very finely folded surface. The cerebellum consists of two lateral cerebellar hemispheres and the midline vermis. The vermis is phylogenetically the oldest area of the cerebellum. It is responsible for the fine regulation of movement, such as enabling one to reach accurately for an object on a direct line without meandering. An individual with damage to the vermis shows motor function similar to a very intoxicated person in that reaching efforts are grossly imprecise, e.g., as in repeated finger to nose pointing. The cerebellar hemispheres are less well understood. They are each arranged in a series of lobules which have a variety of functions. Portions of this region appear to govern fine regulation of timing for movement, such as when to reach to catch a moving object. The lateral regions of the cerebellar hemispheres are phylogenetically newer structures, showing a pattern of development in size similar to the frontal lobe of the cerebral hemispheres.

The *brain stem* is the phylogenetically oldest part of the brain. The brain stem consists of small clusters of neurons (called nuclei) that serve very specific functions. It is the gateway for motor commands from the brain to the rest of the body; it regulates respiration and heart rate; and it has an important role in primitive functions such as sleep. The brain stem also contains the vast majority of the cell bodies for neurons that release modulatory neurotransmitters—the complex chemical messages that modulate the activity of excitatory and inhibitory pathways. These cell bodies project to points all over the cerebral cortex in order to deliver these vital chemicals.

1.2. Studying the Brain

Historically, the study of the brain has progressed at two levels, the macro-level "systems" approach and the micro-level "cellular" approach. The systems approach focuses on the workings of entire brain systems and their relation to observed behavior. Detailed analysis of postmortem brains provides the opportunity to establish clinico-pathological associations by identifying observational correlations between particular clinical/behavioral dysfunctions and the location of lesions or other brain damage. A more advanced method is electrical stimulation of the brain, where an electrical impulse is applied at a particular location either to the surface of the brain or deeper and the corresponding behavioral response recorded. Stimulation has been performed on humans prior to brain surgery and more commonly using animal subjects. Empirical observations of behavior (esp. perception, memory, and problem-solving) provide another tool. This is the basis of cognitive psychology, which attempts to model the function of the brain at an abstract level and understand the capabilities required for high-level cognition to emerge. Results from cognitive psychology set constraints on how the brain can operate. Studies of human development offer further constraints.

The cellular approach focuses on the workings of individual neurons and specialized neural circuits. It involves invasive techniques that continue to be very important for understanding brain mechanisms. The most basic of these techniques is the evaluation of postmortem brain tissue using a variety of dyes and stains that differentially mark various types of cells or neuronal parts. The detailed cellular organization of the brain, particularly the cerebral cortex, has been revealed in this way. Direct cellular (or single unit) recording allows much greater precision but is so invasive that it cannot be used on human subjects. In this technique, a probe electrode is inserted into the brain in close proximity to a single neuron. The electrode records the electrical activity of that neuron over time, and its degree of task-related activity at any time is indicated by its frequency of firing. This method gives extraordinarily detailed information about single neurons, and by painstaking work, the function of many specific circuits has been revealed. Multiple neurons can be recorded simultaneously, but it is currently infeasible to do so if the neurons are widely separated. Hence, this method cannot measure interactions among the distributed neural circuits that are responsible for high-level functions.

In recent years, a synthesis has been forming between the systems and cellular and systems approaches. Structural neuroimaging technologies offered the first non-invasive way to identify, diagnose, and track the development of lesions, tumors, and other brain pathologies. The first of these methods to be widely used is Computed Tomography (CT, the "cat" scan) in which focused beams of x-rays are transmitted through the subjects head from many angles. Because the rate of x-ray transmission through tissue depends on the tissue density, it is possible (with sufficiently many views) to reconstruct an image of tissue density within the brain. CT imaging revolutionized diagnostic radiology, and it was the first radiological technique that required non-trivial statistical methodology to construct the images. Magnetic Resonance Imaging (MRI) measures the density of particular atomic nuclei in tissue by careful manipulation of magnetic fields (see Section 3). Structural MR images can be acquired at very high spatial resolution and provide excellent contrast for clinical radiology.

The advent of functional neuroimaging took this a step further with the non-invasive study of the *active* brain, allowing direct connections to be made between the systems and cellular levels of inquiry. Functional neuroimaging differs from structural neuroimaging in that the technique must be sensitive to physiological changes in the brained that are correlated with varying task demands. Event Related Potentials (ERP) measure functional changes directly from the electrical activity of the brain. This can be based on variation in electric fields, Electro-Encephalography (EEG), or magnetic fields, Magneto-Encephalography (MEG). The electrical activity of neurons firing during the task causes measurable changes in these fields. In ERP research, field variations are measured on several "channels" as the subject repeatedly performs some behavioral task (e.g., viewing blinking lights), and under some assumptions, features of the neural activity can be reconstructed from the measured signal. In practice, ERP's provide excellent temporal resolution, but because the underlying reconstruction problems are terribly ill-posed, it offers very poor spatial resolution.

Positron Emission Tomography (PET) images are acquired by injecting a radioactive tracer into the subject's bloodstream [65]. As the tracer decays, it emits positrons that collide with electrons in the tissue after moving a short distance; the collision gives off a pair of photons (gamma rays) that move in opposite directions. A detector ring around the subject's head can count these photon pairs, and the tracer density in the tissue is estimated from these counts. PET can be used to obtain images that describe a variety of brain processes, including a measure of cerebral blood flow that relates to brain activity [43]. PET images have moderate spatial resolution, but because of limits on the amount of radioactive material to which a subject can be exposed, it is common to average PET images across subjects, further blurring the results. Moreover, a dose of the tracer must decay sufficiently (often over several minutes) before the next image can be obtained, and each image represents an average over this decay period. Subject movement in the scanner becomes a critical concern during this period. PET has tremendous potential as a way of tracing other metabolic processes, such as regional glucose metabolism, that are inaccessible to other methods, but its role in functional neuroimaging is being largely subsumed by fMRI.

Among current neuroimaging methods, fMRI offers a superior combination of spatial and temporal resolution, and because the technique involves no known toxicity, subjects can be imaged as often as needed. Functional MRI is having a powerful impact on scientific exploration in many fields and promises to play a vital role in our efforts to understand the brain. We discuss the details of fMRI acquisition in Section 3.

Figure 4 compares the spatial and temporal resolution among all these techniques. The potential for functional neuroimaging, and fMRI in particular, lies in its ability to study higher cognitive processes in humans and to reveal the processing throughout the brain simultaneously.

2. The Oculomotor System: Importance and Function

Of all the senses, vision plays the largest role in helping us interact with our environment. As any computer vision specialist can testify, the challenge of making sense of the myriad visual stimuli we receive is immense [40]. We perceive objects, but the eye sees only gradations of intensity and color in the incident light. We perceive depth and distance, but the eye sees only a projection. We perceive a stable world around us even as we move, but the eye sees only a set of disparate images from different angles and perspectives. Vision requires more than a passive recording of information; it requires a coordinated, integrated decomposition of sensory input at many levels. This is the task of the visual system.

The windows to this system are the eyes. The retina, the light-sensitive region on the back of the eye, does substantial processing of visual input and then passes the results to the brain's visual system via the optic nerve. The center of the retina, called the fovea, is the part of the retina that is most sensitive to color and fine-scale features. Because the fovea is relatively small, eye movements are a critical part of visual processing; in primates, eye movement is the primary mechanism by which the fovea can be aimed at objects of interest. Through eye movements, the brain controls what visual input is processed, and thus the oculomotor system is very much involved in the processes of perception. The development of the fovea and the primary use of eye movements (rather than head movements) to focus the eyes on different targets are phylogenetically late developments present only in higher mammals. It is an advance that supports many high-level perceptual activities in humans such as reading text [32].

There are two main systems—saccades and pursuit—by which the brain positions the fovea on objects of interest. Saccades are the gross re-positionings of eye fixations by rapid shifting of gaze and attention from one point to another. Saccadic eye movements are very fast (lasting 10's of milliseconds depending on their size) and precise (capable of shifting the eyes a large distance to within a degree of angle of the intended location) [3], and they typically occur without explicit control or awareness. Pursuit eye movements keep the eyes focused on slowly moving targets by matching the angular velocity of the eyes to that of a target of interest [9].

The adaptive value of these eye movements is clear: if you want to find and catch dinner for yourself and not be dinner for someone else, it is helpful to be able to scan your surroundings quickly and precisely. In daily life, the eyes scan the key features of the environment approximately 3-4 times per second with some combination of involuntary saccades and pursuit. If there is potentially important information at a location away from the current focus of the eyes, the oculomotor system is engaged to shift the eyes to that new location. Hence, by determining what visual information is processed, the oculomotor system controls the perceptual evaluation of the visual environment.

The process underlying control of eye movements is more complicated than it may initially appear. Maintaining a stable visual perception through the blur of information that occurs during a saccade or pursuit requires holding internal representation of the environment that is separate from the immediate visual input. In addition, whenever the eyes move, the brain must transform its internal coordinate systems, the spatial representations by which it integrates different views into a coherent sense of what and where [13]. The transformation of these spatial representations are intimately involved in the planning and coordination of motor actions. Eye movements are also modulated by the higher-order processes controlling visual attention and working memory [58, 24]. These interactions indicate that the oculomotor system is embedded into many levels of cognitive processing. Consequently, the study of eye movements offers insights into a variety of brain systems, with implications for both basic and clinical neuroscience. In this section, we describe in some detail the central questions and potential gains that motivate our study of the oculomotor system.

2.1. Layout of the System

The oculomotor system is controlled by the complex interaction of several functional components that are distributed throughout the brain. Before examining the underlying scientific issues, it will be helpful to paint a schematic of the system as it is currently understood.

The main areas involved in the oculomotor network are depicted in Figures 2 and 3. The basic functions of these areas are as follows.

The Eyes. As described above, the eyes control and process the input into the system. Six muscles surrounding the eye precisely control its positioning and movement via three cranial nerves.

Thalamus. The thalamus is a complex relay station located superiorally to the brain stem; it serves as the gateway for almost all sensory input to the cortex. The lateral geniculate nucleus in the thalamus receives visual input from the eyes and sends it on to the visual areas in the occipital lobe of the brain. The lateral geniculate is where the initial sorting of visual input takes place, including early differentiation of signals related to color and form (to be processed by the "what" pathway) from those related to location and motion (to be processed by the "where" pathway).

Superior Colliculus. The superior colliculus (SC) lies at the superior and posterior aspects of the brain stem. It is divided into two layers, superior and inferior. The superior layer of the SC receives and processes sensory input and passes information to the inferior layer. This system provides a mechanism for quick-and-dirty sensory processing for tasks or situations where eye movements need to be initiated very quickly. The inferior layer of the SC sends eye movement commands down to oculomotor nuclei in the brain stem. It is arranged in a coarse representation of the visual field that is used to control the path of eye movements [70]. It integrates input from the superior layer and from several cortical regions. When an eye movement is initiated, a burst of activity occurs in the inferior SC at the represented vector difference between the *desired* and current positions, and as the eyes are moved, a wave of activity proceeds back along that vector to the "fixation zone" representing the origin of the coordinate system [44]. In the fixation zone, sustained activity is resumed to help hold the eyes fixed in place between saccades.

V1, V3, V4 and MT/V5. The visual system in the occipital lobe is partitioned into a number of distinct areas, typically labeled with a V followed by a number. V1 represents the most basic and initial area for cortical processing of visual input. It contains a detailed representation of what is currently in the visual field. This information is then passed through more specialized areas that decompose and recombine the input to isolate particular features. For instance, V3 and V4 appear to process color and form, while V5 (a.k.a. Area MT) is primarily sensitive to motion. V5 is particularly involved in the visual pursuit eye movements, providing information about the direction and velocity of moving targets [64].

Frontal and Supplementary Eye Fields. The frontal eye fields (FEF) and supplementary eye fields (SEF) are the output regions in the cerebral cortex for the generation of eye movements [8, 53]. The frontal eye fields are anterior to the primary motor area, while the supplementary eye fields are anterior to the supplementary motor area [58]. (The primary and supplementary motor areas control other aspects of voluntary body movement, such as the hands and legs.) Currently, there is little evidence to distinguish the function of the SEF from the FEF. There has been some suggestion in the recent literature that the SEF is more heavily involved with handling *sequences* of eye movements, which involve different logistics and planning [25]. Differentiation between the two areas nevertheless still remains an open question.

Posterior Parietal Cortex. The Posterior Parietal Cortex (PPC), particularly what is called the Intraparietal Sulcus, lies along the "where" pathway and is responsible for re-mapping the internal coordinate system as the eyes move. The computational burden here is rather stunning: whenever the eyes move, all of the neurons sensitive to particular locations in the visual field must re-map their sensitivity to the new part of the image. The brain appears to maintain multiple coordinate atlases simultaneously and the re-mapping may be made feasible by encoding coordinate transformations that are more easily changed [13].

Dorsilateral Prefrontal Cortex. This area (DLPFC) on the middle frontal gyrus is associated with maintaining spatial information over time to guide movements that are expected to be initiated in the near future (on the order of seconds) [57]. It is also believed to provide contextual information that modulates the commands of the eye movement system.

