SPATIAL UPDATING IN HUMAN CORTEX

by

Elisha P. Merriam

B.A., Vassar College, 1997

Submitted to the Graduate Faculty of

the The School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2005

UNIVERSITY OF PITTSBURGH

THE SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Elisha P. Merriam

It was defended on

November 18, 2005

and approved by

Carol L. Colby, Ph.D., Dept. of Neuroscience, Pitt

Carl R. Olson, Ph.D., Dept. of Neuroscience, Pitt

Christopher R. Genovese, Ph.D., Dept. of Statistics, CMU

Mark D'Esposito, M.D., Dept. of Psychology, U.C. Berkeley

Marlene Behrmann, Ph.D., Dept. of Psychology, CMU

Seong-Gi Kim, Ph.D., Dept. of Neurobiology, Pitt

Dissertation Director: Carol L. Colby, Ph.D., Dept. of Neuroscience, Pitt

ABSTRACT

SPATIAL UPDATING IN HUMAN CORTEX

Elisha P. Merriam, BA

University of Pittsburgh, 2005

Single neurons in several cortical areas in monkeys update visual information in conjunction with eye movements. This remapping of stimulus representations is thought to contribute to spatial constancy. The central hypothesis here is that spatial updating also occurs in humans and that it can be visualized with functional MRI.

In Chapter 2, we describe experiments in which we tested the role of human parietal cortex in spatial updating. We scanned subjects during a task that involved remapping of visual signals across hemifields. This task is directly analogous to the single-step saccade task used to test spatial updating in monkeys. We observed an initial response in the hemisphere contralateral to the visual stimulus, followed by a remapped response in the hemisphere ipsilateral to the stimulus. Our results demonstrate that updating of visual information occurs in human parietal cortex and can be visualized with fMRI.

The experiments in Chapter 2 show that updated visual responses have a characteristic latency and response shape. Chapter 3 describes a statistical model for estimating these parameters. The method is based on a nonlinear, fully Bayesian, hierarchical model that decomposes the fMRI time series data into baseline, smooth drift, activation signal, and noise. This chapter shows that this model performs well relative to commonly-used general linear models.

In Chapter 4, we use the statistical method described in Chapter 3 to test for the presence of spatial updating activity in human extrastriate visual cortex. We identified the borders of several retinotopically defined visual areas in the occipital lobe. We then tested for spatial

updating using the single step saccade task. We found a roughly monotonic relationship between the strength of updating activity and position in the visual area hierarchy. We observed the strongest responses in area V4, and the weakest response in V1. We conclude that updating is not restricted to brain regions involved primarily in attention and the generation of eye movements, but rather, is present in occipital lobe visual areas as well.

Keywords: Spatial Constancy, Parietal Cortex, Extrastriate Cortex, Functional Magnetic Resonance Imaging, Bayesian Statistics.

TABLE OF CONTENTS

LIST OF FIGURES

1.0 GENERAL INTRODUCTION

1.1 SPATIAL CONSTANCY

Our perception of the visual world is extremely precise. We are able to discriminate spatial frequencies of up to 50 cycles/degree — about the size of the letters on this page from a distance of 5.5 m. High acuity enables us to detect edges and textures, to recognize patterns and objects, and ultimately, to make sense of a rich and complex visual world. However, our visual system is only capable of this resolution at the very center of the visual field, in the region of space corresponding to the fovea. Acuity drops off precipitously in the peripheral visual field. As an illustration of the low resolution of peripheral vision, fixate the chapter number at the top of the page and attempt to identify the number at the bottom of the page. We make eye movements in order to overcome the limited spatial range of high-acuity vision — we have to move our eyes in order to fully *see* the world.

Humans make an average of three eye movements every second. These eye movements, called saccades, direct the fovea onto locations of interest. Saccades are completely effortless, so much so that we are typically unaware when we are moving our eyes — you are making saccades right now as you read this text. Yet the saccadic system is critical for our perception of the world around us. Saccades represent a maximally efficient strategy for sampling visual information (Land and Nilsson, 2001), and they are tightly yoked to attention and memory systems (Kowler et al., 1995). Saccades also pose a major computational challenge to the visual system. Each eye movement leads to the displacement of images on the retina (Figure 1). Thus, the retinal input consists of a series of disjointed snapshots. How does the visual system make sense of this chaotic sensory information? The brain must construct a representation of the visual scene by combining information about the size and direction of eye movements with input from the eyes. This is an active process that occurs in real time, from saccade to saccade. The end product of this process is the stable and coherent visual perception with which we are all familiar. This phenomenon, referred to as *spatial constancy*, is the focus of the current experiments.

There are at least three sources of information that the visual system could use to update representations of space in conjunction with eye movements (see reviews in Bridgeman et al., 1994; Donaldson, 2000). First, the retinal image itself informs the visual system about the

Figure 1. Saccades lead to the displacement of stable objects on the retina. A visual scene before (A) and after (B) a 10° horizontal saccade. The rightward saccade causes the leftward displacement of the retinal image. How does the visual system distinguish between image displacement caused by eye movements from image displacement caused by world movement? (Images adapted from Wexler, 2005)

size and direction of an eye movement. Gibson argued that the "ambient optic array" contains sufficient information to reconstruct the position and actions of the observer. For example, displacement of a portion of the retinal mosaic may be interpreted by the visual system as movement of an object in the world, whereas simultaneous movement of the entire mosaic may be interpreted as the effect of an eye movement (Gibson, 1950; 1966; 1979). Bridgeman et al. (1994) pointed out that retinal information can only be used by the visual system if there is a clearly visible background. This limitation indicates that the visual system must rely on additional sources of information regarding eye movements.

Proprioceptive afferents may provide a second source of information regarding eye movements. Stretch receptors in the extraocular muscles are activated each time the eyes move. Sherrington (1898, 1918) suggested that these signals could be used by the visual system to keep track of eye movements. Experimental evidence suggests that proprioception does play a role in spatial constancy (Gauthier et al., 1990; Lewis and Zee, 1993; Bridgeman and Stark, 1991; Donaldson, 2000). However, it is unlikely to be the only source of eye movement information. Afferents from stretch receptors provide a relatively late signal, beginning only 20–30 ms after saccade onset. This signal is so slow that most eye movements are finished by the time proprioceptive information reaches the visual system. Furthermore, psychophysical

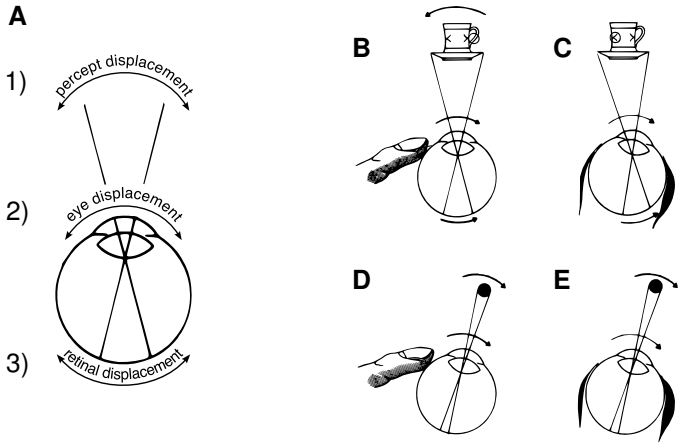

Figure 2. Schematic of Helmholtz's eye displacement experiment. (A) Three events are depicted in each panel: (1) perceptual displacement, (2) eye rotation, (3) displacement of objects on the retina. (B-C) Afterimages move with the eye. Neither finger press (B) or saccades (C) result in the displacement of afterimages on the retina. Therefore, our perception of afterimages moves when our eyes move. (D) Perception of stable objects in the world are displaced when the eyeball is manually rotated. (E) Perception of stable objects in the world *are not* displaced when the eyeball is rotated by a saccade. The retinal displacement depicted in (D) and (E) is identical. The difference in perception in these two cases is attributed to corollary discharge associated with the voluntary saccade in (E). (Modified from Grüsser, 1994)

experiments have demonstrated saccade-related perceptual effects as early as 50 ms *prior* to the saccade (Ross et al., 2001). Other signals besides proprioception must be used by the visual system to compensate for saccades.

Both proprioceptive and retinal information reach the visual system through *neural inflow* — information passes from the periphery (muscle or retina) to the central nervous system. A third source of information regarding eye movements, first proposed by Helmholtz (1911), relies on *neural outflow*. By this account, the visual system receives a copy of the signal that the motor system sends out to the eye muscles. The visual system then uses this motor command to adjust for the displacement of images on the retina. In support of this account, Helmholtz observed that perception of the world is displaced when the side of the eye is lightly tapped (Figure 2B). In contrast, no displacement is perceived when the eyes are moved by making a saccade (Figure 2C). Further support for the centrality of neural outflow comes from the effects of eye movements on stabilized retinal images (images that are fixed to the retina, such as afterimages). In the case of afterimages, there is neither

retinal image motion nor a willful effort to rotate the eye. As predicted by outflow theory, rotating the eyeball manually does not make the afterimage appear to move: the afterimage *stays with* the eye as it moves (Figure 2D-E). Lastly, the effects of oculomotor paralysis support outflow theory. If the eye is immobilized, either pharmacologically or mechanically, subjects report that the visual world appears to move when they attempt to initiate an eye movement (Mach, 1959/1914; Brindley and Merton, 1960). Helmholtz described the effects of eye muscle paralysis on visual perception:

> ...in those cases where certain muscles have suddenly been paralyzed, when the patient tries to turn his eye in a direction in which it's powerless to move any longer, apparent motions are seen, producing double images if the other eye happens to be open at the time. (Helmholtz, 1911, p. 245)

Helmholtz's observations imply that the brain adjusts the representation of visual input based on the intention to move the eyes. When the eyes cannot physically move, these adjustments cause the perception of apparent motion. Helmholtz reasoned that the act of making an eye movement was critical for spatial constancy. Helmholtz used the term *Willensantrengung*, or "effort of will," to emphasize that it was the voluntary nature of self-initiated movements that enabled the visual system to compensate for retinal displacement. Modern scientific evidence favors Helmholtz's proposal. Sperry (1950) proposed that a copy of the motor command, termed a "corollary discharge", was sent to the visual system to enable integration of visual and motor information. At around the same time, von Holst and Mittelstaedt (1950; 1971) described a similar mechanism, calling it *Efferenzkopie*, or "efference copy." A mechanism that makes use of neural outflow is a strong candidate for contributing to stable spatial perception. Neural outflow eliminates the timing issue discussed above because eye movement information is made available to the visual system during the initial stages of motor planning, even before the eyes have begun to move. Furthermore, mechanisms based on neural outflow explain the perceptual disruption in Helmholtz's eye displacement experiment, whereas inflow theories do not (Figure 2).

In the following sections we review the psychophysical literature on spatial updating (Section 1.2). Evidence from psychophysical studies suggests that the visual system uses information about eye movements to create an updated representation of space across saccades.

We discuss the neural basis for spatial updating in Section 1.3. Studies of patients with cortical lesions indicate a central role for the parietal lobe in the integration of visual and oculomotor signals. In the past fifteen years, neurophysiological recordings in alert behaving monkeys have begun to uncover the sequence of physiological events that accompany voluntary eye movements. In Section 1.4, we argue that the phenomenon of visual remapping is the physiological correlate of spatial updating. We discuss the neural circuitry involved in remapping in Section 1.5. Finally, in Section 1.6, we discuss functional imaging studies in humans suggesting that parietal cortex exhibits many of the visual, motor, and cognitive responses properties that are thought to be involved in spatial updating. We argue that these physiological similarities imply, that remapping also occurs in humans. The present experiments directly test the hypothesis that remapping occurs in humans.

1.2 PSYCHOPHYSICS OF SPATIAL CONSTANCY

Our subjective experience suggests that the brain successfully takes our eye movements into account in order to construct a stable percept of the visual world. This subjective observation has been tested experimentally. In a classic study, Hallett and Lightstone (1976) found that subjects are able to direct their eyes accurately to the location of a flashed stimulus presented at the onset of a saccade. The behavioral paradigm used by Hallett and Lightstone (1976), known as the double-step saccade task, is illustrated in Figure 3. The task begins with fixation of a visual target. Two stimuli flash in rapid succession. The fixation cross then disappears and the subject makes a sequence of two eye movements to the flashed target locations. The first saccade can be correctly computed using retinal information alone. In contrast, the second saccade must take into account the size and direction of the first saccade: the second target in this example appears up and to the right, while the required saccade to it is up and to the left (Figure 3D). Accurate performance in this task demonstrates that the visual system is able to compensate for the current position of the eyes. The brain computes the vector of the second saccade using a corollary discharge signal to update the location of the second target.

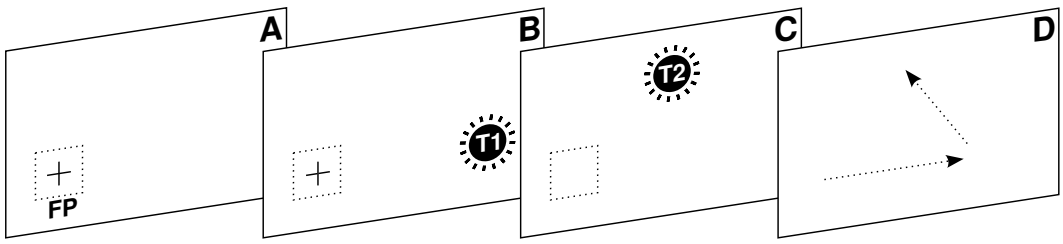

Figure 3. Performance on double-step saccade task requires spatial updating. (A) Subject fixates FP at the start of the trial. Eye position is indicated by the dotted square. (B) A saccade target, T1, flashes briefly. (C) The fixation cross is extinguished, triggering a saccade to T1. Before the eyes begin to move, a second stimulus, T2, is flashed. (D) The subject makes a sequence of two eye movements. The first eye movement to T1 can be computed using retinal information alone. However, T2 was seen up and to the right, but the correct saccade is up and to the left. Thus the second eye movement to T2 requires an updated representation of the stimulus location in eye-centered coordinates. Subjects would make an eye movement up and to the right if no spatial updating occurred.

Accurate performance on the double step task provides behavioral evidence for spatial updating. Despite this behavioral accuracy, there is substantial evidence that perception itself can be inaccurate (see Figure 4). Contrary to our subjective experience, stimuli flashed around the time of an eye movement are perceived at an illusory location (Dassonville et al., 1992; Honda, 1990; Matin and Pearce, 1965; Honda, 1989; Bockisch and Miller, 1999; Cai et al., 1997; Ross et al., 1997; Ross et al., 2001; Schlag and Schlag-Rey, 2002). The magnitude of this perceptual illusion is dependent on both the time of stimulus appearance relative to saccade onset and on the position of the stimulus relative to the saccade target (Mach, 1885; Matin and Pearce, 1965). Perceptual mislocalization has been observed in studies of both horizontal and vertical saccades. Importantly, stimuli are mislocalized regardless of whether the background is illuminated, indicating that the mechanisms responsible for mislocalization are not dependent on retinal inflow (Honda, 1989, 1991; Schlag and Schlag-Rey, 1995).

Morrone et al. (1997) proposed a model of perisaccadic mislocalization consisting of three components: (1) a transient shift in the coordinate system, (2) a compression of the metric of space, and (3) eye position at the time the stimulus appeared. It is likely that the coordinate system shift is caused by a mismatch between actual eye position during saccades (the third component) and a prediction of eye position from an internal corollary discharge signal (Matin et al., 1972; Honda, 1989). The second component consists of a shrinkage of distances

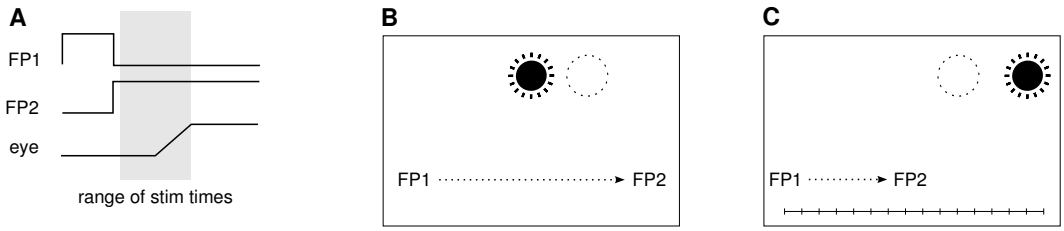

Figure 4. Stimuli presented around the time of a saccade are mislocalized. (A) Time course of events. Trial begins with fixation of FP1. Simultaneously, FP1 is extinguished and FP2 appears, triggering a visually guided saccade. (B) Stimuli presented just prior to saccade onset are mislocalized in the direction of the saccade. The stimulus is indicated by the black star burst. The perceived stimulus location is indicated by the dotted circle. This effect is not dependent on ambient light levels or visual referents, suggesting that it is related to corollary discharge signals. (C) Stimuli presented beyond the saccade are mislocalized against the direction of the saccade, toward the initial fixation point. This effect only occurs in the presence of visual referents, such as a ruler. This effect occurs even if visual referents appear after the saccade, suggesting that mislocalization against the saccade direction is due, in part, to processing of retinal information.

between stimuli, such that locations are compressed toward the saccade target (Ross et al., 1997; Morrone et al., 1997; Lappe et al., 2000). The origin of this compression component has not been fully resolved. It is likely, however, that compression is related to visual reafference rather than corollary discharge. Compression is strongest when visible referents, such as a ruler, are present during the experiment. Visible referents influence compression even if they appear after the saccade, indicating that compression is related to the processing of retinal information, not corollary discharge signals (Lappe et al., 2000; Awater and Lappe, 2004). The existence of perisaccadic mislocalization presents a fundamental paradox: humans and monkeys are able to make accurate saccades in the double-step task, yet are unable to perceive accurately the location of a perisaccadic stimulus. The underlying source of this discrepancy between perceptual and behavioral ability is not currently understood.

1.3 NEUROPSYCHOLOGY OF SPATIAL CONSTANCY

The psychophysical results discussed above indicate that our representation of the visual world is modified around the time of an eye movement. Corollary discharge signals combine

with visual signals to produce a stable percept. We now focus on evidence for a direct neural representation of updated spatial information. One way to address this issue is to determine whether specific brain regions are necessary for spatially accurate oculomotor behavior. In this section we review neuropsychological evidence from monkeys and humans demonstrating that the parietal lobe is critical for accurate spatial representations across saccades.

As discussed above, the double-step task is a direct behavioral test of spatial updating. Do lesions of parietal cortex interfere with performance on this task? To address, this question, Duhamel et al. (1992b) tested a patient with a large right hemisphere lesion that included all of the right parietal lobe and much of the posterior frontal lobe as well. The patient had mild visual deficits and hemispatial neglect, but was able to perform well on a simple visually-guided saccade task. The patient was also unimpaired on trials of the double-step task when the first target appeared in the right (ipsilesional) visual field and the second target appeared in the left (contralesional) visual field, indicating that she could update visual information from the good to the bad hemisphere. Critically, this patient was unable to perform the task accurately when the sequence of saccades was reversed. In these cases, she was able to make an accurate saccade to the first target in the left visual field, even though it was directed to contralesional space. She was unable to make an accurate saccade to the second target, even though the second target was located in the *good*, ipsilesional hemifield. The presumed reason for this deficit is that the visual system was unable to compensate for the saccade made into the contralesional visual field. In other words, the second target could be initially encoded but it could not be accurately updated in conjunction with the first saccade, because the first saccade was directed to contralesional space. The specificity of the behavioral impairment to the second saccade suggests a deficit in spatial updating, and not simply visual or motor processing. The conclusion was further supported by the patient's performance in a control condition. The patient was unimpaired on a visually-guided saccade task that involved the same sequence of eye movements. The conclusion of this study is that parietal cortex is necessary for updating space in conjunction with saccades.

Are the effects of cortical lesions on spatial updating specific to the parietal lobe? To address this question, Heide et al. (1995) studied a large group of 35 patients with lesions

to either parietal, frontal, or dorsomedial frontal cortex. Each patient was tested on four conditions of the double-step task. In two conditions, both the first and the second saccade were directed to the left or right hemifields relative to the initial fixation spot (R-R and L-L conditions). In another two conditions, the two stimuli were located in opposite visual fields (R-L and L-R conditions). In a series of control conditions, all subjects were tested on visually-guided versions of the same sequences . In the visually-guided control tasks, the two stimuli were presented for a longer duration so that the second stimulus was still present on the screen at the time when the eyes had reached the first stimulus. These control conditions were critical in distinguishing an updating deficit from a visual or motor deficit.

Heide et al. (1995) found that although patients with right prefrontal lesions produced large errors in the double-step task, these errors were also significant in the visually-guided control task. Only the patients with parietal lobe damage exhibited a deficit that was specific to the double-step task, confirming the results of Duhamel et al. (1992b). Importantly, the study of Heide et al. (1995) permits comparison of the effects of right and left hemisphere lesions. Both right and left parietal lesions caused errors in the double-step task on trials in which the second saccade crossed the midline (i.e., L-R and R-L cross-hemifield conditions). This indicates that both the right and left parietal lobes are involved in updating the location of contralateral targets. Patients with right parietal damage did exhibit a larger deficit than patients with left parietal damage when both saccade targets were located in the same hemifield. This result indicates that there may be a difference between the hemispheres in the number of targets that can be updated in a sequence of saccades.

These neuropsychological studies indicate that parietal cortex is necessary for spatially accurate behavior. This conclusion has been further supported by a reversible inactivation study in monkeys. Li and Andersen (2001) injected muscimol into the lateral intraparietal area (LIP) of monkeys and then tested them on the double-step saccade task. LIP inactivation in this study interfered with the calculation of the second saccade target. The central conclusion of these studies is that parietal cortex is necessary for updating visual representations in conjunction with saccades.

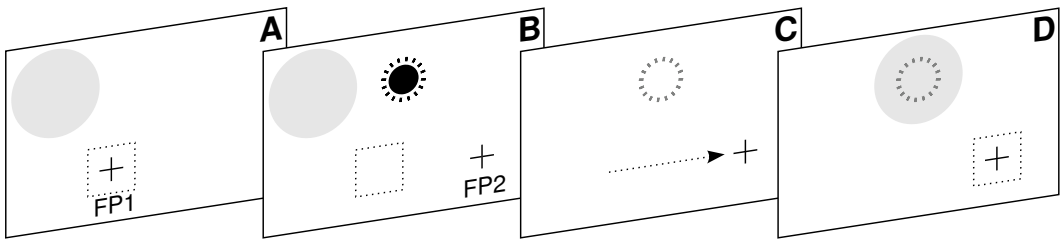

Figure 5. The single-step task used to observe remapping in monkeys. (A) The trial begins with fixation at FP1. Shaded circle represents the receptive field of the neuron being recorded. Dotted square represents the position of the eyes. (B) Three sensory events take place simultaneously: The initial fixation cross (FP1) disappears, a new fixation cross (FP2) appears, and a stimulus (black star burst) is flashed briefly in the upper right visual field. (C) The monkey makes a rightward saccade to FP2. (D) The saccade moves the receptive field onto the location of the recently-flashed stimulus, indicated by the gray star burst. Critically, the stimulus has been extinguished by the time the monkey initiates the saccade – the stimulus was never physically present in the neuron's receptive field at any point in the trial.

1.4 NEUROPHYSIOLOGY OF SPATIAL CONSTANCY

In the last fifteen years, neurophysiological studies have begun to uncover the neural mechanisms that contribute to spatial constancy (Colby and Goldberg, 1999; Merriam and Colby, 2005). Single-unit recording studies in awake behaving monkeys have shown that neurons in the lateral intraparietal area (area LIP) update representations when the eyes move (Duhamel et al., 1992a). Area LIP is a cortical area located in the posterior half of the intraparietal sulcus (Andersen et al., 1990a; Colby et al., 1988). The first demonstration of spatial updating in LIP was made while monkeys performed a simple single-step saccade task schematized in Figure 5. In this task, the monkey is instructed to make an eye movement from one fixation point (FP1) to another (FP2). The cue to make a saccade is the offset of FP1. Simultaneous with the offset of FP1, a stimulus is flashed outside of the receptive field of the neuron being recorded (Figure 5A). The location of FP2 is chosen by the investigator so that a saccade brings the receptive field of the neuron onto the location of the recently flashed stimulus (Figure 5B). The critical feature of the task is that the stimulus is flashed for a very brief period (50 ms). It is thus extinguished before the eyes actually begin to move — the stimulus is never physically present in the receptive field.

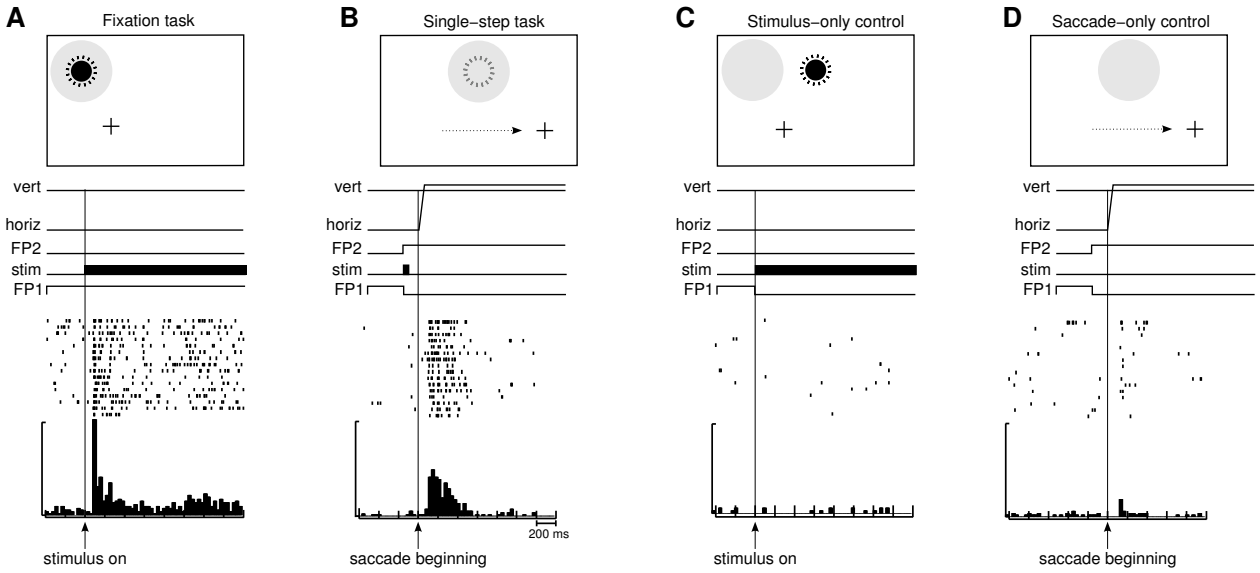

Figure 6. Neural activity in the single-step task and control conditions. In each series, the upper panel represents the spatial configuration used during the task; lower panel represents data from a single neuron. (A) Visual response in fixation task. The monkey fixates while a visual stimulus appears inside of its receptive field (shaded circle). The neuron fires briskly in response to the stimulus. (B) Single-step task. The monkey makes a single rightward saccade from FP1 to FP2, moving the neuron's receptive field onto the location where the stimulus had previously appeared. Critically, the stimulus has been extinguished by the time the monkey initiates the saccade to FP2. This means that there is no physical stimulus in the receptive field either before or after the saccade. The neuron responds with a strong burst of activity that begins in conjunction with the eye movement. (C) Stimulus-only control condition. The monkey's task is to maintain fixation at the intial fixation point (FP1) while the stimulus appears outside of the neuron's receptive field. Note that the stimulus location in this control condition matches the stimulus location in the single-step task (B). The neuron is not activated, indicating the stimulus cannot itself account for activation in the single-step task. (D) Saccade control task. The monkey makes the identical saccade to that used in the single-step task, but the stimulus is not presented. Because the saccade is directed away from the response field, the saccade alone does not drive the neuron. The neuron is active only when the stimulus and saccade occur in conjunction with one another, as they do in the single-step task. Histograms are aligned on saccade onset for the single-step and saccade-alone tasks, and on stimulus appearance for the stimulus-alone task. Rasters show activity on individual trials; each tic mark represents a single action potential. Vertical scale bar indicates a firing rate of 100 sp/s. (Modified from Duhamel et al., 1992a).

LIP is a visual area, and neurons in LIP respond strongly when a stimulus appears in the receptive field (Figure 6A). In the single-step task, the stimulus never appears in the receptive field. Thus, LIP neurons would not respond in this task if their responses were similar to those of a simple photoreceptor. Surprisingly, Duhamel et al. (1992a) observed that LIP neurons do respond in the single-step task (Figure 6B). The amplitude of the response is roughly half as large as the response when the stimulus is flashed in the receptive field. The raster plot in Figure 6B is aligned on saccade onset. Note that the neuron begins to respond around the time the eyes reach the new fixation point. This is the point at which the receptive field has moved onto the recently stimulated screen location. This location is referred to as the *new receptive field* to emphasize that the receptive field has shifted with the eye movement. Duhamel et al. (1992a) described activity in the single-step task as *remapping* to emphasize this shift. Remapping activity indicates that visual information has been updated from the coordinates of the initial eye position to the coordinates of the new eye position.

In the single-step task, a visual stimulus appears and the monkey makes an eye movement. Can neuronal activity in this task be attributed to either the stimulus or saccade alone? Responses in two control conditions demonstrate that it cannot. The neuron does not respond in a stimulus-only control condition in which the stimulus appears while the monkey fixates FP1 (Figure 6C). This is because the stimulus is placed in the opposite hemifield, well outside the bounds of the receptive field. Similarly, the neuron does not respond in a saccade-only control condition in which the monkey makes a saccade in the absence of the stimulus (Figure 6D). This is not surprising because the saccade is directed away from the receptive field, and into the ipsiversive hemifield. The neuron responds only if the two events occur together, as they do in the single-step task (Figure 6B). The critical factor is the location of the stimulus relative to the final eye position: it must appear at the retinotopic location of the receptive field following the eye movement. Because the stimulus alone and saccade alone do not generate any response, activity in the single-step task is interpreted as a response to the memory trace of the stimulus, which has been updated in conjunction with the saccade. Throughout this thesis, we refer to activity driven by the stimulus trace as a *remapped response*. We distinguish remapped responses from *visual responses* that are driven by direct retinal stimulation, as in Figure 6A.

The term "remapping" refers to the physiological phenomenon illustrated in Figure 6. Remapping is thought to contribute to both our stable perception of space during saccades and to our ability to generate spatially accurate behavior (Goldberg et al., 1990; Duhamel et al., 1992a; Colby et al., 1995). We typically think of cortical representations of visual space as retinotopic, as in area V1. However, responses to the stimulus trace in the single-step task demonstrate that the spatial representation in area LIP is not simply retinotopic. Rather, because LIP neurons maintain a representation of the updated coordinates of visual stimuli, we refer to the representation in LIP as *oculocentric*, or eye-centered (Colby et al., 1995). The distinction between a retinotopic and an eye-centered representation is as follows. A retinotopic representation codes the location of the stimulus with respect to its projected image on the retina. An eye-centered representation, on the other hand, codes the location of a stimulus relative to the distance and direction from the fovea. Eye-centered representations are therefore updated each time the eyes move. This is significant for spatial constancy because an eye-centered representation can be used by the oculomotor system to guide eye movements toward visual targets.

1.5 NEURAL CIRCUITRY FOR REMAPPING

The neuropsychological and neurophysiological evidence reviewed above indicates that parietal cortex, and area LIP in particular, is central for updating visual information in conjunction with saccades (Colby and Goldberg, 1999; Heide et al., 1995; Quaia et al., 1998). Remapping depends on the conjunction of a visual and a motor signal. In this section, we discuss a plausible scenario for how these signals reach parietal cortex.

LIP has a well characterized visual response (Blatt et al., 1990; Barash et al., 1991b,a; Colby et al., 1996; Bisley et al., 2004). There are at least three routes by which visual information reaches LIP. First, the main visual input to parietal cortex is from extrastriate visual cortex, and in particular, from area V3A. Responses in LIP resemble those in V3A, reflecting the strong anatomical connections between these two areas (Nakamura and Colby, 2000). Area LIP is also anatomically connected to nearly every other extrastriate visual area

(Cavada and Goldman-Rakic, 1989; Andersen et al., 1990a; Blatt et al., 1990; Morel and Bullier, 1990; Baizer et al., 1991), indicating that parietal cortex has access to visual representations at multiple stages of the visual hierarchy.

Second, visual information could reach LIP via intracortical connections with the frontal eye field (FEF). A large population of neurons in FEF have visual responses (Bruce et al., 1985; Goldberg and Bushnell, 1981; Mohler et al., 1973; Pigarev et al., 1979; Pouget et al., 2005; Schall, 1991; Thompson et al., 1996). Visual responses in FEF likely arise via strong connections with striate and extrastriate visual cortex (Jouve et al., 1998; Schall et al., 1995; Stanton et al., 1995). FEF and LIP are heavily interconnected (Petrides and Pandya, 1984; Andersen et al., 1990b; Stanton et al., 1995; Schall et al., 1995), suggesting that input from FEF could contribute to visual responses in LIP.

Third, a less prominent source of visual input to LIP arises from the superficial *visual* layers of the superior colliculus (Clower et al., 2001). Neurons in the superficial layers project to the pulvinar nucleus of the thalamus (Harting et al., 1980; Benevento and Standage, 1983). The dorsal portion of the lateral pulvinar, in turn, projects to area LIP (Asanuma et al., 1985; Hardy and Lynch, 1992; Baizer et al., 1993). Neurons in both the superficial layers of the colliculus (Wurtz and Albano, 1980) and the pulvinar are driven by visual stimuli (Bender, 1981; Petersen et al., 1985), suggesting that this route is a potential source of visual information for LIP.

An analysis of visual response latencies in LIP also suggests that multiple sources of input contribute to visual activity. LIP neurons have a brisk, short-latency (50 ms) *on* response that is followed by a lower-amplitude sustained response (Bisley et al., 2004). Very short latency responses have also been reported in FEF (Pouget et al., 2005; Schmolesky et al., 1998). Early responses in these areas have approximately the same latency as the earliest responses in extrastriate cortex (Schmolesky et al., 1998). It is currently not known whether these early responses are driven by direct input from primary visual cortex or by input from the collicular pathway.

Remapping also depends on information about impending eye movements. How do corollary discharge signals reach LIP? In a series of elegant studies, Sommer and Wurtz (2001, 2002, 2004) have demonstrated that corollary discharge signals arrive in cortex via a pathway

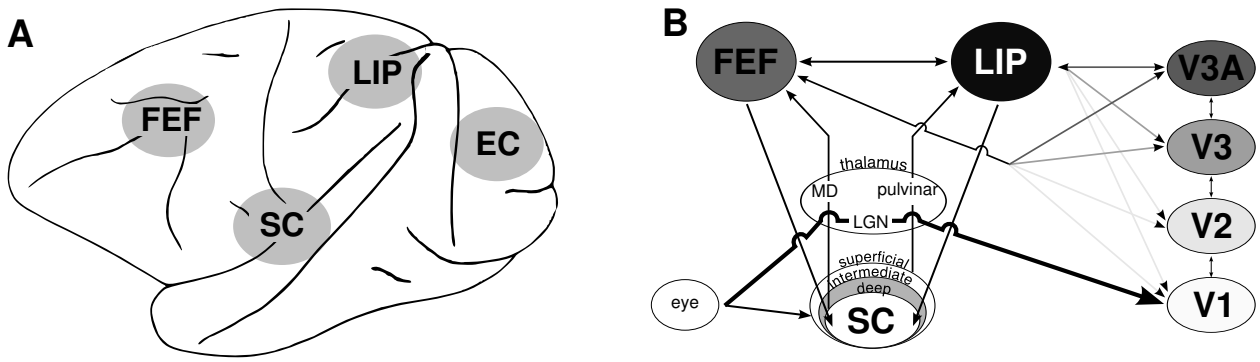

Figure 7. Neural circuity for remapping. (A) Regions in the monkey brain that exhibit remapping. (B) Pattern of anatomical connectivity. Grayscale shading indicates the proportion of neurons in a given area that exhibit remapping. FEF, the frontal eye field; LIP, the lateral intraparietal area; SC, the superior colliculus (deep, intermediate, and superficial layers); EC, extrastriate visual cortex; MD, medial dorsal nucleus; LGN, lateral geniculate nucleus

from the superior colliculus to the FEF. Inactivation of this pathway disrupts both performance on the double-step task (Sommer and Wurtz, 2002) and remapping activity in the frontal eye field (Sommer and Wurtz, 2005). Corollary discharge signals are conveyed from the intermediate oculomotor layers of the colliculus to FEF via projections from the medial dorsal nucleus of the thalamus (Lynch et al., 1994; Sommer and Wurtz, 2002). Corollary discharge signals could in turn be relayed from FEF to LIP through the dense intracortical connections between these two areas (Petrides and Pandya, 1984; Andersen et al., 1990b).

The above discussion assumes that updated visual signals are computed locally in area LIP. However, remapping has been observed in multiple brain regions (Figure 7). A large proportion of visual neurons in the frontal eye field (60%) and the intermediate layers of the superior colliculus (30%) exhibit remapping (Goldberg et al., 1990; Umeno and Goldberg, 1997, 2001; Walker et al., 1995). Remapping has also been observed in several extrastriate visual areas (Nakamura and Colby, 2002). The majority of neurons in V3A (52%) exhibit remapping. Remapping is less prevalent in area V3 (35%), V2 (11%) and V1 (2%). The contribution of each of these areas to spatial updating is largely unresolved. It is unknown whether visual signals are updated independently in each of these areas or whether the computation is carried out in LIP and the results are exported to other areas. Neuropsychological evidence supports the notion that updated signals are computed in LIP. Parietal

lesions in humans and reversible LIP inactivation in monkeys interfere with performance of the double step task, whereas frontal lesions do not, suggesting that LIP is critical for remapping to occur (see Section 1.3; Duhamel et al., 1992a; Heide et al., 1995; Li and Andersen, 2001). Anatomical connectivity suggests that updated signals could be sent from LIP to each of the areas in which remapping has been observed. As described above, LIP projects to the intermediate layers of the superior colliculus, to the frontal eye field, and to extrastriate and striate visual cortex (Lynch et al., 1985; Andersen et al., 1990b; Schall et al., 1995; Stanton et al., 1995). Lastly, the prevalence of remapping in extrastriate cortex reflects the degree of anatomical connectivity with LIP (Stanton et al., 1995), suggesting that remapping in these regions depends on the feedback of updated visual signals from LIP.

1.6 PHYSIOLOGICAL EVIDENCE FOR REMAPPING IN HUMANS

In the previous sections, we discussed lesion and inactivation studies demonstrating that the visual system uses corollary discharge signals to construct a stable representation of the world. We then described a neural mechanism in monkeys — remapping — that contributes to perceptual stability. The present experiments test the hypothesis that remapping occurs in the human brain as well. In this section, we review physiological evidence from humans that supports this hypothesis. We compare cortical responses in humans measured using functional imaging and evoked potentials with neural responses recorded from awake behaving monkeys. In this section, we focus on the neurophysiology of human parietal cortex, as neurophysiological evidence indicates that parietal cortex is critical for spatial updating to occur.

Responses in area LIP are multifaceted: LIP neurons respond in relation to visual stimuli, saccadic eye movements, spatial attention, working memory, and expected reward (Colby et al., 1996; Barash et al., 1991b; Goldberg et al., 2002; Roitman and Shadlen, 2002; Dorris and Glimcher, 2004; Andersen, 1997; Snyder et al., 1997). There is a substantial imaging literature indicating that human parietal cortex exhibits many of these same response properties as monkey LIP. In each of the following sections, we highlight similarities

and differences between monkey and human parietal cortex. Such similarities have been cited by others to argue for cortical area homology (e.g., Sereno et al., 2001; Sereno and Tootell, 2005). Our argument here is that the physiological similarities indicate that human parietal cortex could support remapping activity. We emphasize that this point is orthogonal to the issue of homology. Finally, we review the literature on extrastriate visual cortex and frontal eye field. We argue that human extrastriate cortex is a likely candidate for remapping activity, although human frontal eye field is not.

1.6.1 Visual-responsiveness and remapping

Nearly all LIP neurons respond to simple visual stimuli presented in the receptive field (Colby et al., 1996; Bisley et al., 2004). Human imaging studies have not emphasized the role of parietal cortex in visual processing. Early human imaging studies using positron emission tomography (PET) did not report any visual activity in the parietal lobe (Fox et al., 1986, 1987; Zeki et al., 1991). Instead, these studies focused on visual responses in striate and extrastriate cortical areas. Later functional imaging studies using fMRI also did not report parietal activation in visual stimulation paradigms (Engel et al., 1994, 1997; DeYoe et al., 1996; Sereno et al., 1995), although these studies used surface coils to achieve greater sensitivity near the occipital pole at the expense of sensitivity in the parietal lobes. Does human parietal cortex respond to visual stimuli? Recent functional imaging studies, and the work presented here (Chapter 2), indicate that it does, although the amplitude of visual activity in parietal cortex is smaller than in striate and extrastriate visual cortex (Silver et al., 2005). This is a significant observation because a cortical area that represents updated stimulus traces should also respond to visual stimuli.

1.6.2 Saccade-related activity and remapping

Many LIP neurons respond around the time of an eye movement. Responses can occur before, during, or after the monkey has made a saccade toward the neuron's response field (Barash et al., 1991b,a; Colby et al., 1996). Saccade-related responses in LIP are typically smaller in amplitude than visual responses. Saccade-related responses are also less common:

approximately half of LIP neurons exhibit activity around the time of an eye movement, whereas nearly all LIP neurons respond to visual stimuli.

Is human parietal cortex activated by saccades? Early functional imaging studies in humans revealed multiple activation foci in the frontal lobe and basal ganglia in conjunction with the performance of visually-guided saccades (Melamed and Larsen, 1979; Fox et al., 1985; Petit et al., 1993). Surprisingly, these studies failed to observe any activation in parietal cortex. A number of subsequent functional imaging studies have observed parietal activation in conjunction with visually-guided saccades, however (Luna et al., 1998; Berman et al., 1999; Darby et al., 1996; Merriam et al., 2001; O'Driscoll et al., 1995; Sweeney et al., 1996; Beauchamp et al., 2001; Anderson et al., 1994). The distribution of saccade-related activity in humans has been of considerable recent interest. Koyama et al. (2004) scanned both monkeys and humans during the performance of a visually guided saccade task. Saccades activated area LIP in monkeys (see also Baker et al., 2005), and in the superior parietal lobe in humans. These studies suggest that the parietal lobe in both monkeys and humans is activated during saccades.

Saccade-related activity in LIP raises two important points. First, activity associated with saccades is distinct from activity associated with remapping. In the single-step task, the monkey makes a saccade to a new fixation point (Figure 6B). This is not the optimal saccade for the neuron because it is not directed toward the neuron's receptive field. In fact, this saccade made in the absence of a stimulus does not elicit a response (Figure 6D). Thus, saccade-related activity *per se* does not account for remapping.

Second, activity in a visually-guided saccade task in the human functional imaging literature does not necessarily implicate parietal cortex in the generation of oculomotor behavior. fMRI activation in a visually-guided saccade task may be driven by visual input. In many of the studies listed above, the saccade target was a suddenly-appearing visual stimulus. Thus, the saccade target likely activated visually-responsive parietal neurons. Furthermore, visual reafference associated with fixating the target, or from a visible environment, could drive neurons with receptive fields that encompass the fovea. Lastly, activation in visually-guided saccade tasks may be related to remapping. The saccade target is by its nature a salient stimulus. It is thus likely that the saccade target is remapped in conjunction with the eye

movement. Remapping of the saccade target (or the even of the previous fixation spot) may contribute to activation in the saccade tasks described above. This is an important point, because it is often stated in the literature that parietal activation in saccade-tasks implicates this region in oculomotor processing, which is not necessarily the case.

1.6.3 Spatial attention and remapping

Neurons in area LIP are strongly modulated by the locus of attention (Colby et al., 1996). Attention directed toward the neuron's receptive field causes an increase in both baseline firing rate and visual response magnitude (Colby et al., 1996; Bushnell et al., 1981). Furthermore, temporary inactivation of LIP impairs performance on a covert attention task (Wardak et al., 2004). Attention and remapping are closely linked phenomena (Gottlieb et al., 1998). In this section, we review evidence that human parietal cortex is involved in spatial attention.

Parietal activation in spatial attention tasks was reported in the earliest PET studies (Corbetta, 1993; Nobre et al., 1997; Vandenberghe et al., 1997), and there has been an exponential increase in the number of imaging studies demonstrating attention-related parietal activation since the advent of fMRI in the mid-1990's (see reviews in Behrmann et al., 2004; Kanwisher and Wojciulik, 2000; Culham and Kanwisher, 2001; Corbetta et al., 2002; Pessoa et al., 2003). Some studies have emphasized the similarities between the role of human and monkey parietal cortex in spatial attention. For example, peripheral attention to a visual stimulus in the absence of eye movements activates both LIP neurons in monkeys (Colby et al., 1996) and the medial intraparietal sulcus in humans (Silver et al., 2005). In both monkeys and humans, attention-related responses in parietal cortex follow a rough spatial topography: responses are strongest for attention to contralateral locations.

Other studies have reported a widespread attention-related activation in humans that appears to be dissimilar from that observed in monkeys. For example, several authors have identified multiple foci of activation in spatial attention tasks and have attributed these activations to a fronto-parietal-temporal attention network (review in Corbetta et al., 2002). Current psychological theories attribute specific functions to different nodes in the network.