These are the main players in the oculomotor system, but understanding the functioning of these individual areas is less than half the story. Figure 3 also labels the principal *pathways* along which information flows during the preparation for eye movements. The system consists of two main pathways that process the input at different levels. The first and fastest—from the eyes to the lateral geniculate nucleus to the SC—involves the most basic visual processing and motor operations. This pathway is fast and approximate, using only the most salient features of the input and ignoring many aspects of the environment and their contextual relevance in order to get the fastest possible response to unpredictable visual targets. Think of this as generating the command that shifts attention to the nearby small buzzing object (it's a wasp but we don't know that yet).

The second pathway allows for the cognitive modulation of eye movements. Here, some information from the lateral geniculate is passed to V1; from there it is distributed to V3, V4, MT and then to the intraparietal sulcus and temporal lobe. This information is combined and decomposed in a complex way, enabling recognition of objects and other important associative features of the scene. The output is passed to the frontal eye fields and to the prefrontal cortex. The prefrontal cortex integrates sensory input (e.g., what and where) with information about the current context and feeds higher-order commands to the frontal eye fields to modulate the incipient eye movement command. Both the eye fields and the intraparietal sulcus feed into the inferior SC and directly to the lower brain stem. This pathway integrates a wider variety of information about the environment and current context than subcortical pathways in order to guide adaptive decisions regarding the generation or suppression of eye movements. Think of this as generating the command that, realizing that we are falling from a branch of a tree, decides to ignore the wasp and make an eye movement to the object in the visual field that may be (and is, phew!) a nearby branch with which we can save ourself.

These pathways are quite distinct, even though they involve several common areas. The inferior layer of the SC and other parts of the brain stem integrate the output of these pathways, generating or suppressing the eye movement using the most accurate coordinates available under the time constraints on action. The intraparietal sulcus, prefrontal cortex, and frontal and supplementary eye fields relay information both to the SC and directly into the lower brain stem.

While this description of the oculomotor pathways is necessarily simplified, it does reveal some interesting principles underlying the design of the brain. First, there are multiple partially redundant pathways and connections arranged hierarchically. The redundancy provides robustness against gross loss of vital function in the event of focal trauma or disease. At the same time, the incorporation of distinct functions subserved by similar regions allows the system to function more efficiently than if the components were purely redundant. Thus the first pathway provides a way to make a crude but fast eye movement in the absence of other information or if there is need for quickness, while the other pathways do the same job but in a broader environmental context.

Second, as processing proceeds to higher levels, the neural circuits become more general. For example, the part of the visual field to which a given neuron is sensitive (called its receptive field) is small for low-level processing (e.g., the first node in the cortical pathway) but gets larger as the processing becomes more abstract (e.g., later nodes in the pathway). Similarly, the system integrates more and more disparate kinds of input in higher cortical functions. Only basic motor operations and crude sensory information about target location are involved in the first pathway, but later, motion, form, and color information are extracted, then object recognition occurs, then input from other senses is integrated so we can look to where we touched or heard a sound, then memory and other context information are brought to bear progressively to prioritize and plan possible responses.

Third, attention modulates the functioning of the pathways. In our example, if we are falling off the tree branch in the wind and attending to the need to regain control, the fast moving, loud buzzing, possibly pain inducing wasp does not induce an eye movement. The command is generated by the first pathway, but the context of the situation, the relative costs and benefits of the possible actions (having been learned by association with long experience) in some sense determine the focus of our attention. If we are attending to averting a life-threatening fall, our system tends to down-weight the importance of the wasp, and the eye movement to look toward it, as it settles on an arm, may not be made.

Finally, neural computation appears to proceed by successive and repeated splitting and merging of the input in a highly nonlinear way. Each step in the process sorts the data on different features and integrates other features; the result is passed on to other areas where it will be further sorted and integrated.

2.2. Why Study Eye Movements?

2.2.1. Basic Neuroscience

One of the basic challenges of neuroscience is to determine how the various functions of the brain are implemented. This requires learning not only about the conditions under which individual neurons fire but also about the interactions among different neural circuits. The complexity of the brain makes it necessary, however, to focus this effort as much as possible, for instance on a particular functional system. The more specialized a system, the easier it is to understand its implementation, but the more integrated a system is with other functions, the greater the insight to be gained. The oculomotor system offers a good balance of these features; it has a specific and measurable purpose and numerous distributed connections into other major systems in the brain.

The study of such a neural system can proceed at four different levels:

1. Localization, in which the anatomical location of the system components is identified,
2. Characterization, in which the basic function of each component is determined,
3. Differentiation, in which related but distinct functions within the system are distinguished, and
4. Integration, in which the interactions among the system components that produce the final behavioral responses are understood.

Although questions at all four levels can be addressed concurrently, this ordering tends to represent the path (from the first to the last) along which the science progresses. Note that the second and third levels here are strongly related, but we distinguish them because they represent the underlying steps in an iterative process of identifying functionality. It is also worth noting that all the various tools used to study brain systems—from single-unit recording to fMRI—also tend to develop along the same path. Part of the goal of this paper is to show how questions at the higher levels can begin to be addressed. For the oculomotor system, interesting scientific questions arise at all four levels.

Localization. Most of what is known about the organization of the oculomotor system was obtained using non-human primate models [71]. With single-unit recording, the behavior of neurons at a very specific location can be monitored while a monkey performs a variety of eye movement tasks. In this way, the functions of the areas in Figures 2 and 3 have been slowly revealed by showing that a high percentage of cells in a given region exhibit significant alteration in their firing rates in response to specific eye movement tasks. An important question remains, however: How closely related are the organizations of the oculomotor system in humans and non-human primates? Current evidence suggests that the homology is quite good, but even beyond gross morphological differences, there are several examples in which the correspondence breaks down [39] (see Section 2.3). Another issue for mapping out the locations of components in the system is variability across individuals in the anatomical geometry of the cerebral cortex. An important task is to assess the magnitude and nature of these variations.

Characterization. Consider what is called the visually-guided saccade task: a subject holds her gaze fixed on a marked point in the center of the visual field, and when a point of light appears at a random location in the visual field, she makes a saccade to the location of the light. This is a simple example of an eye movement task that is used to map out the oculomotor system. Increased neuronal activity close to the time of the saccadic response identifies neurons that play a role in generating the saccade. Different neurons will show sensitivity to the task for saccades at specific direction or distance. Although this is potent information, the story becomes more complicated when other tasks are included. It was long thought that brain areas showed tight specialization in their function. For example, the frontal eye fields are involved in voluntary motion of the eyes, an explicitly motor operation. But more recently, it has become clear that there is substantial overlap in the functions carried out by different components in a system. For example, the frontal eye fields also contain cells involved explicitly in cognitive and sensory operations, even though other parts of the oculomotor network seem to carry out these functions as well [58, 55]. How and why this sharing takes place remains to be understood, but it does imply that great care is needed in characterizing functionality, even in low-level sensory-motor systems.

Differentiation. There appears to be a hierarchical organization to functionality in brain systems: at one level, an area will appear to be quite specialized to a particular type of processing, but as we look deeper, we see systematic subdivisions into quite distinct functionality. For example, in the macaque monkey, pursuit and saccade areas within the frontal eye fields are clearly separated [27, 62]. In humans, the separation of these systems in the frontal eye fields is not yet clear [5, 47].

Integration. Having identified the nodes of a functional network like the oculomotor system, the next critical step will be to assess how those nodes interact. How does information flow through the system? How are the functions of the different areas coordinated? How do they share seemingly redundant functionality? What is the temporal sequence of processing? What is the impact on the system if parts of the network are impaired? These questions relate to the functioning of the system as a whole, to the emergence of observable behaviors from distinct functional units. For example, there are two competing views regarding the cognitive control by the prefrontal cortex of the oculomotor and other systems. The first is the "executive" model in which the prefrontal cortex is the central regulator of high-level processing, a sophisticated control system sending commands that coordinate the activity of other systems throughout the brain. The second is the "distributed processing" model in which the coordination is automatic, built into the system through complex pathways and interconnections with higher functions arising as an emergent property of the system integration. In the distributed processing model, the prefrontal cortex is important as a supply of cognitive resources that improves the achievable complexity of the operations but does not control other areas. Both of these views seem reasonable, but they make very different predictions about high-level cognitive processing. This question is concerned completely with the integration of distinct components and with the basic issue of how high-order cognition arises from the brain's architecture. The oculomotor system is an ideal testing ground for studying the various theories, and effective statistical methods are crucial for being able to distinguish fine structure in the brain processes.

2.2.2. Clinical Issues

Eye movement abnormalities are a common neurobehavioral associate for a variety of psychiatric and neurological disorders [37]. For some disorders, the affected patient exhibits such abnormalities, allowing identification of the disturbed neural pathways and enabling physicians to monitor treatment response. In other cases, biological relatives of an individual affected with an illness exhibit the eye-movement abnormalities without expressing the disorder, indicating that the abnormality may be a familial or genetic marker for a predisposition to develop the illness later in life [31]. An improved understanding of the oculomotor system thus promises to yield insights into the changes at the root of these conditions, to help identify risk factors for these disorders, and perhaps to facilitate new diagnostic techniques. Most importantly, oculomotor performance can reflect cortical integrity, enabling the identification of abnormalities in specific cortical systems.

A neurological disorder involves a disturbance in nervous system function. A critical aspect of understanding disorders of the central nervous system is to identify the nature of these disturbances and their precise effect on the brain. While many insights can be gained using animal models or examining postmortem tissue, it would be most desirable to study the functional effects of disturbances in brain tissue as they occur in patients afflicted with a given illness. Functional neuroimaging has made this possible. One basic approach is to study the impact of the disorder on a selected system, and the greatest gains are likely to be realized when the chosen system satisfies the following:

1. Allows accurate measurement of the system's inputs and outputs,
2. Operates at all levels of processing, from basic sensory-motor to abstract cognitive domains,
3. Involves regions that are distributed throughout the brain,
4. Can be studied with relatively simple tasks that can be performed by both patients and healthy individuals over a wide age range.

The oculomotor system has all of these properties. The timing of input stimuli can be precisely controlled and the system output (the eye movement itself) can be precisely measured, both in time and location. Because the visual input is controlled and the motor response readily observable, the dynamics of the task can be monitored. As we have described, the eye movement system integrates all levels of processing in a widely distributed network. Moreover, there is a range of eye movement tasks that exercise the system in different ways (see Section 5), tasks that can be learned and performed by non-human primates, children, and elderly patients. In contrast, complex cognitive tasks can be difficult for patients to grasp and perform and can induce subject movement in the imaging scanner. Eye movement studies also often provide direct linkages to animal models.

The oculomotor system is therefore a powerful tool for studying a number of clinical disorders, and indeed, the study of eye movements has direct implications for many notable psychiatric and neurological diseases, including the following.

Schizophrenia. This severely disabling illness often causes lifelong impairments in occupational and social adjustment. It typically begins in late adolescence or early adulthood and is marked by hallucinations, delusions, thought disorder and losses in emotional responsivity and social interest. It arises in varied forms and resists a tight classification. However, one robust finding in schizophrenia research has been that a high percentage of affected patients and many (perhaps 40%) of their close biological relatives exhibit two eye-movement abnormalities [37, 10]. The first of these, called low-gain pursuit, is a difficulty matching the velocity of the eyes to that of a moving target during visual pursuit, with the pursuit velocity being too slow. The second abnormality is the presence of anticipatory saccades during pursuit where the subject abruptly makes saccadic movements ahead of the target and then must wait for the target to catch up. Anticipatory saccades are less commonly observed than low-gain pursuit but may be more specific to schizophrenia [56, 51].

The presence of these abnormalities offers two distinct paths to better understanding schizophrenia. First, the abnormalities reflect both differences and interactions among the sub-systems that control pursuit and saccades. The low-gain pursuit abnormality arises in a task like the following: a dot to be tracked begins on the left side of a screen then moves right with a constant velocity. In low-gain pursuit, the eyes move to track the target at a lower angular velocity than that at which the target is moving. Anticipatory saccades occur during simple pursuit, in which the subject pursues an object moving slowly back and forth across a screen. The anticipatory saccades appear to be a predictive jump based on the motion of the object, to a point where it is expected rather than to where it is located. This suggests that cognitive information (predictions) from the frontal lobe may be drowning out the motion information from MT. Both of these effects are consistent with the idea that certain pathways to the pursuit system are being disturbed by the disorder. If we can characterize these sub-systems and distinguish their implementation, we may be able to identify some of the specific focal brain abnormalities associated with the disorder.

The second path to understanding schizophrenia is that these eye-movement abnormalities serve to identify families prone to schizophrenia and thus provide more generally homogeneous populations to study. Comparisons within these families can yield insight into how inherited factors influence brain physiology and also identify risk factors that suggest who in the family needs preventative intervention. The genetic basis of schizophrenia is not nearly understood, but the transmission of eye-movement abnormalities to offspring gives information about the genetic basis of the disorder.