For example, top-down attentional mechanisms have been localized to the superior parietal lobule while bottom-up attention mechanisms have been localized to a right-lateralized ventral parietal region, close to the parietal-temporal lobe border (Corbetta et al., 2002). In support of this view, many studies have found that salient stimuli activate a region in ventral parietal cortex (Corbetta et al., 2000; Serences et al., 2005; Arrington et al., 2000; Marois et al., 2000; Downar et al., 2000, 2001, 2002). In contrast, tasks that require the control of attentional resources activate the superior parietal lobule (Kastner et al., 1999a; Shulman et al., 1999; Corbetta et al., 2000; Hopfinger et al., 2000).

We have thus far shown that human parietal cortex is activated by spatial attention. What is the relationship between remapping and spatial attention? Gottlieb et al. (1998) tested the effects of attention on remapping in area LIP. In this experiment, monkeys performed a variant of the single-step task in which a stable array of stimuli remained on the screen for the duration of the experiment. The monkey was cued to make a saccade to the center of the array, bringing the receptive field of the neuron under study onto one of the targets in the stable array. LIP neurons responded in this task only if the stimulus was rendered salient by a brief flash or if the stimulus was behaviorally relevant. The conclusion of this study is that LIP only remaps visual stimuli that are attended by the animal. Attention may be drawn to the stimulus by attributes that engage *bottom-up* attentional mechanisms (e.g., a sudden onset, as in Jonides and Yantis, 1988), or attention may be directed toward the stimulus by *top-down* attentional mechanisms, such as task relevance. That LIP only remaps salient visual stimuli makes sense from a computational perspective — it would be inefficient to remap all visual input. Despite the close link between attention and remapping, the two phenomena are distinct. Covert attention in the absence of an eye movement does not result in remapping of a stimulus trace (Colby, 1996). This observation indicates that a corollary discharge command is necessary for remapping to occur.

1.6.4 Spatial working memory and remapping

About a third of neurons in monkey area LIP carry memory signals: they are active in the delay period between the appearance of a target and the later initiation of a saccade toward

that target (Gnadt and Andersen, 1988). Spatial working memory and remapping are closely intertwined processes. In the single-step task, LIP neurons respond to the updated stimulus trace (Figure 6B). There is no stimulus present on the screen when the neuron responds. Thus, activity in the single-step task is mnemonic in nature. Functional imaging studies of spatial working memory in humans have focused on the role of dorosolateral prefrontal cortex (review in Curtis and D'Esposito, 2003). Many of these experiments have also shown sustained parietal activation during memory delays (Curtis et al., 2004; Pessoa et al., 2002a). Memory-related activity in human parietal cortex indicates an important similarity in function between humans and monkeys. Furthermore, because the stimulus trace response in the single-step task can be considered a mnemonic signal, the presence of delay-period activity in humans suggests that human parietal cortex may also play a role exhibit remapping.

1.6.5 Time course of remapping

Remapping occurs in a relatively wide temporal window around the time of the saccade. In area LIP, a substantial proportion of neurons (44%) respond at a latency that is predictive (Duhamel et al., 1992a). In these neurons, the response to the stimulus trace has a shorter latency than a typical visual response during fixation. A subset of these neurons respond even before the eyes have begun to move. Predictive remapping has been observed in other cortical areas as well, including the intermediate layers of the SC (30%), the FEF (31%), and area V3A (70%), indicating that the computations involved in remapping begin in anticipation of the actual saccade (Nakamura and Colby, 2002; Walker et al., 1995; Umeno and Goldberg, 1997). Anticipatory responses may enable the visual system to represent space in the coordinates of the next fixation, even before new afferent signals have time to reach cortex. This predictive representation could underlie our perception of spatial continuity across saccades.

While some neurons exhibit predictive remapping, a large population of neurons respond relatively late in time relative to saccade onset. Heiser and Colby (*in press*) constructed population histograms of LIP responses in the single-step task. Across the population, responses were clearly evident prior to saccade onset, consistent with predictive remapping. However,

the most robust response at the population level occurred around 100 ms *after* the saccade. This result indicates that remapping is not instantaneous. Rather, the representation of the updated stimulus takes tens of milliseconds to compute.

Are responses in humans consistent with these temporal dynamics? Responses observed by functional imaging are an order of magnitude slower than the time course of neural events that occur in remapping. Because of the sluggishness of the hemodynamic response measured with functional imaging, the temporal dynamics of spatial updating in humans can best be studied using evoked cortical potentials. Most of the ERP work on saccades have focused on the phenomenon of saccadic suppression. Several studies have shown that the visually-evoked potential (VEP) is suppressed around the time of a saccade (Anagnostou et al., 2000; Skrandies and Anagnostou, 1999; Kleiser and Skrandies, 2000; Skrandies and Laschke, 1997), consistent with single-unit studies showing saccadic suppression in early visual cortex. In a recent study, Bellebaum et al. (2005) recorded evoked potentials in humans during performance of the double-step task, and found that the largest response occurred 100 ms after the eye movement, consistent with Heiser and Colby (*in press*). This correspondence indicates that there may be a meaningful relationship between the time course of neural remapping signals in parietal cortex and the modification of signals in early visual areas.

1.7 EXPERIMENTAL AIMS

1.7.1 Remapping in human parietal cortex

Does remapping occur in human parietal cortex? We have discussed evidence in the preceding sections suggesting that it does. Lesions of human parietal cortex result in behavioral deficits on spatial updating tasks. Furthermore, human parietal cortex exhibits many of the physiological response properties — visual, saccade, and cognitive — that are present in monkey LIP. Both of these lines of evidence suggest that human parietal cortex could support remapping activity. The goal of Chapter 2 is to test the hypothesis that it does, and that it can be visualized with fMRI.

1.7.2 Bayesian inference applied to remapping

Analysis of functional imaging data has typically relied on a frequentist approach to statistical inference. Although the frequentist approach has been successfully applied to data sets collected under a broad range of circumstances, the Bayesian approach can offer several advantages in flexibility and performance. The goal of this chapter is to describe a Bayesian method for analyzing fMRI time series data. We apply the model to data from a remapping experiment, and compare the performance of the model to the performance of a more conventional general linear model.

1.7.3 Remapping in human extrastriate cortex

We have focused our discussion thus far on remapping in area LIP. Neurophysiological studies have identified several cortical and subcortical areas where neurons exhibit remapping. The goal of this chapter is to test the hypothesis that remapping occurs in human extrastriate cortex. Specifically, we use the Bayesian model described in Chapter 2 to determine (i) whether remapping exists in human extrastriate cortex, and (ii) whether there are differences between cortical areas in the strength of remapped responses.

2.0 REMAPPING IN HUMAN PARIETAL CORTEX

2.1 OVERVIEW

The vast majority of neurons in monkey parietal cortex update visual information in conjunction with eye movements. This remapping of stimulus representations is thought to contribute to spatial constancy. In this experiment, we hypothesized that remapping occurs in human parietal cortex and that it can be visualized using functional MRI. Neurophysiological experiments in monkeys have studied remapping using the single-step saccade task. These single-unit experiments require detailed information regarding the spatiotemporal responses of visual neurons. fMRI does not afford equivalent resolution. To overcome this limitation, we adapted the single-step task for use with fMRI. In this task, subjects view a visual stimulus for 1 s while maintaining fixation. The stimulus then disappears and subjects are cued to make a large horizontal eye movement. This saccade brings the stimulated screen location into the opposite visual field. This modified version of the single-step task enabled us to visualize the interhemispheric transfer of the stimulus representation in conjunction with the saccade. We observed an initial response in the hemisphere contralateral to the visual stimulus, followed by a remapped response in the hemisphere ipsilateral to the stimulus. We ruled out the possibility that this remapped response resulted from either eye movements or visual stimuli alone. Rather, we found that remapping depended on the conjunction of a visual stimulus and a saccade. Our results demonstrate that updating of visual information occurs in human parietal cortex.

These results have been published in abstract form (Merriam and Colby, 2001) and in a full length journal article (Merriam, Genovese, and Colby, 2003).

2.2 INTRODUCTION

The idea that visual perception is an active process has attracted considerable interest. Nowhere is the active nature of vision more evident than in the construction of a stable image of the world. As we move our eyes, new images are constantly presented to the brain, yet we perceive the world as remaining still. This spatial constancy suggests that the act of making an eye movement changes the brain's representation of visual information. Visual and motor signals interact to construct an internal representation that is constantly updated and spatially stable (Colby and Goldberg, 1999).

Neurons in monkey parietal cortex participate in updating visual signals. Parietal neurons have retinotopic receptive fields that move with the eyes. In the remapping paradigm, when the eyes move so that the receptive field of a neuron lands on a previously stimulated screen location, the neuron fires even though the stimulus is no longer present (Duhamel et al., 1992a). This response to the trace of the stimulus indicates that a transformation in the neural representation has occurred: neurons that initially encoded the location of the visual stimulus have transferred their information to the neurons that will represent the location of the stimulus after the eye movement. The representation of visuospatial information is thereby remapped, or updated, in conjunction with the eye movement. A corollary discharge of the eye movement command is thought to trigger remapping. Updating creates a stable representation of space by compensating for the displacement of objects on the retina.

We hypothesized that spatial updating also occurs in humans. Behavioral results in humans and nonhuman primates have shown that they have similar abilities in double-step eye movement tasks that require the use of updated visual information (Baizer and Bender, 1989; Hallett and Lightstone, 1976). Moreover, the parietal lobe is critical for task performance. Humans with parietal lobe damage are unable to perform double-step tasks (Duhamel et al., 1992b; Heide et al., 1995), and parietal neurons in monkeys are specifically active in these tasks (Goldberg et al., 1990). We thus hypothesized that updating would produce physiological activity in human parietal cortex and that we would be able to visualize it using fMRI. We were encouraged in this endeavor by a previous human fMRI study that tested

subjects in a conceptually related triple-step eye movement task (Heide et al., 2001). While these authors found activation in multiple cortical sites, they suggested that parietal cortex in particular was related to the updating component of the task. For these reasons, we focused our data acquisition and analysis on parietal cortex.

We scanned subjects while they performed an eye movement task that is directly analogous to the task used to demonstrate remapping in monkeys. The sequence of events in our spatial updating task is shown in Figure 8A. At the beginning of the trial, the subject fixated a stable target (right cross). A visual stimulus appeared at the center of the screen and remained on for 2 seconds. The stimulus disappeared at the same time that a tone cued the subject to make a leftward saccade to a stable target (left cross). This eye movement brought the previously stimulated screen location into the right visual field. No physical stimulus was present in the right visual field at any time during the trial.

Previous studies have shown that visual stimuli activate extensive regions of contralateral visual and parietal cortex (DeYoe et al., 1996; Engel et al., 1997; Sereno et al., 1995). We hypothesized that ipsilateral parietal cortex would also become activated when an eye movement brought the stimulus location into the opposite visual field, even though the stimulus had already disappeared by the time of the eye movement. In the example shown in Figure 8A, we expected activation to shift from the right to the left hemisphere when the eyes moved, despite the fact that no physical stimulus ever appeared in the right visual field. In addition, we predicted that the remapped activation would appear later than the visually evoked activation, because the eye movement was cued 2 seconds after the onset of the visual stimulus (Figure 8B). Based on previous single-neuron studies in monkey (Duhamel et al., 1992a), we also expected that activation due to remapping would be smaller in amplitude than the visually evoked activation. Finally, we predicted the opposite pattern of activation when a right visual field stimulus was followed by a rightward saccade (Figures 8C and 8D).

2.3 METHODS

2.3.1 Stimuli and task conditions

Subjects were scanned during three experimental conditions:

2.3.1.1 Spatial updating condition. In the spatial updating condition, subjects fixated one of two stable crosses located at $+8°$ and $-8°$. A flickering stimulus appeared at the center of the display screen. The stimulus was a white spot, the size and flicker frequency of which changed randomly (ranging from $0°$ to $2°$, and from 6 to 10 Hz). The stimulus fell in either the right or left visual hemifield ($8°$) and remained on the screen for 2 seconds. At the time of stimulus offset, a binaurally presented tone cued the subject to make a saccade to the fixation cross located on the opposite side of the screen. This eye movement caused the screen location where the stimulus had appeared to enter the opposite visual hemifield. There were two types of trials in the spatial updating condition: trials in which the stimulus appeared in the right visual field and was followed by a rightward saccade, and trials in which the stimulus appeared in the left visual field and was followed by a leftward saccade.

2.3.1.2 Stimulus-only condition. In the stimulus-only condition, the subject maintained fixation on either the right or left cross, and the stimulus appeared for 2 s at the center of the screen. There were two types of trials in the stimulus-only condition: trials in which the stimulus appeared in the right visual field, and trials in which the stimulus appeared in the left visual field.

2.3.1.3 Saccade-only condition. In the saccade-only condition, the subject made saccades to stable crosses when prompted by a binaurally presented auditory cue, but no visual stimulus appeared on the screen. There were two types of trials in the saccade-only condition: trials in which subjects made a rightward saccade, and trials in which subjects made a leftward saccade.

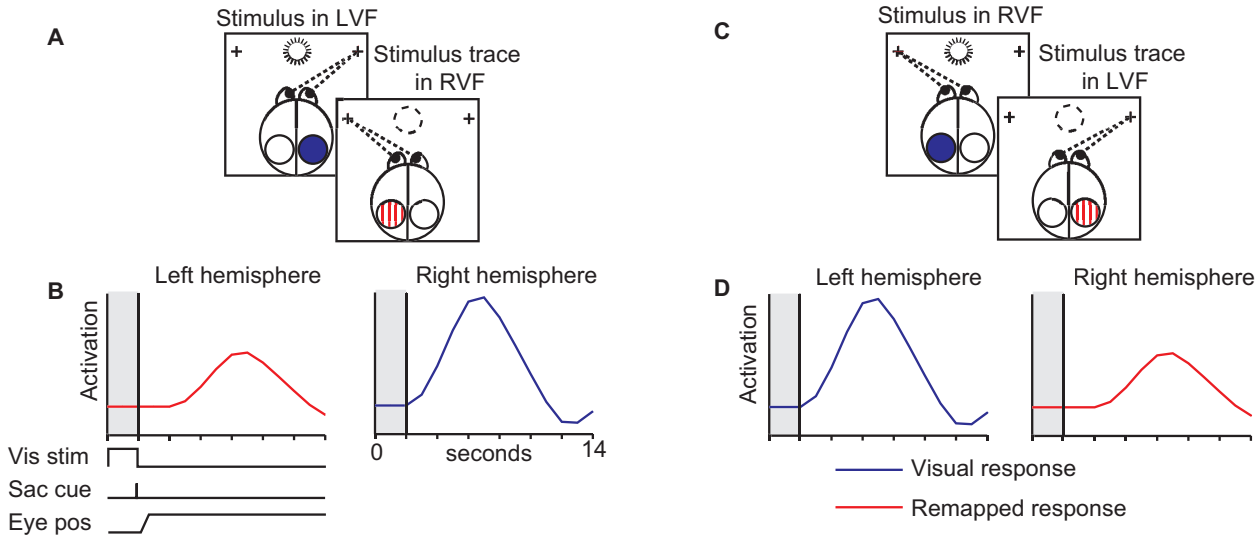

Figure 8. fMRI version of the single-step task and predicted results. (A) Sequence of events in the spatial updating condition. The stimulus appears in the left visual field at the beginning of the trial and remains on the screen for 2 seconds. We expected the stimulus to activate right hemisphere occipital and parietal cortex (blue circle). Simultaneously, the stimulus disappears and a tone cues the subject to make a leftward eye movement. This saccade brings the screen location of the now-extinguished stimulus (dotted circle) into the right visual field. We predicted that remapping of the stimulus trace would cause activation to shift from the right to the left hemisphere (red hatched circle). (B) Predicted time course of activation. The shaded region indicates the time that the stimulus is on, and the vertical line at 2 seconds indicates the time of the auditory cue to make an eye movement. Activation in the right hemisphere, due to the stimulus, was expected to follow the standard hemodynamic time course (blue curve). Activation in the left hemisphere, due to the remapped stimulus trace, was expected to occur with a similar time course but shifted by 2 seconds because the cue to make an eye movement occurs 2 seconds after stimulus onset. We also expected the remapped response to be smaller in amplitude than the visual response. (C and D) The spatial updating condition and predicted results on trials in which the stimulus appears in the right visual field and is followed by a rightward eye movement. Note that the expected pattern of activation is a mirror reflection of that described in (A) and (B). In this and all subsequent figures, shades of blue represent visual responses and shades of red represent remapped responses.

Visual and auditory stimuli were presented using CORTEX software running on a PC (`http://www.cortex.salk.edu`). Visual stimuli were back projected onto a Lucite screen using a 3-panel LCD projector. Auditory stimuli were presented using pneumatic headphones.

2.3.2 Experimental design

The particular ordering and spacing of trials in fMRI experiments is known to have a profound impact on both the nature of the evoked MR response and the analysis methods used to analyze the data (Aguirre and D'Esposito, 2000). There are costs and benefits associated with different design options. We used two different experimental designs to enhance our ability to detect remapped response, to control cognitive factors, and to estimate the fine temporal aspects of the responses.

Theoretical work indicates that randomization of both trial order and interstimulus interval enhances the ability to estimate the shape of the MR response (Birn et al., 2002; Liu et al., 2001). We tested half of the subjects (n = 4) using a randomized design in which three conditions occurred in an interleaved fashion. On a given trial, subjects either viewed a peripheral stimulus while maintaining fixation (stimulus-only condition), or made an eye movement in the absence of a stimulus (saccade-only condition), or made an eye movement after the stimulus disappeared (updating condition). We randomized both the trial order and the interstimulus interval within the ordering constraints of the task (e.g., subjects never performed two successive eye movements in the same direction). Within each 480 second run, subjects performed an average of eight trials of each condition and the interstimulus interval varied from 6 to 15 seconds. The interleaved design ensured that cognitive factors and motor preparation were held constant across the different conditions. This design also ensured that variables such as slow drifts in the MR signal and head motion were evenly distributed across task conditions. Our analysis allowed us to calculate an impulse-response function for each condition independently even though the MR responses overlapped in time.

It was also critical for our study to be able to detect small temporal differences between the visual and remapped responses. In periodic designs, the power of the signal is concen-

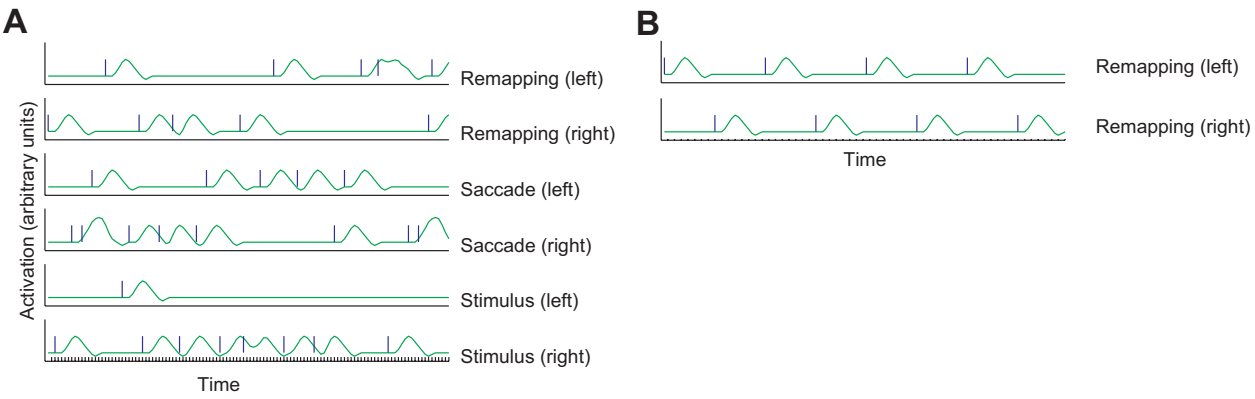

Figure 9. Experimental design. The x-axis represents time (seconds); the y-axis represents MR response amplitude. Vertical lines indicate trial times. Green curves represent hypothetical BOLD responses. (A) Interleaved design in which the six trial types were intermixed. The order of trial types was constrained by the requirement that two successive trials did not involve two eye movements in the same direction. (B) Periodic design. Trials were separated by 15 s, and alternated between remapping right and remapping left. Both the stimulus-only and saccade-alone conditions (not shown) were also tested using periodic designs.

trated around the fundamental frequency of the task, rather than across the entire spectrum, as in randomized designs (Aguirre and D'Esposito, 2000). This type of design allows the use of Fourier analysis to estimate the time of the MR signal evoked by each response type (Saad et al., 2001; Sereno et al., 1995). Accordingly, we tested the other four subjects using a fixed interval, periodic design. Within each 551 second run, subjects fixated during the first 11 seconds. They then performed 36 trials. One trial occurred every 15 seconds. This task frequency confers optimal sensitivity for detecting brief events (Bandettini and Cox, 2000). In two control runs, subjects viewed a peripheral stimulus while maintaining fixation (stimulus-only condition). In another two control runs, subjects made eye movements in the absence of a stimulus (saccade-only condition). In an additional two to four runs, subjects made an eye movement after the stimulus disappeared (updating condition).

2.3.3 MR data collection

Eight right-handed subjects, ages 23–30, participated. All subjects had been scanned previously for other studies and were highly experienced in performing oculomotor tasks in an

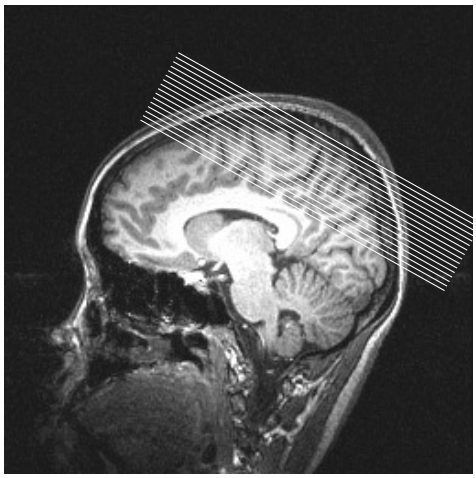

Figure 10. Slice prescription. Slices were angled 30° from the horizontal (axial) plane so as to cover the entire occipital lobe.

MR environment. The protocol for this study was approved by the IRB at the University of Pittsburgh and informed consent was obtained from all subjects. MR images were collected on a GE Signa 3 Tesla scanner. A high-resolution SPGR anatomical sequence was collected at the beginning of each scanning session: TR, 24 ms; flip angle, 30°; number of slices, 124; in-plane voxel dimensions, 0.98×0.98 mm; slice thickness, 1.0–1.5 mm, depending on head size.

We used two different pulse sequences to collect BOLD-sensitive images. For subjects scanned using the interleaved design, we collected images using an echo-planer pulse sequence developed by GE (EPI-RT) with a 2000 ms time to repetition (TR) and a 30 ms echo time (TE). Twenty oblique slices ($3.125 \times 3.125 \times 3.0$ mm, with a 1 mm gap) covered the entire parietal lobe, as well as superior portions of the occipital lobe, and posterior portions of the frontal lobes. In each run, we collected 240 images per slice. Subjects were scanned for six to eight runs each, resulting in at least 28,800 images per subject. For subjects scanned using the periodic design, we used a locally developed spiral scanning sequence to achieve higher temporal resolution (TR 1 s; TE 18 ms). Eighteen oblique slices ($3.125 \times 3.125 \times 3.0$ mm, with no gap) were positioned to cover the entire parietal lobe. In each run, we collected 551 images per slice. Subjects were scanned for six to eight runs each, resulting in least 59,508 images per subject. Subjects lay quietly in the dark 1–2 min between runs and were instructed to keep their eyes closed.

2.3.4 Analysis of fMRI data

Our approach to data analysis for both experiments involved three stages: (1) preprocessing and noise reduction; (2) voxel selection; and (3) response parameter estimation:

2.3.4.1 Preprocessing and noise reduction The data were preprocessed using locally developed FIASCO software (available at `http://www.stat.cmu.edu/~fiasco`, Eddy et al., 1996). Preprocessing steps included a correction for fluctuations in mean intensity; motion correction of the raw, complex-valued k-space data; image reconstruction; linear detrending; and outlier correction using a Windsor filter. Outliers were defined as data points farther than ten times the interquartile range from the median. The image data were not spatially smoothed.

2.3.4.2 Voxel selection The goal of this second stage of data analysis was to identify voxels that were candidates for exhibiting remapping activity. We used two criteria to select voxels. First, we applied anatomical criteria to define a region of interest in parietal cortex in each hemisphere. Second, we applied a functional criterion by selecting voxels that showed a significant response in at least one of the six conditions (contra or ipsilateral visual stimuli, saccades, or spatial updating). Voxels that fulfilled these criteria were considered to be task related.

We drew anatomically defined regions-of-interest (ROIs) on the high-resolution structural scans from each hemisphere using tools included in the FreeSurfer software package (Dale et al., 1999; Fischl et al., 1999). ROIs included the entire extent of the intraparietal sulcus. The ROI extended anteriorly to the segment of the sulcus that joins the postcentral sulcus, and extended posteriorly toward the occipital lobe, stopping at the "T" junction formed by the IPS and the transverse occipital sulcus (Figures 12A and 12B). The ROI extended outside the sulcus to include the adjacent gyral surface, as well as both medial and lateral side branches. ROIs were then resampled to the resolution of the functional data. This method allowed us to define the ROIs using the detailed anatomical information from the structural scans without resampling or distorting the functional data in any way.

We used multiple regression to detect voxels that showed significant responses to any task condition. Multiple regression was performed using software from the AFNI analysis package (`http://afni.nimh.nih.gov/afni`, Cox and Hyde, 1997). The detailed methods for this procedure have been described elsewhere (Ward, 1998). In short, the regression model included a separate regressor to model the signal at each time point in a 15 second window following the start of each trial. Activation maps were created by using a partial F-test that compared the variance accounted for by the regressors associated with each condition to the full model fit to the data. Activation maps were thresholded using a false discovery rate procedure that controls the rate of false positives while obviating the need for explicit correction for multiple comparisons (Genovese et al., 2002).

2.3.4.3 Response parameter estimation Once we identified a population of task-related voxels in each hemisphere, we carried out several analyses aimed at describing the properties of the visual and remapped responses. Raw time series data were analyzed using custom software written in MATLAB (MathWorks, Inc.) and in the S programming language implemented in the R computing environment (R Development Core Team, 2005).

BOLD-image rasters. We created BOLD-image raster plots (Duann et al., 2002) in order to display the variability of responses across trials (Figures 13B and 13C). Each line in the raster plot represents the response in a single trial, averaged over all task-related voxels in the ROI. Data were then converted to units of percent signal change by dividing by the mean. Each time point of the 15 second hemodynamic response was smoothed with a 2 second wide moving window. Data were not smoothed across trials. The time course data for each trial were represented as a series of horizontal bars, stacked sequentially in trial order, and pseudocolor coded for activation level. Conventional time series plots were created by averaging across the trials represented in the BOLD-image plots using the unsmoothed data (Figures 13A and 14).

Response Magnitude Analysis. We used the F-values from the multiple regression analysis described above to determine the significance of the visual and remapped responses. We

converted the F-values to Z-scores and then displayed the Z-score values for each hemisphere using standard boxplots (Figure 15). The population response in a given hemisphere was considered significant if the median Z-score exceeded a threshold of $p < 0.05$ (horizontal dotted line, Figure 15). We chose to use the median of the population because it is both interpretable and robust to outliers.

In order to compare the size of the remapped responses to the size of the other response types, we first calculated the minimum and maximum percent signal change in each voxel for each condition. We then created a single amplitude measure by calculating the difference between the minimum and maximum values. We compared the amplitudes of the remapped responses to the amplitudes of the visual responses, to the responses in the saccade-only condition, and to the responses in the stimulus-only condition (Figure 16). Comparisons were made for the same saccade direction and for the same stimulus location. We performed two analyses to determine the significance of the comparisons. In the first, within-hemisphere analysis, we subtracted the amplitudes elicited by the control conditions from the amplitudes elicited by the remapped responses. We then calculated the median and a 95% confidence interval of the median of these difference values. Within-hemisphere population comparisons were considered significant if the 95% confidence interval of the median did not intersect zero. In the second, across-hemisphere analysis, the median amplitudes of the remapped responses in each hemisphere were plotted against the median amplitudes of the visual responses (Figure 16A), ipsiversive saccade responses (Figure 16B), and ipsiversive stimulus responses (Figure 16C). We used conventional t-tests to test the significance of the comparisons across the population of hemispheres.

Fourier Analysis. It was critical to measure the latency of the visual and remapped responses from the same population of voxels because response latency is known to vary considerably from voxel to voxel in the brain (Saad et al., 2001). Temporal variability appears to be substantial even within a given visual area. This variability likely arises from differences in vascular physiology rather than from neuronal sources, and highlights the importance of making within-voxel comparisons. To measure the latency of the response, we took the FFT of the time series from each voxel. We then extracted the phase at the funda-

mental frequency of the task (one trial every 15 seconds) from the unsmoothed periodogram. To illustrate the differences in phase between the two conditions, we made scatter plots of all voxels in the ROI in polar coordinates, simultaneously representing both the phase (θ) and the magnitude (ρ) of the response (Figures 17A and 16B). The phase spectrum is a good estimate of the phase difference between two time series (Brockwell and Davis, 1991). We calculated the phase spectrum of the two signals and performed a t-test, comparing the phase spectra values of the population of voxels in each hemisphere to zero phase.

2.3.5 Eye movement recording

We monitored eye position with a video-based eye tracking system (ISCAN, Boston, MA) and analyzed the data using ILAB software (`http://groups.yahoo.com/group/ilab/`; Gitelman, 2002). We used the saccade finder utilities in ILAB to detect saccades and calculate their latencies. A saccade was registered if eye velocity exceeded 40°/s and the eyes moved more than 1.5°. We recorded saccade latency by calculating the time when the eyes reached 15% of their maximal velocity.

2.4 RESULTS

2.4.1 Eye movement analysis

To ensure that subjects performed the updating task accurately, we tested subjects outside the scanner and measured their eye position. Three issues were addressed by this test. First, it was important that subjects be able to maintain fixation while the stimulus flashed in the periphery. We tested subjects on either 36 or 72 trials, as in the scanned experiment. We calculated the mean and standard deviation of the eye position over all trials from all subjects (Figure 11A). There was no deviation in gaze caused by the presence of the stimulus. This assured us that the physical stimulus did not enter the opposite hemifield. Second, it was crucial that subjects made an eye movement at the appropriate time in response to the auditory cue. Average saccade latency was 279 ± 85 ms relative to the auditory cue. Subjects

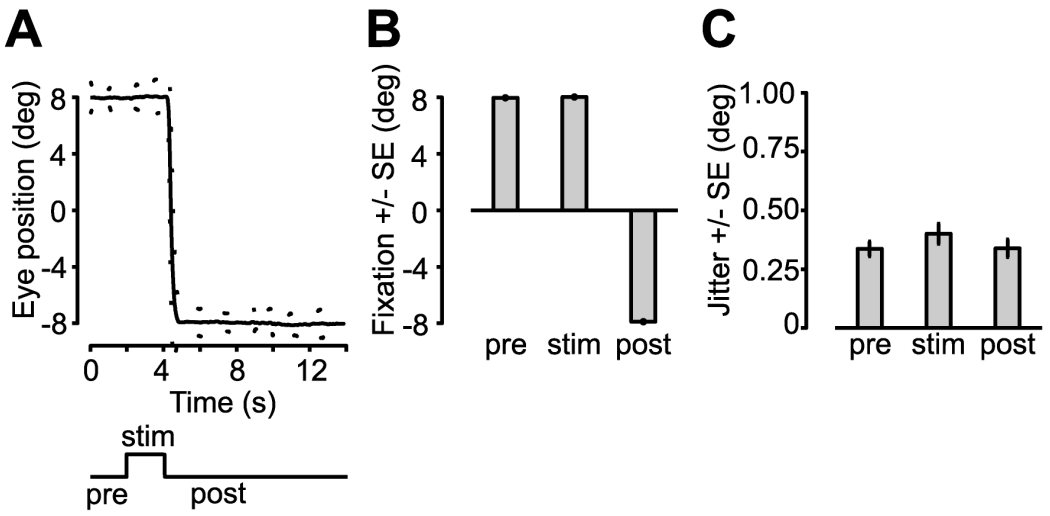

Figure 11. Eye position during single-step task performance. (A) Horizontal eye position over time averaged across all 450 trials from eight subjects. The dotted line indicates 1 SD of the mean. Traces from leftward eye movement were flipped and averaged with traces from rightward eye movements. (B) Mean eye position during each epoch averaged across all trials and subjects. (C) Mean jitter from the fixation point during each epoch averaged across all trials and subjects. Error bars indicate 1 SEM in (B) and (C)

never made a saccade prior to the auditory cue. Third, it was important that subjects be able to fixate after the eye movement for the remainder of the trial. As is shown in Figure 11A, subjects maintained gaze after the saccade.

In order to quantify these data, we divided the task into three epochs, each lasting 2 seconds: a "prestimulus" epoch consisted of the period prior to the onset of the visual stimulus; a "stimulus" epoch consisted of the period in which the stimulus was present, and a "postsaccade" epoch consisted of the 2 s period beginning 1 s after the auditory cue to make the saccade. Figure 11B shows that there was no difference in eye position between the prestimulus and stimulus epochs, indicating that subjects were able to maintain accurate fixation despite the presence of the peripheral stimulus. We also tested the amount of eye position jitter (the standard deviation of eye position across each epoch). While there was slightly more jitter in the stimulus epoch (Figure 11C), this difference was not significant ($t[318] = 1.54$, n.s.). These data indicate that subjects were able to perform this simple oculomotor task with a high degree of spatial and temporal accuracy.

2.4.2 Parietal voxels respond to visual stimuli and updated stimulus traces

We found evidence for remapping in the form of strong and consistent activation in the ipsilateral parietal lobe during the spatial updating task. We used anatomical criteria to define the borders of regions of interest (ROIs) that included the intraparietal sulcus and adjacent gyral cortex (Figures 12A and 12B; see Methods, Section 2.3.4.2). Within this ROI, we selected voxels for inclusion in our analysis if there was a significant response in at least one of the six types of trials: trials of the stimulus-only condition in which the stimulus appeared in the (1) contralateral or (2) ipsilateral visual field; trials of the saccade-only condition in which subjects made (3) contraversive or (4) ipsiversive saccades; and trials of the spatial updating condition in which the stimulus location was remapped to the (5) contralateral or (6) ipsilateral visual field. A large population of task-related voxels was identified in each of the 16 hemispheres tested. Figure 12C shows activation in a single subject on trials of the spatial updating condition in which the stimulus appeared in the left visual field and was followed by a leftward saccade. The left visual field stimulus resulted in strong visual responses in the right hemisphere. The leftward eye movement brought the stimulus location into the right visual field, resulting in remapped responses in left parietal cortex.

We illustrate our main results with data from a single subject who showed responses typical of the group (Figure 13). In a given hemisphere, we measured visual responses on trials of the spatial updating condition in which the stimulus appeared in the contralateral visual field and was then followed by a contraversive saccade. In the example illustrated in Figure 13A, visual responses (blue) were measured in the left hemisphere (left panel) on trials in which the stimulus appeared in the right visual field and was followed by a rightward saccade. The visual stimulus elicited a large, biphasic response: the signal first rose to a peak of 1% above baseline at 6 seconds after the onset of the stimulus. The response then quickly dipped to 0.5% below baseline at around 10 seconds poststimulus. The shape and magnitude of the visual response correspond well with previous reports of BOLD activation evoked by contralateral visual stimuli (Boynton et al., 1996).

The updating task was designed so that this initial visual response would be followed about 2 seconds later by a remapped response in the opposite hemisphere. We measured

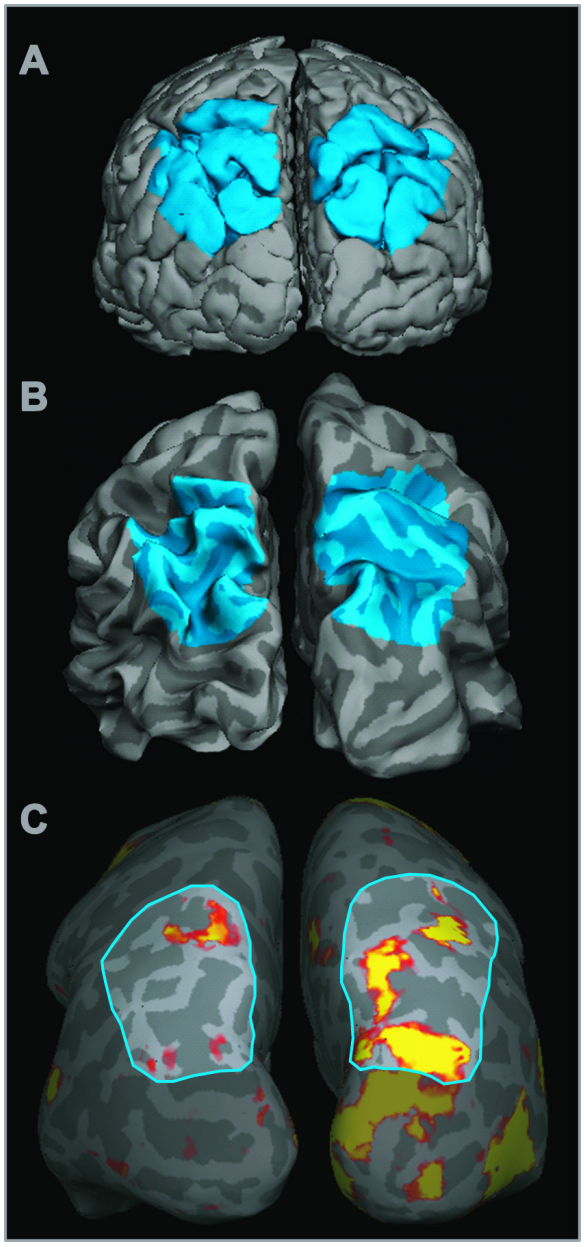

Figure 12. Region of interest in parietal cortex. (A) Posterior view of both hemispheres of a single subject rendered at the outermost layer of gray matter. The regions of interest are shown in blue. (B) Partially unfolded view of the same two hemispheres. Blue shading indicates the location of the ROI. Shades of gray indicate the curvature of the cortical surface: dark gray indicates concave areas, and light gray indicates convex areas. (C) Activation from a single subject on updating trials in which a left visual field stimulus was followed by a leftward saccade. This condition elicited activation in contralateral (right) hemisphere occipital and parietal areas, as expected. Activation was also observed in the ipsilateral (left) parietal lobe, indicating that the visually evoked activation was remapped in conjunction with the eye movement.

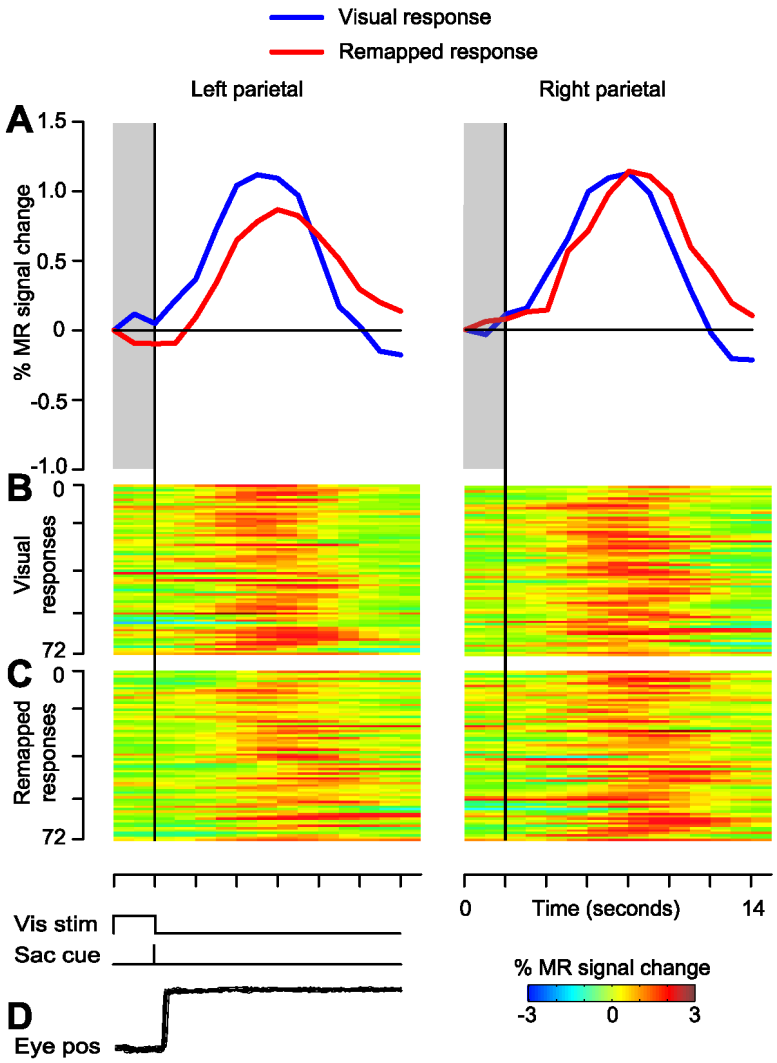

Figure 13. Visual and remapped responses from both hemispheres of a single subject. (A) Time course of activation evoked by the visual and remapped responses from parietal cortex. The MR time course over the 15 second task epoch represents the average of 72 trials from 49 voxels in the left hemisphere and 33 voxels in the right hemisphere. The shaded gray bar indicates when the stimulus was present, and the vertical line at 2 seconds shows the time of the auditory cue to make a saccade. The remapped response (red line) occurs later and is smaller than the visual response (blue line). (B) BOLD-image raster plots of the visual responses from the same hemispheres for 72 successive trials. Activation on individual trials is plotted along the y-axis, with percent signal change represented in pseudocolor plotted over time (x-axis). (C) BOLD-image raster plots of the remapped responses for 72 successive trials. (D) Eye position recorded in 36 trials of the same task performed outside the scanner.

remapped responses on trials of the spatial updating condition in which an ipsilateral stimulus was followed by an ipsiversive saccade. In the example illustrated in Figure 13, remapped responses (red) were measured in the left hemisphere (left panel) on trials in which the stimulus appeared in the left visual field and was followed by a leftward saccade. Remapped responses differed from visual responses in several respects. In this example, the remapped responses had later rise times than the visual responses, as expected. The remapped responses shown here also had lower peak amplitudes and slower returns to baseline. The BOLD-image raster plots indicate that these remapped responses were both present and robust throughout the course of the scanning session (Figures 13B and 13C).

A similar pattern of activation was present in all 16 hemispheres (Figure 14). We observed strong visual responses to stimuli in the contralateral hemifield in every hemisphere. We also observed consistent remapped responses. The response curves in Figure 14 are the average of all task-related voxels in each hemisphere. We used a partial F-test to assess the significance of both response types for each voxel independently (see Experimental Procedures). We converted the F-values to Z-scores and represented the distribution of Z-scores in each hemisphere in a boxplot (Figure 15). Z-scores associated with the visual responses (blue boxes) exceeded the significance threshold ($p = 0.05$; horizontal dotted line) in virtually every individual task-related voxel in every hemisphere. A smaller proportion of voxels had a significant remapped response (red boxes). In two hemispheres, the median Z-score for the remapped responses fell below the statistical threshold of $p = 0.05$. We excluded these two hemispheres from subsequent analyses.