Parkinson's Disease. This disease, which typically appears in patients older than 50 years, disrupts the brain's motor systems, especially the voluntary initiation of movement. Reflexive movements are generally unaffected. In the oculomotor system, the disease interferes with the pathway from the frontal eye fields to the basal ganglia to the superior colliculus (see Figure 3). This results in inaccurate voluntary saccades [6, 38]. Saccadic accuracy in a variety of eye-movement tasks provides an effective measure for monitoring treatment response (e.g., to L-DOPA) of Parkinson's patients.

Childhood Autism. Autism, or Pervasive Developmental Disorder, is a complex neurologic disorder associated with abnormalities in language skills, social interaction, and higher order cognitive processes. Its onset occurs in the earliest years of life. An ongoing debate about the pathophysiology of this disorder surrounds the question of whether autism arises from pathology in the cerebral cortex, primarily in frontal lobe systems, or in the cerebellum, particularly the posterior lobules of the cerebellar vermis (Lobules VI and VII) [41]. Structural neuroimaging studies are not conclusive on this point, in part because they have been confounded by including patients with a number of complicating factors such as fetal rubella and mental retardation that on their own can cause neurological deficits. Little data exist at present directly linking functional deficits with focal anatomic disturbances in this illness.

Studies of eye movements provide a valuable opportunity to test the competing models of frontal and cerebellar origins of autism, because these brain regions are known to play different specific roles in the regulation of saccadic eye movement activity. The control of dynamic aspects of saccades, particularly their accuracy, is known to be regulated by the specific regions in cerebellar vermis (posterior lobules VI and VII) which are believed to be abnormal in autism. The frontal lobe regions believed to be compromised in autism would be expected to be associated not with a disturbance in basic saccade dynamics, but rather with abnormalities in the higher-order voluntarily control of saccades, such as is required to suppress saccades to compelling visual targets and for making saccades to remembered locations without sensory guidance. While our behavioral studies suggest that disturbances in autism are specific to the higher order control of saccades [42], fMRI studies are needed to directly document and examine the nature of information processing disturbances in cortical areas of interest, and to link functional and anatomic disturbances to establish clinico-pathologic relationships.

Alzheimer's Disease. Alzheimer's disease is the most common dementing condition of late life. It typically begins with memory disturbance, but all cognitive functions become compromised later in the course of the illness. The degeneration of cortical neurons appears to lead to progressively severe inefficiencies in computational activity that are the likely cause of the cognitive manifestations of the disorder. One way to conceptualize such a progression is through a "left-shift" in the relationship between brain activation and task difficulty. In this view, the peak sustained activation would occur during less difficult versions of tasks in Alzheimer's patients than in normal individuals because they need a maximal response from affected regions to perform even simpler cognitive tasks. A second testable prediction is that brain regions (such as the anterior cingulate cortex) in which activation generally parallels increases in task difficulty especially those imposed on prefrontal cortex, would exhibit activation during processing of even simpler tasks under which activation would not be detectable in healthy control subjects. Our preliminary data indicate that activation in the anterior cingulate is robust in Alzheimer's patients performing even a simple visually guided saccade task, while this is not seen in healthy individuals.

Stroke. A stroke, or cerebral vascular accident, is a clinical event in which a reduction of blood flow (e.g., a clot in an artery) disrupts and potentially kills brain neurons because of several factors, including a loss of oxygen needed for cellular metabolism. As the brain recovers from a stroke that permanently destroyed a specific area of the brain, it demonstrates a stunning level of plasticity in its reorganization. This is believed to occur both as areas in the affected hemisphere take over functions that they did not previously serve, and as the unaffected hemisphere takes over functions previously served by the other half of the brain. Understanding the recovery process is important for learning about the plasticity of brain reorganization: to what degree it is reorganizing function and how this reorganization varies with demographic covariates such as gender and age. On the clinical side, this understanding may lead to the development of new treatments to improve recovery from stroke and other localized lesions. Neuroimaging studies of brain reorganization could prove invaluable in developing and validating such novel treatments since they provide a window for directly monitoring physiologic reorganization. Eye movement activation tasks performed in an fMRI environment provide an effective way to track recovery and characterize functional re-mapping because they elicit multiple robust and widely-distributed patterns of activation throughout the brain. Our preliminary data provide striking examples of brain reorganization in the oculomotor system during the months after stroke.

Brain Development and Developmental Disorders. With respect to brain anatomy, one of the most dramatic evolutionary developments in higher primates is the growth of the frontal lobes. The frontal lobes, thought to be critically involved in higher cognitive functions, are especially pronounced in humans, comprising a much higher percentage of brain tissue than in other primates. Recent data indicate that neurobiological development in the frontal cortex continues well into adolescence [50], a pattern that closely parallels the age-related maturation of higher executive cognitive processes thought to be subserved by this brain region. Single-cell recording studies of non-human primates yield sharper results about how neurons in prefrontal cortex control higher-order eye movements than perhaps any other function. Hence, the study of normal developmental changes in prefrontal activation as individuals perform voluntary saccade tasks provides an especially promising tool to discern the interrelationship between brain maturation and cognitive development. Functional MRI is particularly well-suited to such a study because the injection of radioactive tracers may pose an unacceptable health risk for studies of healthy children and adolescents. This approach will address several fundamental questions in cognitive neuroscience and clinical neuropsychiatry. For example, a number of neuropsychiatric illnesses, such as schizophrenia and obsessive compulsive disorder, are believed to result from abnormalities in later development. Functional MRI studies on children and adolescents whose parents are affected with such disorders provide a novel way to find associations between abnormal developmental processes and the various clinical conditions.

2.3. Specific Objectives

Given the importance of pursuit and saccades as the two primary eye-movement mechanisms, a basic question of interest about the oculomotor system is as follows: How do the systems that manage saccadic eye movements and visual pursuit of a target relate to each other functionally? In this paper, we begin to address this question by attempting to differentiate pursuit and saccades, with an eye towards differences across primate species:

> **Target Question.** Are saccades and pursuit implemented by separate sub-systems in both humans and non-human primates?

One of the key distinctions between humans and macaque monkeys is the anatomical classification of the brain areas that subserve pursuit and saccades. In the monkey, the pursuit and saccade areas span the boundary of the prefrontal cortex (one lies in what is called Broadman's area 6 and the other in area 8). On the other hand, both pursuit and saccades appear to be implemented outside the prefrontal cortex in humans (both in Broadman's area 6) [5, 47]. Differentiating the pursuit and saccade areas would provide important information regarding the neurophysiology of the cerebral cortex in humans with some clues to the developmental differences between the two species. Specifically, some cellular properties of prefrontal cortex are different than the rest of the cortex (in cortical layer four, the prefrontal cortex is called "granular" because of the types of synaptic and neural elements it contains). Prefrontal cortex also receives inputs from different nuclei in the thalamus than does the rest of the frontal cortex. If pursuit and saccades are differentiated in humans as they are in the monkey, it would suggest that there was an evolutionary "reallocation" of the frontal eye fields in the human, in which eye movement systems were moved. If the two functions are not differentiated, it will provide a significant example in which functional homology between monkeys and humans fails.

3. Overview of fMRI

An MR scanner is a tube containing a very strong, uniform magnetic field along with some mechanism for manipulating the magnetic field within the tube. At its most basic level, the scanner serves a tool for measuring the density of a single nuclear species (e.g., hydrogen) within a volume of tissue. It is possible to isolate a single nuclear species because of the phenomenon called Nuclear Magnetic Resonance (NMR). By carefully modulating the magnetic field, it is possible to encode spatial information in the measured NMR signal so that an MR image can be constructed, showing the nuclear density in an array of smaller volumes. Functional MRI arises from the sensitivity of these measurements to a blood-flow response in the brain that is related to brain activity. In this section, we describe all of these phenomena to clarify the basic aspects of fMRI data acquisition and to indicate what fMRI can say about brain function. We start at the sub-atomic level.

3.1. Principles of Nuclear Magnetic Resonance

Both protons and neutrons, the basic constituents of atomic nuclei, act and interact like small bar magnets. The strength of their magnetic field is described by a vector quantity called a *magnetic moment*. When particles are grouped together, as in an atomic nucleus, their magnetic moments combine by superposition, and the resulting aggregate acts like a larger bar magnet. Protons and neutrons in an atomic nucleus tend to pair with like particles in such a way that the net magnetic moment of the pair is effectively zero. Thus, only nuclei with an odd number of protons or neutrons will have a non-trivial magnetic moment. Such nuclei can thus be influenced by external magnetic fields.

When a nucleus with a non-zero magnetic moment is placed in a strong, *uniform* magnetic field, the nucleus—acting like a small magnet in a much stronger field—tries to align its magnetic moment either parallel or anti-parallel to the magnetic field. It does so because these are the two orientations of lowest energy among all possibilities, with the parallel orientation slightly lower in energy than the anti-parallel orientation. However, the magnetic moment cannot align itself exactly in either direction because the nucleus has a non-zero angular momentum which must be conserved. Instead, the nuclear magnetic moment *precesses* about the field at some fixed frequency. This angular momentum and precession are the result of quantum mechanical effects rather than any physical rotation, but the conceptual model of a spinning top precessing about the vertical is still a useful way to visualize what is happening here. The precession frequency of the nucleus is determined by two factors: (i) the type of nucleus and (ii) the strength of the uniform field. Specifically, the precession frequency can be expressed as γB_0 where B_0 is the strength of the magnetic field and γ, the gyromagnetic ratio, is a constant characteristic of the nuclear species (e.g., hydrogen, sodium, etc.).

The superconducting magnets typically used for imaging humans have fields in the range of 1.5 to 4 Tesla (T), on the order of one million times the strength of the Earth's magnetic field at the surface. Among the elements whose nuclei contain an odd number of protons or neutrons, there are several that occur naturally in living tissue in sufficient quantity to produce a non-negligible signal: Hydrogen-1, Carbon-13, Sodium-23, and Phosophorus-31. Hydrogen, with a single proton and no neutrons, has the strongest magnetic moment among these and, being a key component of both water and lipids, is far more abundant in the human body than any of the other elements. Consequently, the NMR signal obtainable from hydrogen is about 1000 times stronger than that from any other element, and hydrogen is thus the species to which images are most commonly calibrated. Hydrogen (^1H) has a gyromagnetic ratio of approximately 43 MHz/T, so that its precession frequency lies in the radio band of the electromagnetic spectrum.

Now we step back to consider a mass of material containing such nuclei. In many materials, particularly living tissue, there is no preferred orientation, so the orientations of the nuclear magnetic moments are distributed approximately uniformly on the sphere. The superposition of all these moments consequently sums to essentially zero, and the material shows no magnetic properties. This changes when the material is placed in a strong, uniform magnetic field. The magnetic moments of the nuclei within the material align themselves and precess as described above, clustering about the field lines. Since the parallel orientation is a state of slightly lower energy than the anti-parallel orientation, just over half of these nuclei align with the field, and the net magnetization of the material becomes non-zero, i.e., the material acts like a magnet.

At this point, the nuclear magnetic moments (of a given species such as hydrogen) within the field are all precessing about the uniform field with the *same frequency*, but the phases of this precession are distributed uniformly about the circle. The mass of material thus has a constant magnetic moment. If we could get the nuclear magnetic moments to precess *in phase* with each other, the net magnetic moment of the material itself would then also precess with that frequency. Why would we want that? A varying magnetic field induces a current through any nearby loop of wire. If the mass of material had a precessing magnetic moment, the amplitude of the current induced in a nearby receiver coil—the NMR signal—would be proportional to the density of nuclei (of a given species) within the material.

This is where resonance comes into play. To make the nuclei precess in phase, we expose the material to a radio-frequency (RF) pulse of the exact frequency at which the nuclei are precessing. This pulse *excites* the nuclei by "tipping" their magnetic moments away from the main field; whereas they had been precessing in a small circle about the field, they are now precessing in a large circle about the field and *in phase*. Only those nuclei precessing at the right frequency are affected by the tipping pulse; this is resonance. Moreover, the effect of the tipping decays exponentially with time, in two ways. The first results from the gradual realignment of the tipped magnetic moments with the strong, uniform magnetic field. The second results from a gradual loss of phase coherence, where local magnetic irregularities move the precessing magnetic moments out of phase with each other and reduce the measurable NMR signal. This second source of decay tends to be faster than the first and plays an important role in fMRI, as described below.

3.2. Magnetic Resonance Imaging

The main magnet in the MR scanner is designed and adjusted to generate a magnetic field that is as uniform as possible within the imaging volume. As such, the NMR phenomenon allows an MR scanner to measure the total density of a nuclear species such as hydrogen within a volume of tissue. In order to acquire detailed spatial images of how this density varies in the tissue, however, it is necessary to carefully encode spatial information into the field of precessing nuclei. There are countless variants and details of acquisition that apply in practice, but the following is a simplified conceptual description of MR images acquisition, using hydrogen imaging as an example. (See also [7].)

Introduce a Cartesian coordinate system (x, y, z) where, for definiteness, we take the z-axis to be aligned with the axis of the subject's body. We will consider acquiring an M-slice image, where each slice is an $N \times N$ array of contiguous volume elements, or *voxels*. For the purposes of this discussion, we will take $N = 4$ and $M = 1$. (In typical fMRI studies, N is either 64, 128, or 256; and M ranges from 1 to 14. Typical voxels have sides from 1mm to 3mm in plane with slice thickness ranging from 3mm to 5mm.) Let $\rho_{jk} \geq 0$ for $j = 0, \ldots, N-1$ and $k = 0, \ldots, N-1$ denote the densities of hydrogen nuclei in our 16 voxels, with j varying in the x direction and k varying in the y direction. The task is to use the NMR phenomenon to extract the ρ_{jk}'s. Spatial information is encoded into the magnetic field by introducing gradients in the magnitude, but not the direction, of the main magnetic field. There are three basic parts of the acquisition sequence: (a) slice selection, (b) phase encoding, and (c) frequency encoding.

Although it is possible to obtain fully 3-dimensional images, most current acquisition methods acquire each image as a sequence of distinct, and possibly separated, slices. Let ν_0 be the known precession frequency in the main magnetic field for the target nuclear species (e.g., hydrogen). Because the RF pulse excites only the nuclei that are precessing at frequency ν_0 and because the precession frequency of a nucleus depends linearly on magnetic field strength, we can control which nuclei get excited by adjusting the gradients. Slice selection is accomplished by applying a constant gradient to the magnitude of the main magnetic field in the z-direction, which makes the magnetic field strength vary linearly as a function of z coordinate but remain constant in x and y for fixed z. The gradient causes only those hydrogen nuclei in a target slice of tissue to be precessing at frequency ν_0; those nuclei outside the slice are precessing faster or slower because the magnetic field is larger or smaller there. The selected slice lies in the x-y plane. By changing the gradient, we can select which slice of tissue gets excited by the pulse and how thick that slice is. After the slice selection gradient is applied, the RF excitation pulse is delivered, and because of the NMR phenomenon, only the hydrogen nuclei in the target slice are excited. The slice-selection gradient is then turned off, but the excited nuclei continue to precess in phase at frequency ν_0, generating a time-varying magnetic field. (In practice, there is decay and de-phasing as described earlier, but we ignore that here.)

Figure 5a shows a snapshot of our 16 voxel image just after the slice selection gradient is turned off. The arrows show the in-slice component of net magnetization for each voxel. These are all precessing in phase (indicated by the arrows pointing in the same direction) and at the same frequency (indicated by $\Delta\nu = 0$ within the cells). The length of each arrow is proportional to the hydrogen density ρ_{jk} within the corresponding voxel, and the net magnetization of the entire slice is the vector sum of these individual magnetizations. Since they are in phase and at the same frequency, this time-varying magnetic field will induce a current in a nearby coil of wire which varies as $e^{2\pi i\nu_0 t}\sum_{j,k}\rho_{jk}$, the amplitude of which is the total hydrogen density within the selected slice. This signal is not yet recorded.

Instead, we apply a constant gradient G_{PE} to the magnitude of the main magnetic field in the y-direction for some specified amount of time τ_{PE}. This is called *phase encoding*; to see why, consider Figure 5b which shows a snapshot of our 16 voxel image just before the time τ_{PE} has elapsed. While the phase encoding gradient is being applied, voxels located at higher y positions experience larger magnetic fields, and consequently, the nuclei in those voxels precess at a higher frequency. When the phase encoding gradient is turned off, all of the arrows return to the same precession frequency, but the arrows are now out of phase with each other. All voxels at the same y-coordinate are in phase, but voxels at different y coordinates are out of phase. Again, no signal is yet recorded, but if we were to examine the induced current in a nearby coil of wire, it would vary as $e^{2\pi i\nu_0 t} \sum_{j,k} \rho_{jk} e^{2\pi i k\phi}$, where ϕ is a phase shift that depends on τ_{PE}, G_{PE} and the voxel size. In our example, $\phi_1 = 1/4$.

Next, we apply a constant gradient G_{FE} to the magnitude of the main magnetic field in the x-direction for some specified amount of time τ_{FE}. This is called *frequency encoding.* Figure 5c shows a snapshot of our 16 voxel image during this period. The effect of the gradient is to alter the precession frequencies of the voxel magnetizations; voxels at larger x value precess faster than those at smaller x values. Notice in the figure that all 16 voxels are now distinguished by phase and/or frequency. During frequency encoding, the NMR signal (i.e., the current in the coil of wire) is recorded. This signal is varies as $e^{2\pi i\nu_0 t} \sum_{j,k} \rho_{jk} e^{2\pi i(j\nu_1 t + k\phi)}$, where ν_1 is a frequency increment that depends on τ_{FE}, G_{FE}, and the voxel size. In our example, $\nu_1 = 10$MHz.

Thus after one pass of slice selection, phase encoding, and frequency encoding, we have recorded a complex linear combination of the quantities ρ_{jk} which we wish to recover. Note that we can completely describe this recorded signal by saving the values at times $t = 0, \ldots, N-1$. To completely reconstruct the ρ_{jk}'s, we need to execute the above process a total of N times, using phase shifts $s\phi$ for $s = 0, \ldots, N-1$. Figure 5d shows a snapshot of the image using a phase shift value of 3/4.

The resulting measurements are a matrix R_{st} of the form

$$R_{st} = e^{2\pi i\nu_0 t} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \rho_{jk}\, e^{2\pi i(j\nu_1 t + k\phi s)},$$

for $s = 0, \ldots, N-1$ and $t = 0, \ldots, N-1$. In words, after a demodulation at known frequency ν_0, the matrix R is just the Fourier Transform of the matrix ρ. To reconstruct ρ, we just apply the inverse transform to the demodulated measurements.

3.3. Functional Magnetic Resonance Imaging

3.3.1. The BOLD Effect

Among the first studies to use MRI to assess functional neural activity in humans are [4, 34, 45]. The latter two introduced the Blood Oxygenation Level Dependent (BOLD) effect for characterizing activation in the brain. Most fMRI experiments are based on the BOLD effect because it gives roughly an order of magnitude better sensitivity to task-related changes than do the available alternatives.

The BOLD effect owes its existence to two distinct phenomena, one chemical, one biological. First, the presence of atoms like iron with magnetic properties can speed the decay of the MR signal by de-phasing the precession of nearby excited nuclei. This is an irregularity, an artifact, where randomly scattered iron atoms exert a differential magnetic force on the excited nuclei, pulling them out of phase more quickly than in homogeneous tissue. The net result is that the local strength of the MR signal is reduced. Iron is an important component of hemoglobin, the molecule in blood responsible for carrying oxygen. When hemoglobin is oxygenated, the attached oxygen molecule shields the magnetic properties of the iron, eliminating its impact on the MR signal. In de-oxygenated hemoglobin, the iron is unshielded. Consequently, as predicted by Thulborn *et al.* [61], oxygenated blood gives a larger MR signal than de-oxygenated blood with the same hydrogen density. The second phenomenon underlying the BOLD effect is that the brain responds to concentrated activity with an in-flow of oxygenated blood to the active region. This *hemodynamic response* leads to an increase in the relative volume of oxygenated to de-oxygenated blood in the active region. The result is that a voxel tends to show a slightly (about 1%-5%) larger signal during periods of activity than during periods of inactivity. This is the BOLD effect.

3.3.2. fMRI Experiments

During an fMRI experiment, the subject performs a specific sequence of behavioral tasks in the MR scanner while images of the subject's brain are acquired at regular intervals. These tasks are designed to exercise specific behavioral processes and can range from active reasoning and cognition through passive attention to sensory stimuli. The observed pattern of responses in the brain offers insight into how the processes of interest are implemented. Each distinct task defines a single experimental *condition*. The experimental run is typically divided into blocks of time in which a single condition applies, with the corresponding task repeated as necessary to fill the block. These blocks of time are called task *epochs*, and within the experiment, there are multiple epochs nested within each condition distributed throughout the experimental run.

Figure 6 illustrates the design of a simple fMRI experiment, in which two conditions alternate in epochs of equal length. In one condition, labeled "Task", the subject performs a task that involves a specific behavioral process of interest. It also necessarily involves other process of lesser interest. Consider for example a simple task for studying eye movements: moving one's eyes from a mark at the center of the presentation screen to the location of a light flash. This requires attending to the mark at the center, recognizing the light flash, and moving the eyes to the desired location, only the last two of which are of direct interest. In order to isolate these processes of interest, our simple design also includes a condition, labeled "Control", that (ideally) involves all of the processes in the "Task" condition except those explicitly being studied. Hence, in the example just described, the control condition may involve fixating one's gaze at the mark in the center of the screen. Any observed differences in the brain's response to the two conditions may be attributable to the isolated processes of interest. This approach to fMRI design requires that some strong assumptions be met regarding the set of processes involved in the two conditions; nevertheless, it is pervasive in the field and generalizes to much more complicated designs.

The data from an fMRI experiment is a time-series of three-dimensional images, where each image is composed of an array of volume elements, or voxels, ranging in volume from roughly 3 to 30 mm^3. Figure 7 shows a schematic of an fMRI data set. There is a trade-off in image acquisition between spatial and temporal resolution, where the cost of finer spatial resolution is longer acquisition time. Figure 8 shows two image slices, one with very fine resolution ($\sim 4\text{mm}^3$) voxels but acquired slowly and the other with coarser ($\sim 27\text{mm}^3$) voxels but acquired rapidly.

What makes *functional* MRI distinct from structural MRI is that the primary interest centers not on the images themselves, but on subtle changes in the measured signal across time, changes caused by physiological effects (i.e., the hemodynamic response) related to neural activity. If we consider the measurements associated with a particular voxel, we obtain a time-series describing the density of hydrogen (or some other nuclear species) within the volume of the voxel. See Figures 9 and 10. However, these density measurements are distorted by a variety of magnetic effects, including the BOLD effect discussed above. Consequently, the hemodynamic response to a period of neural activity perturbs the measured hydrogen density in a systematic way. Figure 11 shows a voxel time series exhibiting a large BOLD perturbation that is correlated with performance of the "Task" condition. The scientific challenge in fMRI is to determine from such task-related changes how the brain subserves the processes under study. The statistical challenge in fMRI is to infer the pattern of task-related changes, if any, from the voxel time-series.

3.3.3. Statistical Techniques

If the neurons within a voxel respond strongly to one condition but not another, then the resulting hemodynamic response in the brain will lead to a detectable perturbation in that voxel's time series through the BOLD effect, cf. Figure 11. This suggests a statistical approach for identifying such "active" voxels: Compare the average signal during each condition, and declare a voxel "active" if the difference between these averages is sufficiently large. Most current methods for the statistical analysis of fMRI data generalize this idea and are based on classification of voxels, typically via statistical hypothesis tests. For example, the two-sample t-test is perhaps the most commonly used method to compare the average signal in two conditions. Others in current use include correlations with a fixed reference function (e.g., a lagged sinusoid) [2], non-parametric tests (e.g., Kolmogorov-Smirnov) [1, 36], spectral analysis of periodic designs [68, 22], split-sample t-tests [21], cluster-size modulated t-tests [19], and F-tests based on the general linear model [11, 12, 23, 69]. More general classification methods have also been brought to bear, including principal components, cluster analysis, and neural networks [66, 48]. The implicit goal underlying all of these methods is to find the location of the activity associated with the processes of interest.

While these classification approaches can provide reasonable identifications of active voxels, they suffer from two basic limitations: statistical limitations in the underlying models and inferential limitations in the scope of questions that can be addressed. First, the assumptions underlying many of the testing and classification methods are typically simplistic regarding both the noise process and the structure of the activity-induced signal. With respect to the latter, the hemodynamic-response that gives rise to the BOLD perturbation does not have a simple, uniform structure. Variations in the shape of this response can be highly informative and can significantly affect the quality of fit. Nonetheless, most current methods based on hypothesis tests assume a fixed, and often simplistic, response shape, while the available nonparametric methods offer no way to systematically constrain the response shape. See [26, 35, 16] for further discussion of these points. With respect to modeling the noise process, fMRI data is subject to diverse sources of variation, but it can be very difficult to account for these complex features of the noise. For example, recent work [67] suggests that non-local, anisotropic correlations between voxels can occur. Typically, fMRI data is passed through several stages of pre-processing before analysis to "correct" for nuisance sources of variation one at a time. The distortion to the data is estimated and residuals are passed to the next stage. This serial estimation detracts from the quality of fit, and because the uncertainties at each stage are not propagated, it makes the final assessment of uncertainty in the inferences optimistic.

The second and much more important limitation of classification-based approaches is that they can primarily address a single basic question: where did the activity occur? While finding the location of activity is an important part of an analysis, it is for many scientific applications of fMRI only a first step. For cognitive psychologists, the primary interest centers on theories that describe the integrated function of the brain, and their goal is to test the predictions of these theories. This requires the ability to use fMRI data to answer questions about relationships among the task responses both within voxel and across regions of the brain. Examples of such questions are described in detail in [26]. Many of these cannot be addressed directly with classification methods. Investigators using classification methods must then rely on *ad hoc* analysis without an effect assessment of uncertainty.

Our approach is to move away from testing and classification towards estimates of meaningful parameters, to make inferences about the components of variation simultaneously, and to provide effective assessments of uncertainty in the results. We use a Bayesian hierarchical model that accounts for the detailed structure of the underlying processes; it is built on substantive information and is extendable as new information comes to light. Inferences under the model are based on the posterior distribution of functions on the parameter space. This offers both flexibility and interpretability and makes a wide range of scientific questions directly accessible to statistical methods.