The data shown in Figures 14 and 15 indicate that the remapped responses were typically smaller than the visual responses. To compare response amplitudes directly, we calculated the amplitude of both response types by taking the difference between the minimum and the maximum of the response curve (Figure 16). We then subtracted the remapped response amplitudes from the visual response amplitudes. Large values indicate an amplitude difference between response types; a value of zero indicates no difference. The 95% confidence interval of the median was greater than zero in 8 out of 14 hemispheres, indicating that there was a significant amplitude difference in just over half of the hemispheres (Figure 16A, filled dots). A t-test on the median values from each hemisphere revealed that the difference,

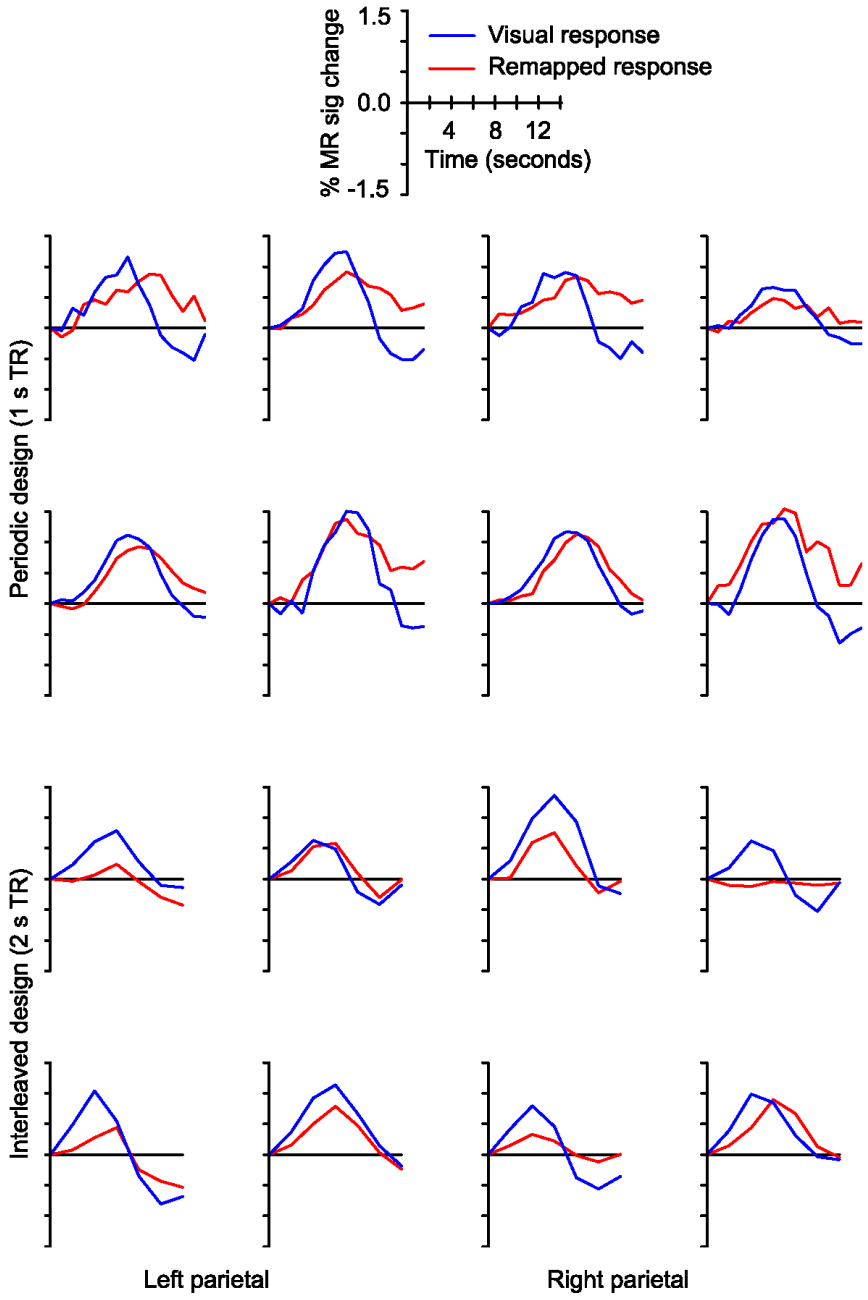

Figure 14. Visual and remapped response from all 16 hemispheres. Time course of activation evoked by the visual and remapped responses averaged over all task-related voxels in each of 16 hemispheres. Same format as in Figure 13.

while small, was significant across the population ($t[13] = 3.96$, $p < 0.01$).

We interpret the existence of remapped responses as evidence of a spatial updating signal in human parietal cortex. In the following sections, we rule out two alternate explanations for these results. First, we show that the remapped responses are not due to the ipsiversive eye movements *per se*. Second, we demonstrate that remapped responses cannot be attributed to the retinal impact of the stimulus in the ipsilateral visual field.

2.4.3 Saccades alone do not account for the remapped responses

One possible explanation of the present results is that saccades by themselves activate ipsilateral parietal cortex. To rule out this possibility, we tested all subjects on a saccade-only condition that was identical to the spatial updating condition in terms of oculomotor requirements. However, in the saccade-only condition, no stimulus appeared before the eye movement. Remapped responses were larger than the responses to ipsiversive saccades in this control condition (Figure 16B). This difference was significant across the population of 14 hemispheres ($t[13] = 5.44$, $p < 0.001$). A within-hemisphere analysis revealed that this difference was also significant in 11 out of the 14 individual hemispheres (Figure 16B, filled circles).

2.4.4 Ipsilateral stimuli alone do not account for the remapped responses

A second possible explanation of the present results is that visual stimuli activate ipsilateral parietal cortex directly. To rule out this possibility, we tested all subjects on a stimulus-only condition that was identical to the spatial updating condition in terms of retinal input but that did not require an eye movement. Remapped responses were larger than responses to ipsilateral stimuli (Figure 16C). This difference was significant across the population of 14 hemispheres ($t[13] = 5.44$, $p < 0.001$). A within-hemisphere analysis revealed that this difference was also significant in 11 out of the 14 hemispheres (Figure 16C, filled circles).

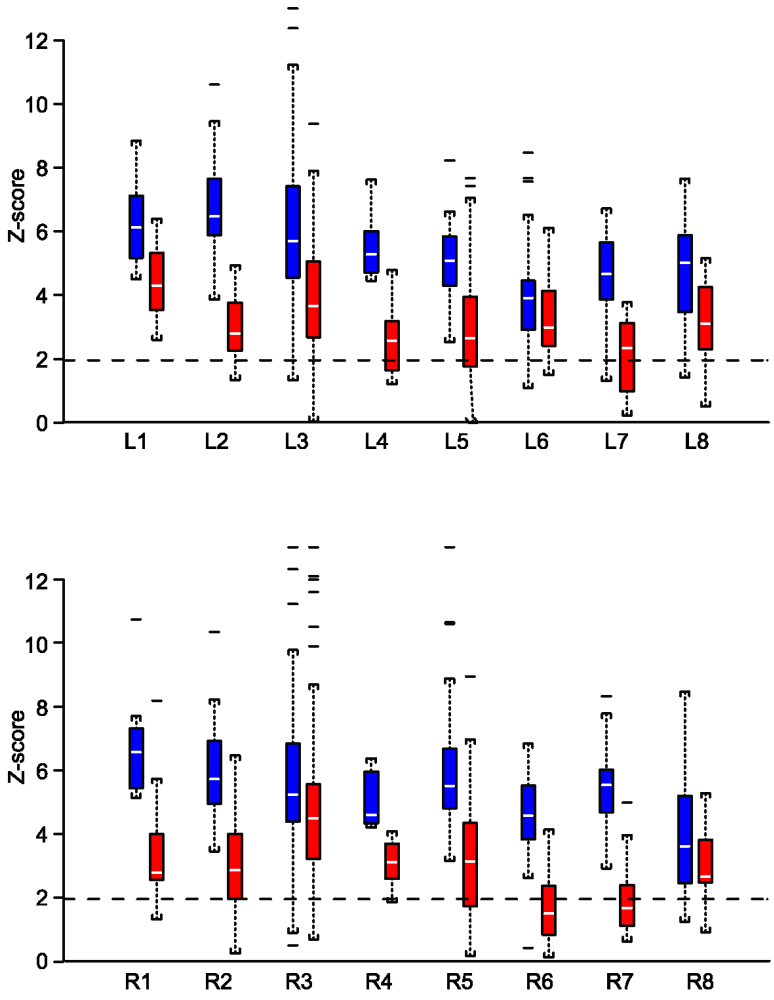

Figure 15. Significance of the visual and remapped responses. Boxplots of Z-scores representing the significance of visual and remapped responses for left (L1–L8) and right (R1–R8) hemispheres. For each box, the white horizontal line indicates the median Z-score for the hemisphere. The top and bottom of the box indicate the first and third quartile. The vertical lines ("whiskers") above and below the box indicate the upper and lower range of the data for that hemisphere, and the floating horizontal bars indicate outliers. The response in a given hemisphere was significant if the median exceeded the statistical threshold ($p = 0.05$; indicated by dotted horizontal line). All 16 hemispheres had a significant visual response, with the first quartile of voxels falling well above the significance threshold (blue boxes). In each hemisphere, the remapped response was associated with smaller Z-scores (red boxes). The median remapped response was not significant in two hemispheres.

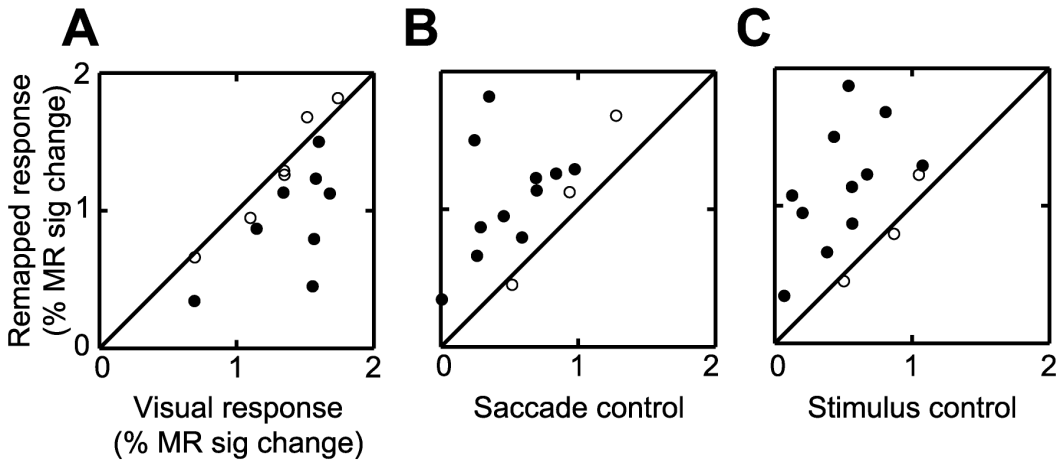

Figure 16. Comparison of remapped responses with control conditions. (A) Amplitude of the remapped responses (y-axis) compared with responses to contralateral stimuli (x-axis) for the 14 hemispheres that had significant remapped responses. Each dot represents the median response amplitude from a single hemisphere. Filled dots indicate that the 95% confidence limits of the median response for a single hemisphere did not intersect the unity line, indicating a significant effect across the population of voxels in that hemisphere. Contralateral stimuli elicited a larger MR response than the updated trace of the stimulus in 8 out of 14 hemispheres. A t-test on the medians confirmed that the visual responses were consistently larger than the remapped responses across hemispheres ($t[13] = 3.96$, $p < 0.01$). (B) The remapped responses were significantly larger than the responses elicited by ipsiversive eye movements in the saccade alone control condition in 11 out 14 hemispheres. This difference was significant across the group of 14 hemispheres ($t[13] = 5.44$, $p < 0.001$). (C) The remapped response was significantly larger than the response in the stimulus alone control condition in 11 out of 14 hemispheres. This difference was significant across the group of 14 hemispheres ($t[13] = 5.44$, $p < 0.001$). All amplitude values represent the maximum percent MR signal change minus the minimum.

2.4.5 Remapped responses occur later than the visual responses

2.4.5.1 Fourier Analysis. The timing of the remapped responses also indicates that they were not caused directly by the ipsilateral stimuli. The time series plots in Figures 13 and 14 show that the remapped responses occur later than the visual responses, which would not have been the case if they had been elicited by the stimulus itself. We used Fourier analysis to estimate the temporal shift between the visual and remapped responses in the eight hemispheres that were tested using a periodic design (Figure 17B). We first applied a Fast Fourier Transform (FFT) to the time series from each voxel. We then plotted the phase and magnitude of the FFT at the task frequency (1/15 s) in polar coordinates (Figures 17A and 17B). Each voxel is plotted twice: once for the signal elicited by the visual stimulus, and again for the signal elicited by the updated stimulus trace. The polar plot of voxels from the left hemisphere of a single subject reveals a clear segregation in phase based on response type (Figure 17A). The visual responses cluster in quadrant II, while the remapped responses cluster in quadrant III, indicating a phase shift across the population of voxels. A similar pattern was observed when voxels from all eight hemispheres were pooled together (Figure 17B).

This Fourier analysis shows that voxels in parietal cortex exhibit a remapped response that occurs later than the visual response in the same voxels. To test the significance of this observation, we calculated the phase spectrum of the visual and remapped responses. The distribution of phase spectra across all voxels and hemispheres is summarized in a boxplot in Figure 17C. Large phase spectrum values indicate a large difference in the time of the responses. We calculated a 95% confidence interval of the median phase spectra value in each hemisphere. In each case, the 95% confidence interval was greater than zero, indicating a significance phase shift between the two signals.

2.4.5.2 Cross Correlation Analysis. We asked whether the temporal shift between the visual and remapped responses was evident on a trial by trial basis (Figure 18). To address this issue, we cross-correlated visual responses on each individual trial with remapped responses on each other trial, resulting in a series of cross-correlograms. Each cross-correlation

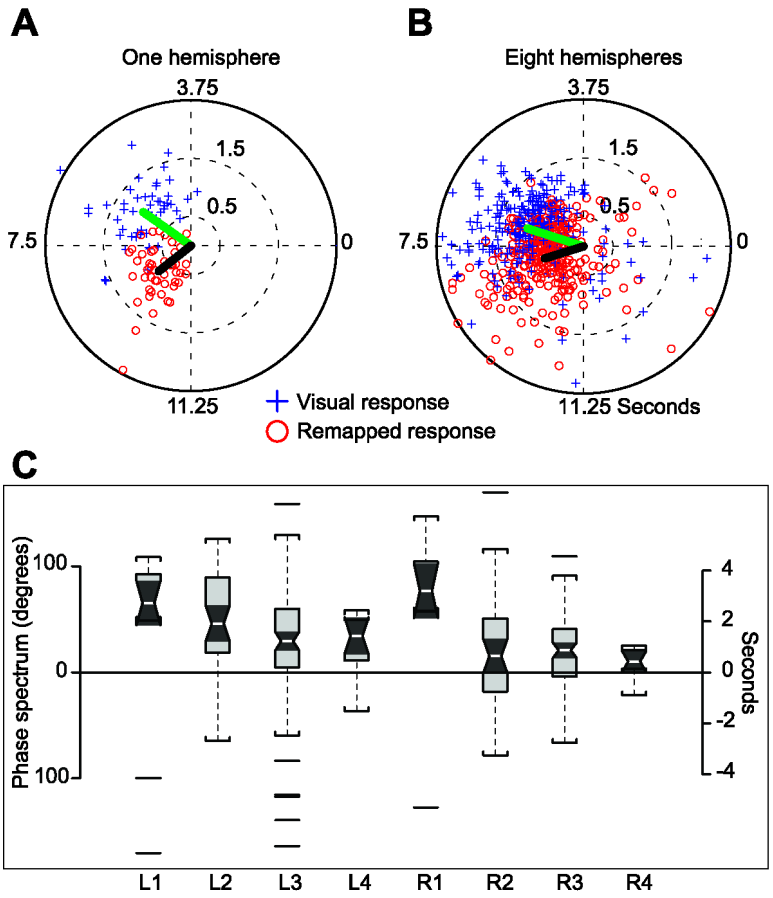

Figure 17. Fourier analysis of visual and remapped responses (A) The magnitude (ρ) and phase (θ) of the FFT at the task frequency are plotted in polar coordinates for a single subject. The vertical and horizontal dotted lines indicate 0, 3.75, 7.5, and 11.25 seconds. Each voxel is plotted twice: once for the visual response (blue, +), and again for the remapped response (red, ∘). The average vectors for the two response types are shown as green and black solid lines. The magnitudes of the visual responses were normalized to 1; the dotted circles represent magnitude values proportional to the average magnitude of the visual responses. Remapped responses tended to have later phases and smaller magnitudes than the visual responses. (B) Phase and magnitude of all voxels from all eight hemispheres scanned using a periodic design. (C) Boxplot of phase spectra representing the difference in time between the visual and remapped responses for eight hemispheres. Phase is represented on the y-axis in both degrees (left) and seconds (right). For each box, the white horizontal lines indicate the median phase. The notched black region around the median indicates the 95% confidence limits for the median. The bottom of the lower gray portion of the box indicates the first quartile; the top of the upper gray portion of the box indicates the third quartile. The vertical lines (*whiskers*) above and below the box indicate the upper and lower range of the data for that hemisphere, and the floating horizontal bars indicate outlier voxels. The notched black region is above the zero line in each hemisphere, indicating a significant phase shift between the visual and remapped responses across the population of voxels.

was calculated from −15 s to +15 s time lags. This analysis was performed on every combination of visual response ($n = 72$) and remapped responses ($n = 72$), resulting in a series of 5,184 visual-remapped cross-correlograms per voxel. We then calculated the peak of each cross-correlogram. The peak in the cross-correlogram is a measure of the temporal offset between the particular pair of visual and remapped responses. In order to derive a null distribution, we performed the same procedure, except we cross correlated visual responses with other visual responses, resulting in a series of visual-visual cross-correlations. Finally, we plotted the proportion of correlations that had a peak at each time lag (Figure 18A). The distribution of peaks at each lag for the visual-visual correlations was, as expected, symmetrically distributed around zero seconds (blue bars). In contrast, the distribution of peaks for the visual-remapped correlations was skewed to the right (red bars).

We repeated the analysis described above for all voxels from each hemisphere independently, and found that the difference between these two distributions was evident in each hemisphere. We averaged the results across hemispheres (Figure 18B). The distribution of peaks for visual-remapped correlations was significantly skewed to the right (KS-test, $p < 0.05$). These analyses indicate that the response to the updated representation of the stimulus occurred later than the response to a visual stimulus on the majority of individual trials.

2.5 DISCUSSION

We found that visual information is updated in human parietal cortex when the eyes move. Stimulus evoked activation in one hemisphere is remapped to the opposite hemisphere. This activation is a response to the trace of the stimulus at a particular spatial location. The stimulus had already disappeared by the time the eye movement was cued, so it could not be a direct visual response. This response to the stimulus trace indicates that neurons in human parietal cortex maintain a representation of space that is dynamically updated in conjunction with eye movements

We were able to rule out several alternative explanations for our findings. First, remapped

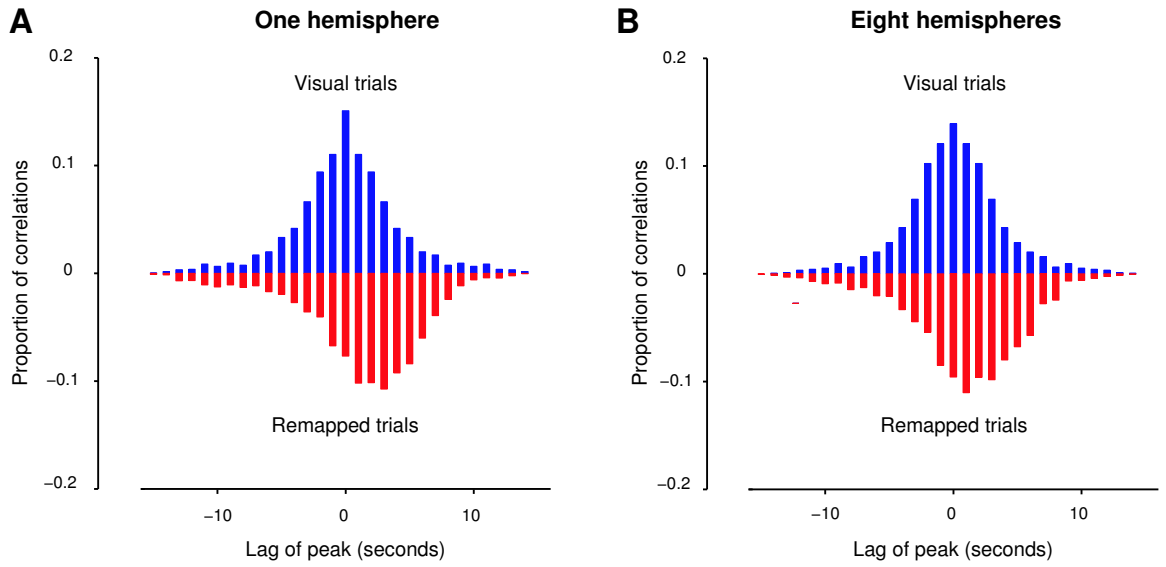

Figure 18. Cross-correlation analysis of visual and remapped responses. (A) The visual responses from each of 72 trials were correlated with each other visual response and correlated with each remapped response (data from Figure 13A, left hemisphere). The blue bars show the proportion of visual-visual correlations in which the peak in the correlogram fell at a particular lag. The red bars show the proportion of visual-remapped correlations in which the peak fell at each lag. The distribution of peaks for the remapped trials is skewed rightward relative to the visual trials, indicating that the remapped responses were later than the visual responses on the majority of trials. (B) Results of the correlation analysis combined over all subjects and hemispheres.

responses were not due to the ipsiversive eye movements *per se*. Some neurons in monkey LIP fire when a saccade is made toward the visual receptive field, which is almost always located in contralateral space (Colby et al., 1996). As expected from these single-unit data, human parietal cortex is also activated by saccades to contralateral targets (Sereno et al., 2001). While these and other studies suggest that parietal cortex should not be activated by ipsiversive eye movements, it was important to consider the possibility. The saccade-only condition was thus crucial in demonstrating that the remapped responses we observed were not due to a motor response associated with the saccade itself.

Second, remapped responses were not due to direct visual stimulation. Neurons in monkey area LIP have large, contralateral receptive fields that increase in size with eccentricity from the fovea (Barash et al., 1991b; Ben Hamed et al., 2001). Some LIP neurons have receptive fields near the fovea, and a subset of these neurons have receptive fields that extend across the vertical meridian. The representation of the ipsilateral visual field is not extensive

in LIP, but some cells do respond to stimuli located as far as 5° in the ipsilateral visual field. Human imaging studies have reported both positive and negative activations in early visual areas in response to ipsilateral, contrast-modulated checkerboard stimuli (Tootell et al., 1998), and it is possible that these ipsilateral responses might extend into parietal cortex. Nonetheless, ipsilateral responses were not observed in our stimulus-only condition. Moreover, the stimuli in our experiment were located far from the fovea (8°) along the horizontal meridian, decreasing the likelihood of activating receptive fields in ipsilateral cortex. Finally, the remapped responses were time locked to the eye movement, not the onset of the visual stimulus. If the remapped responses were in fact retinal in origin, they would have been concurrent with the visual responses.

The third alternate explanation that we considered was whether the remapped responses might be related to the auditory stimulus that was the instruction cue to initiate a saccade. A small number of cells in monkey LIP respond to auditory stimuli if the monkey has been trained to make saccades to a spatial location cued by an auditory stimulus (Grunewald et al., 1999; Linden et al., 1999; Mazzoni et al., 1996). In our experiment, a tone cued subjects to initiate an eye movement, but the saccade itself was directed toward a stable visual target. While it is possible that the tone activated parietal cortex, it was not the source of the responses we observed. The saccade-only condition also involved making eye movements in response to the same auditory stimuli, yet we did not observe the same pattern of activation in that condition. Furthermore, the auditory stimulus was presented binaurally, so we would not have observed the temporal shift between the visual and remapped responses had the auditory stimulus itself driven the activation in parietal cortex.

2.5.1 Cognitive factors and spatial updating

We considered the degree to which the remapped responses might be related to cognitive factors, such as anticipation, attention, and memory. Half of the subjects were scanned using a randomized interleaved design that prevented the subjects from predicting which condition would be tested on upcoming trials, and from anticipating when the next trial would occur. This experimental design helped ensure that additional variables, such as anticipation, were equally present in the both the experimental and control conditions.

Spatial updating and attention are fundamentally similar phenomena at the neuronal

level. The spatial response properties of LIP neurons change when the eyes move, and can do so in advance of the eye movement (Duhamel et al., 1992a). This predictive remapping follows the same time course as the behaviorally measured attention shift that occurs prior to eye movements (Hikosaka et al., 1996; Kowler et al., 1995). Remapping of receptive fields may thus be a neural instantiation of an attention shift. However, as shown by single-unit recording studies, an attention shift by itself is not sufficient to induce remapping (Colby, 1996; Duhamel et al., 1992a). The attention shift that occurs prior to the saccade must by followed by an actual eye movement, or remapping does not occur. This is evident in our experiment because remapped responses did not occur in the stimulus-only condition, which had the same attention demands, but did not require an actual eye movement.

We also considered the relationship between spatial working memory and the parietal activation we observed. At the beginning of the trial, the stimulus activates a set of neurons that maintain a representation of the stimulus location. At the time of the eye movement, those neurons transfer their information to a second set of neurons that represent the new retinotopic location of the stimulus. This updating involves memory since the neurons are responding to a stimulus that is no longer physically present. Many previous imaging studies have reported frontal, parietal, and extrastriate activation in tasks that involve the active maintenance of spatial information (Berman and Colby, 2002; Courtney et al., 1998; Heide et al., 2001; Sereno et al., 2001; Sweeney et al., 1996). Two studies in particular have demonstrated the importance of memory-related activation in parietal cortex during eye movements. Heide et al. (2001) found preferential activation in parietal cortex when subjects performed memorized sequences of eye movements, and (Sereno et al., 2001) demonstrated a spatial map in parietal cortex when subjects made memory-guided saccades to retinotopic targets. These studies indicate that memory traces exist in human parietal cortex, in accord with the present findings.

2.5.2 Conclusions

In the present study, an eye movement brings a previously stimulated screen location into the opposite visual field. In single-unit studies in monkeys, the eye movement brings the

spatial location of the vanished stimulus into the receptive field of the neuron being recorded. Both cases are demonstrations of a spatially specific response in the absence of direct visual stimulation.

The results from our study suggest functional similarities between monkey and human parietal cortex. The physiological response properties of spatial updating were comparable in two important ways. First, in our study, both the visual and remapped responses occurred at about the same time relative to the hypothesized arrival of visual information in cortex. In the updating condition, the visual stimulus appeared 2 seconds prior to the cue to make an eye movement. The remapped response was therefore shifted in time relative to the visual response, reflecting the period between the onset of the visual stimulus and the cue to make an eye movement. In monkeys, neurons exhibit a remapped response that occurs coincident with or even before the eye movement (Duhamel et al., 1992a; Nakamura and Colby, 2002; Umeno and Goldberg, 1997). We found that some voxels responded earlier than the time predicted by the cue to make the eye movement. Early neuronal activity may reflect a predictive response that allows the brain to represent the updated location before the end of the saccade.

Second, we found that the remapped response was smaller in magnitude than the visual response. This finding is consistent with physiological studies in monkeys in which cells have remapped responses that are on average half as large as the responses to stimuli in the receptive field (Duhamel et al., 1992a). These similarities suggest that the mechanisms governing remapping may be similar in humans and monkeys.

3.0 BAYESIAN INFERENCE APPLIED TO REMAPPING

3.1 OVERVIEW

In Chapter 2, we used fMRI to test the hypothesis that remapping occurs in human parietal cortex. Our conclusions were based on standard statistical procedures which revealed that (i) the remapping task activated parietal cortex, (ii) remapped responses were larger than responses in control conditions, and (iii) remapped responses had a longer latency than visual responses. We estimated BOLD activation by signal averaging and Fourier transforms, and used F-tests and *t*-tests to compare responses estimated in different conditions. This is the prevailing approach, used in the majority of fMRI studies. The results in Chapter 2 demonstrate that this approach is adequate for observing remapping. In the current chapter, we introduce an alternative method of analyzing fMRI time-series data that is based on the Bayesian statistical paradigm. We discuss the relative advantages of Bayesian methods and argue that they are useful for studying remapping.

Our argument is structured into four sections. First, we discuss a set of common inferential challenges that are inherent in most functional imaging experiments. For each challenge, we discuss a set of statistical approaches that are currently used (Section 3.2). Second, we introduce Bayesian methods (Section 3.3). Third, we describe a specific Bayesian model that is capable of addressing each of the major inferential challenges (Section 3.4). Finally, we apply the Bayesian model to remapping data (Section 3.5). We compare results obtained using a standard general linear model to responses estimated with the Bayesian model. We conclude that the Bayesian model yields comparable results to the GLM, both in terms of detecting active voxels and in describing the time course of evoked responses.

The comparison of the Bayesian model to the GLM serves a second, more practical function. The results in Chapter 2 were obtained using standard methods, yet the results reported in Chapter 4 were obtained using the Bayesian method. In the current chapter, we show that the two methods produce virtually identical results. The primary advantage of the Bayesian method in the present context is the logic of the Bayesian paradigm and the inferential flexibility it affords. The conclusions we reach are not dependent on statistical a particular methodology, however.

3.2 INFERENTIAL CHALLENGES IN FUNCTIONAL NEUROIMAGING

The goal of fMRI data analysis is to identify and characterize spatiotemporal patterns of activation associated with a given task. The underlying assumption is that the temporal dynamics of the BOLD signal and neuronal activity are themselves related (Duong et al., 2000). A myriad of statistical approaches have been adopted in the literature (see reviews in Lange, 2000; Petersson et al., 1999b,a). One of the simplest and most commonly used methods is a two sample t-test comparing activity in an experimental condition to activity in a control condition. Other frequently used tests include the Kolmogorov-Smirnov test (Aguirre et al., 1998a), the split sample t-test (Friston et al., 1995), cross-correlation between the time-series and either a fixed reference curve (Bandettini et al., 1993), or a time-shifted stimulus sequence (Hansen et al., 2004), t and F-tests in ANOVA (Cohen and Bookheimer, 1994), more general linear models (Worsley and Friston, 1995), linear deconvolution (Glover, 1999), selective averaging (Dale and Buckner, 1997), permutation tests (Bullmore et al., 1999; Belmonte and Yurgelun-Todd, 2001; Nichols and Holmes, 2002), and for periodic designs, tests for large power in a frequency band (Weisskoff et al. 1993; Friston et al. 1994). The wide diversity of methods indicates that none is ideal for all purposes. Rather, investigators have tailored their analysis strategies to address particular sets of scientific goals. In this section, we discuss inferential challenges that are present in all imaging experiments, and describe existing statistical approaches used to address these challenges. Our goal here is not to provide a comprehensive review of statistical methodologies in fMRI, but rather, to highlight the key statistical issues that arise in the analysis of imaging data.

3.2.1 Detection versus characterization

The challenge. Most data analysis strategies focus on identifying the set of brain regions that become active in response to some stimulus or in the context of some task. This is fundamentally a detection problem. Typically, voxels are classified as either *active* or *inactive* with respect to some contrast between two conditions. A standard statistical test, such as a t-test, is used to compare mean responses in two conditions. This comparison

results in a map of statistical values, which when appropriately thresholded yields an image of active brain regions. Activation maps are referred to as *statistical parametric maps*, or SPMs (not to be confused with the popular software package of the same name). Detection-based approaches are currently the prevailing analysis paradigm.

While localization is an important step in fMRI data analysis, there are features of interest in the pattern of activation that go beyond detecting which voxels are active. Neurophysiological processes are both spatially and temporally complex phenomena. For example, during the course of a given task, sets of brain regions may co-activate in a coherent fashion (Curtis et al., 2005; Sun et al., 2005), while other areas grow silent (Raichle et al., 2001), some regions exhibit subtle task-related modulations, increasing or decreasing gradually over multiple trials (Avidan et al., 2002; Gardner et al., 2005), while yet other areas exhibit sustained responses (Cohen et al., 1997). Moreover, critical information may be embedded in the precise timing and sequence of activations, as successive sets of areas become active during task performance (Formisano et al., 2002; Formisano and Goebel, 2003). The challenge is to capture this spatio-temporal complexity.

Solutions. Several recent methodological developments have enhanced the ability to characterize the spatio-temporal dynamics of BOLD activity. One of the most promising approaches is a multivariate method based on phase coherence (Genovese, 2000; Sun et al., 2005). Coherence is mathematically related to cross-correlation, but it is computed in the frequency domain, rather than in the time domain. Both techniques measure of the similarity between two signals. In neurophysiology, cross-correlation has been used to identify anatomical connections between pairs of neurons (Bruno and Simons, 2002; Kara and Reid, 2003). In human functional imaging, on the other hand, coherence has been used to identify functional connections among sets of brain regions. For example, using this approach, Curtis et al. (2004) identified distinct cortical networks involved in working memory. Independent components analysis (ICA) is another method that has been used to characterize spatiotemporal cortical responses (Esposito et al., 2002; McKeown et al., 1998b,a). While the use of ICA in fMRI is still in the early stages of development, several studies have demonstrated its ability to resolve temporal sequences of activation during the performance of cognitive tasks (Formisano et al., 2002; Formisano and Goebel, 2003).

3.2.2 Unit of measurement versus unit of inference

The challenge. In fMRI, the very unit of measurement, the voxel, poses an inferential challenge. Voxels make up an arbitrary coordinate system, imposed by scanning technology, rather than by physiology. As such, voxels do not have any relationship to neuroscientifically meaningful units of interest, such as brain regions, cortical columns, or layers. The methodology thus imposes a disconnect between what is measured on one hand and the target of inference on the other. While most statistical tools make inferences at the individual voxel level, the specific challenge is to relate measurements across groups of voxels to underlying physiological processes.

There are many examples in which testing an hypothesis hinges on making accurate inferences regarding groups of voxels. In one of the most common scenarios, two different experimental conditions activate clusters of voxels that lie in close proximity. The question is whether the two regions are in fact distinct, or alternatively, whether they lie within the spatial variability inherent in the signal. For example, pursuit and saccadic eye movements activate neurons in distinct regions of macaque frontal eye field (Bruce et al., 1985), and there has been some debate regarding whether such a distinction is present in humans as well (Berman et al., 1999; Petit and Haxby, 1999; Rosano et al., 2002). Traditional voxel-wise analyses can test whether a given voxel is selective for pursuit or saccades. It is difficult, however, to test whether the spatial extent of activation associated with the two conditions is distinct. What appears to be spatially discrete clusters of voxels could in fact result from a spatially variable process. Traditional statistical parametric maps do not provide an estimate of the spatial uncertainty associated with loci of activation, which makes it impossible to assess the confidence with which localization is determined. In other words, making a spatial inference on the basis of an activation map is like making an inference about a population mean without knowing the standard deviation.

Solutions. Most statistical methods for making inferences regarding groups of voxels come at the expense of spatial resolution, which is undesirable from a neuroscientific perspective. Forman et al. (1995) described a cluster-size method for identifying significant activation. This procedure is widely used, as it provides a reliable method for controlling false positives

resulting from multiple test comparisons, but it also biases results toward larger regions of activity, because true activations in small anatomical structures do not pass the cluster threshold. A related approach uses random field theory to assess the spatial uncertainty of an active cluster of voxels (Worsley et al., 1996; Ma et al., 1999). This approach, too, results in a loss of effective spatial resolution as substantial smoothing is required.

3.2.3 Linear versus nonlinear models

The challenge. Most analysis methods are based on linear statistical models, though the responses they aim to characterize are inherently nonlinear. Three nonlinearities are especially prominent. The first and most obvious nonlinearity is the shape of the hemodynamic response function (HRF) itself. The BOLD response is characterized by a complex waveform. It begins with a small signal decrease (the *initial dip*), followed by a larger bi-phasic response with a sharp positive component and a prolonged overshoot (Duong et al., 2000; Boynton et al., 1996). Second, responses can combine nonlinearly under certain conditions. The product of two responses depends on their temporal juxtaposition, their duration, their relative magnitude, and their spatial proximity to other activation loci (Wager et al., 2005; Vazquez and Noll, 1998; Nadell and Heeger, 2003, but see Hansen et al. 2004; Boynton et al. 1996). Finally, the underlying signal, in the absence of any activation, is subject to nonlinear drifts that are unrelated to neuronal activity (Smith et al., 1999). The signal drift can take a complex profile over time. It can appear stable for several tens of images, follow a linear increase or decrease, and then suddenly dip and rise in unpredictable patterns. Finally, this drift is not spatially homogeneous, but rather can vary widely across voxels (Smith et al., 1999). The challenge is to characterize BOLD responses accurately in the face of these nonlinearities.

Solutions. The most common solution to this problem is to approximate the nonlinear response using linear methods. Linear models are ubiquitous in fMRI analysis because (i) they work well in a wide range of applied areas, (ii) they are well understood from a theoretical perspective, and (iii) they are computationally tractable. With regard to the nonlinearity of the HRF, one method is to convolve a fixed nonlinear response function (see

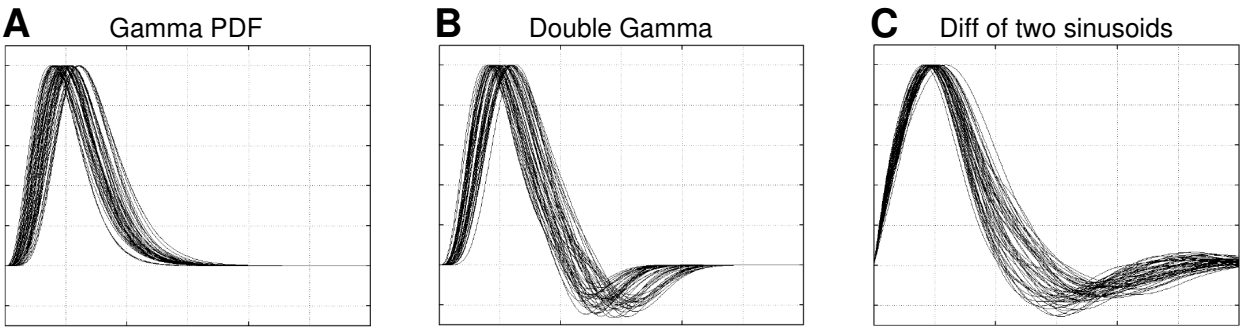

Figure 19. Theoretical temporal impulse response functions used in the statistical analysis of the BOLD response. These curves are intended to represent the BOLD response to brief stimuli of unit amplitude. (A) Gamma probability density function used as a model hemodynamic response function by Boynton et al. (1996); (B) Difference of two gamma PDFs used by Friston et al.; (C) Difference of two half-sinusoids used by Lee et al. (2005). The distributions of curves were created by randomly perturbing the parameters in the respective formulae by 10%.

Figure 19) with a boxcar model of the stimulus sequence (Cohen, 1997). The convolved waveform is then used in a general linear model that scales the convolved waveform to fit the data. The one free parameter in this analysis corresponds to the magnitude of the HRF — the parameters that govern response shape are fixed. This approach is effective in evaluating the extent to which the time series in a given voxel corresponds to an expected HRF. The method also has the advantage of being easy to implement.

Using a single canonical HRF, however, introduces a substantial degree of bias because any mismatch in shape between the HRF and the evoked response results in a poor model fit (Figure 20B). A mismatch could arise from at least three sources. First, the parameters for the canonical response functions are typically derived from BOLD responses in a particular type of stimulus (e.g., responses to flickering checkerboard stimuli). These parameters may not generalize to other stimuli. Second, the shape of the hemodynamic response is known to vary widely, both across individuals (Aguirre et al., 1998b) and across brain regions (Schacter et al., 1997). Thus using a single waveform may bias results toward a particular set of brain regions or individuals. Third, the temporal dynamics of BOLD responses are known to be modulated by task condition, as we observed in the remapping data in Chapter 2. Because the shape of the canonical HRF is fixed, any task-related modulation in response shape could bias estimates toward one condition or another.

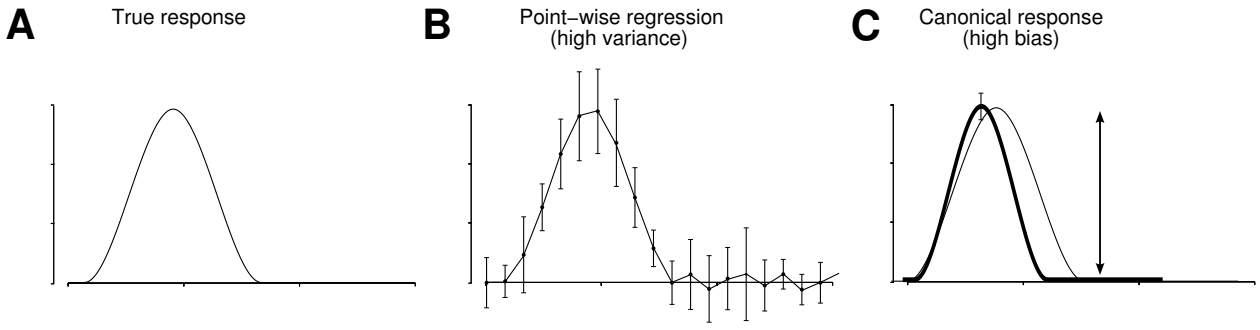

Figure 20. Bias/variance tradeoff introduced by approximating a nonlinear function with a series of linear estimators. (A) The underlying hemodynamic response is a nonlinear function. We simplify it here by ignoring the initial dip and post-response overshoot. (B) The response is often characterized using a series of linear estimators. Each point along the curve represents a separate parameter in general linear model. The model's fit to the data increases with the number of parameters. However, variance increases as well, making this a suboptimal approach. (C) The hemodynamic response is often characterized by fitting a single magnitude parameter using a *canonical response function* (thick line). This approach decreases variance. It also can introduce large bias because there is often a mismatch in shape between the canonical function and the underlying responses (thin line).

An alternative approach to approximating the nonlinear response is to use a point-wise regression. This method yields a series of linear estimates for several point along the hemodynamic response curve (Figure 20). Several closely-related strategies for accomplishing this have been described, including linear deconvolution (Glover, 1999; Ward, 1998; Serences et al., 2004), selective response averaging (Dale and Buckner, 1997; Dale et al., 1999; Hinrichs et al., 2000; Burock and Dale, 2000), and reverse correlation (Hansen et al., 2004; Hasson et al., 2004). These methods make no assumptions about the shape of the hemodynamic response. They allow each voxel to assume a different response profile. These methods have become the standard in the literature: there are implemented in virtually every software package, and they are used in the vast majority of current functional imaging studies.

The property that makes this approach so attractive — its flexibility — also make it suboptimal in three important ways. First, using a point-wise regression results in a high degree of variance in the response estimate. Typically, between one and two dozen independent parameters are used to describe a single response curve. Hence, the total variance associated with the response estimate is amplified considerably. One common approach for dealing with

the variance increase is simply to ignore it. Thus, it is common for imaging studies to report response curves without also reporting the standard error. Another strategy of dealing with the high variance is to blur the data or to average responses across voxels in an ROI, and in many cases, even across subjects. This has the obvious consequence of throwing away hard-earned temporal and spatial resolution. Second, using a point-wise regression places no constraints on the shape of the estimated response. Some flexibility is desirable, as response shape does vary across voxels and as a consequence of different task conditions. Yet, we know from the biophysics of the BOLD response that responses fall within a coherent family. Unconstrained flexibility needlessly increases the variance in the estimate. Finally, the results of a point-wise regression are difficult to interpret. It is difficult to derive quantities of interest — such as latency, response width, magnitude — from the parameters of a point-wise regression because the individual parameters have not direct meaning.

The problems associated with approximating a nonlinear function with linear methods invokes the familiar bias/variance tradeoff: If a statistical procedure is too simple it is biased and makes systematic prediction errors. If it is too complex, it will make unreliable, highly variable predictions. A better solution would be to use a nonlinear method to characterize the response. Nonlinear statistical methods are fraught with computational complexities, however, including numerical convergence issues (e.g., local minima, dependence on starting point) and exorbitant computation time. For many purposes, nonlinear regression is simply not practical.

3.2.4 Technique-driven questions versus question-driven technique

The challenge. Most analyses strategies are bounded by the specific models, the parameters of which do not directly relate to physiological quantities of interest. For example, inferences regarding response onset time were critical to the main finding in the remapping experiment described in Chapter 2. The response parameters that the analysis were based on were an indirect measures of the response onset time (see Figure 14). Instead, the point-wise regression that we used yielded parameter estimates for the height of the response at each point in time, and response onset was then inferred from these parameters. (note, however,

that are use of response phase in Figure 17 is not subject to this criticism). There is a disconnect, therefore, between what is estimated given the constraints of the technique, and the quantity one wants to know given the scientific question.

Solutions. The problem is frequently addressed using a two-stage solution. For example, several methods first model the response using a canonical function, as described above (3.2.3), and then estimate the specific properties of the HRF using non-linear regression (Lange and Zeger, 1997; Purdon et al., 2001). In practice, this works reasonably well (e.g., Lee et al., 2005), but the method has very little statistical power because it provides no estimate of uncertainty. A better solution is to use a model in which the estimated paramteres are directly related to the physiological variables under investigation. We described such a method for estimating the HRF in Section 3.4.

We have identified a set of inferential challenges that are common to all functional imaging studies. Several methods exist for addressing each of these challenges listed above. This diversity indicates the complexity of the problems at hand. Research on all of these problems is active and ongoing, indicating that the challenges remain real and significant.

3.3 MOTIVATION FOR A BAYESIAN APPROACH

In this Chapter, we describe a Bayesian model of fMRI time series data. It is useful to first consider the differences between Bayesian statistics and standard statistical approaches (which are referred to *frequentist* approaches). The key distinction between the Bayesian and the frequentist tradition is the definition of probability. For the frequentist, probability represents limiting relative frequencies over a series of (hypothetical) replications. Probabilities are objectively defined quantities. This definition of probability has several implications. First, probabilities can only be stated for observations that are in principle replicable (although they need not actually be replicated). Second, any parameters describing the probability distribution of a random quantity do not vary across replications. As such, no useful probability statements can be made about them. For example, we cannot speak of the

probability that a hypothesis is true. Third, statistical procedures should be chosen to have good long-run frequency performance. For example, a 95 percent confidence interval should cover the true value with limiting frequency 95 percent while being as short as possible on average.

For the Bayesian, probability represents a degree of belief. Probabilities are subjectively defined quantities. This definition of probability has several implications. First, probabilities can be stated for essentially any event, and relating to any quantity, random or non-random. We are using the classical calculus of probability without requiring "physical" randomness. In particular, we can make probability statements about parameters of interest in a statistical model. Statistical inference is a process by which the observed data update our beliefs. Our beliefs before seeing the data are described by a prior distribution; our beliefs after seeing the data are described by the posterior distribution. Bayesian inferences derive entirely from the posterior.

Bayesian methods have become an increasingly popular alternative to classical statistical approaches (Genovese, 2000; Gössl et al., 2001; Marrelec et al., 2003, 2004; Friston et al., 2002a,b; Friston and Penny, 2003; Penny et al., 2003, 2005; Smith et al., 2003; Woolrich et al., 2004a,b). There are at three primary advantages to Bayesian approaches:

Performance. One of the most important advantage conferred by Bayesian methods is their performance. While this statement sounds somewhat tautological, this point is at the heart of the argument for using Bayesian methods. Bayesian estimators derived from the mean of the posterior distribution have good frequentist performance. In other words, even in the case of a simple model, Bayesian methods provide superior parameter estimates relative to those derived from standard statistical procedures, such as the GLM. This reason alone is enough to motivate the use of Bayesian methods.

Flexibility. Bayesian methods confer greater inferential flexibility than traditional statistical methods. In order to use a statistic for inference in the standard approach, it is necessary to compute a statistic's sampling distribution and relate that to the parameters of interest. Computing a traditional test statistic can be difficult when the necessary assumptions regarding the distributions from which the parameters are to be drawn are not necessarily valid. For example, the assumptions inherent in the frequentist model are violated

in cases in which parameters have bounded ranges, multiple discrete model components, or spatial clusters. All three occur in the context of fMRI. These situations are not problematic for Bayesian methods, however. Bayesian models only require one to be able to compute the likelihood and the prior. Bayesian models do not require computing a test statistic for the purpose of comparison to a probability distribution.

Logic. The Bayesian paradigm provides a unified approach for addressing a wide range of disparate questions. The logical coherence of the Bayesian paradigm enables a single approach to apply to all inferences. This has direct application in neuroimaging. Because the data are used to obtain direct quantitative inferences about the parameters, uncertainty rather than randomness is the central focus. Thus, rather than use a separate statistical procedure depending on the distribution of the statistic, one can use the inference method in all cases. One is only concerned with the data in hand, not with other possible data sets that could have been observed but were not observed.

3.4 A BAYESIAN MODEL OF FMRI TIME SERIES DATA

3.4.1 Components of the model

In this section, we describe a Bayesian statistical model that we have applied to fMRI data collected in the context of a remapping experiment. The details of this model have been published elsewhere (Genovese, 1998a,b, 2000). Our purpose here is to provide an outline of the essential components of the model. The model is nonlinear, fully Bayesian, and hierarchical. It is in some ways similar in structure to the general linear models commonly used in the literature, but whereas the GLM assumes a rigid structure for convenience or necessity, this model confers greater flexibility. The model decomposes the fMRI data at each voxel into four additive components, each of which describe a different source of variation:

μ	Baseline level of signal in the absence of activity
Drift	Coefficients of drift profile in current basis ($\boldsymbol{\theta}^{\text{Drift}}$)
Response	Amplitude of response within an epoch ($\boldsymbol{\theta}^{\text{Response}}_{c,k}$)
	Average amplitude of response in a condition ($\boldsymbol{\theta}^{\text{Response}}_{c}$)
	Shape of response curve (2–8 parameters $\boldsymbol{\theta}^{\text{Shape}}$)
Noise	Noise precision

The components combine to form the likelihood as follows. Let $Y(t)$ be the observed MR signal at time t from a specific voxel, where $t = 0, \Delta, \ldots, (T-1)\Delta$ for $\Delta > 0$. The model decomposes this time series into four discrete components

$$Y(t) \;=\; \mu \;+\; d(t; \boldsymbol{\theta}^{\text{Drift}}) \;+\; a(t; \boldsymbol{\theta}^{\text{Response}}_{c(t),k(t)}, \boldsymbol{\theta}^{\text{Shape}}, \mu) \;+\; \epsilon(t; \boldsymbol{\theta}^{\text{Noise}}), \tag{1}$$

where μ, $\boldsymbol{\theta}^{\text{Shape}}$, and $\boldsymbol{\theta}^{\text{Response}}$ and the function $d(\cdot)$ are model parameters and ϵ is a parameterized noise process with mean 0 and variance 1. Once the distribution of ϵ is specified, this equation determines the likelihood for the model.

The baseline (parameter μ) is a constant that reflects the mean signal in the absence of task-related signal changes. The drift is a smooth cubic spline with ten regularly-spaced knots. The drift spline is weighted toward linearity by a prior that exponentially penalizes the integrated squared second derivative. The noise component is simple white noise with unknown variance, although the model could in principle be extended to include autoregressive noise models.

From a physiological point of view, the most critical feature of the model is the shape of the response, $\boldsymbol{\theta}^{\text{Shape}}$, which we refer to as a bell function (this is the HRF model). The Bayesian model is general, and any parameterized function can be "plugged in" for the response shape. The function chosen should, ideally, be motivated by both the known biophysics of the BOLD response and the scientific questions of interest. The activation model can thus be tailored to the specific experiments and it can incorporate new information about the biophysical properties of the BOLD response. In the original model, Genovese (1998a) used a bell function with eight parameters (Figure 21A). We have explored a range of bell functions and have found that, for experiments with event-related designs, a simple three-parameter polynomial bell works well (Figure 21B). The parameters for this bell are: γ^{lag}, which specifies the time between onset and signal increase, γ^{attack}, which specifies the time between signal increase and peak, and γ^{decay}, which is the time from peak to baseline.

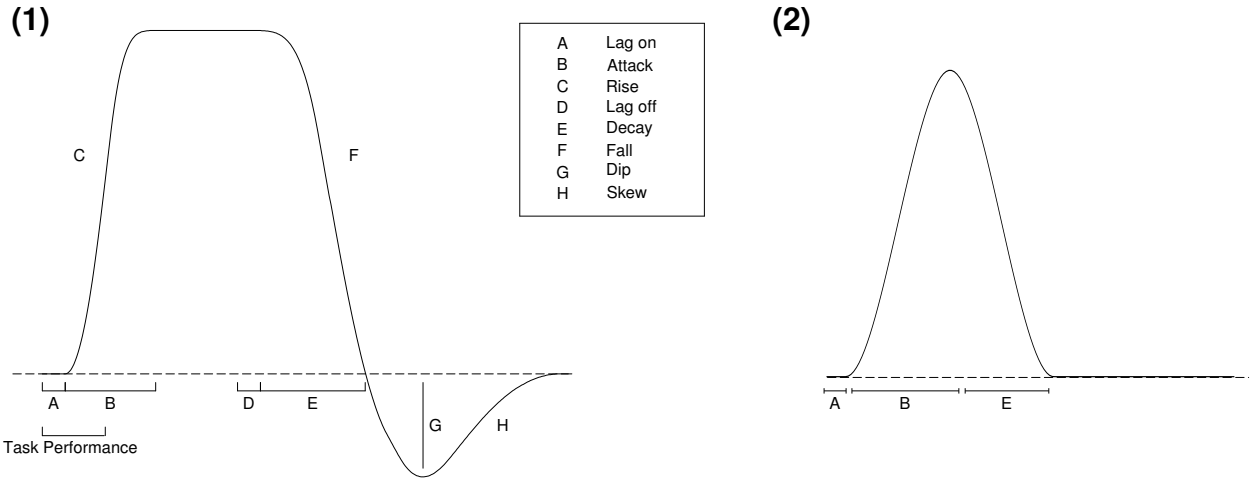

Figure 21. Parameterized activation profile, $a(t)$. (1) Activation bell with eight response parameters, as in the original model description (figure adapted from Genovese, 1998a). (2) Three-parameter bell, with shape parameters that characterize response latency (*lag*), rise time (*attack*), time-to-baseline (*decay*). Our subsequent analyses use the three-parameter bell in (2).

The bell is the product of two piecewise polynomials. The first polynomial is an *up* ramp that models the latency of response onset (*lag*) and the rise of the response from onset to peak (*attack*). The *up* ramp, $U(t)$, over the interval $[0, 1]$, is defined as follows:

$$\begin{aligned} U(t) &= 0, t \leq 0 \\ &= u(t), 0 < t < 1 \\ &= 1, t \geq 1, \end{aligned} \tag{2}$$

where $u(t)$ is a monotonically increasing and smooth function over $[0, 1]$ such that $u(1/2+t) = 1 - u(1/2 - t)$ for $0 \leq t \leq 1/2$. The second polynomial is a *down* ramp that models the time the MR signals takes to return to baseline from the peak. This *down* ramp, $D(t)$, over the interval $[0, 1]$, is simply

$$D(t) = 1 - U(t). \tag{3}$$

The full response curve, b, is simply the difference between the up ramp and the down ramp

$$b(t; s; \{L_r, r, f\}) = U\left(\frac{t - L_r}{r}\right) D\left(\frac{t - L_r - r}{f}\right). \tag{4}$$

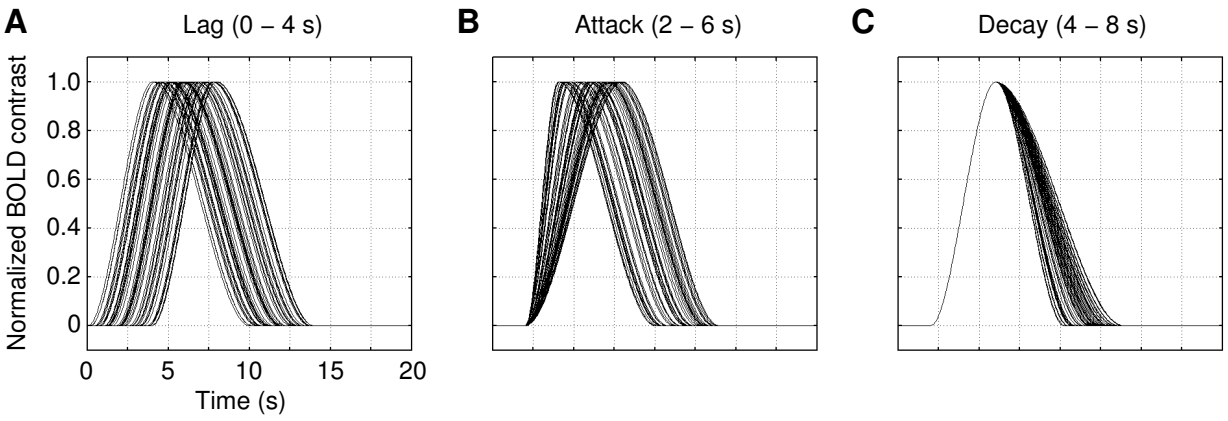

Figure 22. Variation in the shape of the three-parameter bell function. Each of the three parameters, lag, attack, and decay, was independently varied over a range of values, illustrating the range of shapes that the activation profile, $a(t)$ and take. Note the similarity in shape between this bell and the gamma probability function shown in Figure 19 (Cohen, 1997).

The shape of the bell is shown in Figure 22. This polynomial bell is similar in shape to the gamma PDF widely used to model the hemodynamic response (Boynton et al., 1996). This bell function has two advantages over the gamma function, however. First, the parameters of the bell (lag, attack, and decay) are directly related to hemodynamic events of interest. This is not true for the gamma function, which has parameters for shape and scale. Second, this bell is very flexible and can assume a wider range of shapes than the gamma function. Thus, it is more likely to capture the dynamic range of hemodynamic responses. We fit the bell to each voxel independently. Thus, our method yields an estimate of each of the three response parameters for each voxel.

The full activation profile is specified as follows. The response functions were each shifted to the onset time of the event (at sub-TR resolution). The task-related signal component consists of a superposition of responses for each event.

$$a(t) = \mu \sum_{c} \gamma_c^{\text{resp}}\, b\left(t - t_{0,c};\, \gamma_c^{\text{lag}}, \gamma_c^{\text{attack}}, \gamma_c^{\text{decay}}\right). \tag{5}$$

Multiplying by the baseline μ simply scales the γ^{resp}. This scaling has the consequence that γ^{resp} is expressed in units of proportional signal change relative to baseline.

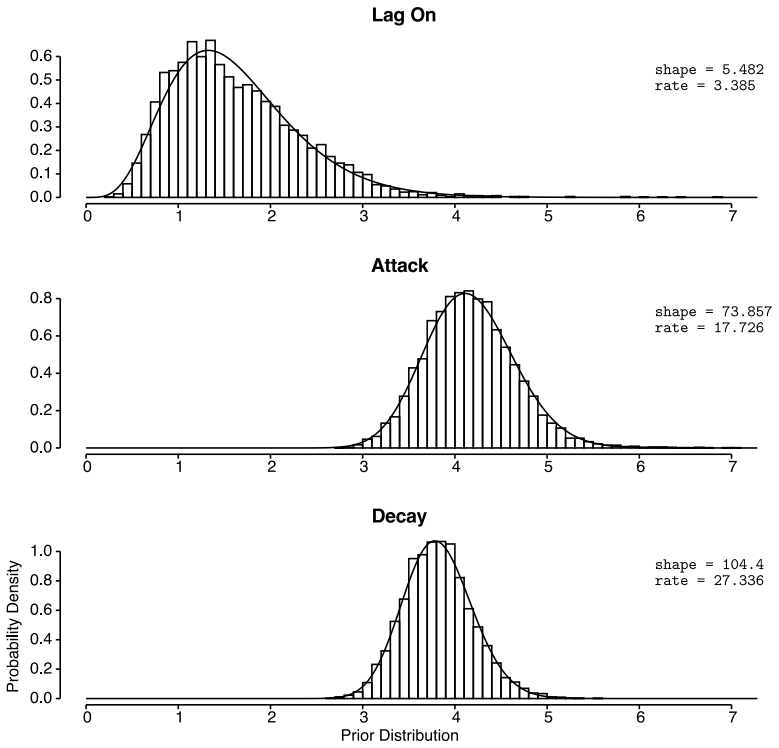

Figure 23. Histogram of shape parameters used to determine the prior in the Bayesian analysis.

3.4.2 Prior specification

One of the principal advantages of the Bayesian model is the use of prior information. To illustrate why this is an advantage, consider the case of the point-wise regression discussed above. Responses estimated in this way can take a wide range of shapes, many of which are physiologically meaningless. We now have a good understanding of the general form of the BOLD response, and our analysis method should be able to use this *prior* knowledge to inform our statistical approach. The Bayesian model that we used (Genovese, 2000) takes into account priors for each of the components discussed above: baseline, drift, responsiveness, shape, and noise.

In applying the model to our remapping data, we based our priors on previous results. We analyzed all active voxels in a subset of subjects in Merriam et al. (2003), deriving estimates of lag, attack, and decay from a simple nonlinear least squares fit to HRFs estimated using point-wise regression (see Figure 14). We then fit gamma probability density functions to the

distribution of shape parameters (Figure 23). The best fit gamma PDFs were then used as *priors* for the Bayesian model. Hence, the Bayesian approach takes advantage of information that we have gleaned from earlier studies.

The hierarchical structure of the priors is as follows:

$$\begin{aligned}
\boldsymbol{\theta}^{\text{Drift}}|\lambda, \boldsymbol{\theta}^{\text{Response}}, \boldsymbol{\theta}^{\text{Shape}}, \mu, \sigma &\sim A \exp\left(-\frac{1}{2\lambda}\left[\rho_{\text{nc}} \int |d(t)|^2 + \int |d''(t)|^2\right]\right) \\
\lambda|\boldsymbol{\theta}^{\text{Response}}, \boldsymbol{\theta}^{\text{Shape}}, \mu, \sigma &\sim \mathsf{Exponential}(\theta^{\text{Noise}}/\lambda_0) \\
\theta_c^{\text{Response}}|\boldsymbol{\theta}_{-c}^{\text{Response}}, \boldsymbol{\theta}^{\text{Shape}}, \mu, \sigma &\sim \text{Gamma/Point-Mass Mixture} \\
\boldsymbol{\theta}^{\text{Shape}}|\mu, \sigma &\sim \text{Gamma} \quad \text{[Independent Components]} \\
\mu|\sigma &\sim t_1(\mu_0), \\
\sigma &\sim \text{Inverse Gamma} \quad \text{[Proper and diffuse].}
\end{aligned}$$

Lower levels in the model were conditional on the fit at high levels of the hierarchy.

3.4.3 Calculating the posterior

The basic Bayesian method is the same in all cases. The first step is to select a probability model, $f(y \mid \theta)$ that reflects our beliefs about the data y for each value of the parameters. Note that this likelihood is now considered a conditional probability distribution, not just an index set of distributions. The second step is to select a prior distribution, $f(\theta)$ for the parameter. The third step is to combine these to form a posterior distribution via Bayes Theorem,

$$f(\theta \mid y) = \frac{f(y \mid \theta) f(\theta)}{\int f(y \mid \theta') f(\theta') d\theta'}. \tag{6}$$

These three steps are applied in every situation and produce a posterior distribution.

3.4.4 Making inferences based on the posterior

The primary output of the Bayesian model fit is the posterior distribution of the parameters given the data, $\mathsf{P}\{\gamma \mid Y\}$. From this distribution we derive point estimates, such as posterior means on a given parameter. The most relevant measure of what we can infer about a comparison between conditions is embodied in posterior probabilities. We base our inferences on the probabilities of several specific events. The first is the probability that there is a nonzero response in a task condition. For condition "c", this is denoted by $\mathsf{P}\{\gamma_c^{\text{resp}} > 0 \mid Y\}$. Because our hierarchical model allows for a nonzero probability on the discrete value 0, this probability indicates the strength of evidence for any particular response. A second event that we consider is the probability that the response in one condition is greater than the response in another, denoted by $\mathsf{P}\{\gamma_c^{\text{resp}} > \gamma_{c'}^{\text{resp}} \mid Y\}$. We make similar comparisons for shape parameters, such as when comparing response onset times across conditions. Third, we compute the posterior probability that the remapped response is larger than the maximum of the saccade and stimulus-only responses, denoted by $\mathsf{P}\{\gamma_{\text{remap}}^{\text{resp}} > \max(\gamma_{\text{sac}}^{\text{resp}}, \gamma_{\text{stim}}^{\text{resp}}) \mid Y\}$. Finally, we consider the posterior probabilites of more complicated events such as the monotonicity in the group mean response parameters and the population probabilities of nonzero response. We have used two different approaches for estimating the probabilities. The probabilities were computed both analytically, and to arbitrary precision by Monte Carlo simulation. All probability statements in subsequent analyses were calculated using Monte Carlo simulations because this method is computationally simpler and is sufficiently accurate.

3.5 MODEL VALIDATION

In order to validate the model, we analyzed the same data set twice: first with a general linear model, and again with the Bayesian model. We then compared results obtained with the two methods. This section is organized as follows. First, we describe the general linear model in some detail (Section 3.5.1). Second, we describe a data set in which the subject performs the same remapping task described in Chapter 2 (Section 3.5.2). Finally,

we present results from the data analyzed using the two methods. We show that the two methods provide approximately equal fits to the full fMRI time series data, approximately equal HRF estimates, and approximately equal spatial activation maps (Section 3.5.3).

3.5.1 Details of general linear model

The use of the general linear model is widespread in fMRI. Its ubiquity makes it a good standard against which to evaluate the results from the Bayesian model. In many event-related fMRI paradigms, the BOLD signal is modeled at a given time (t) using a GLM of the form:

$$s(t) = b_0 + x_1(t)b_1 + x_2(t)b_2 \ldots + x_{n-1}(t)b_{n-1}(t)b_{n-1} + \epsilon(t). \tag{7}$$

where n is the number of variables to be estimated, x is some pre-specified stimulus-response function for a given variable, and b is a scaler beta weight that minimizes the error between $x_n(t)$ and $s(t)$. This expression is typically written in matrix form as

$$\boldsymbol{s} = \boldsymbol{X}\boldsymbol{b} + \epsilon \tag{8}$$

where $\boldsymbol{s}$ is the observed BOLD time series vector, $\boldsymbol{X}$ is the *design matrix* that describes the relationship between task events and the observed data, and $\boldsymbol{b}$ is a vector of model parameters that describes the magnitude of the relationship between $\boldsymbol{X}$ and $\boldsymbol{s}$. The design matrix $\boldsymbol{X}$ has a row for each time point in the BOLD signal ($\boldsymbol{s}$) and a column for each estimated variable. The vector $\boldsymbol{b}$ contains a scalar value for each corresponding column in $\boldsymbol{X}$. The least squares solution that minimizes the error between $\boldsymbol{X}$ and $\boldsymbol{s}$ is found by inverting the design matrix:

$$\boldsymbol{s} = (\boldsymbol{X}^{\mathrm{T}}\boldsymbol{X}^{-1})\boldsymbol{X}_S^{\mathrm{T}} \tag{9}$$

As discusses earlier, one common approach is to estimate the magnitude of the response using a fixed canonical response function (Figure 24A). This is done by convolving the function with a boxcar model of the stimulus sequence, yielding a single predictor waveform. This predictor is then included as a column in the design matrix, $\boldsymbol{X}$. The result is a scalar value in the corresponding row of vector $\boldsymbol{b}$ that scales the canonical input function to minimize the error between $\boldsymbol{X}$ and $\boldsymbol{s}$.

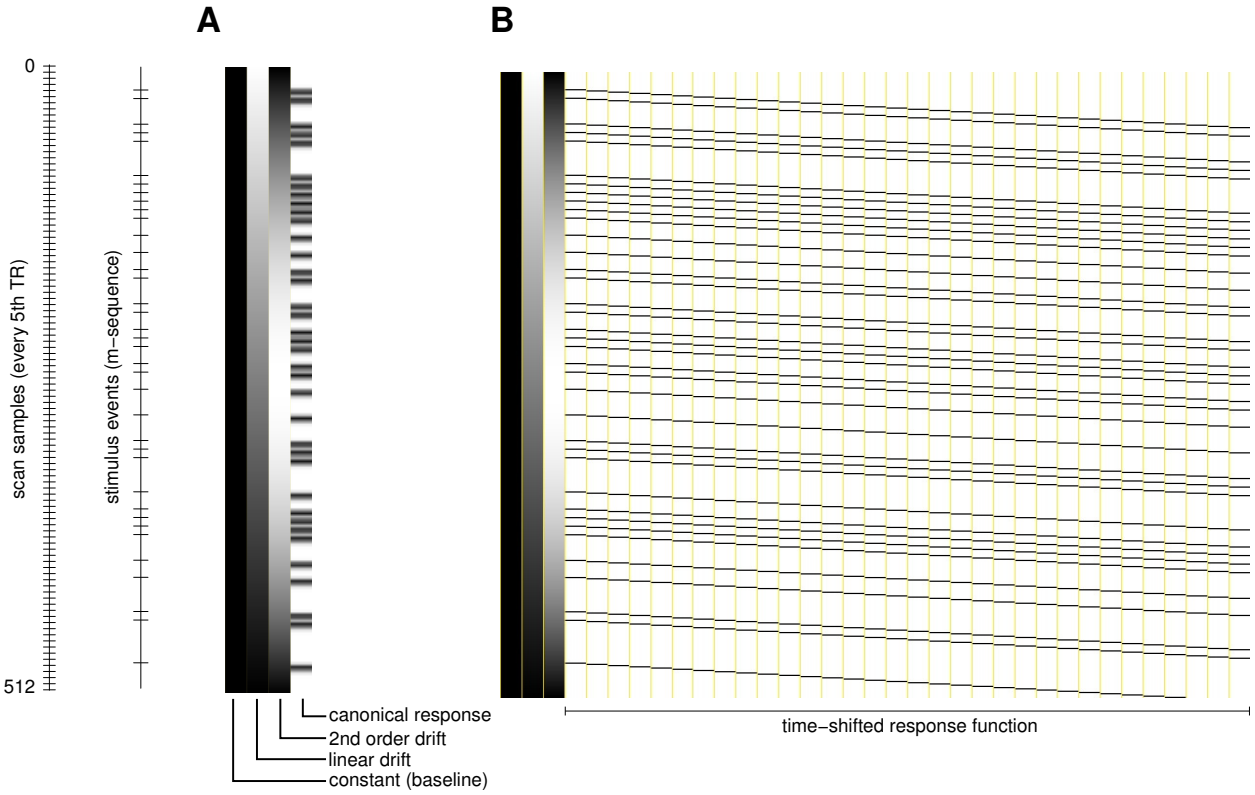

Figure 24. Design matrices for two general linear models. (A) Design matrix in which the response is modeled with a canonical function. (B) Design matrix in which the response is modeled with a separate regressor for each point in time in a 32 s post-stimulus window. Rows in the matrices represent time (in TRs); columns represent individual regressors. Scan number is shown in the far left (every 5^{th} TR is shown).

There are several disadvantages to using a canonical response function. The approach currently favored is to estimate the BOLD response using point-wise regression. There are a number of ways of doing this, but for the present purposes, we will use the "deconvolution" approach (Glover, 1999; Serences et al., 2004), as implemented in AFNI software (Ward, 1998). The details of the approach are as follows. Instead of a single column in $\boldsymbol{X}$ representing scaler values of the expected response to a given event over the entire time series, a separate column of $\boldsymbol{X}$ is created for each point in time along the expected event-related BOLD response for each event (Figure 24B). Each column in $\boldsymbol{X}$ consists of a binary array of 1's for points in time when an event has occurred and 0's at all other points in time. Inverting the design matrix specified in this way results in an estimate, $\boldsymbol{b}$, for each point along the BOLD response for each variable entered into the model. The series of estimates

for a single event are called the system's *impulse-response*.

The design matrix also contains columns to factor out sources of unexplained variance. For example, it is typical to include a column of 1's to factor out a constant (the baseline response level), and another column with a linear term to factor out signal drift. Sometimes additional columns are added to account for higher-order drift components, head motion estimates, or potential physiological confounds, such as heart rate and respiration.

3.5.2 Example data set

In this comparison, we use data from two subjects performing the remapping task. The task was similar to that described in Chapter 2 (see Figure 8). The data for this example were collected in runs in which subjects performed only one task condition (the remapping task), and in only one direction. Thus, there is only a single response to estimate. This simplifies matters considerably, and makes it easier to evaluate and compare the two models. Finally, it is worth noting that the timing for the task events was determined by a special class of pseudorandom sequences known as *m-sequences* (Sutter, 2001; Reid et al., 1997; Buracas and Boynton, 2002). Experimental designs based on m-sequences are thought to have several advantages over the periodic and random event-related designed used in Chapter 2 (see Figure 9, Liu, 2004; Liu and Frank, 2004).

3.5.3 Results

We first asked whether the two models, the GLM-based model and the Bayesian model, provide equally reasonable fits to the data. We plot the model fits for six different voxels in Figure 25. Subjectively, the GLM and the Bayesian model appeared to provide comparable fits to the MR time series. It is evident that the GLM was more sensitive to noise, however. The shape of the HRF estimated by both models for 25 different voxels is shown in Figure 26. The HRFs estimated by the Bayesian model (shown in blue) closely corresponded to an HRF that we would expect to see based on our understanding of the biophysics underlying the BOLD response. The fits obtained by the Bayesian model were relatively insensitive to noise, because the model was constrained to conform to our prior expectations about the shape

of the HRF. In contrast, the GLM provided a much noisier estimate of the HRF (shown in red). The raggedness of these HFR estimates is attributable to the fact that the GLM model contained 32 parameters as opposed to 3 parameters, and thus contained degrees of freedom which allowed for over-fitting. A comparison between the HRF estimated by both models provides an intuitive account of the advantages of the Bayesian approach: by taking into account our knowledge of of the likely shape of the hemodynamic response, we can avoid over-fitting our estimate to random perturbations in the MR signal that are probably not of physiological interest.

To compare the performance of the two models in identifying which voxels in the brain had task-related activation, we plotted the SPMs obtained using the GLM in Figure 27A. Figure 27B shows the posterior probability maps obtained for the corresponding data sets using the Bayesian model. The dark squares are voxels for which the GLM returns a high F-statistic, indicating statistically significant levels of activation. These voxels in the SPM maps appear to correspond closely to the voxels for which the Bayesian model indicates a high posterior probability of a nonzero response (i.e., the dark squares in panel B).

To formally compare the results of the two models, it is necessary to apply some significance threshold to the SPMs. Figure 28A shows the results from three different thresholding methods. In this panel, the dark squares are voxels that were declared significant at an uncorrected threshold of $p < 0.05$, by using the false discovery rate (FDR, Genovese et al., 2002), and by using the Bonferroni correction. Of the three activation maps, the map thresholded with FDR is most similar to the posterior probability map obtained from the same data set 28B. Subjecting probability maps to a 95% threshold means that surviving voxels have a 5% probability of not exceeding the default threshold. In other words, if we declared these voxels to be "activated", we should expect 5% of the voxels thus declared to be false positives. This fact points to a fundamental similarity to the probability map and the false discovery rate (Friston and Penny, 2003). The FDR is the proportion of voxels that are declared significant but are not. It should be noted that many voxels will have a posterior probability that is more than 95%. Therefore, the 5% is an upper bound on the FDR.

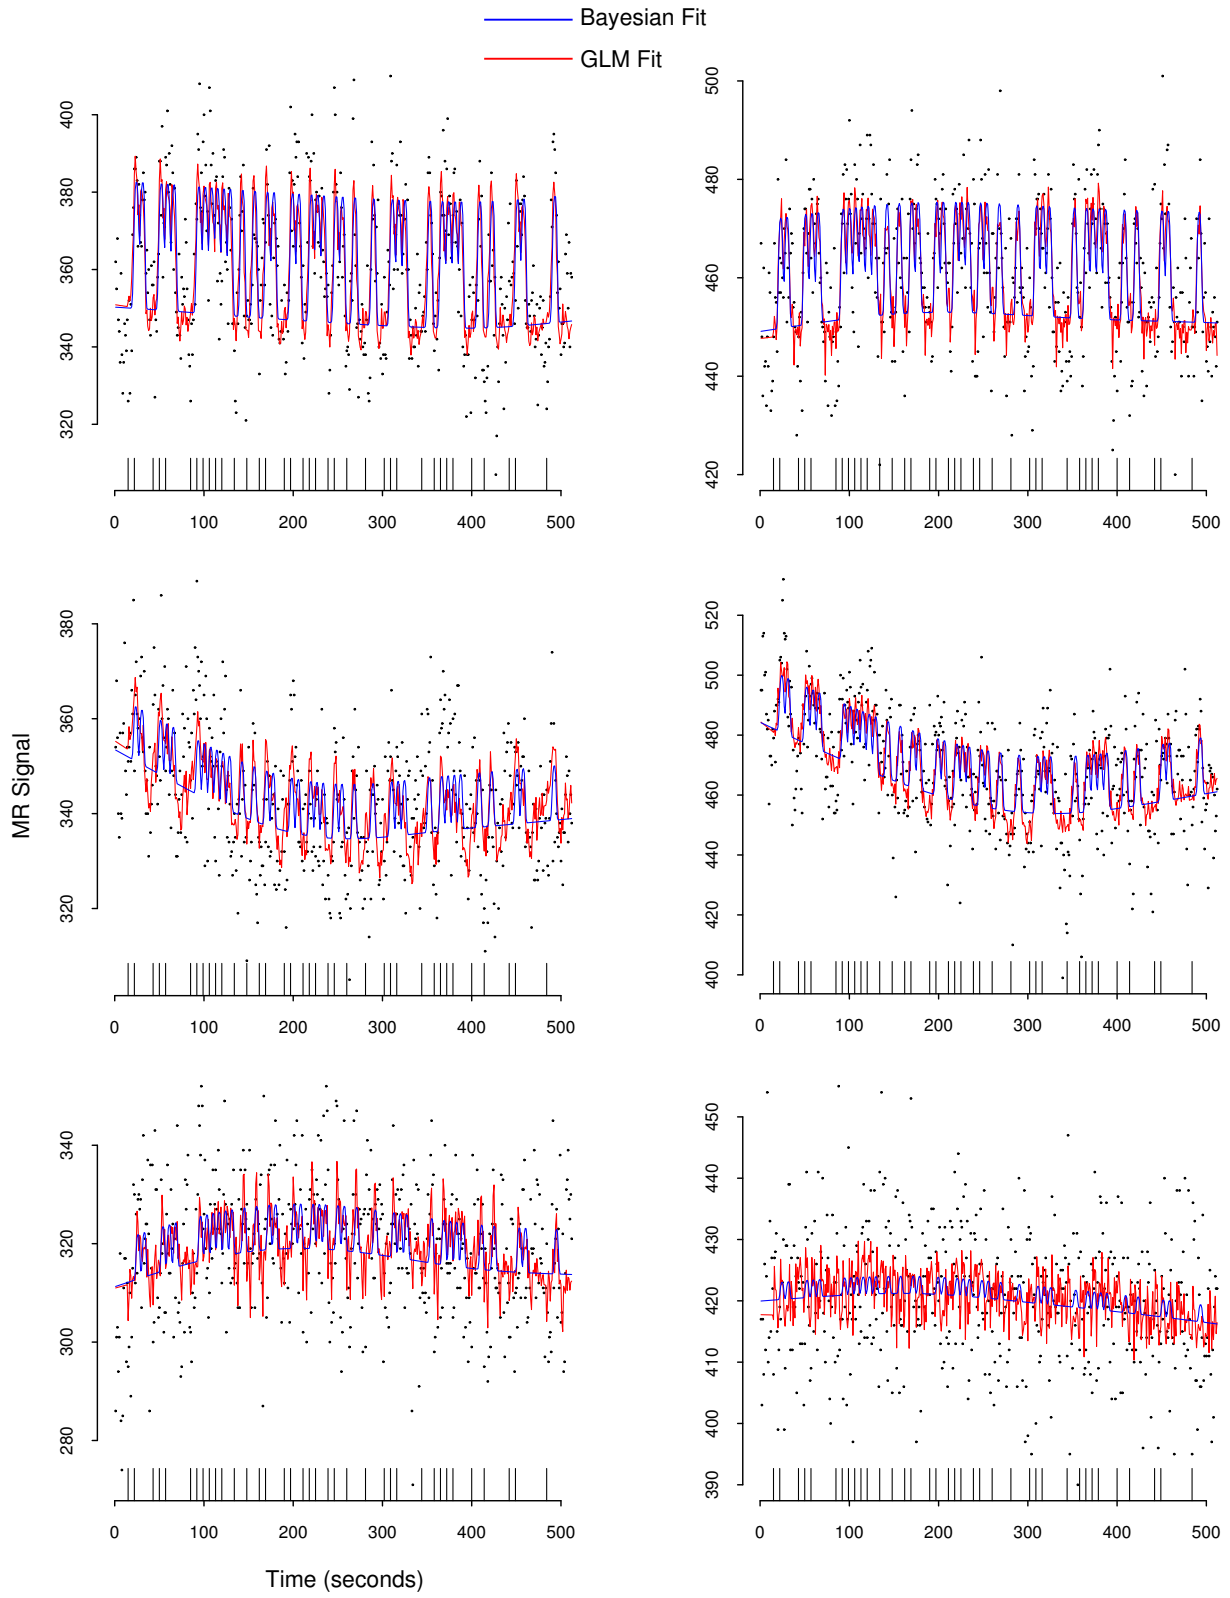

Figure 25. The general linear model (red) and the Bayesian model (blue) both provide reasonable fits to the data. Each panel represents the time series from a single voxel (black dots). In each case, the Bayesian model provides a slightly more conservative estimate of response amplitude — the range of the blue lines is smaller than the range of the red lines.

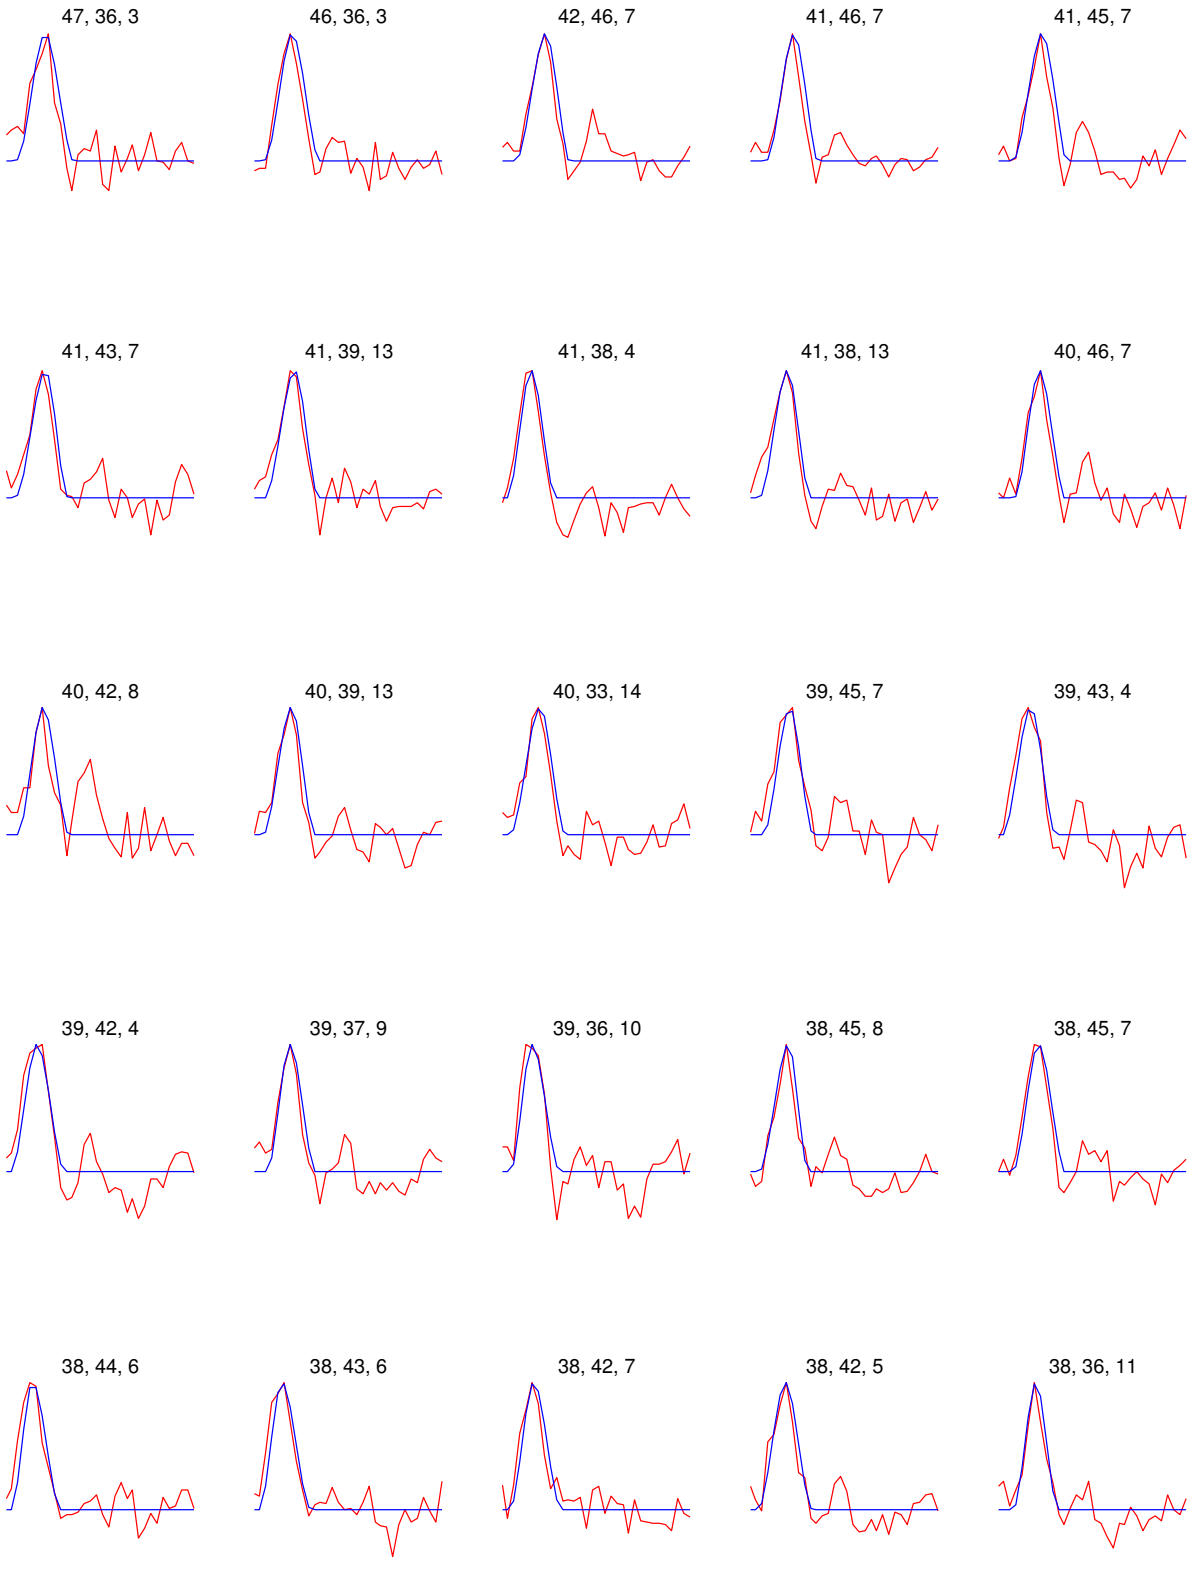

Figure 26. Good agreement between GLM (red) and Bayesian (blue) response estimates. Responses are shown over a 32 s time window following trial onset. Response magnitudes were normalized to 1, hence all responses have the same height. Note that the responses estimated with the GLM contain 32 parameters, while the responses estimated with the Bayesian model have only 3 parameters.

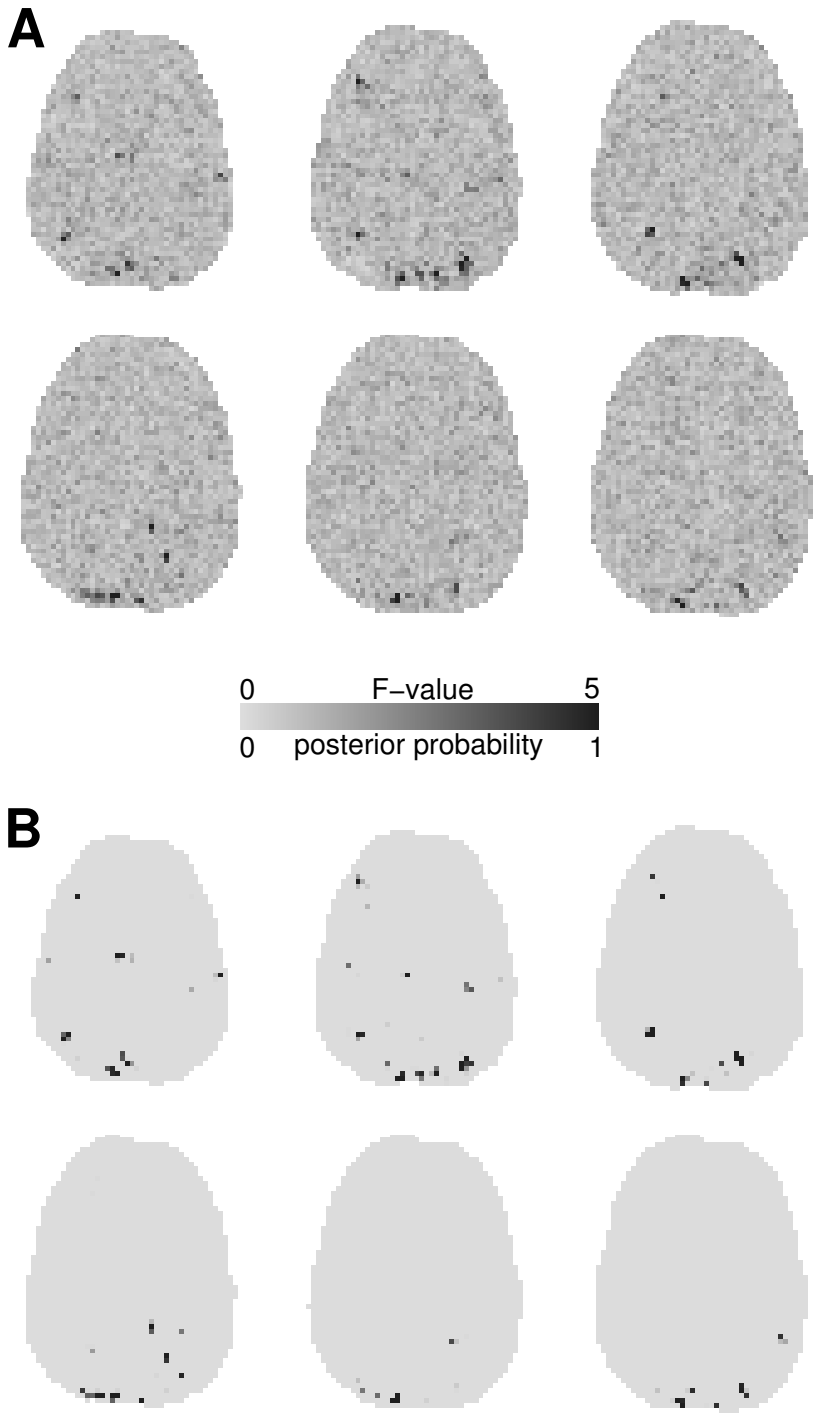

Figure 27. Unthresholded activation maps. (A) Statistical F-value maps generated from the general linear model. (B) Posterior probability maps. Data are from six different slices from the same subject.

Figure 28. Posterior probability map corresponds most closely to the F-map thresholded using the false discovery rate. (A) Activation maps generated with the GLM with three different thresholds. *Uncorrected:* thresholded at $p \leq 0.05$, not corrected for multiple comparisons. *False discover rate:* thresholded at an FDR of 0.05. *Bonferroni:* thresholded at $p \leq 0.05$ after correcting for multiple comparisons using Bonferroni correction. Black voxels exceed threshold. (B) Unthresholded posterior probability map. Maps agree, but F-map *requires* thresholding, while the posterior probability map does not.