4. Modeling fMRI Data

The data from an fMRI experiment result from a sequence of complex processes with biological and technological sources of variation intertwined. The goal of the analysis is to make inferences about the slight, activation-induced perturbations to the signal that enable scientists to gain insight into the intricacies of the brain. The statistical methods for analyzing fMRI data must enable a wide range of inferences and give appropriate assessments of uncertainty. To support inference with such complicated spatio-temporal data, a detailed and flexible model is needed. Three principal challenges arise in modeling these data. First, the underlying noise process is composed of several distinct sources of variation, many of which have an intricate structure. Subject movement during the experiment can cause drastic and irregular signal changes. Much of this movement is non-rigid and three-dimensional and thus is difficult to capture. The subject's physiological cycles—respiration and heart beat—cause large fluctuations in the voxel time series that may be nearly confounded with the task design. These fluctuations are enhanced near tissue boundaries, where much of the interesting activation takes place. Outliers are a persistent problem, and a variety of scanner instabilities and calibration errors can give rise to localized anomalies. The measured MR signal exhibits nonlinear and inhomogeneous drifts that can be large in magnitude relative to the background noise. These drifts can take diverse forms and are spatially irregular; for example, Figure 12 shows a sample of time courses with a variety of drift profiles.

The second principal challenge that arises in modeling fMRI data lies in capturing the structure of the hemodynamic response. The signal generated by the hemodynamic response in the brain arises from a complex mechanism that is only incompletely understood. The shape of these signal changes varies across the tissue and depends on the distribution of vasculature within a given voxel. Moreover, the magnitudes of the task-related changes are small relative to the noise level and vary over time even in response to the same task. The importance in capturing the shape of the hemodynamic response accurately is that it provides a rich source of information regarding the timing and structure of the underlying activation [60].

A third and much more difficult challenge is the complicated spatial structure across the brain in both the noise components and the task-related response. Taking advantage of these spatial relationships poses a number of statistical problems but promises to improve the precision of inferences markedly. This is a direction of development for our research, but for now we focus on the temporal aspect alone.

The model we use for fMRI data is described in detail in [26], but we review its structure here. Currently, the model takes each voxel to be independent both structurally and stochastically. We are consequently treating each voxel time series as an independent data set with its own parameter values. The model is designed to account for the different sources of variation that affect the voxel time series; the goal is to depict the underlying structure as accurately as possible while being extensible in light of new information.

4.1. Structure of the Model

Amidst the components of variation that make up the voxel time series, we wish to make inferences about the task-related signal changes. Because all these components can take diverse forms, we require substantive constraints to separate them. The model is consequently cast as a Bayesian hierarchical model, with substantive constraints reflected both in the parameterization and by prior distributions to be described below.

Let experimental conditions be indexed by $c = 1, \ldots, C$ and each condition be divided into a number of epochs indexed by $k = 1, \ldots, K_c$. We will use $c(t)$ and $k(t)$ to denote the condition and epoch being performed at time t. The model parameters are arranged in blocks where each block relates to a single source of variation in the voxel time series. Within a single level of the hierarchy, these blocks are conditionally independent, but they are intertwined in general.

At the top level of the hierarchy is the likelihood, which describes the gross structure of the data at a single voxel. The current model distinguishes four interpretable components: the baseline signal, the signal drift, the activation profile, and the noise.

$$Y(t) \ = \ \mu \ + \ d(t; \boldsymbol{\delta}) \ + \ a(t; \tilde{\gamma}_{c(t),k(t)}, \boldsymbol{\theta}, \mu) \ + \ \epsilon(t; \boldsymbol{\phi}) \tag{1}$$

The baseline signal, denoted by the real-valued parameter μ, is the mean level of the signal in the absence of activation and drift, the underlying measure of hydrogen density in the voxel. The drift profile $d(t)$ is modeled as the equivalent of a smoothing spline with respect to a fixed B-spline [14] or orthogonalized spline basis (where the latter is derived by orthogonalizing the power basis). For both bases, $d(t)$ is explicitly constructed to be orthogonal to constants with respect to the empirical inner product, i.e., $\sum_i d(t_i) = 0$ where the sum is over image acquisition times. The parameters in $\boldsymbol{\delta}$ represent the coefficients of the profile in the corresponding basis. The activation profile $a(t)$ represents the task-related signal change as a function of time; it is a parameterized family of smooth curves, with the parameters describing the magnitude of the response for each condition and the shape with which the response manifests itself in the signal. The precise structure of this curve depends explicitly on the known experimental design through the times at which task performance begins and ends. The parameter $\gamma_{c,k}$ denotes the response amplitude for epoch k within condition c as a proportion of baseline signal; the amplitudes for epochs within a condition c are centered on the average response amplitude γ_c, as described below. The parameter vector $\boldsymbol{\theta}$ describes the shape of the response function for a single epoch of task performance, where each component specifies one feature of the response function. Finally, $\epsilon(t)$ denotes the temporal noise process parameterized by the vector $\boldsymbol{\phi}$. In the model fits presented in this paper, $\boldsymbol{\phi} = (\sigma^2)$ for $\sigma^2 > 0$ and the ϵ's are taken as independent $\text{Normal}(0, \sigma^2)$. In more general versions of the model, this term can include autocorrelations, physiological fluctuations, and heavy-tailed components.

The additive combination of components in the likelihood would not be identified without strong constraints on the components. These constraints are not arbitrary but reflect the underlying structure as accurately as possible. For example, a great deal is known about the basic shape of the hemodynamic response function, and the signal drifts tend to be very smooth. The priors used in the model are based on substantive information—theory, experiment, and observation about the basic processes generating fMRI data—as described in [26].

The next level of the hierarchy describes variation in the response magnitude across epochs but within conditions. Specifically,

$$\widetilde{\gamma}_{c,k} \mid \widetilde{\boldsymbol{\gamma}}_{-(c,k)}, \boldsymbol{\gamma}, \ldots \sim \begin{cases} N_+(\gamma_c, \tau_0^2) & \text{if } \gamma_c > 0 \\ \text{point mass at } 0 & \text{o.w.,} \end{cases} \tag{2}$$

where N_+ is a normal distribution truncated at 0, $\boldsymbol{\gamma}$ is the vector of average response magnitudes for all conditions and $\widetilde{\boldsymbol{\gamma}}_{-(c,k)}$ refers to all of the response magnitude parameters except that associated with the specified epoch. All of these parameters are non-negative and can take on the value 0, and the variance hyperparameter τ_0^2 is fixed (e.g., 2.5×10^{-5}). This level of the hierarchy can be excluded by setting $\widetilde{\gamma}_{c,k} \equiv \gamma_c$ for all conditions and epochs; in this case, $\gamma_{c(t)}$ is used in the likelihood above.

The remaining levels of the hierarchy enforce constraints on the variations of the components within a voxel. The drift is represented by a high-dimensional spline basis, but without constraint it would be grossly over-parameterized. The prior for the drift regularizes the profile $d(t)$ by penalizing both its magnitude and curvature, as follows:

$$\boldsymbol{\delta} \mid \sigma^2, \lambda, \boldsymbol{\theta} \ \sim \ A \exp(-Q(\boldsymbol{\delta})/2\lambda)), \tag{3}$$

where A is a normalizing constant. Here, Q is a quadratic form defined by

$$Q(\boldsymbol{\delta}) \propto \left[a_n \, \|d\|_2^2 + a_c \, \|d''\|_2^2\right],$$

where $\|f\|_2 = \int |f(t)|^2$ is the $\mathcal{L}^2$ norm. The constants $a_n \geq 0, a_c > 0$ are fixed, and specify the trade-off between the norm and curvature penalties; typical values are 0.01 and 1 respectively. Since d is explicitly orthogonal to constants, a very large a_n prevents any drift, whereas a very large a_c forces the drift to be linear. The smoothing parameter λ controls the overall degree of smoothing. This is is a hyper-parameter of the model with prior given by

$$\lambda \mid \sigma^2, \boldsymbol{\gamma}, \boldsymbol{\theta} \ \sim \ \text{Exponential}(1/\sigma^2\lambda_0). \tag{4}$$

The mean of this distribution depends on the fixed hyper-parameter λ_0 and is targeted to achieve a desired "degrees of freedom" for the drift profile, as measured by the trace of the corresponding smoothing matrix [29]. The dependence of this distribution on the noise level is a technical rather than substantive feature, since the smoothness of the drift is not known to depend explicitly on the noise level. However, this noise dependence makes the drift smoothing more adaptive and provides an interpretable way to control the impact of the prior. (It is easier to think in terms of degrees of freedom than in terms of values of the weights.) The other important feature of this distribution is that it shrinks strongly towards zero. When $\lambda = 0$, there is no drift (or only linear drift if $a_n = 0$), so the prior encourages simple drift structure unless sufficient likelihood gain can be achieved with a more complex profile.

The prior for the average response magnitude parameters takes the form

$$\gamma_c \mid \boldsymbol{\gamma}_{-c}, \boldsymbol{\theta}, \sigma^2 \sim \eta \times \text{ point mass at } 0 + (1-\eta)\,\text{Gamma}(\rho_{1c}, \rho_{2c}) \tag{5}$$

where $0 < \eta < 1$, ρ_{1c}, and ρ_{2c} are fixed hyperparameters. Typical values are $\eta = 0.99$, $\rho_{1c} = 1$, and $\rho_{2c} = 50$. The Gamma component of the prior is chosen to cover the interval from 1%–5% of the baseline signal which is roughly the expected range of task-related signal changes for an "active" voxel. The vast majority of voxels will show no response at all to a particular task condition, and the point-mass at zero reflects this possibility. Including such atomic components in the prior is equivalent to specifying a number of distinct sub-models in which various combinations of conditions are excluded. All inferences are generated by averaging over these sub-models. See [26] for a further discussion.

The shape of the response is more complicated. Observations have shown that the response manifests itself in the signal roughly according to the following basic template: a lag between beginning of task performance and signal change, a signal rise over approximately 6 to 8 seconds, signal plateau while the task continues, another lag after task performance ends, followed by a slow decay over 8 to 20 seconds with a possible dip below baseline before returning. The shape parameters $\boldsymbol{\theta}$ each capture one of these features, and the model parameterizes the response by a family of piecewise polynomials adapted to this template [26]. The priors for the shape parameters take the components as independent and are thus of the form

$$\boldsymbol{\theta} \mid \sigma^2 \sim \prod_s \text{Gamma}(\nu_{1s}, \nu_{2s}) \tag{6}$$

where the product is over the component indices of $\boldsymbol{\theta}$ and the parameters ν_{1s} and ν_{2s} are chosen based on current understanding of the blood-flow response. Typical values are $\boldsymbol{\nu}_1 = (2, 4, 2, 4)$ and $\boldsymbol{\nu}_2 = (4.3, 1, 4.3, 0.43)$ for the four parameter response model. The choice of hyperparameters depends in part on the experimental design, as the responses to very short periods of task performance may differ qualitatively from the responses to longer (and more common) periods.

Finally, the noise model used in this paper requires only a parameter to describe the noise variance. The noise level depends on the acquisition parameters, in particular the voxel size and the rate of image acquisition. This dependence can be expressed theoretically for a particular acquisition scheme, and it can also be measured empirically using recorded measurements of "signal-to-noise" tracked by the scanner technicians for quality control. Our prior for the noise level is a proper and diffuse distribution of the form

$$\sigma^2 \sim \text{Inverse Gamma}(\xi_1, \xi_2) \tag{7}$$

where the distribution is centered on the predicted noise level for the given scanner and acquisition scheme. Typical values for the hyperparameters are $\xi_1 = 1.6$ and $\xi_2 = 100$. The spread of the distribution is kept rather diffuse to account for the inevitable variations across voxels in the tissue content and physiological effects. Even more precise information can be brought to bear voxel by voxel if pilot data are acquired prior to the experiment, using the same subject, scanner, and voxel coordinates.

4.2. Fitting the Model

We use two different approaches for fitting the model to fMRI data: direct posterior maximization and posterior sampling using Markov Chain Monte Carlo (MCMC). In the first approach, the location of the posterior mode is used to derive parameter estimates, and standard errors are derived from the normal approximation about the mode with the inverse observed Fisher information matrix. This approach is computationally efficient even for large data sets and is sufficient for most common inferential objectives. Comparisons with the full posterior sampling suggest that it gives good approximations in practice. For numerical optimization to be efficient, it is necessary that the posterior be smooth, so the response amplitude priors with atomic components must be represented as a set of distinct sub-models. Each sub-model corresponds to a constraint that a subset of γ_c parameters are zero. The posterior is maximized separately for each sub-model, and the Bayes Factors are estimated with a form of the Laplace approximation [15] or with the Schwarz criterion [54, 46]. From this, we derive the posterior probabilities of each sub-model, and average results across models using these probabilities as weights.