3.6 CONCLUSIONS

In this chapter, we have discussed many of the inferential challenges present in all functional imaging studies. We then highlighted some of the statistical approaches that are currently used to address these challenges. Our goal was not to provide an exhaustive review of statistical methodology in fMRI. Rather, our aim was to emphasize that none of the current solutions are entirely satisfactory. Thus, there is a strong incentive for development of more powerful statistical approaches. We then described a flexible statistical model of the BOLD response (Genovese, 2000), and discussed how the model may ameliorate many of the problems that current methods face. Finally, we tested the model by analyzing remapping data using both the Bayesian model and a standard general linear model. We found that the two methods yield comparable estimates: they both reveal similar response shapes across voxels, equivalent spread in noise level within the brain, and a similar classification of voxels as either active or in active.

While the estimates of response shape derived with the two models are comparable, results derived under the Bayesian model confer several advantages. First, access to the full posterior distribution enables inferences that are otherwise difficult to make. The analysis in the subsequent chapter provides concrete examples of such inferences. Second, the posterior probability maps obviate the need for statistical thresholding and multiple comparisons correction. In practice, it is sometimes convenient to apply a threshold to the probability maps (as we do in Chapter 4), but unlike F- or t-maps, there is no hypothesis testing, and hence no threshold is needed. Finally, estimates of response shape derived from the nonlinear fit of the bell function to the data is useful in certain contexts, such as when response latency is a critical factor.

One drawback of our approach is the computation resources necessary to compute the posterior. The difference in computational demands between the two models is substantial. The general linear model that we have described takes about 30 seconds to run on a data set from an 8 minute scan. On the same computer, it takes about 24 hours to run the same data set through the Bayesian model, and it can take many times longer if there are multiple conditions per scanning run. This disadvantage, however, will soon disappear, as computers become faster and more efficient algorithms are developed.

4.0 REMAPPING IN HUMAN EXTRASTRIATE CORTEX

4.1 OVERVIEW

In Chapter 2, we described a functional imaging experiment in which we demonstrated that remapping occurs in humans and that it can be observed with functional MRI. Furthermore, we showed that remapping is present in human parietal cortex in an area that is thought to be functionally analogous to monkey area LIP. In Chapter 3, we described a nonlinear and fully Bayesian approach to analyzing event-related functional imaging data. We showed that the Bayesian method can be used to observe remapping in human parietal cortex, and that it outperforms more traditional analysis methods based on the general linear model, In the current Chapter, we combine the remapping task described in Chapter 2 with the methodology described in Chapter 3 to test the hypothesis that human striate and extrastriate cortex exhibit remapping. The questions addressed in this chapter are (i) whether remapping is present in early and intermediate-level visual areas in the occipital lobe, and (ii) whether the strength of remapping varies systematically across occipital cortex. To address these questions, we identify the borders of several retinotopically defined visual areas. We then use the single-step saccade task to determine whether these areas exhibit remapping. The main finding of this experiment is that the remapping activity is widespread throughout occipital cortex. Furthermore, the strength of remapping is monotonically related to position within the visual area hierarchy. Remapping is strongest in areas V4 and V3A, and it is weakest in V1. These results indicate that remapping is present in visual areas that are directly linked to visual perception.

4.2 INTRODUCTION

Remapping is thought to arise from a neural circuit that includes the lateral intraparietal area (LIP), the frontal eye field (FEF), and the superior colliculus (SC). Neurons in all three areas have spatially selective visual and perisaccadic responses, are modulated by spatial attention, and exhibit remapping in the single-step saccade task (Duhamel et al., 1992a; Walker et al., 1995; Colby et al., 1996; Umeno and Goldberg, 1997, 2001). We asked whether remapping is also present in visual cortex. If remapping is important for perceptual constancy, remapping should not be limited to brain regions with attentional and oculomotor functions. Rather, updated spatial information should reach visual areas that play a more direct role in visual perception.

At the single neuron level, remapping has been demonstrated in several extrastriate regions (Nakamura and Colby, 2002). In the present study, we focus on humans, where we have previously shown evidence for remapping in parietal cortex (Merriam et al., 2003). The current experiment tests the hypothesis that remapping also occurs in human extrastriate cortex. We were encouraged in this endeavor by previous human fMRI studies demonstrating strong *top down* effects throughout occipital cortex. Multiple striate and extrastriate areas are activated in tasks that involve spatial attention (Brefczynski and DeYoe, 1999; Gandhi et al., 1999; Kastner et al., 1999a; McMains and Somers, 2004; Ress et al., 2000; Tootell et al., 1998). Furthermore, many of these areas are modulated by oculomotor signals (Sylvester et al., 2005; DeSouza et al., 2002), suggesting that visual cortex may have access to corollary discharge signals important for remapping.

To test this hypothesis, we scanned the occipital lobes of 12 healthy human participants during the performance of an fMRI version of the single-step task. This task is based on the remapping paradigm used previously to observe remapping in human parietal cortex (Merriam et al., 2003), and is conceptually-related to the paradigm used to observe remapping in monkeys (Duhamel et al., 1992a). Using this task, we found strong evidence for remapping in each extrastriate visual area. Furthermore, we found that remapping was strongest in extrastriate areas V3A and hV4. These results indicate that remapping is present in visual areas that are directly linked to visual perception.

4.3 METHODS

4.3.1 Subjects

We studied a total of fourteen healthy participants (7 female, aged 25-35) participated in this study. All subjects had extensive prior experience with both fMRI and psychophysical experiments. Informed written consent was obtained in accordance with the University of Pittsburgh IRB. All subjects had normal or corrected vision. Data from two subjects were discarded because of technical problems.

4.3.2 Behavioral paradigms

Visual stimuli were generated on a PC computer using the Psychophysics Toolbox (Pelli, 1997) running in MATLAB 6.5 (Mathworks, Natick, MA). Stimuli were presented via an LCD projector and long-throw optics onto a back-projection screen in the bore of the MR scanner. Subjects viewed the projected stimuli through an angled mirror, resulting in a 10° vertical $\times$ 20° horizontal field of view. We measured fMRI activation while subjects performed three tasks, as described below.

4.3.2.1 Single-step task. Subjects fixated one of two crosses at the beginning of the trial. After a variable fixation period (1000 ± 200 ms), a visual stimulus appeared at the center of the screen, 3° above the horizontal axis in the upper quadrant of the right or left visual field. After 1 s, the stimulus was extinguished and a tone cued the subject to make an eye movement to the opposite fixation cross. This saccade brought the screen location of the now-extinguished stimulus into the opposite visual field. The trial ended after a variable period of fixation (2000 ± 200 ms) when a second tone instructed the subject to make a return saccade back to the initial fixation cross.

4.3.2.2 Stimulus-only fixation task. This task is identical to the single-step task, except that there is no auditory cue and subjects do not make a saccade. Two crosses were located 8° to the left and right of screen center. After a variable period (1000 ± 200 ms),

a visual stimulus appeared at the center of the screen, 3° above the horizontal meridian in the upper quadrant of the right or left visual field. Subjects continued to fixate while the stimulus flickered on the screen for 1 s. The trial ended 2000 ± 200 ms after the stimulus was extinguished. Thus, subjects maintained fixation for the entire duration of the trial. Trials of the stimulus-only fixation task in which the stimulus was located in the contralateral visual field were used to identify visually-responsive voxels. Trials in which the stimulus appeared in the ipsilateral visual field were used as a control condition in several analyses (Section 4.4.5).

4.3.2.3 Saccade-only task. This task is identical to the single-step task, except that no salient visual stimulus appears prior to the eye movement. Subjects fixated one of two crosses at the beginning of the trial. After a variable period (2000 ± 200 ms), a tone cued the subject to make an eye movement to the opposite fixation cross. Subjects again maintained fixation for a variable period (2000 ± 200 ms) until a second tone instructed them to make a return saccade back to the initial fixation cross. The timing of the task was identical to the single-step saccade task.

Trials of each condition lasted an average of 4000 ms. In each condition, periods of fixation were variable in duration. In pilot behavioral studies, we found that jittering fixation duration reduced anticipatory saccades. Because of the jittering, the start time of each trial was not yoked to the scanner TR, as is typically the case in fMRI experiments. We recorded trial timing information and scanner sync pulses for every run. We used this timing information in our data analysis to estimate responses at a sub-TR temporal resolution.

4.3.3 Experimental design

Each subject participated in at least two scanning sessions, one for the main experiment, and another for retinotopic mapping. A single scanning session lasted about 1.5 hr. A single session consisted of 6-10 runs. Each run lasted 512 s. On half (3-4) of the runs, all trials began with fixation on the left cross, and on the other runs, trials began with fixation on the right cross. Thus, each of the three tasks was performed in both directions over the course

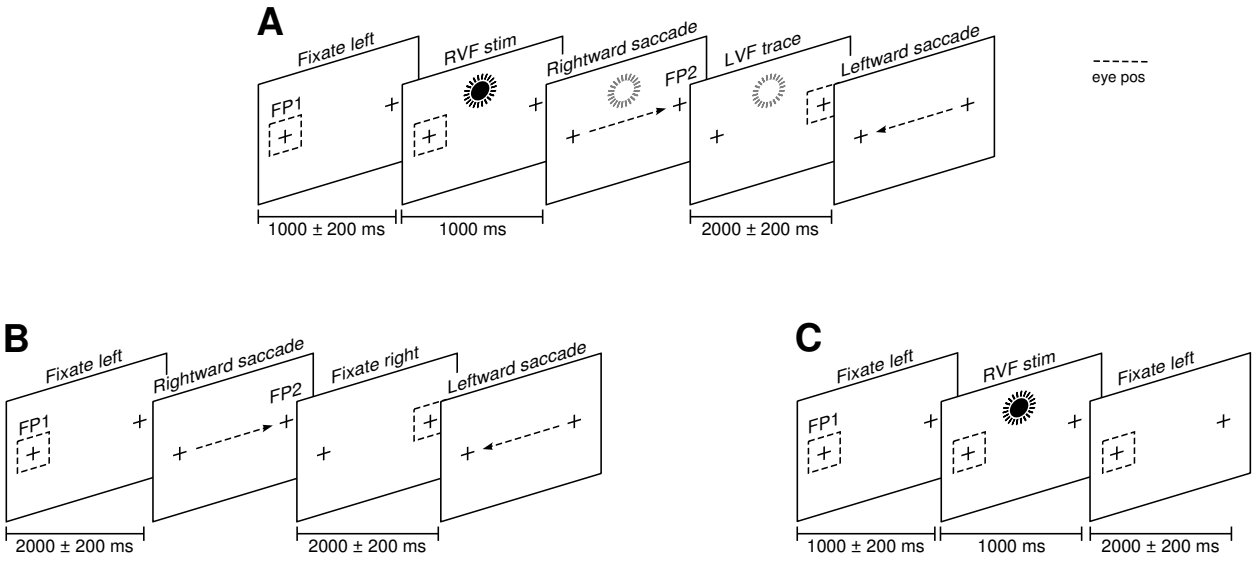

Figure 29. Three task conditions. (A) Single-step task. Subject fixates a cross (FP1) located 8° to the left of screen center. After a variable period (1000 ± 200 ms), a salient stimulus appears in the right visual field and flickers on the screen for 1 s. The stimulus (closed circle) activates contralateral (left) hemisphere visual cortex. Next, the stimulus is extinguished and a tone cues the subject to make a rightward eye movement to FP2. This saccade brings the screen location of the now-extinguished stimulus (open circle) into the left visual field. After a variable period (2000 ± 200 ms), a tone instructs the subject to make a return saccade back to FP1. (B) Saccade-only control task. Subject fixates FP1. After a variable period (2000±200 ms), a tone cues subjects to make an eye movement to FP2. Subjects fixate for a variable period (2000 ± 200 ms) until a second tone cues them to make a return saccade. (C) Stimulus-only fixation task. Subject fixate FP1. After a variable period (1000 ± 200 ms), a visual stimulus appears in the periphery and flickers. The stimulus is extinguished after 1000 ms. The trial ends after a variable fixation period (2000 ± 200 ms). Dashed squares and arrows indicate the location of the eyes.

of the session. This was a critical feature of the experimental design because it enabled us to measure responses in each hemisphere when the stimulus was located in either the contra or ipsilateral visual field. Two directions were never mixed within a run. For example, on runs in which the stimulus appeared in the left visual field, the stimulus never also appeared in the right visual field. Over the course of a single session, subject performed 128-256 repetitions (trials) of each task.

We used a rapid event-related design in which *null* trials were interspersed with the experimental trials. On null trials, subjects simply maintained fixation on the initial cross for the duration of the trial. Thus, on runs in which trials began with fixation on the left cross, null trials consisted simply of fixation on the left cross for 4000 ms, and *vice versa*.

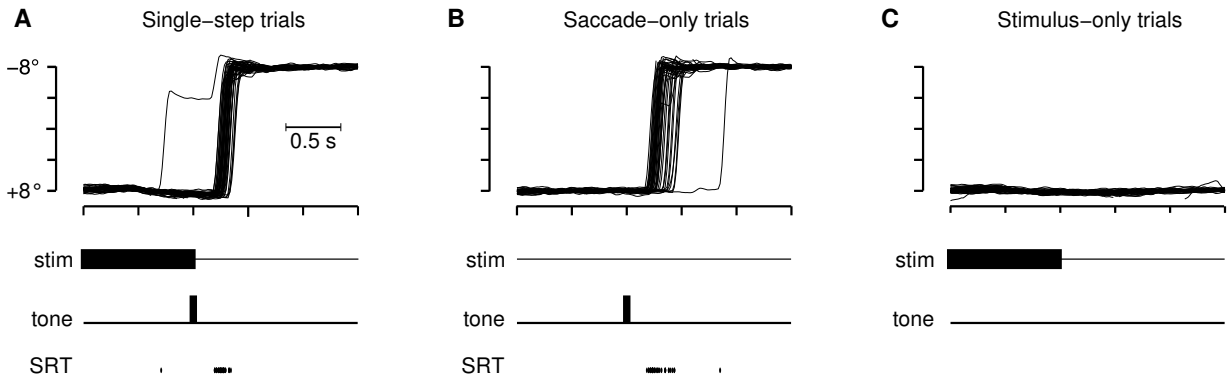

Figure 30. Analysis of eye position data. Eye traces from 64 trials of (A) the single-step task, (B) the saccade-only control task, and (C) the stimulus-only control task. Calculated saccadic reaction time is indicated by tick marks in (A) and (B). This subject made one anticipatory saccade in the single-step task and one late saccade in the saccade-only task. Error trials were excluded from the analysis of fMRI data.

Null trials were matched with experimental trials for orbital position, duration, and frequency. The ordering of experimental and null trials was determined by a special class of pseudorandom sequences known as *m-sequences* (Sutter, 2001; Reid et al., 1997). (sc Matlab functions for generating m-sequences were provided by Thomas Liu, Radiology, UCSD `http://cfmri.ucsd.edu/ttliu/ttliu_software.html`) Randomly generated stimulus sequences often have temporal autocorrelations. These autocorrelations create dependencies in the MR signal that can interfere with response estimation (Liu, 2004; Liu and Frank, 2004). In contrast, m-sequences have a flat autocorrelation function. Thus, m-sequences are perfectly counterbalanced n-trials back, so that trials from each condition, including the null condition, are preceded equally often by trials for each of the other conditions. This property makes m-sequences well suited for use in fMRI experiments (Buracas and Boynton, 2002). Several recent functional imaging studies have demonstrated the advantages of using m-sequences (e.g., Hansen et al., 2004; Murray and Wojciulik, 2004; Kellman et al., 2003).

4.3.4 Eye position recording

We monitored eye position during each fMRI session using a video-based eye tracker (ASL, Boston, Mass). The eye tracker has a temporal resolution of 60 Hz. Stability of eye data was typically better than 2°, as determined by the SD of the data during periods of stable fixation.

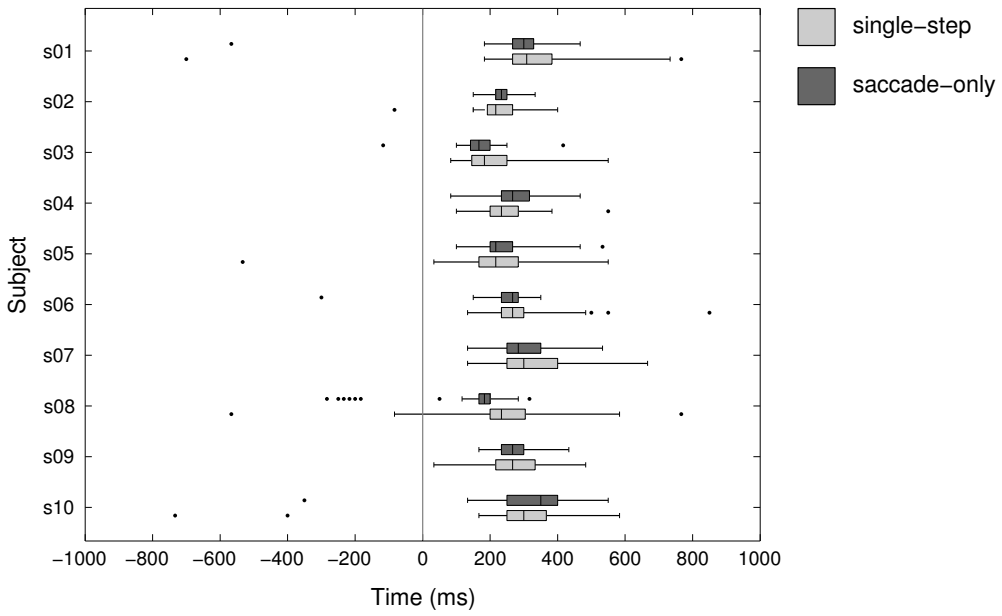

Figure 31. No difference in saccade reaction time in trials of the single-step task (light gray bars) and saccade-only control task (dark gray bars). Outliers are indicated by tick marks.

Analysis of the eye position data ensured that subjects (1) maintained stable fixation within a 2° window on FP1 during the 1 s of visual stimulation, (2) made horizontal eye movements in a 500 ms temporal window following the auditory cue, and (3) made accurate saccades (to within 2°) on single-step and saccade-only trials. We analyzed each subject's eye position data on a trial-by-trial basis and discarded trials in which subjects failed to meet this set of performance criteria. All subjects made errors on fewer than 5% of trials.

We measured fixation stability by calculating the median absolute deviation (MAD) of the horizontal eye position. There were no differences in MAD values between any of the three conditions during the first 1 s of the trial. During this period, a salient visual stimulus was present on the single-step and stimulus-only trials, but not on the saccade-only trials. The analysis of fixation stability during the period of visual stimulation indicates that fixation was not affected by the presence of the stimulus.

We used an automatic saccade finding algorithm to search for saccades in the first two seconds of each trial (ILAB software, `http://groups.yahoo.com/group/ilab`; Gitelman, 2002). The software identified the occurrence of a saccade if eye velocity exceeded 50°/s and the eyes moved more than 2°. Subjects were expected to have made saccades in response

to the auditory cue on single-step and saccade-only trials, but not on stimulus-only trials. Mean saccadic reaction time was 255 ms (128 ms SD) and 269 ms (136 ms SD), respectively, in the single-step and saccade-only conditions. The difference in SRT was not significant (t-test, $p > 0.05$). On occasional trials, subjects made a saccade prior to the cue. These anticipatory saccades appear as tick marks (outliers) in the boxplots of saccadic reaction time in Figure 31. These trials were discarded from the analysis of the fMRI data.

4.3.5 MRI data acquisition and preprocessing

4.3.5.1 Structural data. We acquired a 3D anatomical volume in each scanning session using an MPRAGE pulse sequence (30 ms TE, 8° flip angle). Anatomical volumes consisted of 192 sagittal slices at a 1 mm^3 voxel resolution. We combined multiple (2-3) structural scans across sessions to increase the signal to noise ratio. Gray and white matter were segmented from the anatomical volume using FreeSurfer software (Fischl et al., 1999; Dale et al., 1999, `http://surfer.nmr.mgh.harvard.edu`). The structural data were subject to several preprocessing steps prior to brain segmentation. Steps included motion correction, nonlinear intensity normalization to correct for inhomogeneity, averaging, and skull stripping. Brain segmentation was followed by tessellation of the cortical surface and the creation of 3D surface models representing the border between white and gray matter.

We visualized activation on 2D flattened representations of the cortical surface (Figure 32). Flat maps were created using custom software adapted from Wandell et al. (2000) (`http://white.stanford.edu/software`) and implemented in MATLAB. This flattening procedure does not introduce topological errors (twists or folds) and it minimizes metric distortions. Pilot analyses from our lab indicate that this method outperforms other commonly used brain-flattening methods implemented in FreeSurfer and BrainVoyager™ software. The primary disadvantage of this method is that it does not perform well on large patches of cortex. Consequently, we limited the flattened cortical patches in this study to a radius of 50 mm.

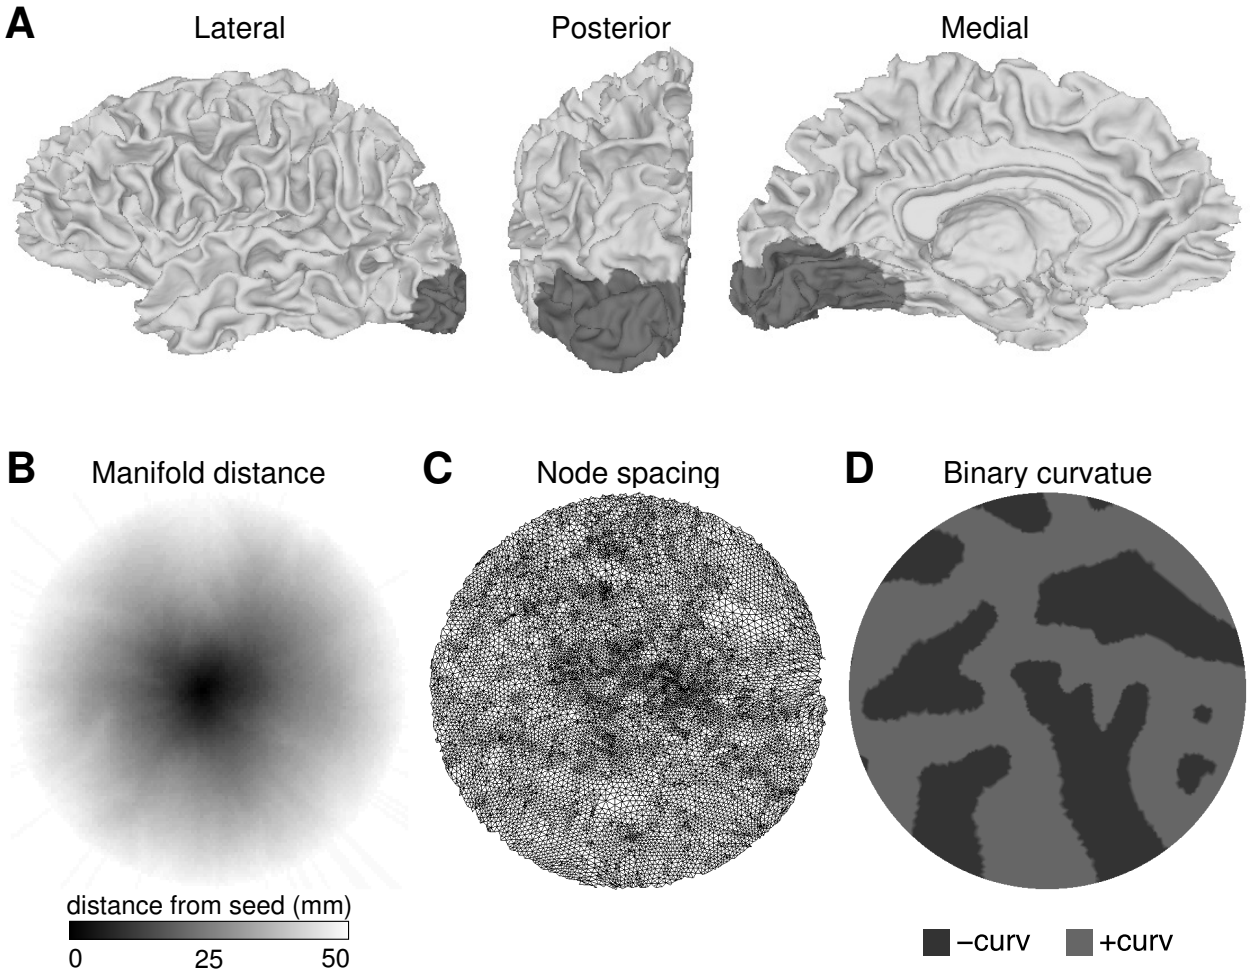

Figure 32. Occipital lobe patch. (A) Three-dimensional cortical manifold. A seed voxel was selected at the occipital pole. The patch (gray shading) was defined used a region-growing algorithm that selects a perimeter of nodes that have equal distance from the seed. (A) Distance (mm) between the seed node and all other nodes in the flattened cortical patch. The distance map is radially symmetrical, indicating that the flattening procedure did not introduce large distortions. (B) Node locations in the flattened map. Nodes are distributed evenly, with a slight increase in density at the center. Even spacing indicates that the flattening procedure did not introduce folds or twists. (C) Binary curvature map. Faces in the tessellated surface were assigned one of two grayscale values depending on the local curvature of the 3D manifold.

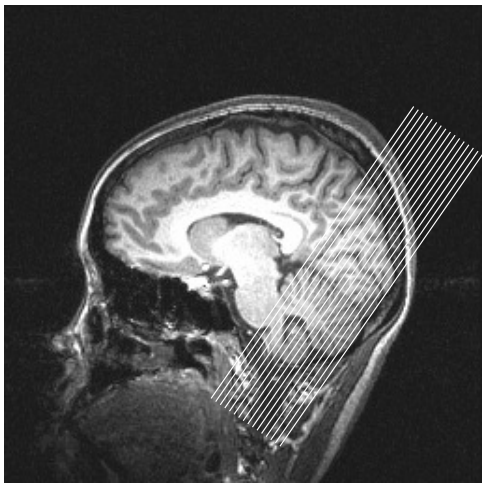

Figure 33. Slice prescription. Sixteen slices were oriented perpendicular to the calcarine sulcus. Slices were 3 mm thick and covered the entire occipital lobe.

4.3.5.2 Functional data. We used functional magnetic resonance imaging at 3T (Allegra, Siemens, Erlangen) and a T2*-sensitive EPI pulse sequence to measure changes in BOLD activity. Scan parameters were as follows: TR = 1000 ms; TE = 30 ms; flip angle = 65°. We collected 16 slices (3 mm^3 voxels, 192 mm FOV) in each volume. Slices were oriented perpendicular to the calcarine sulcus in order to cover the entire occipital lobe. We collected a time-series of 512 volumes in each functional run. Subjects participated in multiple sessions. We scanned subjects in 8-10 runs in each retinotopy session, and in 6-8 runs in each session for the remapping experiment.

We preprocessed the functional data using FIASCO software (McNamee and Eddy, 2001, `http://stat.cmu.edu/~fiasco`). Preprocessing steps included correction for fluctuations in mean intensity; motion correction of the raw, complex-valued k-space data (Eddy et al., 1996); image reconstruction; and outlier correction using a Windsor filter. Outliers were defined as data points farther than ten times the interquartile range from the median. The image data were not smoothed, temporally filtered, or spatially normalized.

4.3.5.3 Functional and structural alignment. Functional data for each subject were registered to the subjects whole-brain 3D anatomical images using a fully automated algorithm implemented in FIASCO software. We identified functional voxels that intersected the

white/gray matter boundary using tools that are part of the AFNI software package. This procedure is thought to sample neuronal activity preferentially from deeper cortical layers (Tootell et al., 1997a). We further selected voxels based on inclusion within the boundaries of predefined cortical visual areas. Visual area ROI's were defined during retinotopic mapping experiments. We considered only voxels that were contained within a single ROI for a given visual area.

4.3.6 Retinotopic mapping

We used phase-encoded visual stimuli to identify the borders between occipital lobe visual areas V1, V2, V3, V3A, and hV4 (Engel et al., 1994, 1997; DeYoe et al., 1996; Sereno et al., 1995). Retinotopic mapping was carried out in a separate scan session in which subjects underwent 6-8 runs of phase-encoded eccentricity and polar angle mapping. Stimuli were counter-phase, color-modulated flickering checkerboards that took the shape of rotating wedges and expanding/contracting annuli. The spatial frequency of the checkerboards was scaled to accommodate larger receptive fields in the periphery. Stimulus movement was periodic, with a frequency of 1/64 s. The stimulus completed 8.5 cycles per run. We used the phase cancellation technique described by Kalatsky and Stryker (2003) in which the direction of stimulus movement was reversed on successive runs. We summed the complex-valued data prior to calculating the phase and magnitude of the response. This procedure removes the hemodynamic delay associated with the BOLD response, thereby yielding more accurate estimates of the stimulus position that elicited the maximal response.

Visual area boundaries were defined using a conjunction of polar angle and eccentricity maps according to the following three criteria, as described by Dougherty et al. (2003). First, each area was bounded by phase reversals in the angular component of the retinotopic map. Second, a given area had to be activated by both the wedge and annulus stimuli. Third, the phase gradient in the angular and eccentricity maps had to run orthogonal to one another. Both the dorsal and ventral portions of areas V1, V2, and V3 were easily identified using these criteria.

We identified area hV4 using the criteria described by Tootell and Hadjikhani (2001).

Area hV4 is the ventral retinotopic area that continues laterally from ventral V3. Area hV4 has a full hemifield representation and is located proximal to the medial lip of the collateral sulcus. We refer this area as "hV4" because the degree of functional coorespondence between this region and monkey area V4 has not been fully resolved (Tootell and Hadjikhani, 2001; Brewer et al., 2005). It is possible that the cortical area that we labeled hV4 also contains additional subdivisions (e.g., areas VO-1 and VO-2; Brewer et al., 2005).

We identified area V3A using the criteria described by Tootell et al. (1997b). Area V3A is the dorsal retinotopic area that continues anteriorly from dorsal V3. Area V3A contains a full hemifield representation and is located proximal to the transverse occipital sulcus, at the base of the intraparietal sulcus. It is likely that the cortical area that we identified as V3A also contained other cortical areas (e.g., V3B and LO2; Press et al., 2001; Larsson and Heeger, 2005). Area V7 was not reliably identified in our data, and was therefore not included in this study.

Of the 24 hemispheres studies, all but two had clean retinotopic maps in which the borders between visual areas could be discerned. In two hemispheres (one left, one right) from two different subjects, retinotopic maps were ambiguous. These hemispheres were not included in the analysis. Hence, all analyses report results from $n = 22$ hemispheres.

4.3.7 Statistical modeling

We used a fully Bayesian approach to analyzing the MR data, the details of which have been described in Chapter 3. The primary output of the Bayesian model fit is the posterior distribution of the parameters given the data, $\mathsf{P}\{\gamma \mid Y\}$. From this distribution, we derived point estimates, such as posterior means and standard deviations on a given parameter. The statistical inferences in this study are based on these Bayesian estimates. We describe three such inferences below.

First, we calculated the probabilities of several specific events, such as the occurrence of a visual or a remapped response. For example, in Sections 4.4.3 and 4.4.4, we calculate the probability that there is a nonzero response to the visual stimulus in the stimulus-only fixation task, or to a stimulus trace in the single-step task. For condition "c", this is denoted

by $\mathsf{P}\{\gamma_c^{\text{resp}} > 0 \mid Y\}$. Because the hierarchical model allows for a nonzero probability on the discrete value 0, this probability indicates the strength of evidence for any response. Second, we consider the probability that the response in one condition is greater than the response in another. For example, in Section 4.4.3, we calculate the probability that responses to contralateral visual stimuli are larger than responses to ipsilateral visual stimuli. This is denoted by $\mathsf{P}\{\gamma_c^{\text{resp}} > \gamma_{c'}^{\text{resp}} \mid Y\}$. We make similar comparisons for shape parameters in Section 4.4.7 when comparing response onset times across conditions. Third, we compute the posterior probability of more complex events, such as the probability that remapped responses are larger than the maximum of the two control conditions (Section 4.4.5). This is denoted by $\mathsf{P}\{\gamma_{\text{remap}}^{\text{resp}} > \max(\gamma_{\text{sac}}^{\text{resp}}, \gamma_{\text{stim}}^{\text{resp}}) \mid Y\}$.

One of the key advantages of the Bayesian framework is that the same analytic procedure can be used to address a wide range of questions (see Chapter 3). In order to compute each of the values described above, we performed Monte Carlo simulations to determine the probabilities of interest given the posterior distribution of the relevant parameters. We used this approach to determine probabilities for individual voxels, collections of voxels in individual hemispheres, and the population of hemispheres in the group of subjects. For example, the group analysis in Figure 39 is based on the following analysis. Within hemisphere j, let $\overline{\gamma}_{j,\text{remap}}^{\text{resp}}$ denote the average of the remapped response magnitudes (proportional signal change) for all voxels that show a visual response. If we further assume that the responses in different hemispheres are independent, then we do a similar combination to determine the posterior distribution for the population averaged response. All of these posterior probabilities were calculated to arbitrary precision by Monte Carlo simulation. Note that we report the posterior probability, p, which should not to be confused with a p-value from a standard statistical test. Large probabilities (i.e., approaching 1) indicate strong evidence for the particular statement. A probability of 0.5 indicates that the response could have been observed by chance.

4.4 RESULTS

Our central finding is that the single-step task elicits responses in the hemisphere that was initially ipsilateral to the stimulus. This activation cannot be attributed to either the stimulus or to the saccade. We argue that this activation is a response to the updated stimulus trace. We begin by illustrating this result by describing responses from a single right hemisphere hV4 voxel that exhibits remapping (Section 4.4.1). This single-voxel is representative of the larger population of voxels from 12 subjects. In subsequent sections, we combine data across hemispheres and subjects to show that the main findings illustrated in this example voxel are characteristic of the larger population.

4.4.1 Single voxel example

4.4.1.1 Responses estimated with general linear model. We begin by illustrating our main findings with responses from a single right hemisphere hV4 voxel that exhibits remapping (Figure 34). Each curve represents the averaged response from 64 trials per condition. For purposes of this illustration only, the responses were estimated using standard analysis procedures based on the general linear model (Dale and Buckner, 1997). This voxel is strongly activated by the appearance of a stimulus in the contralateral visual field during the fixation task (Figure 34A). The MR response in this condition is referred to as the *visual response* because it is driven by direct visual stimulation. In contrast, the voxel did not respond when a stimulus was presented in the ipsilateral hemifield without any saccade (Figure 34C) nor did it respond in conjunction with a saccade from FP1 to FP2 when no stimulus was present (Figure 34D). Nevertheless, this voxel did respond after the saccade brought the screen location where the stimulus had recently appeared into the contralateral visual field (Figure 34B). We interpret this as a response to the trace of the stimulus. The representation of the stimulus trace has been updated in conjunction with the saccade.

This interpretation is supported by analysis of eye position data. We calculated saccade latency on each of 64 trials of the single-step task from eye position data recorded during scanning (Figure 34B, eye traces). This subject had an average saccadic reaction time of

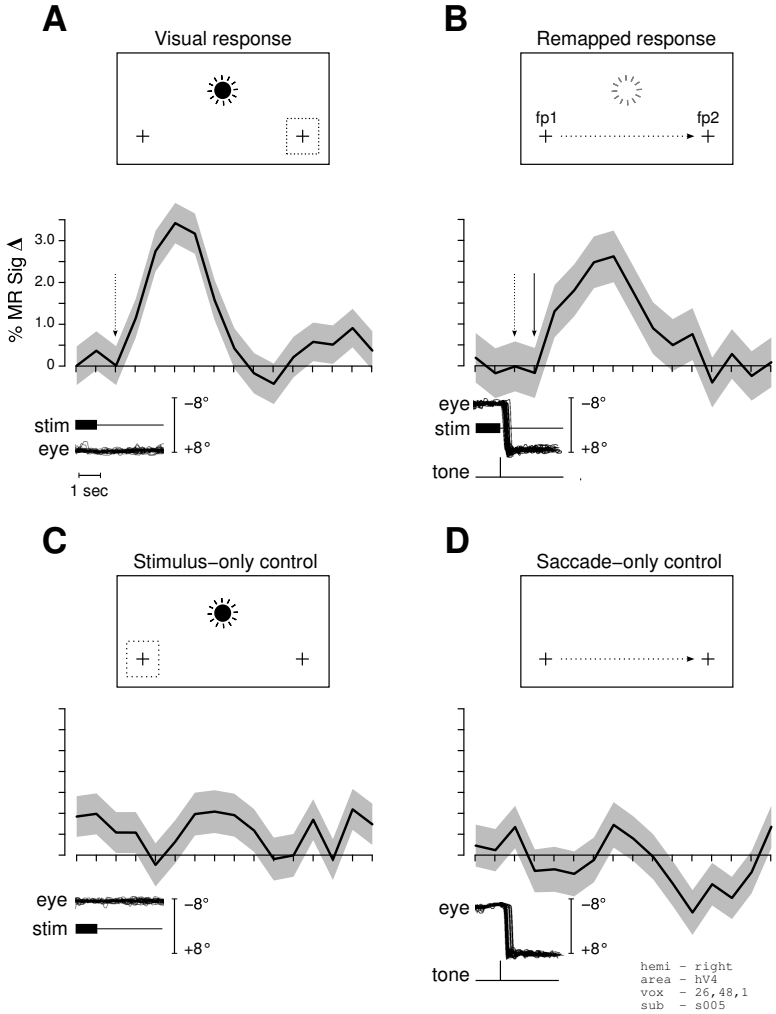

Figure 34. A single right hemisphere hV4 voxel that exhibits remapping. The cartoon shows the location of the stimuli on the screen. Horizontal eye position and timing of stimulus events are shown below (calibration bar, 16°). (A) Visual response in fixation task. A contralateral stimulus during fixation elicits a strong response. (B) Remapped response in single-step task. The subject fixates FP1 as a stimulus flickers in the ipsilateral visual field for 1 s. After 1 s, the stimulus is extinguished and a tone cues the subject to make a saccade to FP2. The saccade brings the screen location of the extinguished stimulus into the contralateral visual field. The remapped trace of the stimulus elicits a response. (C) Stimulus-only control. Presentation of the stimulus in the ipsilateral visual field does not elicit a strong response in the absence of a saccade. (D) Saccade-only control. The saccade alone does not elicit a strong response in the absence of a stimulus. Each curve was estimated from responses on 64 trials.

257 ms (55 ms SD). Thus, the eyes began to move only after the stimulus had already been extinguished. Because the stimulus was never physically present in the contralateral visual field, we conclude that this voxel responded to the remapped *trace* of the stimulus.

The temporal profile of the response on single-step trials also indicates that the response is driven by the remapped trace of the stimulus rather than the stimulus itself. In the single-step task, the stimulus appears and stays on the screen for 1 s prior to the cue to make an eye movement. The response to the remapped stimulus trace should thus occur about 1 s later than the response to the visual stimulus. The example voxel in Figure 34 indicates that this is in fact the case. The visual response in Figure 34A begins to rise at about 2 s after the onset of the stimulus, consistent with the time course of visually-driven hemodynamic response curves (Boynton et al., 1996). In contrast, the remapped response in Figure 34B begins to rise about 3 s after the onset of the stimulus. The latency difference between these two curves is about 1 s, corresponding to the period between the onset of the stimulus and the cue to initiate a saccade.

The example voxel in Figure 34 illustrates four response properties that characterize remapping. This voxel (1) responded to the contralateral visual stimulus in the fixation task; (2) responded to the remapped stimulus trace in the single-step task; (3) did not respond strongly in either of two control conditions; and (4) responded in the single-step task at a latency predicted by the timing of the task. Note that each of these properties directly parallels observations made at the single-cell (see Figure 6).

4.4.1.2 Responses estimated with Bayesian model. We used the Bayesian statistical model described in Chapter 3 to evaluate each of these four properties. In order to illustrate how the model parameters relate to each response characteristic, we show the same voxel illustrated in Figure 34, but here the response curves were generated from the posterior distributions of the response profile, estimated under the Bayesian model (Figure 35). Each of the four properties described above are evident in Figure 35. First, this voxel responded strongly when the stimulus was located in the contralateral visual field (Figure 35A). The estimated magnitude of this visual response is 3.42% (relative to baseline). There is a high probability ($p \geq 0.95$) that this visual response is greater than zero.

Second, this voxel responded strongly in the single-step task when the stimulus appeared in the ipsilateral visual field and an eye movement brought the recently-stimulated screen location into the contralateral visual field (Figure 35B). The estimated magnitude of this remapped response is 2.62% of baseline, 77% of the visual response magnitude. There is a high probability ($p \geq 0.95$) that this remapped response is greater than zero.

Third, this voxel did not respond strongly in either of the control conditions. The estimated response magnitude in the stimulus-only control condition was 1.09% baseline, 32% of the visual response (Figure 35C). The estimated responses magnitude in the saccade-only condition was 0.67% baseline, 20% of the visual response (Figure 35D). The posterior probabilities that the responses in the stimulus-only and saccade-only conditions are greater than zero is $p = 0.80$ and $p = 0.02$, respectively. This voxel's response to the ipsilateral stimulus is greater than chance, indicating that the ipsilateral stimulus did elicit a response. Can the response in the single-step task be attributed to the ipsilateral stimulus? A direct comparison between the response magnitudes in Figure 35B and 35C suggests that it cannot. There was a high probability ($p \geq 0.95$) that the response in the single-step task is larger than the response in the stimulus-only control condition. We conclude that the remapped response can not be attributed to the ipsilateral visual stimulus or to saccades alone.

Lastly, as discussed above, the remapped response should begin about one second later than the visual response. The estimated response latency (i.e., the *lag* parameter) is 1.93 s for the visual response and 2.62 s for the remapped response. The difference between the two lag estimates is 0.69 s. There is a high probability the two lag estimates differ ($p \geq 0.95$). The difference in latency indicates that the response in Figure 35B was driven by the updated stimulus trace rather than the actual visual stimulus. The difference in lag between the two conditions is about 300 ms less than predicted by the timing of the task. This may indicate that remapping was predictive.

4.4.2 Roadmap

We have described four response properties indicating that activity in the single-step task is related to remapping. In the following sections, we characterize each of these properties

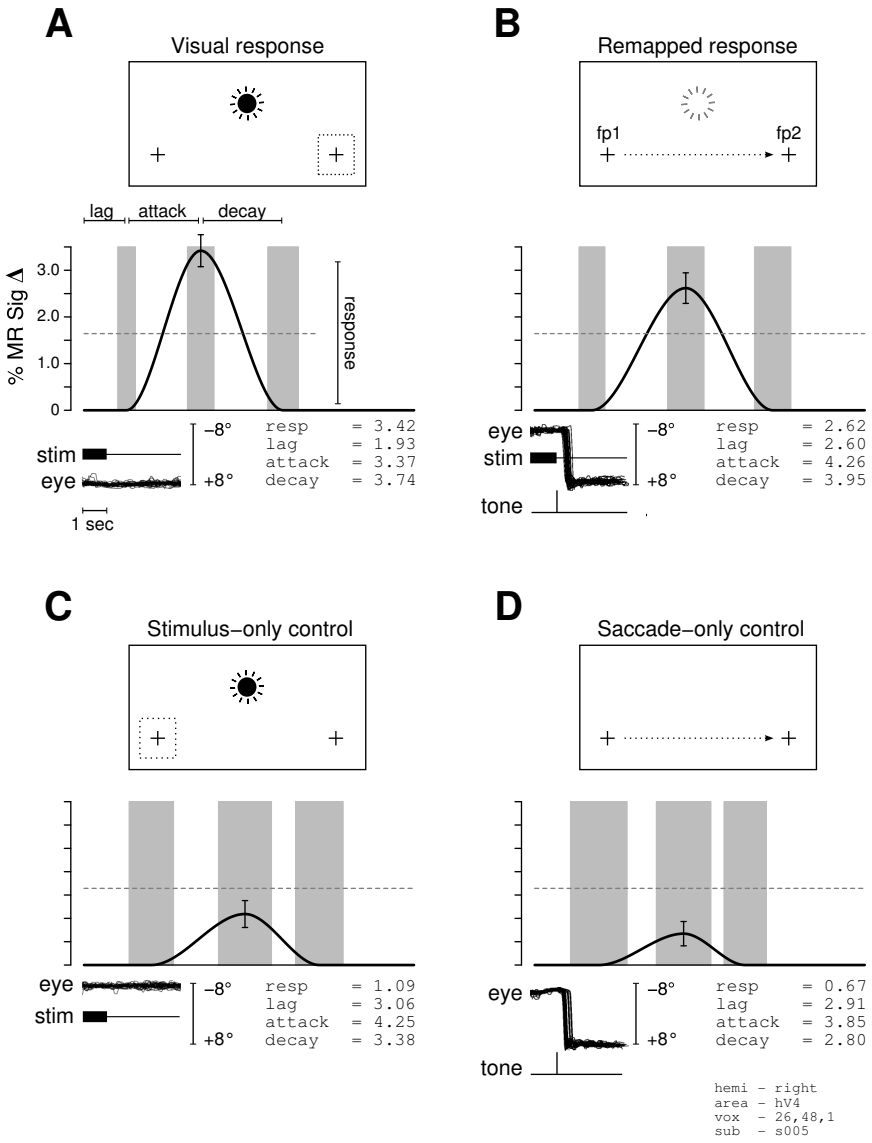

Figure 35. Response estimated with the Bayesian model rather than with the GLM. Same voxel shown in Figure 34. Response curves reflect the posterior distribution of four model parameters. *Response* is the magnitude of activation, expressed in units of percent signal change baseline. *Lag* is the time between the onset of the stimulus and the onset of the response. *Attack* is the time from response onset to response peak. *Decay* is the time from response peak to the point at which the response has returned to baseline. Vertical gray bars indicate the posterior standard deviation around each parameter estimates (e.g., wide bars indicate low confidence in the mean). Vertical error bars indicate the standard deviation of the response parameter. Dotted horizontal line indicates the estimated noise level.

in a larger population of voxels from 12 subjects. Our approach is as follows. We first applied the Bayesian model at the individual voxel level in order to derive joint posterior distributions of each model parameter for each voxel for each condition. We then combined parameter estimates across voxels in a given visual area (for each hemisphere) and then across hemispheres, yielding group-level posterior probabilities. The use of Bayesian statistics is not widespread in the imaging literature. For many of the key inferences in this experiment, we also used standard statistical tests to determine whether a given effect is *significant* in the classical sense. In all cases, results from the two approaches agree.