The second approach, MCMC, is used when more refined inferences are required, for example to study the distribution of complex functions of the parameters or to elucidate the detailed structure of the posterior. The MCMC approach is essential for making inferences about detailed spatial or temporal relationships in the responses across conditions. The sampling strategy is a mix of Metropolis and Gibbs steps [63], depending on the priors used. Sub-models are handled using the reversible jumps of [28], which yields direct estimates of the posterior probabilities across sub-models. Sampling occurs in three stages: a pre-scan where the initial Metropolis jumping distributions are adjusted, a period of burn-in where no output is recorded and the chain equilibrates, and the final sampling. The MCMC approach is computationally intensive and quite time consuming for large data sets, but it allows variations in the qualitative structure of the model and provides a way to tackle questions that cannot easily be addressed otherwise.

The data analysis process begins with image reconstruction and pre-processing using the FIASCO (Functional Image Analysis Software, Computational *Olio*) software package [17]. This includes an algorithm [18] to align successive images that reduces movement artifacts. Maximum posterior estimates are then computed, and if MCMC sampling is to be done, these estimates are used to generate initial values and Metropolis jumping distributions.

4.3. Using the Results

The result of the model fit is an estimate of the posterior distribution of the parameters given the data at *every* voxel. Of primary interest in most cases are the response amplitude parameters $\boldsymbol{\gamma}$ which relate to the level of neural activity involved in processing the corresponding tasks. However, other model parameters, particularly those describing the response shape, can be highly informative regarding particular questions. There are two challenges in using these results effectively: we must construct comparisons that address the scientific questions of interest, and we must present these comparisons in an interpretable way. The latter issue is complicated by the need to communicate the results within the fMRI and scientific communities, many of which are unfamiliar with how to interpret a posterior distribution.

In the fMRI literature, statistical results are often reported as maps showing the value of some estimate or test statistic for every voxel, perhaps averaged over subjects. These maps are usually thresholded to classify which voxels are "active", and the thresholded maps are usually overlayed on an anatomical image to associate the observed pattern activity with anatomical structures.

To the degree that familiarity in form is a benefit to communication, we can use the posterior from our model fit to produce similar but more sensitive maps. However, we eschew direct subject averaging as it is currently implemented: rescaling the brains relative to a select few gross measurements [59] and then mapping classifications onto a common coordinate system. Such across-subject averaging can significantly blur results. We present single subject results here. The problem of integrating results from multiple subjects is an ongoing area of research. See [26] for a discussion of other options.

The simplest way to summarize the fit is to map of the marginal probability that $\gamma_c > 0$ for each condition c. These probabilities reflect the strength of evidence in the data that there is some response in the voxel to the corresponding tasks. These probabilities tend to be very small over most of the voxels yielding an unthresholded map with only a few prominent features. If it is desired to overlay this map on anatomical images or if pure classification is of interest, then these probabilities can be thresholded. (The Bayes rule for thresholding with a binary loss function is a reasonable way to do this [52], although it demands some attention to the relative costs.) Maps of the estimated amplitudes $\widehat{\gamma}_c$ are also useful, but unlike the probability maps, these provide no graphical indication of uncertainty.

The unfamiliarity of posterior probabilities raises some concerns from scientists interpreting these images. (It is worth noting however, that scientists often improperly interpret the p-values as posterior probabilities when using test statistics such as the t-test.) A more familiar alternative is to use what we call normalized contrast maps. These are maps whose pixels give the values of a t-like statistic measuring the magnitude of a contrast among conditions relative to the uncertainty:

$$NC_{c,c'} = \frac{\widehat{\gamma}_c - \widehat{\gamma}_{c'}}{\text{Standard Error}(\widehat{\gamma}_c - \widehat{\gamma}_{c'})}. \tag{8}$$

Here, both the estimates and standard errors include the variation over sub-models. Under the maximization regime, these standard errors are computed from the estimated covariance matrices. These maps allow us to consider particular comparisons among the conditions; there is of course no need to restrict to pairwise contrasts, any linear contrasts of the amplitude parameters are acceptable.

All of the maps thus far discussed are useful for identifying and localizing active voxels, but this formalism can address many more complex questions expressed as functions on the global parameter space. Many more complex questions can be addressed using this formalism, which is a big advantage of posterior inference in this context. Questions of relationships among the conditions, of variation in the response shape, of spatial relationships in specified regions can all be addressed with a measure of uncertainty and without selection biases. Examples and discussion of this point is a main theme of [26]. In the study described here, the inferential goals are more modest, and we focus on the response amplitudes at each location.

5. Experimental Methods

In this section, we describe three fMRI experiments designed to help address our target question: are the pursuit and saccade systems differentiated. Our analysis of the data obtained from the first two experiments is presented in the next section. These two experiments were run relatively early in our work when the experimental designs were relatively simple, as largely dictated by the limitations of the statistical tools then available. In this paper, we re-analyze these data using the model and methods discussed above. As an indication of how methods in the field have developed, we also describe a third experiments with a more sophisticated design that makes more subtle comparisons possible. The analyses of this later experiments will be presented in a future paper. All of the data discussed herein were acquired on the 3T GE Signa Scanner at the University of Pittsburgh Medical Center Magnetic Resonance Research Center.

There are several ways to differentiate the brain systems that subserve pursuit and saccadic eye movements with fMRI data. The most direct is to determine if the observed location of activity ascribed to these two processes are physically separated. Specifically, we ask if the voxels classified as active (relative to a suitable control) under the pursuit and saccade tasks are different, and if these differences in location correspond to reasonable anatomic distinctions. Experiments 1 and 2, described below, are motivated by this approach. Such "location dissociations" require careful interpretation because the changes detected by fMRI involve blood-flow responses which act only as a proxy for neural activity. Nevertheless, given what is currently known about anatomical and functional structure of the oculomotor system, we believe that location dissociations can provide strong evidence for addressing our target question.

A second way to differentiate the neural controls for pursuit and saccades is to design an experiment that manifests different temporal patterns of activation from the two types of tasks. For example, both pursuit and saccade are known to activate the opposite hemispheres for a given direction of eye movement [37]. Pursuit to the left is primarily controlled by the left hemisphere (ipsilateral), and saccades to the left are controlled by the right hemisphere (contralateral). Similarly, the size of a saccade (and possibly the speed of pursuit) change the location of activity within the associated brain area in a very specific way. Experiment 3 is motivated by these considerations.

Experiments 1 and 2 both involve the same simple, two-condition design: the subject alternates between 42 second epochs of a control task and an eye-movement task. In both experiments, the control task, called fixation, requires the subject to hold her gaze on a marked point projected onto the center of her visual field during the entire 42 second epoch.

The eye-movement task differs between the two experiments. Experiment 1 uses the Visually-Guided Saccade (VGS) task, and Experiment 2 uses the Pursuit task, both described below. Each of these tasks is repeated continually throughout each 42 second eye-movement epoch.

In the VGS task, the subject is instructed to make saccades to targets that step three degrees of visual angle either to the left or right of the current fixation point. The target starts in the center of the screen at the beginning of each VGS epoch, and the direction of each step is chosen randomly. For example, from the center of the screen, the target would jump randomly either 3 degrees to the left or right, and the subject would shift her gaze to that new target. The target would then jump randomly 3 degrees to the left or right of that new position for the next saccade and so forth. Target jumps occurred every 750 milliseconds. When the target reaches 9 degrees to the left or right of center, it moves back toward center to remain in the subject's field of view from within the scanner.

In the Pursuit task, the subject is instructed to visually follow a target that moves horizontally across the projection screen at a slow speed. When the target reaches 9 degrees to the left or right of center, it returns in the other direction. The target begins the epoch in the center at the fixation point. The target velocity is cosinusoidal as a function of angular distance from the center—it speeds up toward the center of the screen and slows down to zero velocity at the edges—to help the subject reverse pursuit direction easily. This appears to the subject as a pendulum with motion only in the horizontal plane. The speed of pursuit is timed so the target returns to the central fixation point at the end of the epoch.

The data for Experiments 1 and 2 are acquired using two-shot Echo-Planar Imaging (TR=4.2s, TE=25) to obtain very high-resolution voxels of size 1.6mm $\times$ 0.8mm $\times$ 3mm. The images consist of 6 slices separated by 1mm; these were prescribed by hand using scout images obtained before the experiment and selected to include the region of the FEF within the field of view. Images were acquired every 8.4 s.

The experiments just described involve pursuit and saccades in both directions and thus activity in both hemispheres, making it harder to distinguish the two processes. In

Experiment 3, we integrate both the VGS and Pursuit tasks and fixation into the same design with careful control of the timing and direction of eye movements separate the pursuit and saccade activation by hemisphere. Specifically, the experiment consists of a single 30 second fixation period followed by a sequence of identical 24.17 second blocks each of which consists of four epochs, followed by a single 30 second fixation at the end of the experiment.

An important issue is that at the behavioral level, pursuit and saccades are not completely separable. Failure to keep up with the target during pursuit initiates what are called "catch-up" saccades, small saccades that rapidly bring the eyes back to the target during pursuit. The propensity for catch-up saccades is directly related to the speed of pursuit, with a fast moving target requiring a greater number of saccades. However, catch-up saccades do occur during pursuit at any target velocity. Note that the activation from pursuit and the catch-up saccades is expected to lie in different hemispheres since the eye movements are in the same direction. Throughout the task epochs in Experiment 3, pursuit and saccadic eye movements in *opposite* directions are always separated in time enough so that their temporal patterns of activity will be completely out of phase with each other.

Each 24.17 second block in Experiment 3 is divided into four epochs: Pursuit-R, Fixation-R, Pursuit-L, Fixation-L. The block begins with the subjects gaze focused on a fixation point at the left side of the visual field. A target appears there and moves across the screen to the right for 5.8 seconds. This is Pursuit-R, and the subject is instructed to pursue the target as it moves. Catch-up saccades occur throughout this pursuit period, but the activation of saccade and pursuit show up in opposite hemispheres. When the target reaches the other edge of the screen it stops, and the subject is instructed to fixate on the target. This is Fixation-R, and it lasts for 5.8 seconds. The target then moves horizontally left across the screen for 5.8 seconds. This is Pursuit-L. Finally, the target stops at the left side of the screen, and the subject holds fixation for 6.77 seconds. This is Fixation-L.

The timing of the four epochs is important. During 5.8 seconds, a single 6 slice image is acquired. Within a block, the signal measured within one slice occurs at the same point in the task for both the L and R sections. The extra 0.97 seconds of fixation at the end of the block rotates the acquisition by exactly 1 slice, so that each slice is acquired in the same time interval in the Pursuit-L and Pursuit-R epochs within a block. Thus, in the first block of the experiment, slice 1 records what happens at the beginning of pursuit for both L and R, but in the second block, slice 2 records the beginning pursuit and slice 1 records the end. This aspect of the design controls for variation in activity during task performance when making comparisons across hemispheres. The design ensures that simultaneous saccade and pursuit activation occurs in opposite hemispheres. The combined pattern of activity provides an uncorrupted picture in each hemisphere of where pursuit and saccade activation take place.

The data for this experiment are again acquired with two-shot Echo-Planar Imaging (TR=5.8,TE=25) with the high resolution (1.6mm $\times$ 0.8mm $\times$ 3mm) voxels and 6 slices chosen anatomically. One practical issue is that the speed of pursuit must be carefully selected. Pursuit must be fast enough to engage the saccade system (i.e., manifest an adequately high catch-up saccade rate), which we can test outside the scanner using behavioral data measured in the laboratory [20]. On the other hand, pursuit must be slow enough so that target movement across the screen allows enough time to build up a detectable fMRI signal change. This movement is limited by the angular view of the subject, for once the target reaches the edge of the presentation screen, the only options are (i) to switch to fixation, (ii) move in the other direction, or (iii) jump the target across the screen and continue motion in the same direction. The second option here makes anti-directional pursuit epochs adjacent, and the latter introduces a large saccade in the direction opposite the most recent pursuit. Note that the eye-movement task conditions here are lasting for a much shorter time than in Experiments 1 and 2. This is required if we are to separate activation associated with movement in the two directions, but it reduces the duration and probably the peak magnitude of the BOLD response. We attempt to compensate for this reduction by running a large number of blocks. To avoid tiring the subject, we include several periods of rest during which images are acquired while the subject holds fixation. We also take pains to reduce head motion during acquisition by using a carefully trained and thoroughly instructed subject, as well as dense padding around the head. The slow pursuit speed helps reduce movement as well.

6. Results and Discussion

The basic question in this study is the degree to which the brain's systems for implementing saccadic and pursuit eye-movements differ. This question is interesting methodologically because we know from monkey studies that the regions involved are relatively small and spatially contiguous, so that discriminating these functionally distinct regions represents a significant challenge to the spatial resolution of fMRI. We replicated Experiments 1 and 2 described above on eight subjects and here present the results of our analyses. Because these experiments manipulate the two eye-movement tasks independently and compare them separately to the fixation condition, subtle task-related differences in the temporal patterns of response are washed out by variations across scans. Hence, we use the location of activity as our primary tool for separating these systems with these experiments. Fortunately, the methods described in this paper provide better precision and resolution than standard data analytic methods used in fMRI, and we are able to detect systematic (even striking) differences that were obscured in previous analyses.

In this paper, we focus on two specific inferential measures: posterior probabilities that a voxel is responsive to a particular task and normalized contrasts between the task and control conditions. In general, the modeling approach presented above enables more refined comparisons; indeed, the studies in this paper represent the initial stage of work in progress. All of the figures in this section show the results for the same subject. This subject was chosen, quite literally, at random from among the eight. Our results appear to be consistent across all our subjects, and we highlight any notable differences in the discussion.

Figures 13a and 13b shows thresholded posterior probability maps of activation over six brain slices for the VGS (Experiment 1) and Pursuit (Experiment 2) tasks respectively. The probabilities shown in the maps are the posterior probabilities that the response amplitude γ_c

is non-zero for the corresponding task condition. The threshold of 0.2 was chosen for visual clarity, not for classification. The probabilities across all voxels divide into two groups, one concentrated near zero and a smaller group with non-trivial values. The value 0.2 lies in the lower tail of the latter cluster, and the vast majority of voxels have a much smaller probability of non-zero response. The patterns of activity in these maps reveal structure where we would expect it for both of these tasks on the basis of previous studies and the animal literature. The activation follows the sulci quite closely, and the Frontal Eye Fields, Supplementary Eye Fields, and the parietal eye fields along the Intraparietal Sulcus all show activation. Even at this level we can see some gross differences between activations present during the pursuit and saccade tasks. One striking feature of these maps is the minimal structure outside the brain. Voxels outside the brain can show no true activity, but naturally, chance fluctuations can lead to significant test statistics. The t-maps we had used for our earlier analyses of these data showed a substantial number of active voxels outside the brain. In contrast, using the methods in this paper, the likelihood gains from fitting these fluctuations are small because the noise fluctuations outside the brain are less likely to conform to the constraints on the activation profile. Consequently, the shrinkage implicit in averaging over sub-models eliminates the vast majority of these spurious structures. This is one example of how the constraints embedded in our model can improve the sensitivity of the inferences.

Three pairs of figures Figures 14ab, 15ab, and 16ab compare activation in FEF and SEF for two specific slices. The first pair shows a slice near the top of the brain for which the activation is concentrated in a small region near the midline. The pattern of activation for VGS (Experiment 1) and Pursuit (Experiment 2) appear to diverge along different sulci (note the different orientations of the activation clusters), showing a structural difference. In the second pair of figures, the most notable difference between VGS and Pursuit is that the dense lateral clusters of activity appear on both sides of the brain whereas for Pursuit they are primarily on one side. Finally, the third pair of figures shows how the activation for both tasks follow the precentral sulcus (see below). In the SEF (the medial areas of activation), Pursuit activation was very modest. VGS activation, when present in the SEF, was most robust at the extreme lateral extent of small sulcal involutions of the midline of the brain.

In the precentral sulcus, the location of the human FEF, the functional organization of the pursuit and saccade regions appears similar to that present in the non-human primate. Along the anterior wall of this sulcus, as it comes down from the cortical surface, most of the activation is primarily associated with saccades. However, deeper down the sulcus, pursuit activation becomes evident immediately beneath the area of saccade-related activation. Figures 17 and 18 illustrate these differences using maps of normalized contrasts (see equation [8]) between saccade or pursuit and fixation, thresholded at 3. In the case presented in the Figure 18, it can be seen that this pursuit-related activation is restricted to the anterior wall of the sulcus, and is not evident as the sulcus wraps back toward the cortical surface on the posterior wall of the sulcus. In contrast, in Figure 17, the saccade-related activity is restricted to the opposite side of the sulcus. In our other subjects, the spatial relationship of these regions are preserved, including the spatial ordering along the precentral sulcus and the adjacency of the saccade and pursuit activations. However, the saccade-related area can appear deeper toward the base of the precentral sulcus with the pursuit area shifted deeper to the base of the sulcus and sometimes wrapping between 1 to 2 centimeters up the deepest-most aspect of the posterior wall of the sulcus. This is consistent with the results of single-unit recording in monkeys [27].

Figure 19 shows classification maps based on thresholded normalized contrasts for part of six slices from the top of the brain to the bottom. In this figure, the Pursuit and VGS areas show systematic and three-dimensional differences, with minimal overlap. Moving from the top of the brain to the bottom, we can see the movement of the pursuit and saccade areas along the posterior midline of the brain, in the region called the precuneus, and at deeper slices onto the lateral cortical surface toward the intraparietal sulcus. As in the FEF, the activation in parietal cortex associated with saccades and pursuit was comprised of two small and adjacent regions that we had found indistinguishable with lower resolution imaging and without the statistical advances described above.

Finally, in several of our subjects, examination of their anatomical images revealed the presence of a small sulcus bridging the precentral sulcus and the central sulcus in some subjects, The central sulcus runs parallel to the precentral sulcus and 1-2 centimeters posterior to it. This is a known anatomical variation in the human brain, and it is particularly interesting for the study of pursuit and saccade because in subjects who did have this anatomic variant it was often the most robustly activated area of the frontal eye fields. Figure 20 shows a pattern we observed in all such subjects. Along this special sulcus, pursuit and saccade are clearly differentiated, with the pursuit and saccadic regions being isolated to different banks of the sulcus.

7. Conclusions

Using new methods for imaging and statistical analysis, we are able to detect systematic differences between pursuit and saccade that were previously unclear with earlier methods. The statistical model we used offers improved sensitivity and greater inferential scope over current methods. This approach demonstrates the scalability of Bayesian methods to large and complex problems.

We are currently developing this work in several directions. On the scientific side, we are analyzing Experiment 3 for data from several subjects; our preliminary results suggest that we are able to distinguish pursuit and saccade more effectively with this design. Our next steps involve studying the higher cognitive involvement of the oculomotor system, using single-trial (as opposed to block-trial) designs to distinguish different temporal patterns of response for cognitive, motor, and sensory sub-processes. On the statistical side, we are attacking three critical problems: incorporating spatial structure into the model, combining inferences across subjects, and using the estimated temporal patterns of activation in the different oculomotor areas to characterize the functional connectivity among the areas. We are also extending the model to handle more complex features of the noise.

References

[1] J. R. Baker, R. M. Weisskoff, C. E. Stern, D. N. Kennedy, A. Jiang, K. K. Kwong, L. B. Kolodny, T. L. Davis, J. L. Boxerman, B. R. Buchbinder, V. J. Weeden, J. W. Belliveau, and B. R. Rosen. Statistical assessment of functional MRI signal change. In *Proceedings of the Society for Magnetic Resonance, Second Annual Meeting*, volume 2, page 626. SMR, 1994.

[2] P. A. Bandettini, A. Jesmanowicz, E. C. Wong, and J. Hyde. Processing strategies for time-course data sets in functional MRI of the human brain. *Magnetic Resonance in Medicine*, 30:161–173, 1993.

[3] W. Becker. Metrics. In R. H. Wurtz and M. E. Goldberg, editors, *The Neurobiology of Saccadic Eye Movements*, pages 13–39. Elsevier, New York, 1989.

[4] J.W. Belliveau, D.N. Kennedy, R.C. McKinstry, B.R. Buchbinder, R.M. Weisskoff, M.S. Cohen, J.M. Vevea, T.J. Brady, and B.R. Rosen. Functional mapping of the human visual cortex by magnetic resonance imaging. *Science*, 254:716–719, 1992.

[5] R. A. Berman, B. Luna, B. J. McCurtain, M. H. Strojwas, J. Voyvodic, K. R. Thulborn, and J. A. Sweeney. fmri studies of human frontal eye fields [abstract]. In *Society for Neuroscience*, volume 22, page 1687, 1996.

[6] A. M. Bronstein and C. Kennard. Predictive ocular motor control in parkinson's disease. *Brain*, 108:925–940, 1985.

[7] M. A. Brown and R. C. Semelka. *MRI: Basic Principles and Applications.* John Wiley and Sons, New York., 1995.

[8] C. J. Bruce and M. E. Goldberg. Primate frontal eye fields. I. single neurons discharging before saccades. *Journal of Neurophysiology*, 53:603–635, 1985.

[9] J. R. Carl and R. S. Gellman. Human smooth pursuit: Stimulus-dependent responses. *Journal of Neurophysiology*, 57:1446–1463, 1987.

[10] B. Clementz and J. A. Sweeney. Is eye movement dysfunction a biological marker for schizophrenia? A methodological review. *Psychology Bulletin*, 108:77–92, 1990.

[11] J.D. Cohen, S.D. Forman, T.S. Braver, B.J. Casey, D. Servan-Schreiber, and D.C. Noll. Activation of prefrontal cortex in a non-spatial working memory task with functional MRI. *Human Brain Mapping*, 1:293–304, 1994.

[12] J.D. Cohen, D.C. Noll, and W. Schneider. Functional magnetic resonance imaging: Overview and methods for psychological research. *Behavior Research Methods, Instruments, Computers*, 25(2):101–113, 1993.

[13] C. L. Colby, J. R. Duhamel, and M. E. Goldberg. Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *Journal of Neurophysiology*, 76(5):2841–2852, 1996.

[14] C. de Boor. *A Practical Guide to Splines*. Springer-Verlag, 1978.

[15] T. J. DiCiccio, R. E. Kass, A. Raftery, and L. Wasserman. Computing Bayes Factors by combining simulation and asymptotic approximations. *J. Amer. Statist. Assoc.*, pages 903–915, 1997.

[16] W. F. Eddy. Comment on Lange and Zeger. *Journal of the Royal Statistical Society C*, 46:19–20, 1997.

[17] W. F. Eddy, M. Fitzgerald, C. R. Genovese, A. Mockus, and D.C. Noll. Functional image analysis software - computational olio. In A. Prat, editor, *Proceedings in Computational Statistics*, volume 12 pp. 39-49. Physica-Verlag, Heidelberg, (1996).

[18] W. F. Eddy, M. Fitzgerald, and D. C. Noll. Improved image registration using Fourier interpolation. *Magn. Reson. Med.*, 36:923–931, 1996.

[19] S. Forman, J. C. Cohen, M. Fitzgerald, W.F. Eddy, M.A. Mintun, and D. C. Noll. Improved assessment of significant change in functional magnetic resonance fMRI: Use of a cluster size threshold. *Magn. Reson. Med.*, 33:636–647, (1995).

[20] L. Friedman, J. A. Jesberger, and H. Y. Meltzer. A model of smooth pursuit performance illustrates the relationship between gain, catch-up saccade rate, and catch-up saccade amplitude in normal controls and patients with schizophrenia. *Biological Psychiatry*, 30:537–556, 1991.

[21] K. J. Friston, C. D. Frith, and R. S. J. Frackowiak. *Human Brain Mapping*, 1:69–79, 1994.

[22] K. J. Friston, P. Jezzard, and R. Turner. Analysis of functional MRI time-series. *Human Brain Mapping*, 1:153–171, 1994.

[23] K.J. Friston, A.P. Holmes, J.B. Poline, P.J. Grasby, S.C.R. Williams, R.S.J. Frackowiak, and R. Turner. Analysis of fMRI time series revisited. *NeuroImage*, 2:45–53, 1995.

[24] S. Funahashi, C. J. Bruce, and P. S. Goldman-Rakic. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, 61:331–349, 1989.

[25] B. Gaymard and C. Pierrot-Deseilligny. Impairment of sequences of memory-guided saccades after supplementary motor area lesions. *Annals of Neurology*, 28:622–626, 1990.

[26] C. R. Genovese. Statistical inference in functional Magnetic Resonance Imaging. Technical Report 674, Department of Statistics, Carnegie Mellon University, 1997.

[27] J. P. Gottlieb, G. MacAvoy, and C. J. Bruce. Neural responses related to smooth-pursuit eye movements and their correspondence with electrically elicited smooth eye movements in the primate frontal eye field. *Journal of Neurophysiology*, 72(4):1634–1653, 1994.

[28] P. J. Green. Reversible jump MCMC computation and Bayesian model determination. *Biometrika*, 82:711–732, 1995.

[29] T.J. Hastie and R.J. Tibshirani. *Generalized Additive Models.* Chapman and Hall, 1990.

[30] J. V. Haxby, C. L. Grady, B. Horwitz, L. G. Ungerleider, M. Mishkin, R. E. Carson, P. Herscovitch, M. B. Schapiro, and S. I. Rapoport. Dissociation of object and spatial visual processing pathways in human extrastriate cortex. *Neurobiology*, 88:1621–1625, 1991.

[31] P. S. Holzman, L. R. Proctor, D. L. Levy, N. J. Yasillo, H. Y. Meltzer, and S. W. Hirt. Eye-tracking dysfunctions in schizophrenic patients and their relatives. *Archives of General Psychiatry*, 31, 1430151 1974.

[32] M. A. Just and P. A. Carpenter. A theory of reading: From eye fixations to comprehension. *Psychological Review*, 87:329–354, 1980.

[33] C. Kennard and F. C. Rose. *Physiological Aspects of Clinical Neuro-ophthalmology.* Chicago Year Book Medical Publishers, Chicago, IL, 1988.

[34] K.K. Kwong, J.W. Belliveau, D.A. Chesler, I.E. Goldberg, R.M. Weisskoff, B.P. Poncelet, D.N. Kennedy, B.E. Hoppel, M.S. Cohen, R. Turner, H. Cheng, T.J. Brady, and B.R. Rosen. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl. Acad. Sci. U.S.A.*, 89:5675, 1992.