In Section 4.4.3, we describe cortical responses to the visual stimulus. These visual responses were measured in areas V1, V2, V3, V3A, and hV4 in the contralateral hemisphere during the stimulus-only condition. We use the results of this analysis to select the subset of voxels in each area that are visually-responsive; all subsequent analyses were performed on this population of voxels. In Section 4.4.4, we describe responses to the stimulus trace. These remapped responses were measured in the single-step task in the hemisphere that was initially ipsilateral to the stimulus. It was critical to ensure that ipsilateral stimuli and saccades alone do not account for activity in the single-step task. In Section 4.4.5, we describe responses in the stimulus-only and saccade-only control conditions. We then directly compare the magnitude of responses in the single-step task to responses in these two control conditions (Section 4.4.6). This comparison is critical to our central finding — that remapping occurs in human visual cortex. We thus make this comparison using three different analysis strategies, each of which is based on a different set of assumptions. Lastly, in Section 4.4.7, we test the hypothesis that remapped responses have a latency consistent with the expected arrival time of updated visual information in cortex.

4.4.3 Responses to visual stimuli

In the version of the single-step task used in neurophsiology experiments, an eye movement brings the receptive field of the neuron onto the recently-stimulated screen location (see Figure 6; Duhamel et al., 1992a). Remapped activity is observed as a response to the stimulus trace, in the absence of any direct visual stimulation. Implicit in the logic of this paradigm is

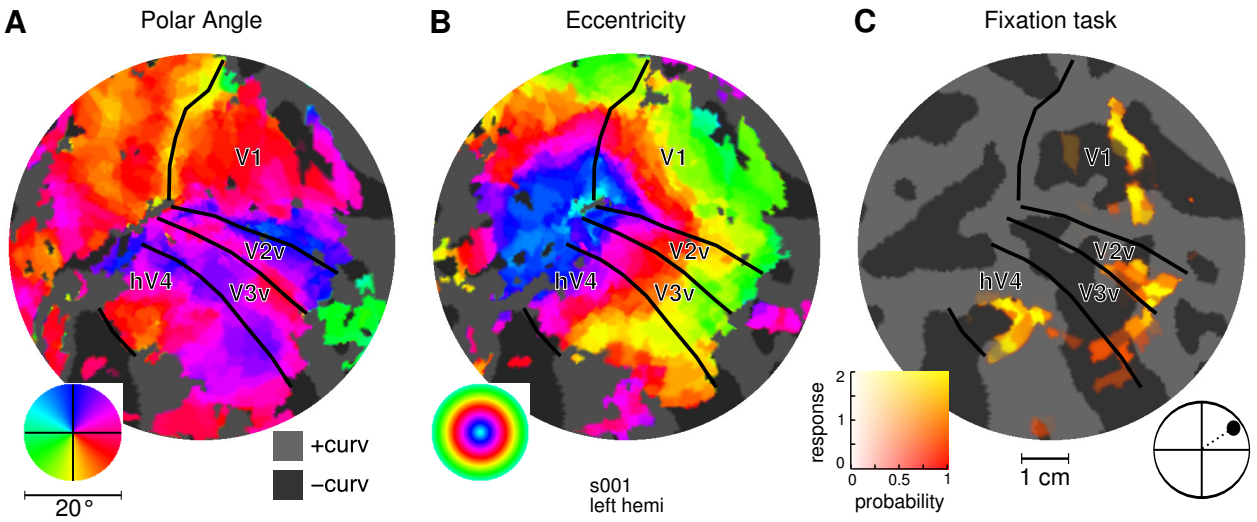

Figure 36. Visual responses in ventral occipital cortex. The field of view includes all of area V1, the ventral portions of areas V2 and V3, and all of area hV4. Each disk shows the same flattened patch of cortex with a 50 mm radius. (A) Representation of polar angle. The stimulus was a rotating 15° checkerboard wedge. (B) Representation of visual eccentricity. The stimulus was an expanding/contracting 15°checkerboard annulus. In both (A) and (B), hue represents the stimulus location that elicited the maximal response. (C) Response to contralateral visual stimulus in fixation task. Hue (red-yellow) represents the magnitude of the response. Color opacity (transparent-opaque) represents the probability that the response to the stimulus was nonzero. Opacity values range from 0 to 1; the data are not thresholded. Gray background represents estimated cortical curvature: (dark gray, concave; light gray, convex). **Note:** retinotopy images in (A) and (B) are optimized for on-screen viewing; they contain RGB colors outside the range of some printers.

the assumption that the neuron fires when an actual stimulus appears inside of its receptive field. The goal of this section is to identify the set of voxels that exhibit visual responses — these are the same voxels which we predict will also exhibit remapping.

We initially screened voxels using large, flickering checkerboard stimuli in a retinotopic mapping experiment. We identified visually-responsive gray matter voxels in each area through a combination of cortical segmentation and standard mapping procedures (Engel et al., 1994, 1997; DeYoe et al., 1996; Sereno et al., 1995). The checkerboard stimuli encompassed a large portion of the visual field (20° horizontal × 15° vertical). These stimuli thus activated a broad region of cortex, extending from the foveal to the peripheral cortical representation. We show the activated cortical regions in two figures, one illustrating activation in ventral occipital cortex, including the ventral portions of areas V1, V2, V3, and

V4 (Figure 36A-B), and another illustrating activation in dorsal occipital cortex, including the dorsal portions of areas V1, V2, V3, V3A, and V7 (Figure 37A-B).

The single-step task involves remapping the representation of a small stimulus located in the contralateral, upper visual field (Figure 29). Thus, only a subset of the voxels activated by the checkerboard stimuli were also activated by the stimulus in the single-step task. We identified the subset of voxels that responded to this stimulus by analyzing activation in the stimulus-only fixation task. For each voxel, we estimated the magnitude of the stimulus-evoked response, $\gamma^{\text{resp}}_{\text{visual}}$, and the posterior probability that this magnitude was greater than zero given the data, $\mathsf{P}\{\gamma^{\text{resp}}_{\text{visual}} > 0 \mid Y\}$. Results from this analysis are plotted on the cortical surface in order to visualize the degree of spatial correspondence between activation in the fixation task and location within the retinotopic maps (ventral cortex, Figure 36C; dorsal cortex, Figure 37C). In these plots, activation magnitude is represented by a red–yellow color scale: voxels with large visual responses are yellow and voxels with smaller visual responses are red. Posterior probability is represented by color opacity, with zero probability being fully transparent. Activation maps were not thresholded. Yellow voxels tend to be more opaque because large responses tend to have a higher probability of being nonzero.

The activation maps illustrate four properties of visual responses. First, activated voxels were located in the appropriate region of the retinotopic maps. In the fixation task, the stimulus appeared 3° above the horizontal axis and 8° from the vertical axis, in the upper quadrant of the right or left visual field. We thus expected the stimulus to activate the upper visual field representation at 8° to 10° eccentricity. This location in cortex is indicated by shades of magenta in the polar angle map and by shades of yellow/green in the eccentricity map. These activation maps indicate that visual responses in the fixation task were located in the subregion of each visual area that correspond to the appropriate location in the retinotopic map.

Second, there is a clear distinction between active and inactive voxels. Posterior probabilities tended to be either high (far greater than chance, $p \gg 0.5$), or low (far less than chance, $p \ll 0.5$). Because of this property, voxels appear as either fully opaque or completely transparent (p approaches 0 in cortical locations in which the underlying grayscale anatomy is clearly visible). The clear classification of voxels as either active or inactive is illustrated

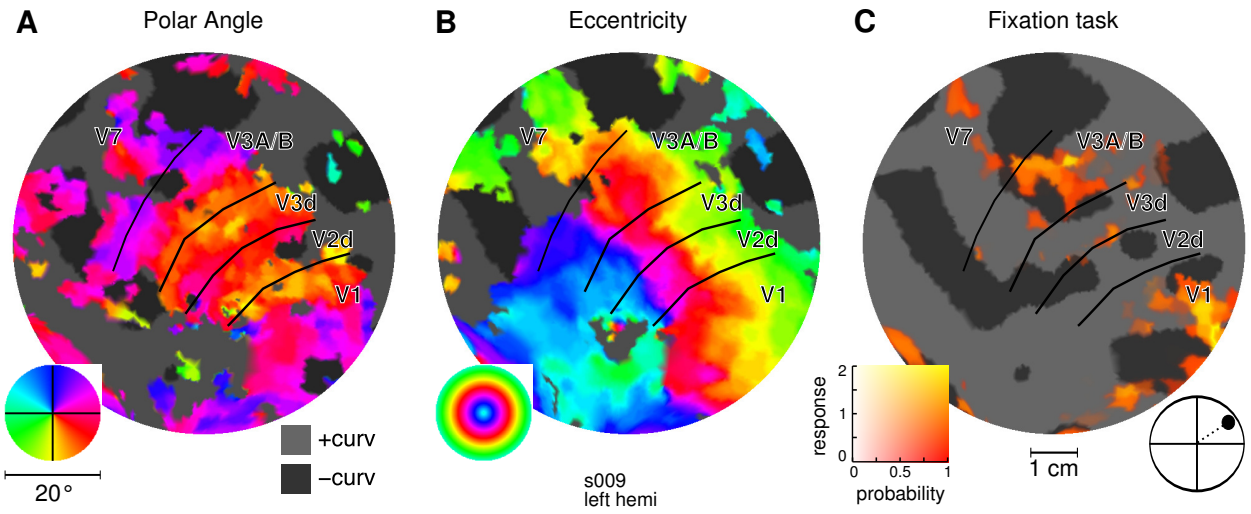

Figure 37. Visual responses in dorsal occipital cortex. Flattened patch of occipital cortex (50 mm radius) with a field of view including the dorsal portions of V1, V2 and V3, and all of V3A and V7. (A) Representation of polar angle. (B) Representation of visual eccentricity. (C) Response to contralateral visual stimulus in stimulus-only fixation task. Conventions are the same as in Figure 36. **Note:** retinotopy images in (A) and (B) are optimized for on-screen viewing; they contain RGB colors outside the range of some printers.

in a strip chart of probability values (Figure 38A). Voxels clustered at both the top ($p = 1$) and bottom ($p = 0$) of the chart — relatively few voxels had intermediate probabilities. All subsequent analyses in this chapter were performed on the subset of voxels in each area that met a $p \geq 0.95$ selection criteria for contralateral visual stimuli (Figure 38A, above dashed horizontal line). The sharp distribution of probability values indicates that the particular threshold did not have a strong impact on which voxels were included in the analysis.

Third, the number of voxels activated by the stimulus in the fixation task decreases slightly at later stages of the visual area hierarchy. Two factors affect the number of voxels that are activated: visual area size and receptive field size. Visual area size decreases at successive stages of the visual area hierarchy (Dougherty et al., 2003). Because there are fewer voxels in area hV4 than in V1, the stimulus may activate a smaller number of voxels in area hV4. On the other hand, receptive field size increases at successive stages of the visual area hierarchy (Gattass et al., 1981, 1988). This increase in receptive field size results in MR responses that are less spatially selective (Kastner et al., 2001; Smith et al., 2001). A decrease in spatial selectivity may result in a larger proportion of hV4 voxels that are

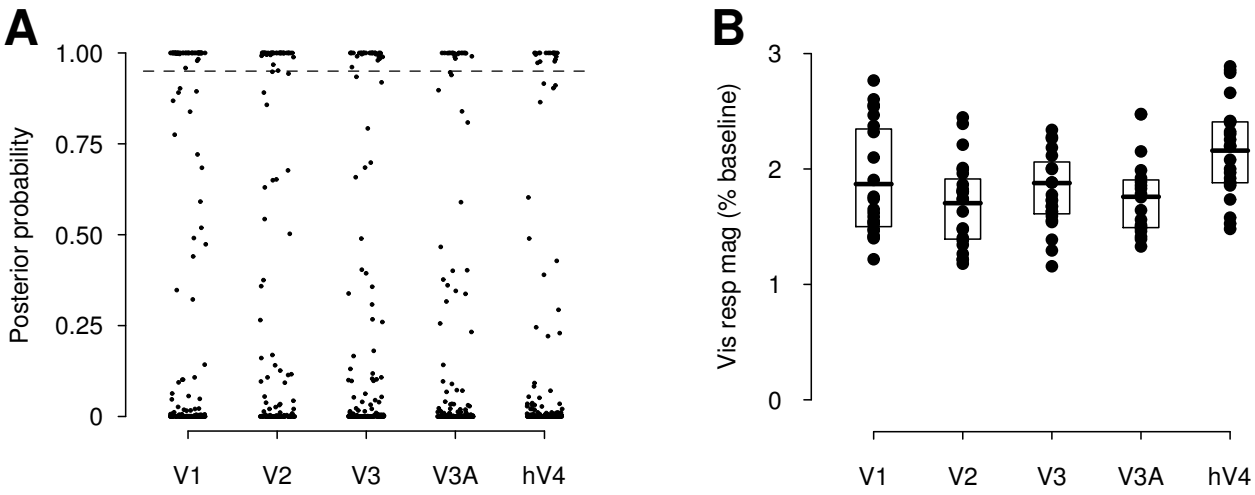

Figure 38. Visual response magnitude. (A) Posterior probability that visual responses in the stimulus-only fixation task are nonzero. Each dot represents a single voxel. Figure shows voxels from a single hemisphere. Note that probabilities cluster at $p = 0$ and $p = 1$, indicating that voxels were classified as either visually-responsive or not. All subsequent analyses were based on the subset of voxels that met a $p \geq 0.95$ selection criteria (dashed line). (B) Magnitude of visual responses for the subset of responses that reached threshold. Each dot represents the response magnitude averaged across all suprathreshold voxels in a single visual area within a single hemisphere. The bottom of the box represents the first (lower) quartile of the population of hemispheres, the thick horizontal line indicates the median, and the top of the box represents the third (upper) quartile. This plotting convention is used throughout this chapter.

responsive to a given stimulus. Thus, the decrease in cortical size in later stages of the visual area hierarchy could be countered by the increase in receptive field size.

We found that the visual stimulus activated slightly fewer voxels in higher-order visual areas. This trend was not evident at the individual-hemisphere level. Voxel counts ranged from 6 to 33 voxels per visual area per hemisphere. A trend was evident, however, when all voxels were pooled together. Total voxel counts pooled across hemisphere were 373 (V1), 305 (V2), 256 (V3), 203 (V3A), and 252 (hV4). These counts indicate that somewhat fewer voxels were activated by the visual stimulus in higher order visual areas. Because there were more visually responsive voxels in V1 and V2, we had more statistical power to detect small effects in these areas.

Fourth, there were differences in visual response magnitude across areas (Figure 38B). Each data point in this figure represents the visual response magnitude, $\gamma_{\text{visual}}^{\text{resp}}$, averaged over all voxels in a single visual area in a single hemisphere in which the posterior probability of a

non-zero visual responses was greater or equal to 0.95. Because we selected high-probability voxels, the responses shown in Figure 38B are large, ranging from 1% to 3% of the baseline. The median response magnitudes, reveal differences in response magnitude across visual areas. Visual responses were largest in hV4 (median of 2.16% of baseline). We used the posterior distribution of response strength to estimate the probability that visual responses in hV4 were larger than responses in each other area. There was a high probability ($p \geq 0.95$) that responses in hV4 were larger than responses in V1, V2, V3, and V3A. The second largest responses were observed in areas V1 and V3 (medians of 1.87% and 1.88% of baseline). The responses in V1 and V3 were not different from each other ($p < 0.50$), but were both larger than responses in V2 and V3A ($p \geq 0.95$). Responses in V2 and V3A did not differ from each other ($p < 0.50$).

It is not clear whether the differences in response magnitude across visual areas (Figure 38B) are due to differences in neural responses strength or to interregional differences in hemodynamics. Logothetis and Wandell (2004) have argued that regional differences in the coupling between neural activity and hemodynamic changes could result in spurious differences in response amplitude across cortical areas. They termed this coupling *hemodynamic response efficiency*, or HRE. It is common to normalize responses by the baseline signal level, thereby expressing activation in units of percent signal change, as we have done in Figure 38B. However, normalizing by the baseline does not account for regional differences in HRE. One solution to this problem is to normalize responses by the magnitude of activation in a second condition that is known to elicit a response (e.g., see Sàenz et al., 2002). Such selectivity measures reflect the proportional increase or decrease in activation in a particular area, given the regional HRE. The data in this experiment lend themselves to this normalization procedure, because voxels were selected based on there being a highly-probable visual response. Based on this line of reasoning, activation magnitudes in all subsequent analyses are normalized by the magnitude of visual responses.

4.4.4 Responses to the remapped stimulus trace

The central question in this experiment is whether early- and intermediate-level visual areas respond to the remapped stimulus trace. In the previous section (4.4.3), we identified voxels that responded to the visual stimulus. In this section, we ask whether these same voxels also respond to the remapped trace of the stimulus. Remapped responses were measured on trials of the single-step task in which the stimulus appeared in the ipsilateral visual field and a subsequent eye movement brought the recently-stimulated screen location into the contralateral visual field. For each voxel, we estimated the magnitude of the responses in the single-step task, $\gamma^{\text{resp}}_{\text{sstep}}$, and the posterior probability that the responses were greater than zero given the data, $\mathsf{P}\{\gamma^{\text{resp}}_{\text{sstep}} > 0 \mid Y\}$. We observed robust responses in each visual area (Figure 39). The majority (60%) of the hV4 voxels that exhibited a visual response also exhibited a response in the single-step task that reached a $p \geq 0.95$ posterior probability threshold (Figure 39A). A substantial proportion of visually responsive voxels in V3A (43%) and V3 (35%) exhibited responses. We observed responses in about a quarter of visually-responsive voxels in V2 (26%) and V1 (22%). These results indicate that the updated visual signals are present in both striate and extrastriate cortex. Furthermore, remapping is more prevalent in higher-order visual areas.

The next question we address is whether the magnitude of responses in the single-step task varies across cortical areas. In order to compare the magnitude of responses across cortical areas, we normalized the response in single-step trials in each hemisphere by the magnitude of the visual response. Normalized values were calculated as follows. We first averaged visual and single-step responses across all visually-responsive voxels in a given visual area. We then took the ratio of the two means,

$$\gamma^{\text{norm}}_{\text{sstep}} = \left(\frac{\overline{\gamma}^{\text{resp}}_{\text{sstep}}}{\overline{\gamma}^{\text{resp}}_{\text{visual}}} \right) \times 100. \tag{10}$$

We found that the strength of remapping increased monotonically as a function of position within the visual area hierarchy (Figure 39B). The largest normalized responses were observed in areas hV4 (median value of 71%) and V3A (61%). Smaller responses were observed in areas V3 (35%), V2 (23%) and V1 (17%).

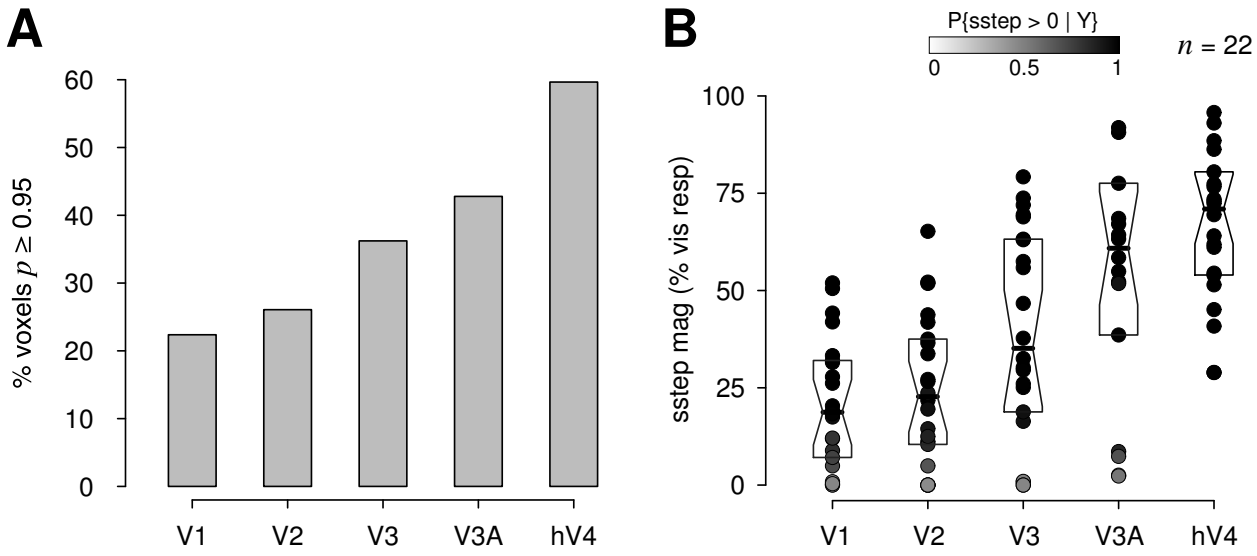

Figure 39. Remapped responses in the single-step task. (A) Proportion of visually-responsive voxels in which the remapped response in the single-step task reached a posterior probability threshold of $p \geq 0.95$. (B) Response magnitude in the single-step task normalized by responses to contralateral visual stimuli. Each dot represents the normalized value of all visually-responsive voxels in a single hemisphere. Grayscale shading corresponds to the posterior probability that the response was greater than zero. The notches in the boxes indicate an estimate of the 95% confidence interval of the median across the population.

This monotonic relationship between response strength and position in the visual area hierarchy (Figure 39B) is demonstrated by a series of pairwise comparisons. Comparisons between visual areas were made under the Bayesian model. We calculated the posterior probability that the response in a given cortical area, a, is larger than the response in another area, a', $\mathsf{P}\{\gamma_{\mathrm{a}}^{\mathrm{sstep}} > \gamma_{\mathrm{a}'}^{\mathrm{sstep}} \mid Y\}$. This analysis yields a Bayesian posterior probability, which we denote as "p". The posterior probability should be confused with a p-value from a classical statistical test (i.e., the probability that the data could be drawn from the population tested given the assumption that the null hypothesis is true).

The largest response was observed in hV4. There was a high probability ($p \geq 0.95$) that the response in hV4 was larger than responses in V1-V3, and there was a marginally high probability ($p = 0.64$) that the response in hV4 is larger than responses in V3A. The next largest response was observed in V3A. There is high probability that this response is larger than responses in V1-V3 ($p \geq 0.91$). The third largest responses was observed in V3. There was high probability ($p \geq 0.95$) that this response is larger than responses in V1 and V2.

Finally, the response in V2 was only marginally larger than the response in V1 ($p = 0.63$). This series of comparisons indicates that the strength of responses in the single-step task increases at each successive stage in the visual area hierarchy.

The use of Bayesian probability is not standard in the functional imaging literature. We confirmed the pattern of results described above using standard statistical tests. We performed each of the comparisons described above using paired t-tests to test for differences in the mean response between visual areas. These tests revealed a similar pattern of results: responses in hV4 were significantly larger than responses in each other area (t-test, $p < 0.01$) except in V3A. Similarly, responses in V3A were significantly larger than responses in each other area (t-test, $p < 0.01$), except V3, the next step down in the hierarchy. Responses in V2 and V1 were not significantly different from each other. These statistical tests and the Bayesian analyses support the same conclusion: the strength of remapped responses in the single-step task increase from V1 to hV4.

4.4.5 Responses in control conditions

As illustrated in the single-voxel example, both ipsilateral stimuli and saccades evoke small responses in extrastriate cortex (Figure 34C-D). It is therefore possible that a portion of activity in the single-step task could be attributed to either the ipsilateral stimulus or to saccades alone, rather than to remapping activity *per se*. In this section, we analyze activity in the two control conditions to determine the degree to which ipsilateral stimuli and saccades contributed to activity in the single-step task.

4.4.5.1 Responses to ipsilateral visual stimuli alone. In the single-step task, a visual stimulus flashes in the ipsilateral visual field. While receptive fields in striate and extrastriate cortex are predominantly contralateral, it is possible that the stimulus itself elicited a response. This consideration is particularly important in areas hV4 and V3A because receptive fields increase in size at later stages of the visual system and some extend into the ipsilateral visual field (Gattass et al., 1981, 1988; Kastner et al., 2001). In the single-step task, it is conceivable that neurons with large receptive fields that extended into the

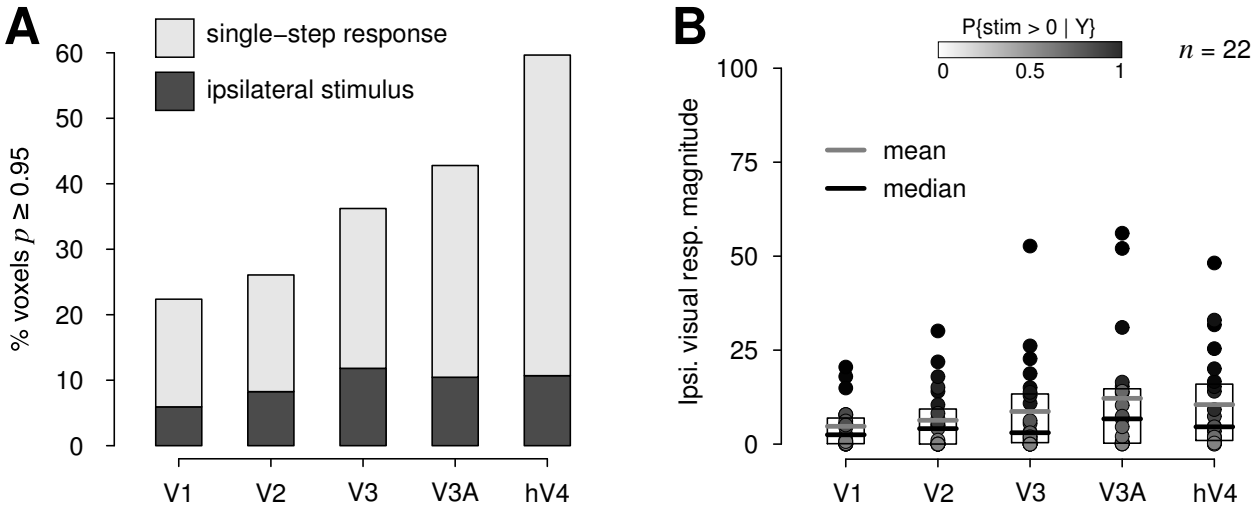

Figure 40. Ipsilateral stimuli elicit small responses. (A) Proportion of visually-responsive voxels in which responses in the stimulus-only condition reached a posterior probability threshold of $p \geq 0.95$. For comparison, we also show the proportion of voxels in the single-step task that reached threshold (light gray bars). (B) Magnitude of responses to the ipsilateral stimuli normalized by response to contralateral visual stimuli. Each dot represents the average normalized value of all visually-responsive voxels in a single hemisphere. Grayscale shading corresponds to the posterior probability that a particular visual area had a nonzero responses to the stimulus. $n = 22$ hemispheres in each area, except for V3A, in which $n = 19$.

ipsilateral visual field could have been driven by the stimulus to a greater degree than V1 neurons that have smaller receptive fields.

We assessed this possibility by measuring responses in a stimulus-only control condition (Figure 29B). In this condition, subjects maintained fixation while a stimulus flickered in the ipsilateral visual field. This condition was balanced with the single-step task for orbital position and visual stimulation. The only difference between the two conditions was the presence or absence of the auditory cue and the resultant saccade. We found that the ipsilateral visual stimulus did activate a small proportion of voxels in each visual area (Figure 40A). Less than 10% of visually-responsive voxels in V1 and V2 responded to the ipsilateral stimulus with a posterior probability that reached a $p \geq 0.95$ threshold. Ipsilateral responses were slightly more prevalent in areas V3, V3A, and hV4, with between 10% and 13% of voxels reaching threshold in each area. These responses in the stimulus-only condition reveal that some voxels in each area are activated by the ipsilateral stimulus. However, such responses do not account for the activation we observed in the single-step task: ipsilateral responses

were far less prevalent in the stimulus-only condition than in the single-step task.

We considered whether there were differences across visual areas in the magnitude of ipsilateral responses (Figure 40B). Areas V3A and hV4 had the largest ipsilateral response (median values of 7% and 5% of the contralateral visual response). The probabilities that either of these responses were larger than responses in any of the other cortical areas were small, though greater than chance ($0.50 < p < 0.75$). A series of t-tests comparing ipsilateral visual responses across cortical areas yielded no significant effects. These comparisons indicate a small and insignificant increase in the magnitude of ipsilateral responses in later stages of the visual area hierarchy.

Two conclusions can be drawn from the analysis of responses in the stimulus-only condition. First, responses to ipsilateral visual stimuli during fixation are small in each visual area. Second, there is a small increase in response strength at later stages of the visual area hierarchy. This is manifest as both a slight increase in the prevalence of active voxels and as a small increase the magnitude of responses. We conclude that ipsilateral responses are too small to account for the relatively large remapped responses observed in the single-step task.

4.4.5.2 Responses to saccades alone. A potential concern is whether saccades activate visual cortex. This issue is particularly important in higher-order visual areas. Neurons in both V4 and V3A fire in relation to saccades directed toward their visual receptive fields (Tolias et al., 2001; Nakamura and Colby, 2000). Furthermore, differences between areas in receptive field size could have increased the chances of observing saccade-related activity in area V4 relative to other visual areas. The logic is as follows. In the single-step task, subjects made 16° saccades. Thus visual responses associated with processing the saccade target were located in the 16° representation in the cortical retinotopic map. Because of differences in receptive field size, the 16° representation in cortex is located closer to the expected site of remapped activation (8°) in area hV4 than in area V1. Thus, saccade-related responses could have overlapped with remapped responses to a greater degree in hV4 and than in V1. It was therefore critical that we determine the extent to which saccades contribute to responses in the single-step task.

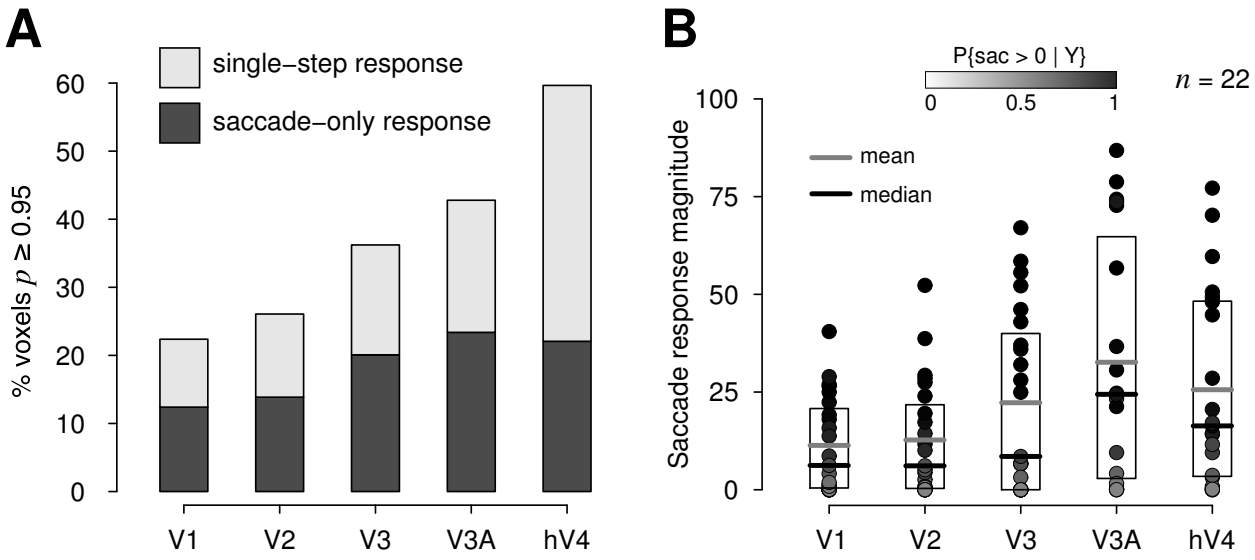

Figure 41. Saccades elicit substantial responses. (A) Proportion of visually-responsive voxels in which responses in the saccade-only condition reached a posterior probability threshold of $p \geq 0.95$. For comparison, we also show the proportion of voxels in the single-step task that reached threshold (light gray bars). (B) Magnitude of responses in the saccade-only control condition normalized by response to contralateral visual stimuli. Each dot represents the average normalized value of all visually-responsive voxels in a single hemisphere. Grayscale shading corresponds to the posterior probability that the averaged response in that hemisphere had response in the saccade-only condition that was larger than zero.

We addressed this issue by testing subjects on a saccade-only control condition (Figure 29C). This condition was balanced with the single-step task for orbital position, auditory stimulation, and number of saccades. The only difference between the two conditions was the presence or absence of the visual stimulus in the 1 s preceding the cue to initiate a saccade. We found that saccades in the absence of the visual stimulus did activate a small proportion of voxels in each visual area (Figure 41A). A minority of visually-responsive voxels in V1 and V2 (12% and 14%) responded in the saccade-only condition with a posterior probability that reached a $p \geq 0.95$ threshold. Saccade-related responses were more prevalent in areas V3 (20%), V3A (24%), and hV4 (22%). This analysis indicates that saccades may have contributed to activity in the single-step task. Moreover, the contribution of saccade-related activity was larger than the contribution of ipsilateral visual responses. However, responses in the saccade-only condition (dark gray bars) were still less prevalent than in the single-step task (light gray bars) indicating that responses in the single-step task cannot be attributed primarily to saccades.

We considered whether there were differences in the magnitude of saccade-related responses across visual areas (Figure 41B). The largest responses in the saccade-only control condition were observed in areas V3A and hV4 (median values of 24% and 16% of the visual response, respectively). The probability that responses in V3A were larger than the responses in hV4 was at chance ($p = 0.53$). Similarly, neither V3A nor hV4 had a high probability of being larger than responses in V3 ($p < 0.80$). However, there was a high probability ($p \geq 0.95$) that responses in V3, V3A, and hV4 were all larger than responses in V1 and V2. Finally, responses in V1 and V2 were not different from each other ($p < 0.5$). These pairwise comparisons indicate that saccades activate high-order visual areas to a greater extent than V1 and V2. Again, we performed standard statistical tests to evaluate the inferences made using Bayesian probability. The difference in response magnitude between V3, V3A, and V4 was not significant (t-test, n.s.). However, all three areas had larger saccade-related responses than V1 and V2 (t-test, $p < 0.01$).

In conclusion, the analysis of responses in the saccade-only control condition revealed that saccades in the absence of a salient stimulus do elicit responses. Furthermore, the pattern of activity is similar to that observed in the single-step task (Figure 41B) in that responses are strongest in higher-order visual areas. Across all visual areas, saccade-related responses were smaller than responses in the single-step task, indicating that saccades alone do not account for remapped responses.

4.4.6 Accounting for both stimuli and saccades

Analysis of responses in the control conditions (Section 4.4.5) indicates that both ipsilateral stimuli and saccades activate striate and extrastriate visual areas to some degree. Throughout this chapter, we have referred to responses in the ipsilateral hemisphere in the single-step task as *remapped responses*. The above analyses suggests that responses in the ipsilateral hemisphere may in fact reflect the combination of remapped responses, ipsilateral visual responses, and saccade-related responses. In this section, we perform three sets of analyses in order to quantify the contribution of remapping in isolation from these other sources of activation:

1. Three-way selectivity analysis. We develop a selectivity measure that quantifies the relative strength of activity in the single-step task, the saccade-only control task, and the stimulus-only fixation task (Section 4.4.6.1). This analysis approach has strong appeal because it gives a precise measure of the proportion of activity that can be attributed to remapping. This approach is closely related to the procedure used to identify remapping in single-unit studies (e.g., Heiser and Colby, *in press*) *2. Linear summation analysis.* We test whether responses in

the single-step task are larger than the sum of visual- and saccade-related responses measured in the two control conditions (Section 4.4.6.2). The results of this analysis demonstrate that responses in single-step trials exceed even this highly conservative criterion, providing strong evidence for remapping in striate and extrastriate visual cortex.

3. Subadditivity analysis. We estimate the degree of subadditivity present in our data by analyzing responses in the *contralateral* hemisphere (4.4.6.3). Because we did not predict any remapping in the contralateral hemisphere, we were able to estimate subadditivity by comparing contralateral responses in the single-step task to the sum of responses to contralateral stimuli and saccades alone. The results of this analysis provide an estimate of the strength of remapping that is independent of activity associated with the control conditions.

In all three analyses, we draw inferences using both Bayesian and non-Bayesian paradigms. All three analyses support the same conclusions: remapping is present throughout visual cortex, and it is strongest in higher order visual areas.

4.4.6.1 Three-way response selectivity. Responses in the single-step task reflect three sources of activity: activity driven by the stimulus in the ipsilateral visual field, activity associated with saccades, and remapping. We calculated a measure of response selectivity, S, in order to quantify the relative contribution of each of these factors. For each of the three conditions, c, we calculated

$$S_c = \frac{\overline{\gamma}_c}{\overline{\gamma}_{\text{sstep}} + \overline{\gamma}_{\text{stim}} + \overline{\gamma}_{\text{sac}}}, \qquad c = \text{sstep, saccade, stim,} \tag{11}$$

where $\overline{\gamma}_c$ is the proportional signal change for a given condition averaged over all visually responsive voxels in a given cortical area. Selectivity values, S_a, sum to 1 ($S_{\text{sstep}}+S_{\text{sac}}+S_{\text{stim}} = 1$) and are all non-negative.

The results from this three-way selectivity measure are represented on triangular simplex plots (Figure 42). Position in the simplex was determined as follows. Let $\mathbf{V}_{\text{sstep}} = (0, \frac{\sqrt{3}}{2})$, $\mathbf{V}_{\text{sac}} = (-\frac{1}{2}, 0)$, $\mathbf{V}_{\text{stim}} = (\frac{1}{2}, 0)$ be the vertices of a triangle. Then the plotted positions of a hemisphere in the simplex are given by $\mathbf{V} = S_{\text{sstep}}\mathbf{V}_{\text{sstep}} + S_{\text{sac}}\mathbf{V}_{\text{sac}} + S_{\text{stim}}\mathbf{V}_{\text{stim}}$. Position in the simplex represents the degree to which the MR response is selective for each of the three conditions. For example, a voxel that responds most strongly on single-step trials relative to the two control conditions would be located in the top sector. A voxel that responds equally strongly in all three task conditions would be represented in the middle of the simplex.

We found that responses from the majority of hemispheres fell within the single-step sector, indicating that responses were selective for the single-step task relative to the control conditions (Figure 42). We performed two sets of analyses in order to determine the degree of response selectivity. In the first analysis, under the Bayesian framework, we estimated the posterior probability that $\mathbf{V}$ was located in the single-step sector of the simplex. Because the use of Bayesian statistics is not widespread in the neuroimaging literature, we perform a second analysis in which we use nonparametric statistical procedures to test whether responses were significantly selective for the single-step task. Both analyses yield the same pattern of results.

Bayesian analysis. The simplex plots in Figure 42 reveal that responses were selective for the single-step task — the majority of hemispheres fell within the top sector. However, there is a degree of uncertainty regarding the position of each point in the simplex, as reflected by the posterior distribution of the response parameters. For example, responses from a given hemisphere would be located in the top sector of the simplex plot if the posterior mean in the single-step task was large relative to the posterior mean in the two control conditions. However, there would be low certainty regarding that hemisphere's position in the simplex if the variance of the posterior distributions was also large. We calculated a series of probabilities that take this uncertainty into account. From the posterior distribution of

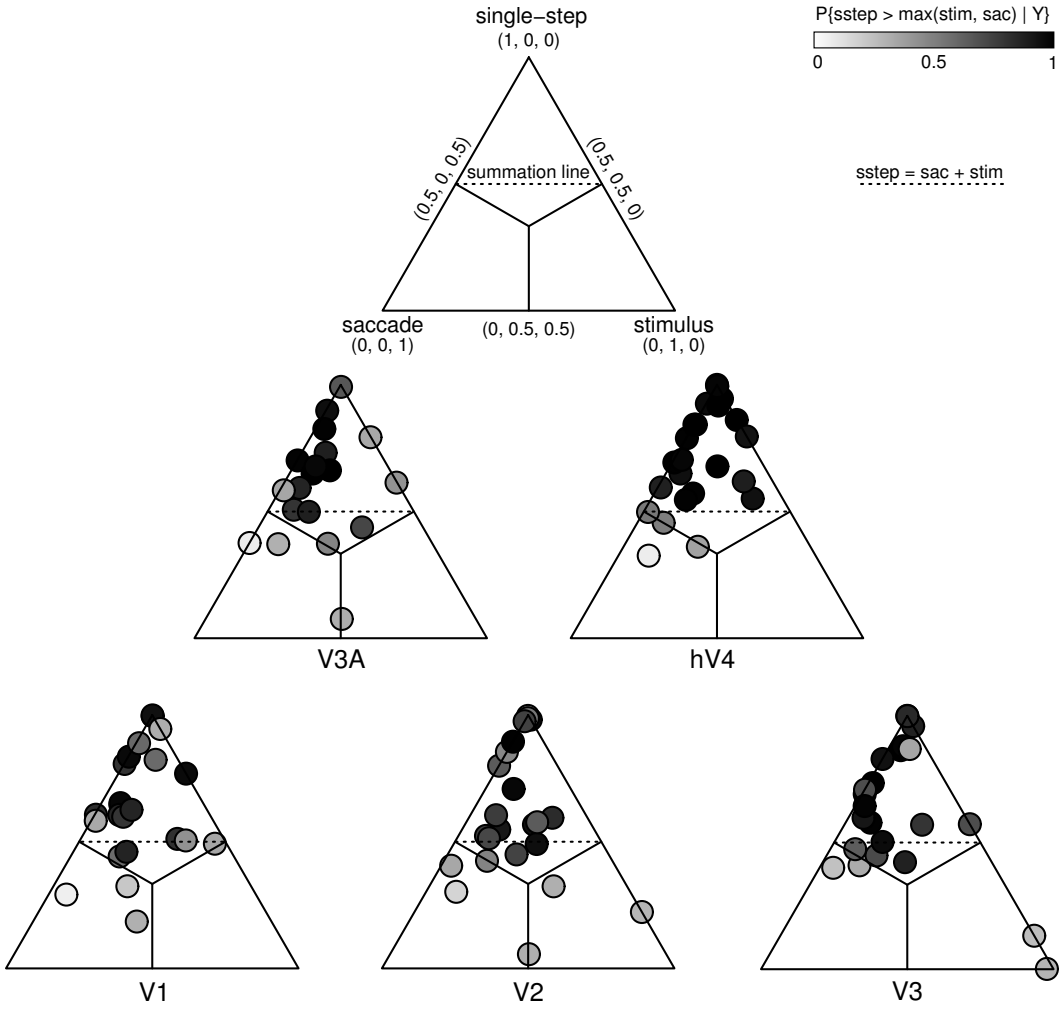

Figure 42. Responses in the single-step task are larger than responses in both control conditions. Ipsilateral responses in the single-step task, stimulus-only control condition, and saccade-only control condition were each averaged within hemisphere, normalized so that they sum to 1, and then plotted on triangular simplex plots. Each dot represents the average response of a single hemisphere. Position of a hemisphere in the simplex represents the selectivity of responses for each of the three conditions. Grayscale shading corresponds to the posterior probability that a particular hemisphere fell within the single-step sector (see Equation 12). Dotted lines indicates the positions in the triangle at which the sum of responses the two control conditions equals the response in the single-step task.

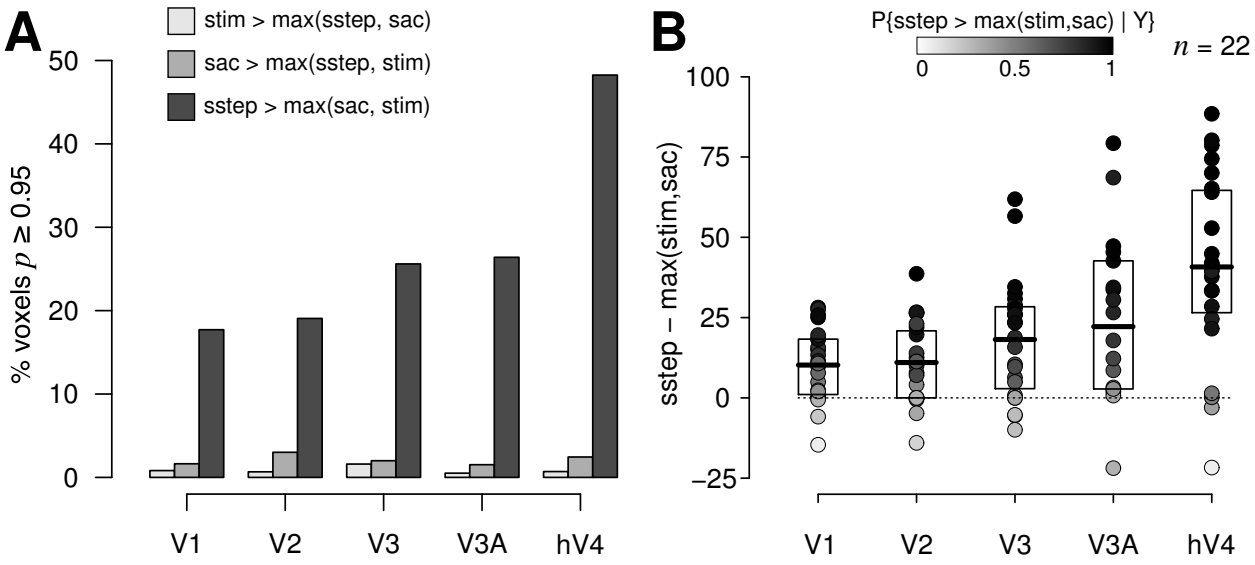

Figure 43. Responses in the single-step task are larger than in either control condition. (A) The proportion of voxels that had a high probability of being located in the stimulus-only sector (light gray), the saccade-only sector (mid-gray), or the single-step sector (dark gray). The y-axis represents the proportion of voxels in each visual area in which the probability reached a $p \geq 0.95$ threshold. A larger proportion of voxels were selective for the single-step task than were selective for either of the other two conditions. (B) Magnitude of responses in the single-step task relative to responses in the two control conditions. Grayscale shading corresponds to the posterior probability that responses in the single-step task were larger than the maximum response in the two control conditions.

the response parameters, γ_{remap}, γ_{sac}, and γ_{stim}, under the Bayesian model, we derived the posterior probability that responses fell within each of the three sectors,

$$\mathsf{P}\{\gamma_{\text{sstep}} > \max(\gamma_{\text{sac}}, \gamma_{\text{stim}}) \mid Y\} \tag{12}$$

$$\mathsf{P}\{\gamma_{\text{sac}} > \max(\gamma_{\text{sstep}}, \gamma_{\text{stim}}) \mid Y\} \tag{13}$$

$$\mathsf{P}\{\gamma_{\text{stim}} > \max(\gamma_{\text{sstep}}, \gamma_{\text{sac}}) \mid Y\}. \tag{14}$$

We calculated these probabilities at the individual voxel level. Fewer than one fifth of voxels in V1 (17%) and V2 (19%) had a high probability ($p \geq 0.95$) of being located within the top sector (Figure 43, dark gray bars). About a quarter of voxels in V3 (26%) and V3A (27%) reached threshold. Voxels in hV4 were most strongly selective for the single-step task; nearly half (49%) of hV4 voxels reached threshold. In contrast, only a small minority ($< 5\%$) of voxels in each area had a high probability of being located in either of the bottom two

sectors (Figure 43, light bars). These results indicate that, in each area, responses in the single-step task are larger than responses in the control conditions. Furthermore, the degree of selectivity increases at later stages of the visual area hierarchy.

Non-Bayesian analysis. We performed a parallel series of standard statistical tests to determine whether responses were selective for the single-step condition. We calculated the proportion of voxels in each area that fell within each sector of the simplex. Voxel counts were skewed toward the top sector in each visual area (Figure 44). To test the significance of this observation, we compared the voxel counts to a χ^2 distribution with two degrees of freedom. The null hypothesis is that data were evenly distributed around the triangle. The alternate hypothesis is that a larger proportion of voxels fell within the top sector. In each area, the test revealed a significant effect (χ^2 test, $p < 0.01$). This result indicates that responses were larger in the single-step task than in either control condition, in agreement with the results of the Bayesian analysis.

In order to compare response selectivity across visual areas, we created a remapping selectivity index (RSI). The RSI is calculated as follows. For each voxel, we normalized responses in each condition by the sum of responses in all three conditions, yielding S_{sstep}, S_{sac}, and S_{stim}, as described above in Equation 11. For each voxel, we then calculated

$$RSI = S_{\text{sstim}} - \max(S_{\text{sac}}, S_{\text{stim}}) \_. \tag{15}$$

The RSI is in the interval $[0, 1]$. It takes a value of 1 for voxels that respond in the single-step task and have no response in either of the two control conditions. It takes a value of 0 for voxels in which the response in either control condition exceeds the response in the single-step task.

We made two sets of inferences based on the RSI. First, we asked whether the distribution of RSI values in each hemisphere differ from that observed by chance. We generated a null distribution of RSI values by permuting the data as follows. For each voxel, we calculated the RSI six times, once for each of the possible permutations of the three conditions. We followed this procedure for each voxel from each visual area, resulting in a single null distribution against which to compare index values from each area (Figure 45; labeled "null").

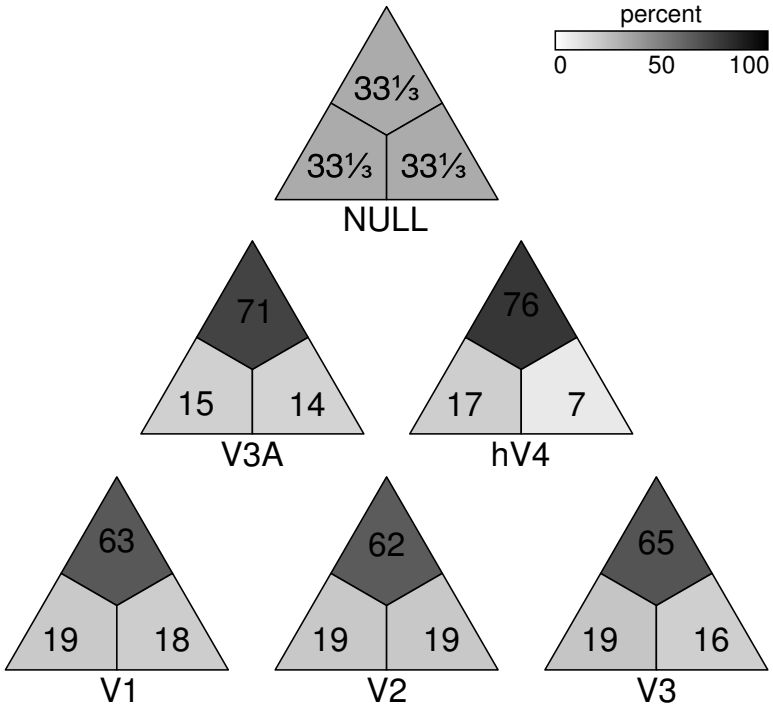

Figure 44. Voxel counts across triangular simplex. Percentage of voxels in each visual area that fell within a given sector of the simplex. Data generated from a random process would be distributed equally around the simplex.

We then performed a series of Kolmogorov-Smirnov tests comparing the distribution of RSI values in each visual area to this null. We used KS-tests because the distribution of RSI values are not normal, but rather are highly skewed toward zero. In each area, the distribution of observed RSI values was greater than the null distribution (Figure 45B), indicating that responses were larger in the single-step task than in either control, consistent with the Bayesian analysis described in Equation 12.

Second, we asked whether the distribution of RSI values differed across visual areas. The distribution of RSI values in area hV4 was significantly greater than that in each other area (KS-test, $p < 0.001$ for comparisons with areas V1, V2, V3, and $p < 0.01$ for the comparison with area V3A). However, none of the other comparisons between visual areas were significant. This analysis confirms that the strongest response selectivity was observed in hV4 but is not sufficiently sensitive to detect the monotonic relationship observed using the Bayesian method (Figure 43).

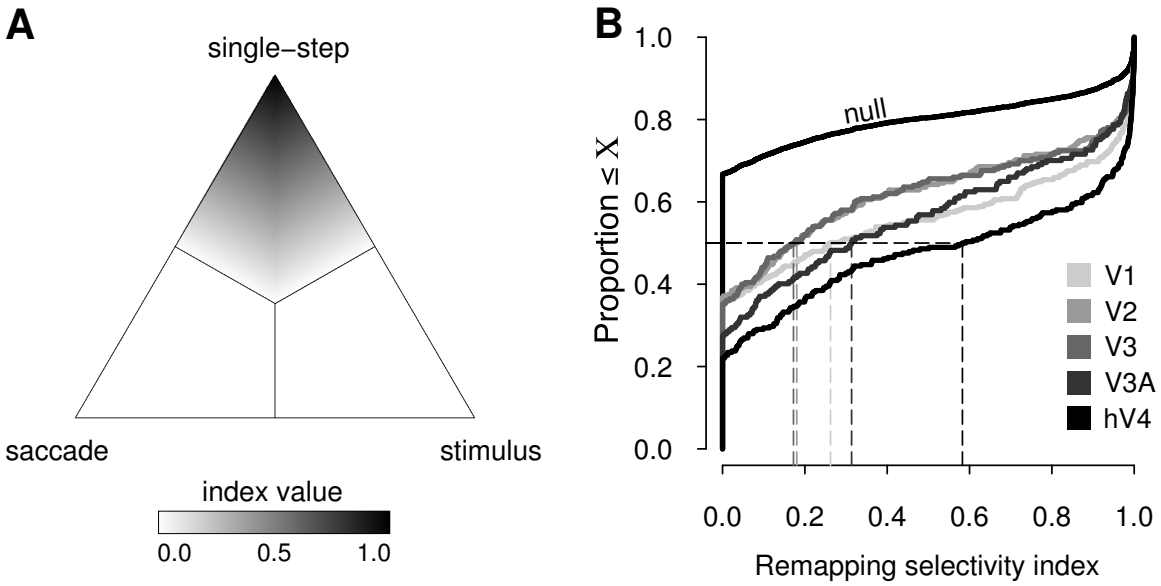

Figure 45. Remapping selectivity index. (A) Correspondence between RSI values and position in the triangular simplex plot. The RSI represents the degree to which MR responses are selective for the single-step task relative to the stimulus-only and saccade-only control tasks. RSI values, shown here in grayscale, range from 0 (white) to 1 (black). Higher values represent increased selectivity. (B) Cumulative distribution functions for the population of voxels from all hemispheres and subjects. The *null* distribution was obtained by permuting the data. The distribution of RSI values in each hemisphere were significantly greater than the null distribution (KS-test, $p < 0.05$).

4.4.6.2 Linear summation of activity in control conditions. In the preceding sections, we found that both the ipsilateral stimuli and saccades themselves elicit small responses (Sections 4.4.5.1 and 4.4.5.2). However, responses in the single-step task were larger than responses in either control condition, indicating that neither ipsilateral stimuli nor saccades alone account for activity in the single-step task (Section 4.4.5). This analysis assumed that responses in the three conditions are independent. That is, we assumed that a random process would result in an equal number of observations in each sector of the simplex. Activation in the single-step task could, however, reflect the sum of responses to ipsilateral stimuli and saccades. In this section, we test this possibility by comparing activation in the single-step task to the sum of responses in the two control conditions. As in the previous section, we address this question using both the Bayesian paradigm and standard statistical tests.

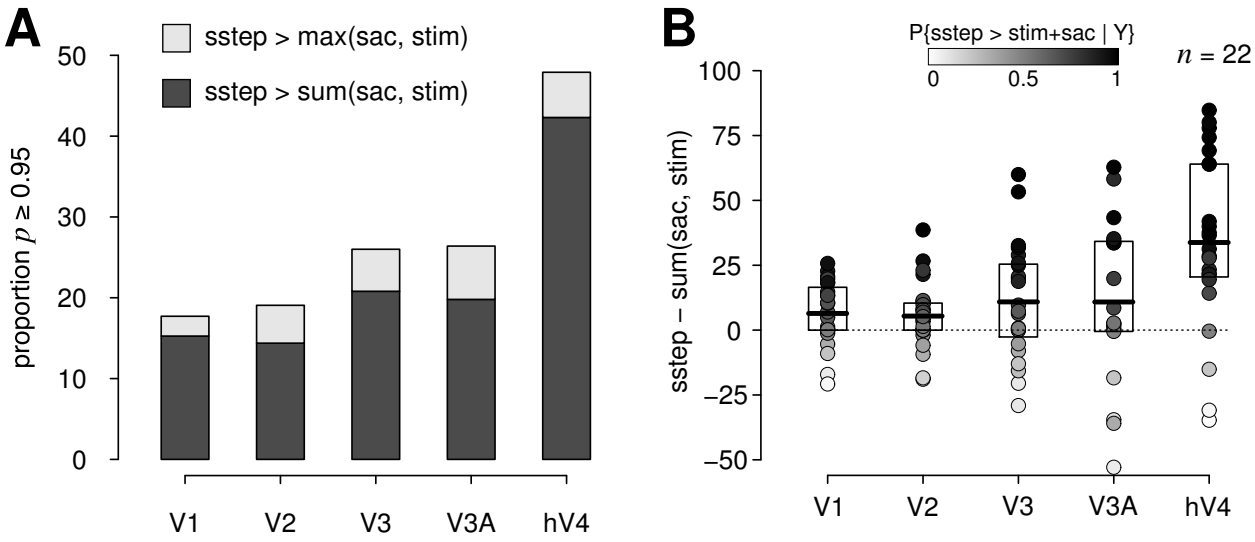

Figure 46. Responses in the single-step task are larger than the sum of responses in the control conditions. (A) The proportion of voxels that had a high probability of being larger than the sum of responses in the two control conditions (dark gray bars). (B) Magnitude of responses in the single-step task relative to the sum of responses in the two control conditions. Grayscale shading corresponds to the posterior probability that responses in the single-step task were larger than sum of control responses.

Bayesian analysis. Under the Bayesian model, we calculated the posterior probability that responses in the single-step task are larger than the sum of responses in the two conditions,

$$\mathsf{P}\{\gamma^{\text{resp}}_{\text{sstep}} > (\gamma^{\text{resp}}_{\text{sac}} + \gamma^{\text{resp}}_{\text{stim}}) \mid Y\}. \tag{16}$$

This analysis is more conservative than the selectivity analysis in the preceding section (4.4.6.1). Accordingly, slightly fewer voxels in each area reached threshold (Figure 46). Fewer than one fifth of voxels in V1 (15%) and V2 (14%) had responses in the single-step task that were larger than the sum of responses in the two control conditions at a probability level of $p \geq 0.95$. About one fifth of voxels in V3 (21%) and V3A (20%) reached threshold. A substantial proportion of voxels in area hV4 (43%) reached threshold. This analysis demonstrates that, for a substantial proportion of voxels, responses in the single-step task are larger than the sum of responses in the two control conditions.

Non-Bayesian analysis. As in the previous section (4.4.6.1), we performed an additional analysis using standard statistical tests. For each voxel, we calculated a summation index

(I_s). This index quantifies the magnitude of responses in the single-step task and relative to the sum of responses in the two control conditions. I_s is calculated as follows:

$$I_s = \frac{\gamma_{\text{remap}} - \gamma_{\text{stim}} - \gamma_{\text{sac}}}{\gamma_{\text{remap}} + \gamma_{\text{stim}} + \gamma_{\text{sac}}}. \tag{17}$$

Index values fall in the interval $[-1, +1]$. At one extreme, values of -1 indicate no response on single-step trials relative to the two control conditions. Values of $+1$ indicate a large response on single-step trials and no responses in either of the two control conditions. Values at 0 indicate that activation in the single-step task is a linear sum of activation in the two control conditions (Figure 47A). I_s provides a more stringent measure of response selectivity than the RSI (Equation 15). The RSI considers responses to be selective for the single-step task if they fell within the top sector of the simplex (Figure 42). I_s, on the other hand, considers responses to be selective only if they fell above the summation line (Figure 42, dotted horizontal line).

We calculated I_s for each individual voxel in each visual area. We then pooled the data across hemispheres and subjects to estimate the distribution of index values. The median I_s in each visual area was larger than 0 (Figure 47B). This observation indicates that responses in the single-step task are greater than the sum of activity in two control conditions.

4.4.6.3 Subadditivity of BOLD responses. In the previous section (4.4.6.2), we found that responses in the single-step task were larger than the sum of responses in the two control conditions, providing strong evidence for the existence of remapping in occipital cortex. In many experimental contexts, it is reasonable to assume that the BOLD response elicited by two events is equal to the sum of responses measured independently — BOLD activity approximates a shift-invariant linear system (Boynton et al., 1996). However, there is increasing evidence that BOLD responses are nonlinear in specific contexts (Birn and Bandettini, 2005; Birn et al., 2001; Friston et al., 2000a, 1998; Huettel and McCarthy, 2001; Vazquez and Noll, 1998). For example, many studies have shown that there is a saturating nonlinearity for closely spaced neural events (review in Wager et al., 2005). Two visual stimuli that occur in rapid succession evoke an MR response that is smaller than would be predicted by the sum of responses to the two stimuli in isolation. In other words, closely

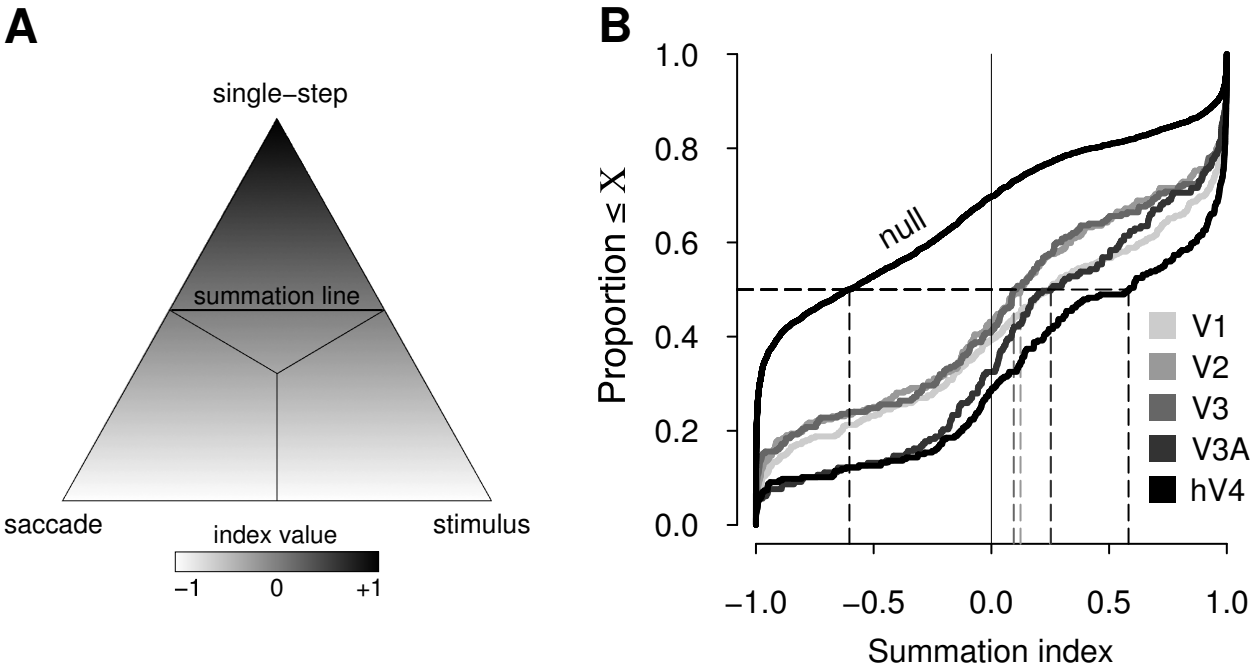

Figure 47. Responses in the single-step task are larger than the *sum* of responses in the two control conditions. (A) Correspondence between summation index (I_s) and position in the triangular simplex. Values at the bottom of the simplex have a value of -1 and values at the apex have a value of $+1$. (B) Distribution of I_s values for all visually-responsive voxels from all hemispheres and subjects. The medians of each distribution are indicated by dotted vertical lines. Each distribution had a positively skewed I_s distribution, with a median value greater than 0. This analysis indicates that responses in the single-step task can not be attributed to the sum of activity in the two control conditions. Note that the distribution of I_s values calculated from random data (labeled "null") has a median that is far less than 0.

spaced events sum *sublinearly*. The nonlinearity can be large, up to 50% of the expected response magnitude (Wager et al., 2005). This nonlinearity is likely due to elastic properties of the microvasculature (e.g., see the *balloon model* of the BOLD response; Buxton et al., 2004; Friston et al., 2000b). The basic idea is that the rate at which the vasculature can accommodate a sudden and rapid influx of blood saturates. This saturation limits the magnitude of the response, causing the response to the second event to be smaller than it would be otherwise. This saturating nonlinearity could have affected responses in the single-step task. In this task, two events, a stimulus and a saccade, occur in rapid temporal succession. The resultant response may thus be smaller than would be predicted by the sum of responses to the stimulus and saccade measured in isolation. This nonlinearity would cause the analysis in Section 4.4.6.2 to be overly conservative.

The issue of sublinear summation is difficult to address directly. The neural phenomenon of remapping is itself nonlinear — the conjunction of a stimulus and a saccade produces a response that is not there if either occur alone. In this section, we model the sublinear summation present in our data by analyzing responses in the *contralateral* hemisphere. Specifically, we analyzed responses in the following three conditions: (1) the single-step task when the stimulus appeared in the contralateral visual field and was followed by a contraversive saccade; (2) the stimulus-only fixation task when the stimulus appeared in the contralateral visual field, and (3) the saccade-only condition in which the saccades matched the saccades on single-step trials. By definition, we did not expect remapping in the contralateral hemisphere: the stimulus trace is always remapped from the contralateral to the ipsilateral hemisphere with the eye movement. In the contralateral hemisphere, responses in the single-step task should therefore reflect the sum of responses in the stimulus-only and saccade-only conditions,

$$\mathrm{sstep_c} = \mathrm{sac_c} + \mathrm{stim_c}. \tag{18}$$

In contrast, we expect that activity in the ipsilateral hemisphere in the single-step task reflects responses associated with each of the two control conditions *and* responses associated with remapping,

$$\mathrm{sstep_i} = \mathrm{sac_i} + \mathrm{stim_i} + \mathrm{remap}. \tag{19}$$

Based on these two assumptions, we calculated a subadditivity parameter, $\hat{u}$, that represents the degree to which responses in the contralateral hemisphere are smaller than would be predicted by linear summation,

$$\hat{u} = \left(\frac{\mathrm{sstep_c}}{\mathrm{sac_c} + \mathrm{stim_c}} \vee 1 \right) \wedge 0. \tag{20}$$

This parameter, $\hat{u}$, is in the interval $[0, 1]$. Values of 0 indicate strong subadditivity; values of 1 indicate that responses in the contralateral hemisphere sum linearly. Constraining $\hat{u}$ to this interval precludes the possibility of superadditivity. We used $\hat{u}$ to derive an estimate of remapping, $\widehat{\mathrm{remap}}$, in the ipsilateral hemisphere,

$$\widehat{\mathrm{remap}} = \mathrm{sstep_i} - \hat{u}(\mathrm{sac_i} + \mathrm{stim_i}) \tag{21}$$

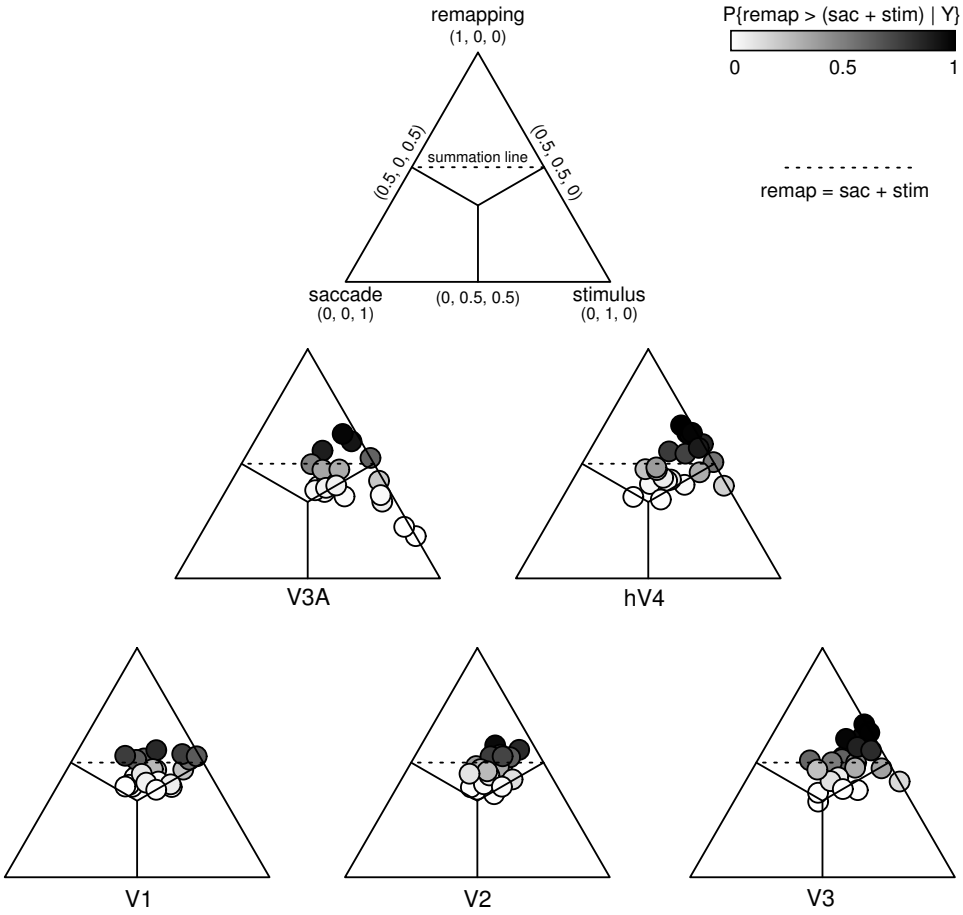

Figure 48. Subadditivity of responses in the contralateral hemisphere. Simplex plots represent the relative strength of responses in the following three conditions: (1) Single-step task when the visual stimulus was located in the *contralateral* visual field and was followed by a *contraversive* saccade. (2) Stimulus-only fixation task when the stimulus was located in the *contralateral* visual field. (3) Saccade-only control condition. Position in the simplex indicated the degree of selectivity of responses for each of the three tasks. Grayscale shading corresponds to the probability that responses in the single-step task are larger than the *sum* of responses in the two control conditions. The critical result is that the majority of responses fell below the summation line (dotted horizontal line) and had low probabilities, indicating subadditivity in the contralateral hemisphere.

Finally, under the Bayesian model, we calculated the posterior probability that $\widehat{\text{remap}}$ is larger than 0 given the subadditivity parameter,

$$\mathsf{P}\{\widehat{\text{remap}} > 0 \mid \hat{u}\} \tag{22}$$

Each of these steps is described in greater detail below.

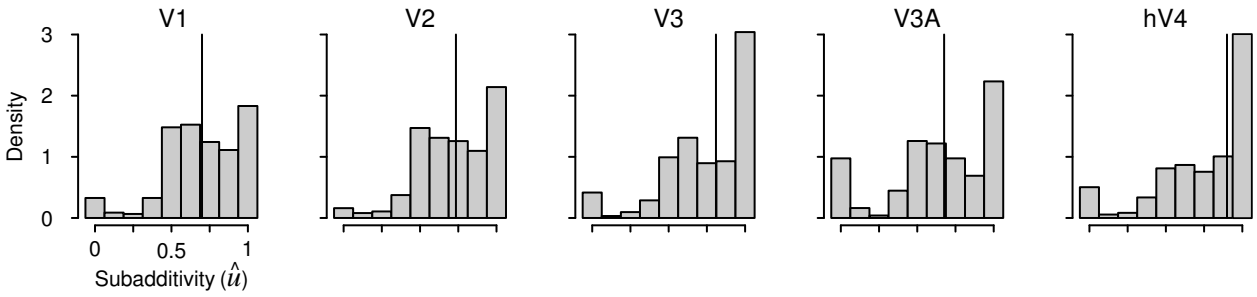

Figure 49. Distribution of subadditivity parameter ($\hat{u}$) calculated from Equation 20. Values of 1 indicate that contralateral responses in the single-step task reflect the linear sum of activity in the stimulus-only and saccade-only control conditions. Values less than 1 indicate that responses in the contralateral hemisphere are smaller than would be predicted by linear summation. Vertical line in each plot indicates the median of the distribution.

We created a series of simplex plots to visualize the degree of subadditivity in the *contralateral* hemisphere (Figure 48.) These simplex plots highlight three points. First, the majority of hemispheres fell well within the top sector. This indicates that voxels in the contralateral hemisphere were more strongly activated by the single-step task than by either visual stimuli or saccades alone. This finding suggests that *both* stimuli and saccades activate the contralateral hemisphere, and that activity in the single step reflects the *joint* activity associated with the two conditions. Second, responses tended to cluster toward the right (stimulus) side of the simplex. This observation indicates that responses to contralateral stimuli were large relative to responses to saccades. This is not surprising given that voxels were selected based on having large responses to contralateral visual stimuli (see Section 4.4.3, and Figure 38). Third, and most importantly, the majority of hemispheres fell at or below the summation line (horizontal dotted line). This observation indicates that visual and saccade responses in the single-step task are *smaller* than would be predicted by taking the sum of responses to stimuli and saccades independently. We interpret this finding to indicate that responses in the contralateral hemisphere sum sublinearly in the single-step task. An important alternative interpretation for this observation is based on response *truncation* rather than sublinear summation. We discusses this issue further in Section 4.5.3.

We quantified the degree of subadditivity in the contralateral hemisphere by calculating the parameter $\hat{u}$ as described above in Equation 20. We found that $\hat{u}$ was less than 1 in

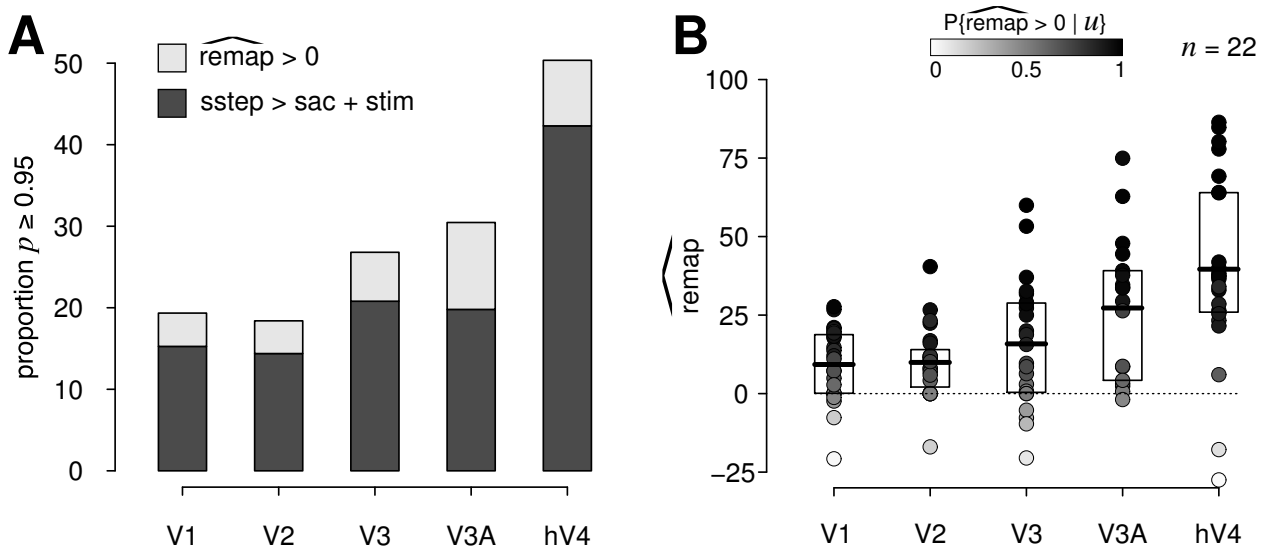

Figure 50. Estimate of remapping, $\widehat{\text{remap}}$, in each visual area. The calculation of $\widehat{\text{remap}}$ takes into account activity in both control conditions *and* estimated subadditivity of responses. (A) The proportion of voxels for which the posterior probability that $\widehat{\text{remap}}$ reached a $p \geq 0.95$ threshold. (B) Response magnitude. Each dot represents the averaged response from a single hemisphere. Grayscale shading corresponds to the posterior probability that $\widehat{\text{remap}}$ was larger than zero, given the subadditivity parameter, $\hat{u}$.

the majority of voxels, indicating that contralateral responses in the single-step task were smaller than predicted by linear summation of activity in the control conditions (Figure 49), consistent with the contralateral simplex plots in Figure 48. Furthermore, we observed differences in $\hat{u}$ across visual area. Subadditivity was most pronounced in area V1 (median $\hat{u} = 0.76$) and was least pronounced in hV4 (0.96). This is notable because the nonlinearity of BOLD responses has been studied most extensively in primary visual cortex. The degree to which hemodynamic nonlinearities are region-specific is a matter of much current interest (Wager et al., 2005).

Finally, based on estimates of $\hat{u}$, we estimate the magnitude of remapping, $\widehat{\text{remap}}$, in the *ipsilateral* hemisphere, according to Equation 22. This estimate is influences to the relative magnitudes of responses in all six trial types (three conditions, two directions). The critical question is there is high probability that $\widehat{\text{remap}}$ was larger than zero, conditional on $\hat{u}$.

We found that a substantial proportion of voxels in each visual area had nonzero $\widehat{\text{remap}}$ values (Figure 50). This estimate of remapping strength is less conservative than the analysis

in Section 46, in which we compare activity in the single-step task to the sum of responses in the two control conditions without accounting for subadditivity. This is because the scaling factor, $\hat{u}$, reduces the size of the summed visual and saccade activity. Thus, as one would expect, the proportion of voxels in which the posterior probability that $\widehat{\text{remap}}$ exceeded a $p \geq 0.95$ threshold was larger than in the linear summation analysis described earlier (Figure 50A, compare light and dark bars). The analysis of sublinear summation provides three important insights. First, is indicates that responses did sum sublinearly, and that nonlinearities are important to consider when measuring response to rapidly occurring events. Second, this analysis further supports the claim that remapping is present throughout occipital cortex. Third, this analysis also suggests a monotonic relationship between the magnitude of remapping and position in the visual area hierarchy.

4.4.7 Time course of visual and remapped responses

Remapping occurs at various points in time relative to saccade initiation. A substantial proportion of neurons in parietal and extrastriate cortex remap predictively, while others begin to respond around the time of the saccade (Duhamel et al., 1992a; Nakamura and Colby, 2002). Predictive responses in single neurons occur at a latency that is shorter than the typical visual response for that neuron. A subset of cells with predictive responses begin to respond to the stimulus trace even before the eyes have moved. We asked whether the timing of remapped responses in human striate and extrastriate cortex is consistent with these observations from monkeys.

In the fMRI version of the single-step task, the stimulus appears and stays on the screen for 1 s prior to the auditory cue to make an eye movement. We expect that remapping occurs around the time of the eye movement. Subjects in this study had a mean saccadic reaction time of 255 ms (128 ms SD). Thus, the onset of the stimulus preceded the onset of the saccade by an average of 1255 ms. Because of this interval, remapped responses driven by the stimulus trace should begin about 1255 ms later than visual responses driven by the visual stimulus. This statement assumes that remapping occurs around the time of the eye movement.

We used the Bayesian estimates of response profile to test this prediction. The three relevant parameters are *lag*, *attack*, and *decay*. We predicted that remapped responses would have a longer lag than visual responses, but that the two response types would not differ in either attack or decay. Our analysis of response profile was performed on the subset of voxels in each visual area in which both visual *and* remapped responses had a high probability of being nonzero ($p \geq 0.95$). This selection criterion was necessary because inferences regarding response profile must be based on actual responses. The analysis described in this section includes a smaller proportion of visually-responsive voxels that the analyses described in previous sections because, as we showed in Figure 39, only a subset of voxels in each visual area responded to the stimulus trace.

4.4.7.1 Analysis of response lag. We predicted that remapped responses in the single-step task would have larger response lags than visual responses. The lag parameter represents the time from stimulus onset to the start of the hemodynamic response (i.e., lag is equivalent to response latency). Across all visual areas, visual responses had lags that were, on average, 1.34 s ± 0.47 s SD (Figure 51A, light bars). In contrast, remapped responses had lags that were, on average, 2.08 s ± 0.64 s SD (Figure 51B, dark bars), indicating that remapped responses did in fact occur later in time than visual responses. We calculated the difference in lag between remapped and visual responses on a voxelwise basis (Figure 51B). Remapped responses had lags that were larger than visual responses by a median of 0.67 s (V1), 0.63 s (V2), 0.90 s (V3), 0.73 s (V3A), and 0.94 s (hV4). These differential lags confirm the prediction that remapped responses have longer latencies relative to visual responses.

We analyzed these data with both Bayesian and with standard statistical methods. First, we calculated the posterior probability that remapped responses had longer lags than visual responses given the data, $\mathsf{P}\{\gamma^{\text{lag}}_{\text{remap}} > \gamma^{\text{lag}}_{\text{visual}} \mid Y\}$. Probabilities were then estimated on a hemisphere basis, taking into account the posterior distribution of the estimated lag parameters. We then combined probabilities across hemispheres, yielding population-level probabilities. In each visual area, there was a high probability that remapped responses had longer lags than visual responses (Figure 52A, grayscale shading). Probabilities were 0.84, 0.80, 0.91, 0.86, 0.92 for areas V1, V2, V3, V3A, and hV4, respectively. This analysis indicates that

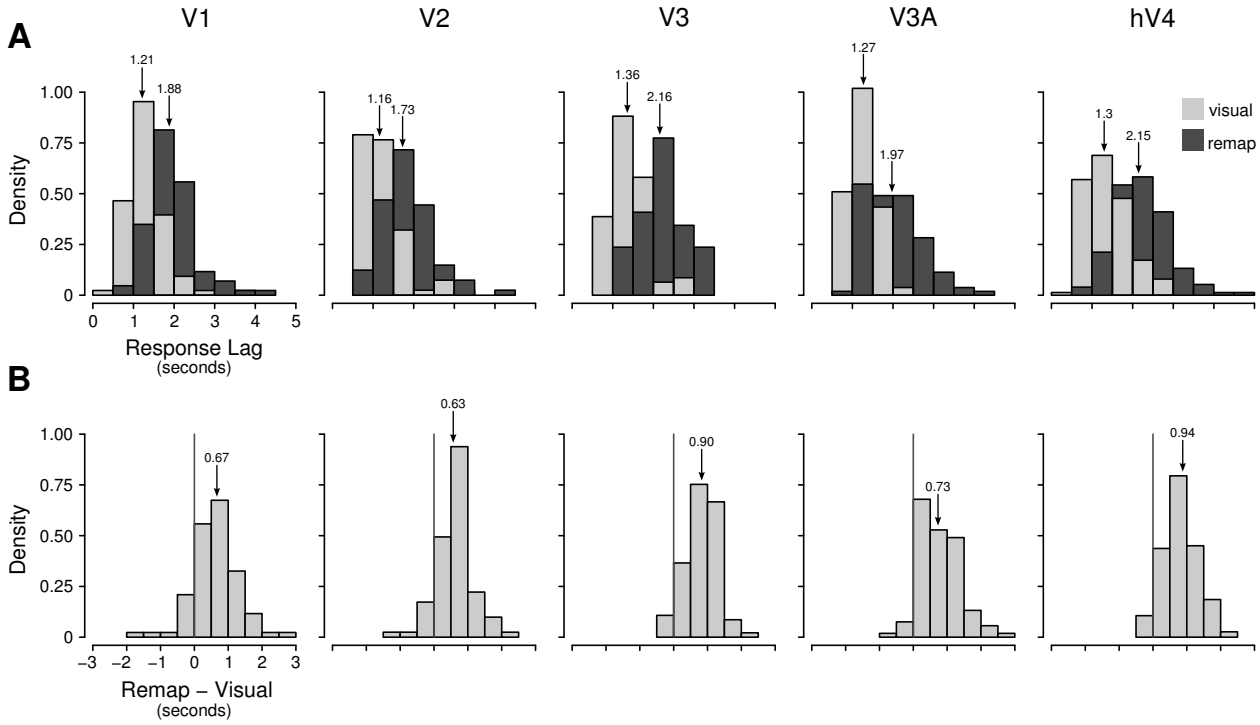

Figure 51. Remapped responses have a longer lags than visual responses. Histograms show distribution of the lag parameter across voxels in which both visual and remapped responses had a high posterior probability of being nonzero ($p \geq 0.95$). (A) Histograms of lag estimates for visual (light gray) and remapped (dark gray) responses. (B) Histogram of the difference between visual and remapped response lags. Downward-pointing arrows indicate the median of the distributions. The lag parameter represents the duration between stimulus onset and the beginning of the fMRI response. Remapped responses in each visual area had longer lags.

remapped responses occur substantially later than visual responses. Second, we performed a series of two-sample Kolmogorov-Smirnov tests on the distribution of visual and remapped response lags to determine if the two distributions differed. We use KS-tests here, rather than t-tests, because the distribution of lags was not normal. The KS-test is distribution free test, but it is known to have low power (Aguirre et al., 1998a). The difference in lag was significant in each visual area (KS-test, $p < 0.01$). Both the Bayesian and the standard analyses provide strong support for the claim that remapped responses occur later than visual responses, as expected. The difference in response lag was smaller than the temporal interval between stimulus and saccade onset, indicating that remapping was predictive.

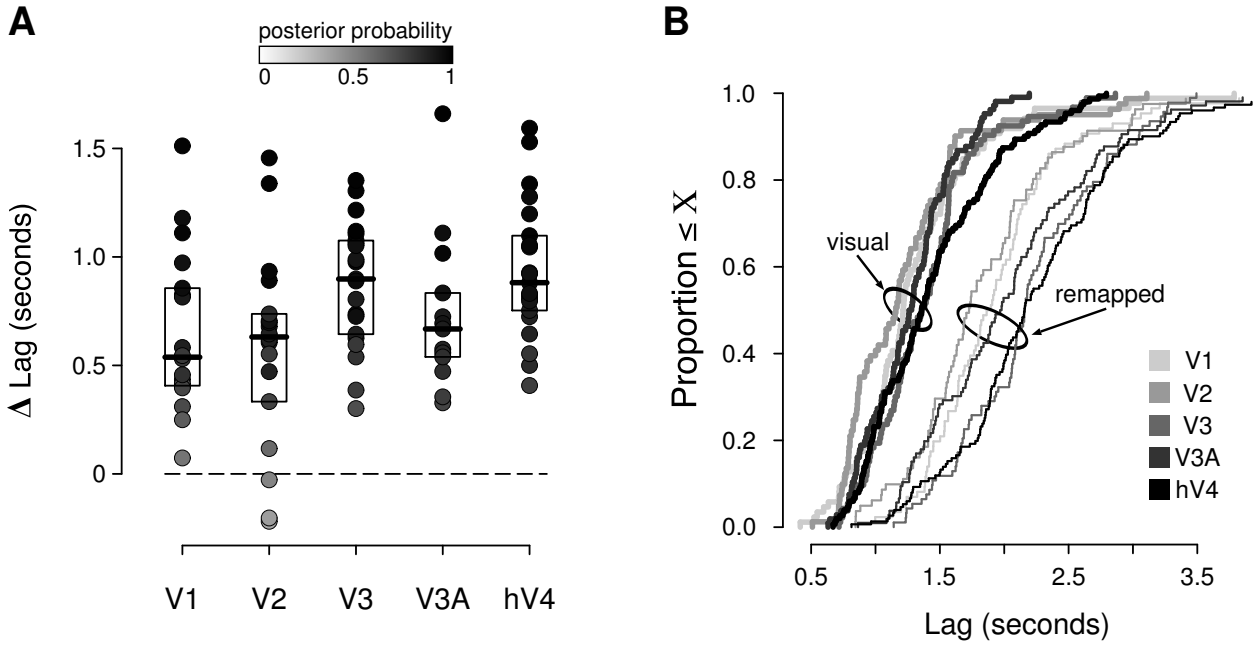

Figure 52. Difference between remapped and visual response lag. Each circle represents the difference in lag averaged over all voxels in a given visual area in a single hemisphere. Grayscale shading represents the posterior probability that remapped responses had longer lags than visual responses. Dotted line at zero represents no difference.

4.4.7.2 Analysis of response attack. The lag parameter measures the point in time at which the BOLD response begins. This parameter may thus be sensitive to the population of neurons that remap early. Single neurons begin responding to the stimulus trace at various points relative to the saccade onset (Nakamura and Colby, 2002). In our study, the neurons that remap later could have influenced the rate at which remapped responses reached peak. This aspect of the BOLD response was characterized by the *attack* parameter. Response attack measures the time between the beginning of the response and response peak. Unlike *lag*, attack is a duration — it is not a measure of time from the start of the trial.