[35] N. Lange and S. Zeger. Non-linear Fourier time series analysis for human brain mapping by functional magnetic resonance imaging. *Journal of the Royal Statistical Society C Applied Statistics*, 46:1–29, 1997.

[36] E. L. Lehmann. *Nonparameterics: Statistical Methods Based on Ranks.* Holden Day, Oakland, CA, (1975).

[37] R. J. Leigh and D. S. Zee. *The Neurology of Eye Movements (2nd ed.).* F. A. Davis, Philadelphia, PA, 1991.

[38] C. J. Lueck, S. Tanyeri, T. J. Crawford, L. Henderson, and C. Kennard. Antisaccades and remembered saccades in parkinson's disease. *Journal of Neurology, Neurosurgery and Psychiatry*, 53:284–288, 1990.

[39] B. Luna, K. R. Thulborn, M. H. Strojwas, B. J. McCurtain, R. A. Berman, C. R. Genovese, and J. A. Sweeney. Dorsal cortical regions subserving visually-guided saccades in humans: An fmri study. *Cerebral Cortex*, 8:40–47, 1998.

[40] D. Marr. *Vision*. W.H. Freeeman and Co., New York, (1982).

[41] N. J. Minshew, J. A. Sweeney, and M. L. Bauman. Neurologic aspects of autism. In D.J. Cohen and F.R. Volkmar, editors, *Handbook of Autism and Pervasive Developmental Disorders*, pages 344–369. John Wiley and Sons, second edition, 1998.

[42] N. J. Minshew, J. A. Sweeney, and J. M. Furman. Evidence for a primary neocortical system abnormality in autism [abstract]. In *Society for Neuroscience Abstracts*, volume 21, page 293.13, 1995.

[43] M. A. Mintun, M. E. Raichle, W. R. W. Martin, and P. Herscovitch. Brain oxygen utilization measured with 0-15 radiotracers and positron emission tomography. *Journal of Nuclear Medicine*, 25:177–187, 1984.

[44] D. P. Munoz and R. H. Wurtz. Role of the rostral superior colliculus in active visual fixation and execution of express saccades. *Journal of Neurophysiology*, 67:1000–1002, 1992.

[45] S. Ogawa, D.W. Tank, D.W. Menon, J.M. Ellermann, S. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping using MRI. *Proc. Natl. Acad. Sci. U.S.A.*, 89:5951–5955, 1992.

[46] D. Pauler. *The Schwarz Criterion for Mixed Effects Models*. PhD thesis, Carnegie Mellon University, 1996.

[47] L. Petit, V. P. Clark, J. Ingeholm, and J. V. Haxby. Dissociation of saccade-related and pursuit-related activation in human frontal eye fields as revealed by fmri. *Journal of Neurophysiology*, 77(6):3386–3390, 1997.

[48] J. B. Poline and B. Mazoyer. Cluster analysis in individual functiona brain images: Some new techniques to enhance the sensitivity of activation detection methods. *Human Brain Mapping*, 2:103–111, 1994.

[49] M. I. Posner and M. E. Raichle. *Images of Mind*. Scientific American Library, 1994.

[50] P. Rakic. The development of the frontal lobe: A view from the rear of the brain. In H. H. Jasper, S. Riggio, and P. S. Goldman-Rakic, editors, *Epilepsy and the Functional Anatomy of the Frontal Lobe*, pages 1–8. Raven Pressn, New York, 1995.

[51] D. R. Rosenberg, J. A. Sweeney, E. Squires-Wheeler, M. S. Keshavan, B. A. Cornblatt, and L. Erlenmeyer-Kimling. Eye-tracking dysfunction in offspring from the new york high-risk project: Diagnostic specificity and the role of attention. *Psychiatry Research*, 66:121–130, 1997.

[52] M. J. Schervish. *Theory of Statistics*. Springer-Verlag, New York, (1995).

[53] J. Schlag and M. Schlag-Rey. Evidence for a supplementary eye field. *Journal of Neurophysiology*, 57:179–200, 1987.

[54] G. Schwarz. Estimating the dimension of a model. *Ann. Stat.*, 6(2):461–464, 1978.

[55] M. A. Sommer and E. J. Tehovnik. Reversible inactivation of macaque frontal eye field. *Experimental Brain Research*, in press.

[56] J. A. Sweeney, B. A. Clementz, G. L. Haas, M. D. Escobar, K. Drake, and A. J. Frances. Eye tracking dysfunction in schizophrenia: Characterization of component eye movement abnormalities, diagnostic specificity, and the role of attention. *Journal of Abnormal Psychology*, 103:222–230, 1994.

[57] J. A. Sweeney, B. Luna, R. A. Berman, B. J. McCurtain, M. H. Strojwas, J. T. Voyvodic, C. R. Genovese, and K. R. Thulborn. fmri studies of spatial working memory [abstract]. In *Society For Neuroscience*, volume 22, page 1688, 1996.

[58] J. A. Sweeney, M. A. Mintun, S. Kwee, M. B. Wiseman, D. L. Brown, D. R. Rosenberg, and J. R. Carl. A positron emission tomography study of voluntary saccadic eye movements and spatial working memory. *Journal of Neurophysiology*, 75(1):454–468, 1996.

[59] J. Talairach and P. Tounoux. *Coplanar Stereotaxic Atlas of the Juman Brain. Three-dimensional Proportional System: An Approach to Cerebral Imaging.* Thieme, 1988.

[60] K. R. Thulborn. personal communication.

[61] K.R. Thulborn, J.C. Waterton, P.M. Matthews, and G.K. Radda. Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. *Biochem. Biophys. Acta*, 714:265–270, 1982.

[62] J. R. Tian and J. C. Lynch. Functionally defined smooth and saccadic eye movement subregions in the frontal eye field of cebus monkeys. *Journal of Neurophysiology*, 76(4):2740–2753, 1996.

[63] L. Tierney. Markov chains for exploring posterior distributions. *The Annals of Statistics*, 22(4):1701–1727, 1994.

[64] R. B. H. Tootell, J. B. Reppas, K. K. Kwong, R. Malach, R. T. Born, T. J. Brady, B. R. Rosen, and J. W. Belliveau. Functional analysis of human mt and related visual cortical areas using magnetic resonance imaging. *Journal of Neuroscience*, 15:3215–3230, 1995.

[65] Y. Vardi, L. A. Shepp, and L. Kaufman. A statistical model for positron emission tomography. *J. Amer. Statist. Assoc.*, 80:8–20, 1985.

[66] J. B. Weaver, A. Y. Saykin, R. B. Burr, H. Riordan, and A. Maerlender. Principal component analysis of functional MRI of memory. In *Proceedings of the Society for Magnetic Resonance, Second Annual Meeting*, page 808. SMR, 1994.

[67] R. Weisskoff. personal communication, July, 1997.

[68] R. M. Weisskoff, J. Baker, J. Belliveau, T. L. Davis, K. K. Kwong, M. S. Cohen, and B. R. Rosen. Power spectrum analysis of functionally-weighted MR data: What's in the noise? In *Proceedings of the Society for Magnetic Resonance in Medicine, Twelfth Annual Meeting*, page 7. MRM, 1993.

[69] K. J. Worsley and K.J. Friston. Analysis of fMRI time series revisited – again. *NeuroImage*, 2:173–181, 1995.

[70] R. H. Wurtz and M. E. Goldberg. Activity of superior colliculus in behaving monkey. III. cells discharging before eye movements. *Journal of Neurophysiology*, 35:575–586, 1972.

[71] R. H. Wurtz and M. E. Goldberg. *The Neurobiology of Saccadic Eye Movements*. Elsevier, New York, 1989.

Figure Captions

Figure 1. A schematic of the brain in external side view, showing a single hemisphere divided into lobes. The arrows indicate the canonical directions, each of which can be referred to by two synonymous terms (e.g., Anterior or Rostral, Posterior or Caudal).

Figure 2. A cut-away side view of the brain at the midline plane between the left and right hemispheres. Various anatomical structures that are relevant to eye movements are labeled.

Figure 3. A schematic of the oculomotor system, including the principal brain regions involved in the system and the principal pathways among them. Here, SEF=Supplementary Eye Fields, FEF=Frontal Eye Fields, PPC=Posterior Parietal Cortex (particularly the Intraparietal sulcus), DLPFC=Dorsilateral Prefrontal Cortex, and SC=Superior Colliculus. MT (a.k.a. V5) is a motion sensitive component of the visual system. The Visual Cortex contains the early visual areas V1, V2, and so forth. V1 is the largest of these.

Figure 4. A heuristic comparison of the spatial and temporal resolutions of various techniques for studying the brain and brain function. The boxes for each technique show the range of resolutions that can be explored with that technique. This display is based on a similar diagram in [49].

Figure 5. Four snapshots of voxel magnetizations at different points during the acquisition of a $4 \times 4 \times 1$ image. The arrows denote the in-slice component of the net voxel magnetization; think of them as rotating clockwise (and clock like) around their flat end. The lengths of the arrows are proportional to the hydrogen densities within the voxel. The directions of the arrows show the relative phase of precession. The value of $\Delta\nu$ for the voxel indicates the difference between the precession frequency for the voxel magnetization (the frequency with which the arrow is rotating) and the base precession frequency ν_0 that would hold in the absence of gradients. The four panels show the configuration in the following cases: (a) just after slice selection, (b) at the end of phase encoding, (c) during frequency encoding, and (d) at the end of phase encoding (using a different phase shift than in (b)).

Figure 6. A representation of a simple alternating fMRI experimental design that indicates which task is being performed at every time throughout the experiment. The horizontal axis shows the corresponding image index, and the heights of the line segments serve to separate the conditions. One image was acquired every 3 seconds. Each horizontal line corresponds to a single task epoch. The task and control conditions in this case are finger-tapping and rest, respectively.

Figure 7. A representation of an fMRI data set. A series of three-dimensional images is acquired at regular intervals. Each of these images (the cubes) consist of an array of values, one value per volume element (voxel). The brain is typically imaged in a series of two-dimensional slices as depicted.

Figure 8. Two slices of functional MR images showing the capability of trading-off temporal and spatial resolution.

Figure 9. A single voxel time series, with time on the horizontal axis (seconds) and MR signal strength on the vertical axis (arbitrary units). The hashed band at the bottom of the figure shows the experimental design, with the task of interest being performed during \\\-hashed epochs and the control task during ///-hashed epochs.

Figure 10. The relationship between the voxel time series and the voxel in the image: each measurement in the time series corresponds to the measured nuclear density in the labeled voxel over time.

Figure 11. This voxel time series shows the correspondence between task performance and the BOLD perturbation. The horizontal axis gives the image index; one image was acquired every 3 seconds. The vertical axis gives the MR signal strength in arbitrary units. The bar along the horizontal axis indicates the experimental design as follows. The \\\-hashed blocks (red) correspond to the "Task" condition where the subject performs the task of interest. The ///-hashed blocks (blue) correspond to the "Control" condition where the subject performs the control task. The lines on the time series curve are styled to indicate the timing of each data point: solid (red) lines correspond to "Task" and dotted (blue) lines correspond to "Control". Note the slight lag in the perturbation relative to task performance.

Figure 12. A collection of voxel time series highlighting some of the heterogeneity of the signal drift. The horizontal axes represent the image index and the vertical axes the strength of the MR signal.

Figure 13. Posterior probabilities of non-zero response amplitude for (a) VGS (Experiment 1) and (b) Pursuit (Experiment 2) overlaid on anatomical images for six slices of the brain.

Figure 14. Close-up view of posterior probability maps in a single slice for (a) VGS (Experiment 1) and (b) Pursuit (Experiment 2). This image highlights activation differences along the midline: saccade and pursuit activation appear to diverge along distinct sulci; note the distinct directions of the central activation clusters.

Figure 15. Close-up view of posterior probability maps in a single slice for (a) VGS (Experiment 1) and (b) Pursuit (Experiment 2). This image highlights activation differences in the Frontal and Supplementary Eye Fields (activity along the precentral sulcus and near the midline respectively).

Figure 16. Close-up view of posterior probability maps in a single slice for (a) VGS (Experiment 1) and (b) Pursuit (Experiment 2) This image highlights activation differences in the Frontal and Supplementary Eye Fields (lateral and medial clusters of activity respectively).

Figure 17. Normalized contrast map for VGS activation (Experiment 1) in a single slice, thresholded at 3 and overlaid on an anatomical image. The area surrounding the precentral sulcus is highlighted and outlined. See equation (8) for definition of normalized contrasts.

Figure 18. Normalized contrast map for Pursuit activation (Experiment 2) in a single slice, thresholded at 3 and overlaid on an anatomical image. The area surrounding the precentral sulcus is highlighted.

Figure 19. Sections of the parietal cortex for several slices arranged from the top to the bottom of the brain. The colored voxels indicate which areas activate in response to each task. Classification based on normalized contrast map thresholded at 3.

Figure 20. Area near the Frontal Eye Fields in three adjacent slices for a single subject. Inset highlights the special sulcus detected in several subjects. The colored voxels indicate which areas activate in response to each task. Classification based on normalized contrast map thresholded at 3.