We found that there was a consistent difference in response attack between visual and remapped responses. Visual responses had attack values that on averaged 3.99 s (0.29 s SD) (Figure 53A, light bars). In contrast, remapped responses had attack values that averaged 4.39 s (0.37 s SD) (Figure 53A, dark bars). We calculated the difference in attack between remapped and visual responses on a voxelwise basis (Figure 53B). Remapped responses had attack values that were longer than visual responses by 0.34 s (V1), 0.37 s (V2), 0.31 s

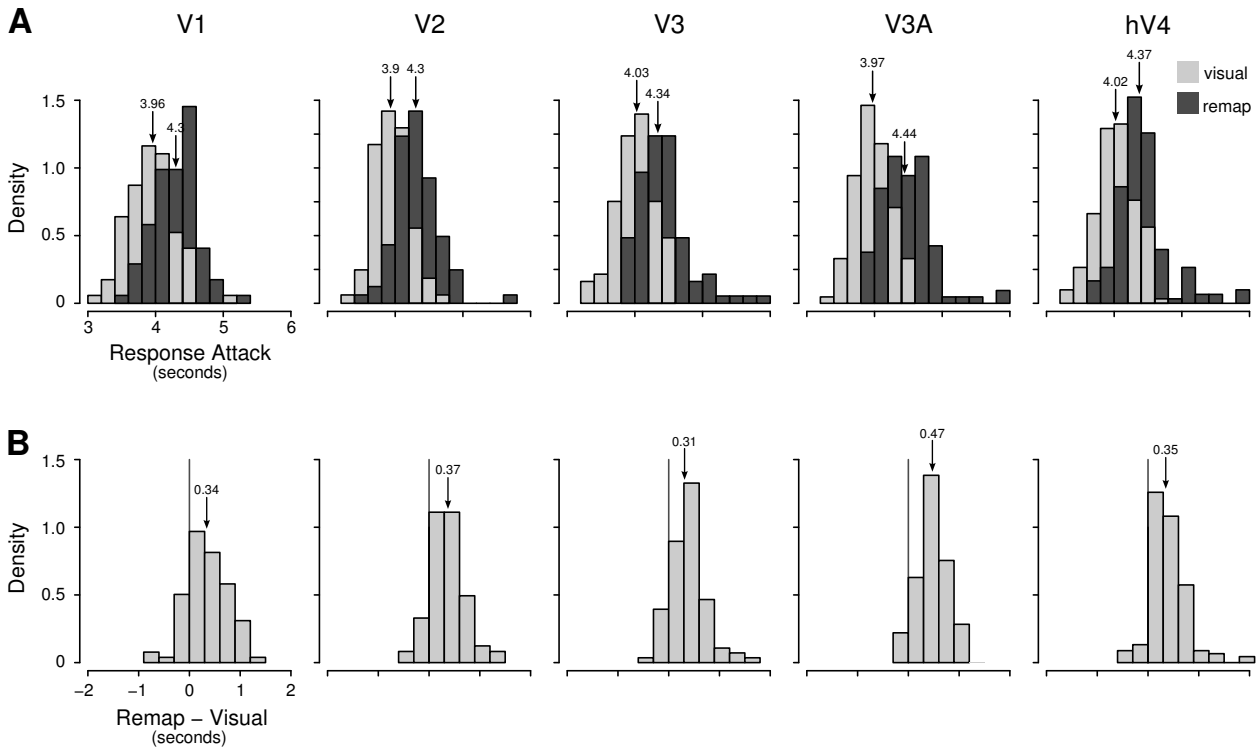

Figure 53. Remapped responses have a longer attack than visual responses. Histograms show the distribution of the *attack* parameter across voxels in which both visual and remapped responses had a high posterior probability ($p \geq 0.95$) of being nonzero. (A) Histograms of attack estimates for visual (light gray) and remapped (dark gray) responses. (B) Histogram of difference between visual and remapped response attack estimates. Downward-pointing arrows indicate the median of the distribution. The attack parameter is the time between the beginning of the fMRI response and the peak of the responses. Remapped responses in each visual area had slightly longer attack values.

(V3), 0.47 s (V3A), and 0.36 s (hV4). These differential attack values show that remapped responses took slightly longer than visual responses to reach peak magnitude.

We used both Bayesian and standard statistical methods to analysis these data. First, we estimated the posterior probability that remapped responses had larger attack values than visual responses given the data, $\mathsf{P}\{\gamma^{\text{attack}}_{\text{remap}} > \gamma^{\text{attack}}_{\text{visual}} \mid Y\}$. Probabilities were estimated on a voxelwise basis, using the same approach described above for response lag. There was a high probability in each visual area that remapped responses had larger attacks than visual responses (Figure 54A, grayscale shading). Probabilities were 0.84, 0.80, 0.91, 0.86, 0.92 for areas V1, V2, V3, V3A, and hV4, respectively. We performed a second standard analysis to test whether remapped responses had significantly larger attacks than visual responses.

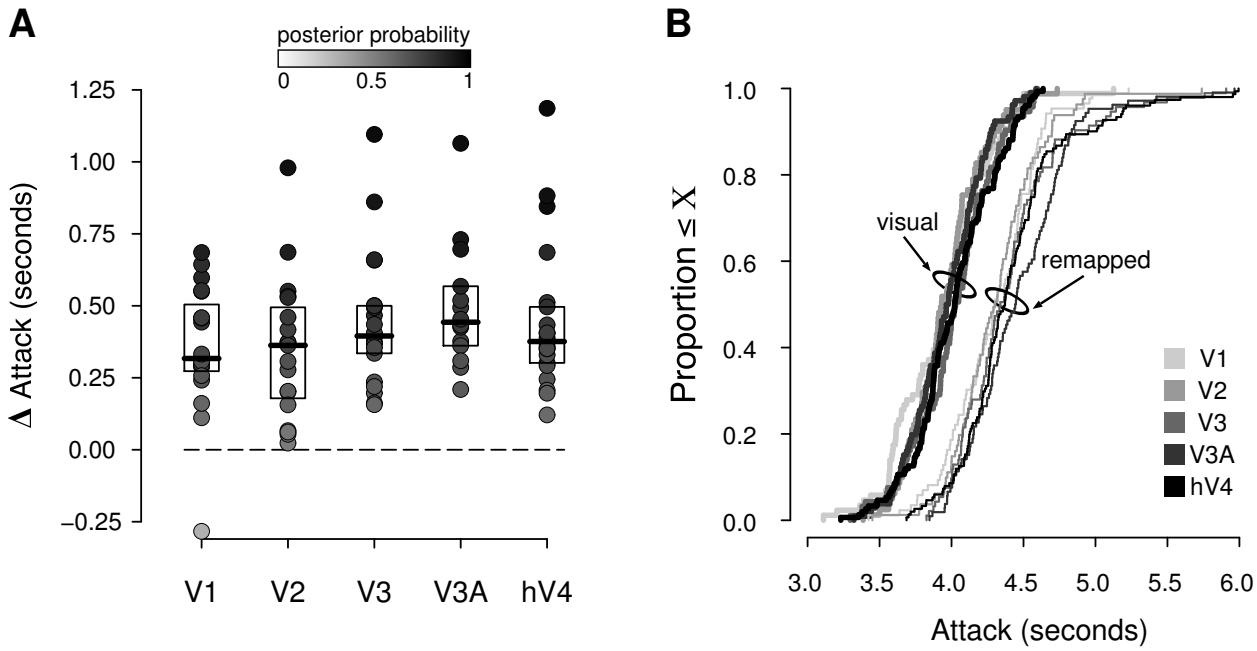

Figure 54. Difference between remapped and visual response attack. Each circle represents the difference in attack averaged over all voxels in a given visual area in a single hemisphere. Grayscale shading represents the posterior probability that remapped responses had larger attacks than visual responses. Dashed line at zero represents no difference.

The difference in attack was significant in each visual area (KS-test, $p < 0.01$). This result indicates that remapped responses take longer to reach their peak than visual responses — remapped responses rise more gradually.

4.4.7.3 Analysis of time-to-peak: lag + attack. The differences in shape parameters for both response lag and attack were smaller than the temporal interval between stimulus and saccade onset. We considered whether the sum of lag and attack would better reflect the time that we expect to observe remapping. The attack parameter is a duration, rather than an absolute time. Combining the lag and attack parameters, however, gives a measure of *time-to-peak* (TTP) relative to the onset of the stimulus. We compared visual and remapped responses to determine whether the difference in TTP corresponded to the interval between the onset of the stimulus and the time at which subjects executed an eye movement.

We first calculated TTP by simply taking the sum of the lag and attack parameters ($\gamma^{\text{lag}} + \gamma^{\text{attack}}$). We then calculated the difference in TTP between remapped and visual

responses on a voxelwise basis. Areas that had relatively small remapped responses had TTP that were slightly earlier than expected: median TTP for V1 and V2 were 0.98 s and 0.97 s, respectively. The three areas that exhibited the strongest remapped response magnitude had TTPs that were consistent with the expected time course of remapping. Median differential TTPs in were 1.20 s in V3, 1.27 s in V3A, and and 1.22 s hV4. These values correspond closely to the interval between the onset of the stimulus and the average time of saccade onset (when we predict remapping to occur). Analysis of TTP indicates that the time at which the BOLD response peaks may contain important information regarding the underlying physiological responses.

4.4.7.4 Analysis of response decay. We next compared response *decay* for the two conditions. Response decay is the length of time that the response takes to return to baseline from the peak. Like the attack parameter, response decay is a duration, not an absolute time relative to the start of the trial. In contrast to differential response lag, we predicted that there would be no difference in response decay between remapped and visual responses.

We found that remapped responses tended to have marginally shorter decay values than visual responses (Figure 55). Visual responses had decay values that averaged 3.96 s (0.22 s SD) (Figure 55A, light bars). Remapped responses had similar decay values that averaged 3.90 s (0.18 s SD) (Figure 55A, dark bars). We calculated the difference in decay between remapped and visual responses on a voxelwise basis (Figure 55B). Remapped responses had decay values that were slightly smaller than visual responses, but in no visual areas was the median difference between the two more than 0.11 s (Figure 55B). These differential decay values show that remapped responses took slightly less time than visual responses to return to baseline from peak magnitude. In each area, the posterior probability that decay for remapped responses was shorter than for visual responses, $\mathsf{P}\{\gamma^{\text{decay}}_{\text{remap}} < \gamma^{\text{decay}}_{\text{visual}} \mid Y\}$, was less than chance ($p < 0.5$). This result indicates that the temporal differences between visual and remapped responses are restricted to response lag and attack; the two curves do not differ in decay.

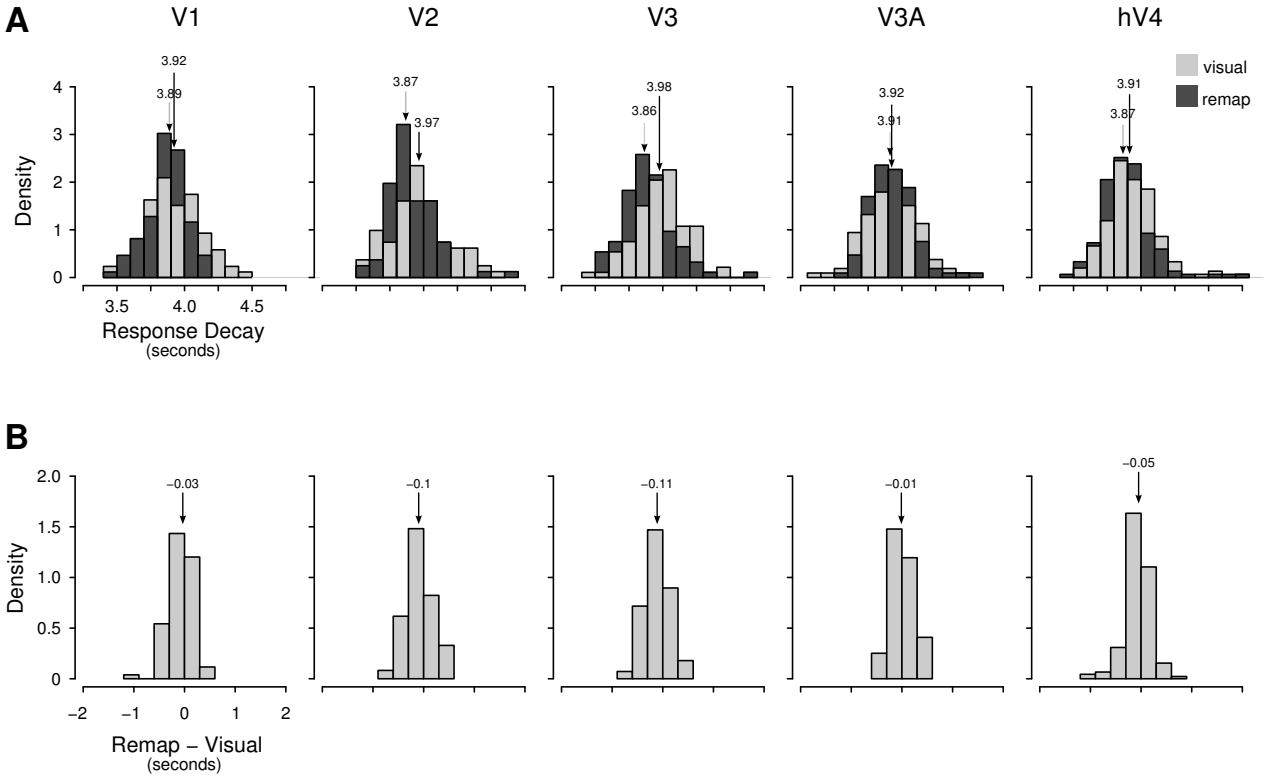

Figure 55. Remapped and visual responses have the same decay duration. Histograms show the distribution of the *decay* parameter across voxels in which both visual and remapped responses had a high posterior probability ($p \geq 0.95$) of being nonzero. (A) Histograms of decay estimates for visual (light gray) and remapped (dark gray) responses. (B) Histogram of the difference between visual and remapped response decay estimates. Downward-pointing arrows indicate the median of the distribution. The decay parameter represents the duration between the peak of the response and the point at which the response returns to baseline. While remapped responses in each visual area had marginally longer decay values, there was only a small probability that this difference is real.

4.5 DISCUSSION

The aim of this study was to determine whether human striate and extrastriate cortex have access to updated spatial information. We found robust remapping activity throughout the occipital lobe. Specifically, we found that the representation of a salient visual stimulus is shifted from one hemisphere to the other in conjunction with a horizontal saccade. Comparison with control conditions confirmed that responses in the single-step task were not due to either stimuli or saccades when they occur alone: remapping depends on the conjunction of these sensory and motor events.

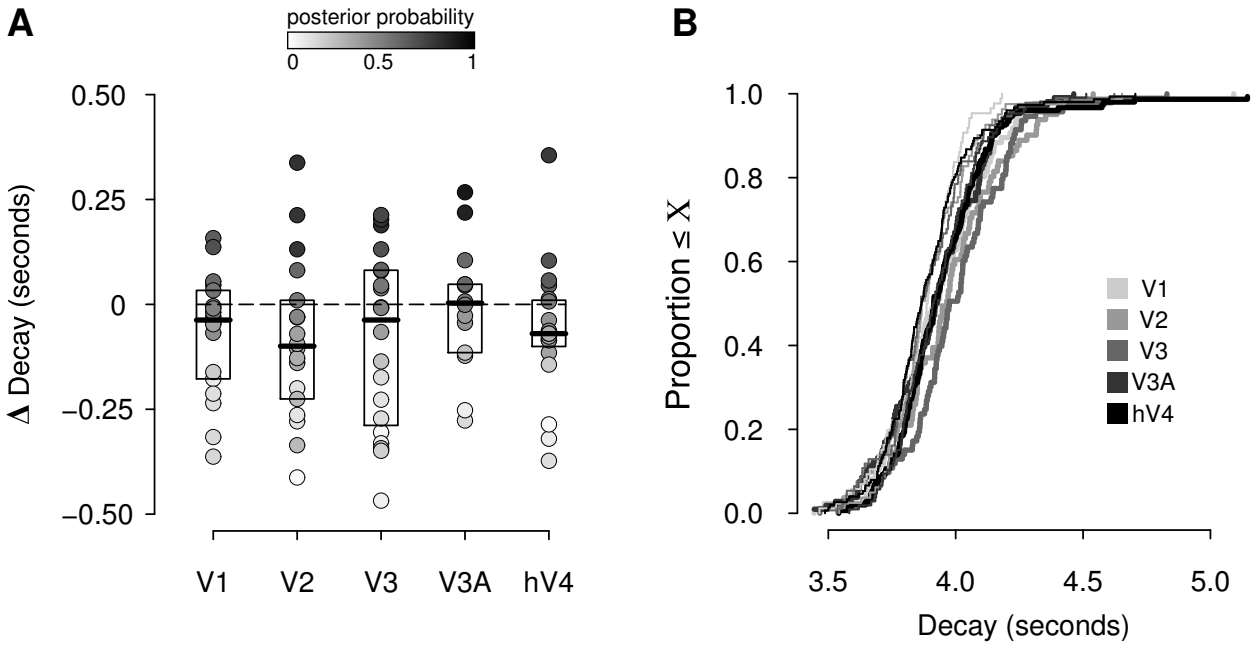

Figure 56. Difference between remapped and visual response decay. Each circle represents the difference in decay averaged over all voxels in a given visual area in a single hemisphere. Grayscale shading represents the posterior probability that remapped responses had larger decays than visual responses. Dotted line at zero represents no difference.

4.5.1 Remapping in areas linked to perception

Previous studies of spatial updating in humans have focused on parietal cortex. In an earlier functional imaging study, we found that human parietal cortex updates the trace of a visual stimulus in conjunction with voluntary eye movements (Merriam et al., 2003, see Chapter 2). Our previous result indicated that human parietal cortex makes use of an eye-centered spatial reference frame. This conclusion has been supported by imaging studies reporting parietal activation in both double- and triple-step saccade paradigms (Heide et al., 2001; Medendorp et al., 2003), and by human lesion studies showing that parietal cortex is necessary for accurate performance on the double-step task (Duhamel et al., 1992b; Heide et al., 1995).

The function of remapping is to maintain a stable representation of the world despite the constant shifting of images on the retina. Neurons in several cortical regions become active or grow silent according to whether a voluntary eye movement is going to place their receptive fields onto or away from the location of a remembered stimulus (Duhamel et al., 1992a;

Nakamura and Colby, 2002; Kusunoki and Goldberg, 2003). Thus activity representing the remembered location passes from one group of neurons to another so as to maintain an accurate retinotopic representation of the remembered location. Remapping enables the visual system to maintain spatially accurate representations across saccades. Remapping has been studied most extensively in cortical areas that are involved in eye movements and attention. If remapping is indeed important for perceptual stability, updated visual information should reach cortical areas that are directly involved in visual perception. We found that this is in fact the case: several regions of human visual cortex exhibit remapping.

4.5.2 Prevalence of remapping and visual area hierarchy

We observed remapping in both lower and higher-order extrastriate areas, consistent with observations in monkey neurophysiological studies (Nakamura and Colby, 2002). In both monkeys and humans, the prevalence of remapping increases monotonically with position in the visual area hierarchy (compare Figure 39B and Figure 57A). We found that remapping is robust in areas V3A and V4, and that the strength of remapping decreases in areas V3, V2, and V1. This gradient in the strength of remapping across visual cortex may reflect the strength of anatomical connections with LIP (Figure 57B). LIP is extensively interconnected with visual cortex, and the density of projections decrease at earlier stages of the hierarchy (see Figure 57; Schall et al., 1995; Stanton et al., 1995).

The neural circuit that produces remapping is unknown. Remapping requires that visual signals be shifted within cortex by corollary discharges of motor commands. Neuropsychological studies indicate that this combination of visual and motor signals is likely to take place in parietal cortex. Damage to parietal cortex in humans results in severe impairment in performance of the double-step saccade task, a task which requires the use of updated stimulus traces (Duhamel et al., 1992b; Heide et al., 1995, see Section 1.4). While visual and corollary discharge signals could potentially be combined in other cortical areas, impairments on the double-step task like those observed after parietal damage have not been demonstrated after other types of lesions, such as frontal eye field lesions (Heide et al., 1995). We conclude that parietal cortex is central to the process of remapping. Given this centrality,

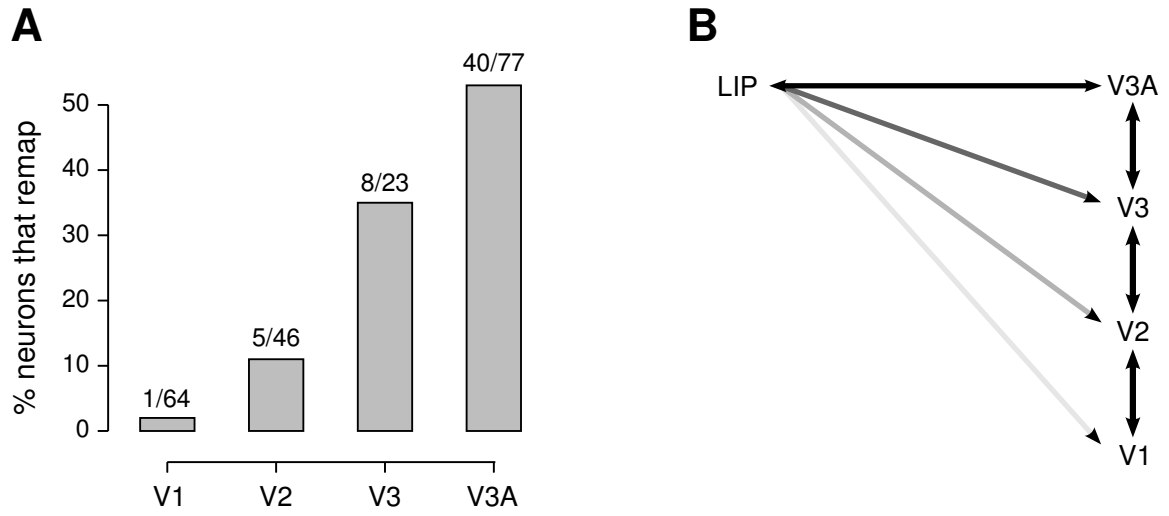

Figure 57. Remapping in monkey striate and extrastriate cortex. (A) At the single unit level, monotonic relationship between the prevalence of remapping and position in the visual hierarchy. (Adapted from Nakamura and Colby, 2002) (B) Prevalence of remapping parallels the strength of anatomical connectivity with area LIP (indicated by grayscale arrows). Remapping has not been studied in monkey area V4.

it is plausible that the visual cortex remapping signals we observed with functional imaging are the result of feedback projections from parietal to extrastriate cortex. The prevalence of remapping in each area is consistent with the relative strength of projections from parietal to extrastriate cortex.

4.5.3 Strength of remapped visual signals

fMRI activity in the single-step task is not a "pure" measurement of remapping: activation associated with stimuli and saccades contributed to the observed responses. In order to assess the magnitude of remapped responses, it is necessary to account for activity generated in the control conditions. This was not an issue in the original single-unit studies of remapping because stimuli and saccade targets were placed so as to ensure that neurons did not respond in either control condition (see Figure 6, Duhamel et al., 1992a). This issue has become increasingly important as studies attempt to identify the neural circuitry involved in remapping in finer detail. For example, a recent study compared the strength of remapping across multiple saccade directions (Heiser and Colby, *in press*). In this experiment, it was

not possible to place the stimulus outside the RF in every condition, nor was it possible to ensure that the saccade was always diverted away from the receptive field. Under these circumstances, LIP neurons commonly responded in at least one of the two control conditions. This problem led Heiser and Colby (*in press*) to develop an index measuring the strength of activity in the single-step task relative to responses in the two control conditions. This index enabled them to quantify the strength of remapping across saccade direction.

Based on the approach adopted by Heiser and Colby (*in press*), we generated an estimate of "pure" remapping activity, which we called $\widehat{\text{remap}}$ (see Section 4.4.6.3). This estimate was calculated by factoring out contributions from the two control conditions. This involved computations beyond a simple subtraction, as we also took into account the expected nonlinearities in the BOLD response. This method for estimating remapping magnitude involves three assumptions which are worth further discussion. First, we assumed that responses associated with visual stimuli and saccades do in fact sum sublinearly in the single-step task. This assumption is supported by empirical work demonstrating subadditivity for when events occur in rapid temporal succession (Wager et al., 2005).

Second, we assumed that responses in the contralateral hemisphere during the single-step task represent only the combined activity associated with contralateral stimuli and saccades. We calculated the subadditivity parameter by determining the degree to which responses in the contralateral hemisphere during the single-step task were smaller than what would be expected by the linear summation of visual and saccade responses measured independently. The validity of this assumption rests on there being no truncation effect. Truncation of a visual responses in LIP was observed by Duhamel et al. (1992a) in conjunction with saccades. LIP neurons grow silent as the receptive field is moved away from a stimulated screen location. Truncation of a visual response could have caused what appears to be the sublinear summation of visual and saccade responses in the contralateral hemisphere. We think this is unlikely, however, because truncation should be largest in areas that exhibit strong remapping — yet we observed the largest subadditivity in V1, the area that exhibited the least remapping.

The third assumption we made is that ipsilateral responses in the single-step task reflect activity associated with the visual stimuli and saccades, both of which are scaled by the sub-

additivity parameter, *and* activity associated with remapping (Equation 21, Section 4.4.6.3). This assumption does not take into account the sublinear summation associated with the addition of remapped responses and the scaled control responses. Our estimate of $\widehat{\text{remap}}$ is therefore likely to be *overly* conservative. The magnitude of the remapped signal may be greater than our estimate of it.

4.5.4 Strength versus prevalence

Comparing the strength of remapping in monkeys and humans raises a broader issue regarding the relationship between the BOLD signal and neural activity (Logothetis and Wandell, 2004; Kim, 2003; Ress and Heeger, 2003). One of the critical questions is whether modulations in neural responses are accompanied by an equivalent modulation in BOLD activity. Several examples indicate a surprisingly close correspondence. Contrast response functions in human V1 measured with BOLD have nearly the same shape as those measured with single units in monkeys (Boynton et al., 1999; Heeger et al., 2000; Logothetis et al., 2001). Similarly, coherence response functions in human area MT measured with BOLD have nearly the same shape as those measured with single-units in monkeys (Rees et al., 2000). BOLD response strength appears to be related to neural response strength.

Our experiment does not speak directly to this issue, as we did not attempt to manipulate remapping strength. We did expect a specific modulation in response strength as a function of visual area, however. Nakamura and Colby (2002) found that the proportion of neurons that exhibit remapping increases with position in the visual hierarchy. Remapping was found to be more prevalent in extrastriate than in striate cortex. The BOLD signal is likely sensitive to both the proportion of cells that exhibit a given effect and the strength of the effect in terms of response magnitude. Consequently, a small population of V1 neurons could drive a large BOLD response if those neurons have large responses (e.g., high firing rate). Nevertheless, our results reveal a pattern of remapping across visual areas that is remarkably similar to the proportion of neurons that exhibit remapping in monkeys. Our results therefore provide additional support to the claim that BOLD activity is sensitive to differences in neural response strength.

4.5.5 Remapping in human striate cortex

We observed remapping in area V1. This finding contrasts with a closely related single-unit study in which only one out of 64 neurons in V1 exhibited remapping (Nakamura and Colby, 2002). There are three potential explanations for this difference. First, it is possible that fMRI measurements of remapping in V1 do not accurately reflect the true degree to which remapping occurs in this region. The BOLD signal is driven by a complex mixture of pooled synaptic activity and spiking neuronal output that is more tightly correlated with the local field potential than with spikes (Logothetis and Wandell, 2004). Local field potentials are thought to reflect mostly somatodendritic events, and may therefore be more sensitive to neuronal inputs (Logothetis et al., 2001). The BOLD signal may thus be influenced by inputs to a given cortical area that are not revealed by classical electrophysiology experiments, which detect neuronal outputs. Our results demonstrating remapping in human striate cortex could be attributable to feedback input to V1 from areas that exhibit robust remapping at the single neuron level, such as extrastriate cortex, LIP, and FEF.

Second, it is possible that Nakamura and Colby (2002) underestimated the prevalence of remapping in monkey striate cortex. Activation in striate cortex is thought to correspond to the perceived, rather than actual, location of a visual stimulus. For example, V1 activation during binocular rivalry reflects the perceived stimulus location (Lee et al., 2005). In the single-step task, a stimulus flashes around the time of the saccade. This timing and sequence of events is known to cause human subjects to mislocalize the stimulus (Ross et al., 1997; Honda, 1999, 1991, 1989). Perisaccadic mislocalization can be several degrees in magnitude for stimuli presented just prior to saccade onset (Krekelberg et al., 2003), as in the monkey version of the single-step task. Perceptual mislocalization raises the possibility that neural signals are mislocalized in cortex. Nakamura and Colby (2002) recorded from neurons that would have received the remapped stimulus trace if remapping were spatially veridical. These neurons would not have responded to the stimulus trace if mislocalization produced a stimulus representation in the "wrong" cortical location. This would be most problematic in striate cortex because it has the smallest receptive fields.

A slight shift in the perceived stimulus location could move the memory trace off of the

receptive field of the neuron under study, thereby interfering with the detection of remapping. Furthermore, striate neurons are arranged in a very precise retinotopic map that occupies a large cortical surface area. Thus a small mislocalization of the stimulus could result in a larger shift in the locus of activity in striate than in extrastriate cortex. Finally, fMRI measurements are more likely to be immune to small cortical mislocalizations, because voxel-level activity reflects the joint activity of a large pool of neurons. This could explain the discrepancy between our observation of remapping in human striate cortex (Chapter 4) and the lack of remapping in striate cortex at the single-unit level (Nakamura and Colby, 2002).

Third, it is possible that our results and those of Nakamura and Colby (2002) differ because of a true species difference. There is a growing consensus that human and monkey visual cortex share many functional and structural properties (Sereno and Tootell, 2005). On the other hand, there is also a growing list of differences. Remapping in human striate cortex may represent yet another difference in the functional organization of macaque and human visual cortex. Our finding of remapping in human V1 may be analogous to the observation of attentional effects in human V1. Attentional effects have been difficult to observe in monkey V1 (McAdams and Maunsell, 1999; Luck et al., 1997) but have been widely reported in the human imaging literature (Pessoa et al., 2003; Kastner et al., 1998, 1999b; Ress et al., 2000; Sàenz et al., 2002; Gandhi et al., 1999). It is still unknown whether this is due to a species difference or to differences in the signal measured by the two techniques (Heeger and Ress, 2002).

4.5.6 Remapping in area V4

We observed the strongest remapping signals in area hV4. It is not known at the single-unit level whether neurons in V4 exhibit remapping. However, physiological studies in area V4 suggest that it has the visual, saccade, and cognitive properties to support remapping activity. Visual responses in V4 are modulated by both covert attention and by eye movements (Tolias et al., 2001). Furthermore, receptive fields in V4 are dynamic. Attention can modulate both the effective size of a neuron's receptive field and its selectivity for object features (McAdams and Maunsell, 1999; Moran and Desimone, 1985; Luck et al., 1997;

Spitzer et al., 1988; Haenny et al., 1988; Desimone, 1995). Attention can even shift the location of the receptive field toward the attended location (Connor et al., 1996, 1997). These single-unit studies suggest that receptive field properties of V4 neurons are not static, but instead change dynamically depending on the locus of attention. Our results indicate the responses in human V4 are dynamic as well.

4.5.7 Active vision

Vision is an active process: sensory and motor systems interact in order to create a coherent perception of the visual world. In both monkeys and humans, visual response properties have typically been studied during fixation. There is growing evidence, however, that eye movements have a profound impact on responses to visual stimuli. Eye movements affect multiple properties of classical receptive fields at several stages of visual information processing, from LGN through extrastriate cortex (Reppas et al., 2002; Sharma et al., 2003; Khayat et al., 2004; Tolias et al., 2001; Sylvester et al., 2005; Sylvester and Rees, 2005). Remapping is a prime example of active vision. In the remapping paradigm, information about the size and direction of each saccade influences the effective location of the neuron's receptive field. Our study demonstrates that saccades have a powerful influence on visual responses in the human visual system as well. The interplay between visual and motor information in visual cortex is likely fundamental to our perception of a stable world.

5.0 GENERAL DISCUSSION

5.1 OVERVIEW

We have demonstrated that remapping occurs in several cortical areas in humans, as it does in monkeys. While remapping is a physiological phenomenon, our underlying assumption is that remapping is related to the neural representation of space, a cognitive function. The experiments in this thesis represent a first step in achieving our overarching goal of understanding this representation in humans. The purpose of this chapter is to discuss our results in light of the larger questions that remain unresolved. In each section, we frame our discussion as a contrast between what we have learned thus far and what remains to be learned. First, we discuss the implications of our findings with regard to the network of brain areas involved in spatial updating in monkey and human (Section 5.2). Second, we discuss the relationship between remapping and accurate spatial perception (Section 5.3). Finally, we discuss the relationship between remapping activity in parietal cortex and the generation of accurate spatial behavior (Section 5.4).

5.2 NETWORK FOR REMAPPING IN MONKEYS AND HUMANS

The circuitry that produces remapping of stimulus traces is a topic of considerable current interest (Berman et al., 2005; Heiser et al., 2005). In monkeys, this circuitry is known to involve multiple cortical areas (LIP, FEF, extrastriate cortex) as well as SC. While a good deal is known about the anatomical connections among these structures, very little is known about their functional connectivity and how they work together to produce a stable framework for spatial perception and action.

Our studies have focused on posterior cortical regions. We assumed that many of the same areas that participate in remapping in monkey cortex would do so in humans as well. Our goal, however, was not to address the issue of homology directly. Rather, we sought to establish that both parietal and extrastriate cortex contribute to remapping in humans. Our results show that they do, but we leave open the question of homology. In the this section, we discuss the functional organization of human parietal cortex, particularly with respect to

the established organization of the monkey brain.

Monkey parietal cortex contains multiple functionally and anatomically defined areas (Colby and Duhamel, 1991). Whether human parietal cortex is divided into a similar set of regions is unknown (for review see Culham et al., 2001). Two interrelated issues regarding parietal organization are currently under debate. The first issue is the degree of functional specialization within parietal cortex. Some imaging studies have reported overlapping regions of parietal activation in tasks with very different cognitive and behavioral requirements, emphasizing the generality of function subserved by human parietal cortex (Wojciulik and Kanwisher, 1999). Other imaging studies have identified multiple, functionally specialized zones within the parietal lobe. For example, a large region in inferior parietal cortex is activated by tasks involving language processing and abstract numerical calculation, while other regions in superior parietal cortex are activated by eye movements, attention, and manual grasping (Corbetta et al., 2002; Simon et al., 2002). Conjunction paradigms (Price and Friston, 1997) have been a successful strategy for identifying functionally specialized cortical zones in the midst of diffuse patterns of activation. In parietal cortex, Bremmer et al. (2001) measured responses to auditory, visual, and somatosensory motion. Large regions are activated by each of these types of stimuli, but only a circumscribed region responds to all three. Similarly, large regions in parietal cortex respond to objects presented in either the visual or tactile modalities, but only a small region is activated by both (Grefkes et al., 2002). Overall, it is likely that human cortex does contain specialized areas comparable in function to those identified in monkey. These sites of multimodal activation in human parietal cortex are consistent with well-established multimodal areas within the primate intraparietal sulcus (Colby and Duhamel, 1991).

The second, related, issue is whether there are clear homologies in humans of cortical areas defined originally in monkeys, such as area LIP. Large portions of human parietal cortex are activated when subjects make saccadic and smooth pursuit eye movements (Berman et al., 1999; Corbetta et al., 1998; Luna et al., 1998; Petit and Haxby, 1999). Parietal cortex is also strongly activated when subjects make saccades to remembered targets (Sweeney et al., 1996). A recent study found two foci of activation near the intraparietal sulcus that encode remembered saccade targets in retinotopic coordinates (Sereno et al., 2001).

One of these regions is located near the middle part of the intraparietal sulcus, at the junction of the superior and inferior branches. The other region is located posteriorly at the base of the sulcus. While further studies are needed to establish the complete set of parietal areas in humans and nonhuman primates (Press et al., 2001; Van Essen et al., 2001), the emerging picture is one of a similar anatomical and physiological organization across species.

The goal of the present study was to demonstrate and characterize updating of visual information in human cortex. Our approach differs considerably from that of the localization studies described above. We asked whether voxels in and around the intraparietal sulcus exhibit response properties that are similar to those of single neurons recorded in area LIP in behaving monkeys. We did this without restricting our analysis by means of a conjunction paradigm. Rather, we began with a large, anatomically-defined ROI and included voxels that responded in any one of our task conditions (visual, saccade, or spatial updating). Once we identified task-related voxels, we performed a series of analyses on the entire population of voxels. This approach allowed us to demonstrate the existence of physiological responses to remapping in humans. The cortical region analyzed in this study likely includes both the anterior and posterior maps described by Sereno et al. (2001). This region presumably includes both the human homologue of area LIP and additional extrastriate and parietal visual areas as well. Further studies are needed to identify the specific regions in which remapping is localized and their correspondence to parietal areas in the monkey.

The network of areas that participates in remapping includes extrastriate cortex. The issue of homology is not straightforward here either. For example, human V4 has a retinotopic organization that is different from its monkey counterpart (Tootell and Hadjikhani, 2001; Brewer et al., 2005). Human area V3A has strong motion responses, though monkey V3A does not (Tootell et al., 1997a). Functional imaging studies have revealed a host of differences between monkey and human in both the anatomical organization and functional selectivity of responses (Larsson and Heeger, 2005; Press et al., 2001; Brewer et al., 2002). A major issue that remains to be addressed is whether extrastriate cortex helps to construct updated spatial representations or whether it simply reflects the outcome of processes instantiated elsewhere.

There are two brain regions known to be involved in remapping that we have not inves-

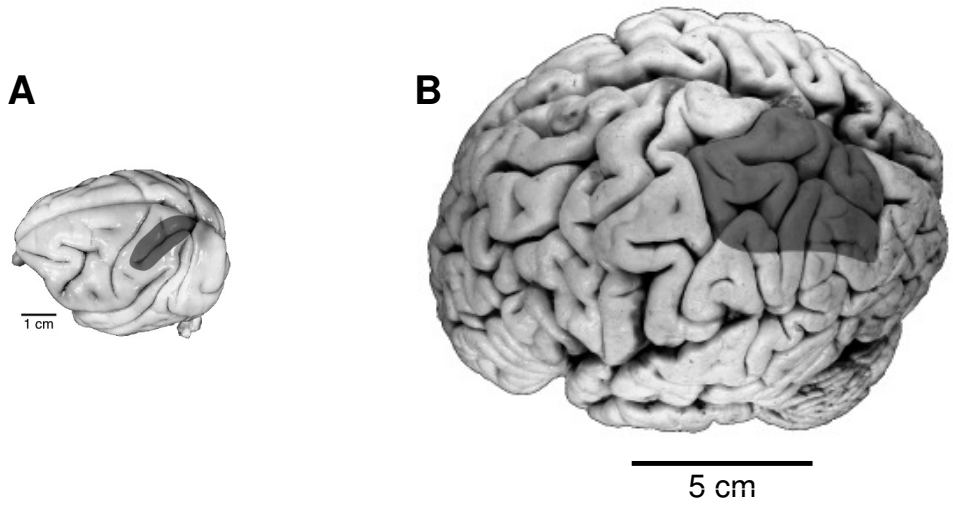

Figure 58. Monkey (A) and human (B) parietal cortex. Dark shaded region indicates the intraparietal sulcus and surrounding gyral surface. Area LIP is located on the lateral bank of the caudal portion of the intraparietal sulcus in monkeys. The human brain is larger and has a more complex folding pattern. Adapted from `http://brainmuseum.org/Specimens/primates`.

tigated in our studies — SC and FEF. A large proportion of neurons (30%) in monkey SC exhibit remapping activity (Walker et al., 1995), and it would be interesting to investigate remapping in human SC. Subcortical structures are notoriously difficult to study using fMRI, however, because of their small size and close proximity to pulsating vascular structures. Technological advances have begun to address these difficulties and several recent studies have demonstrated the feasibility of studying perceptual phenomena in subcortical structures (O'Connor et al., 2002; Schneider et al., 2004; Haynes et al., 2005; Wunderlich et al., 2005). Of particular relevance, one recent fMRI study has described detailed visual topography within human SC, and demonstrated a strong contralateral selectivity, which suggests the possibility of studying remapping in human SC (Schneider and Kastner, 2005).

Several studies have demonstrated that cells in monkey FEF respond to the updated stimulus trace (Goldberg et al., 1990; Umeno and Goldberg, 1997, 2001). Whether human FEF exhibits remapping is unknown. In humans, there is a growing consensus that a region proximal to the anterior wall of the precentral sulcus is the human analogue of monkey area FEF (Luna et al., 1998). This area exhibits many of the physiological response characteristics observed in human parietal cortex, such as activity during spatial working memory

tasks (Curtis et al., 2004; Berman and Colby, 2002; Sweeney et al., 1996), suggesting that this area too may exhibit remapping. However, human FEF shows surprisingly little selectivity for contralateral saccade targets and virtually no response to visual stimuli during fixation (Koyama et al., 2004; Hagler and Sereno, 2003; DeSimone et al., 2005). Both of these factors preclude the possibility of observing remapping activity in FEF with the current paradigms, which require the ability to detect lateralized visual responses.

5.3 REMAPPING AND PERCEPTION

Remapping activity in several brain regions is thought to underlie perceptual stability (Colby et al., 1995). This argument implies a relationship between physiology and visual perception. How strong is the evidence for this relationship? One way to address this question is to identify instances in which subjects misperceive stimulus locations, and then ask whether neural activity reflects the misperception.

Our perception of the location of stimuli presented around the time of an eye movement is not veridical. Rather, the location of perisaccadic stimuli is systematically misperceived. Stimuli presented prior to the saccade are mislocalized in the direction of the eye movement, while stimuli presented after the saccade are mislocalized against the direction of the eye movement (Honda, 1991; Ross et al., 1997). Is remapping activity related to the misperception of space? Consideration of the time course of remapping provides strong evidence for such a relationship. Nakamura and Colby (2002) recorded from extrastriate neurons in a version of the single-step task in which the stimulus was presented at various times relative to saccade onset, as did Kusunoki and Goldberg (2003) in area LIP. This design enabled them to characterize the time at which visual responsivity shifted from the *old* RF to the *new* RF.

Two properties of remapping indicate a correspondence between single-unit activity and perceptual localization. First, both stimulus mislocalization and remapping follow a similar time course. Perceptual mislocalization can begin as early as 200 ms prior to the saccade, but it is maximal as the eyes begin to move. Similarly, some neurons begin to remap

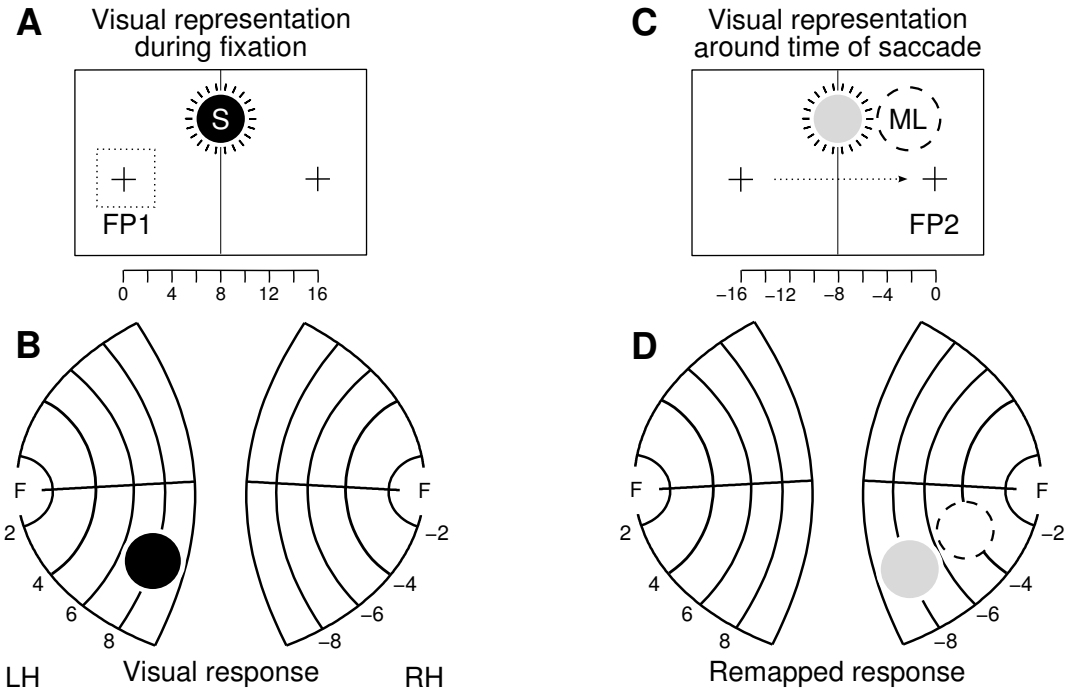

Figure 59. Perisaccadic stimulus mislocalization and cortical representation. (A) Fixation task. Stimulus flashed 8° in the right visual field is accurately perceived. (B) Cortical representation. Stimuli presented during fixation activates the corresponding retinotopic location in left hemisphere V1. (C) Single-step task. Stimulus flashed in the right visual field and a rightward eye movement brings the recently stimulated screen location (dotted circle) into the left visual left. The stimulus location is misperceived by the subject. The subject perceives the stimulus at a location that is shifted in the direction of the saccade. The mislocalized stimulus (ML) is represented by the dotted circle. (D) Cortical representation. Remapped activity in right hemisphere V1 may correspond to the perceived location (dotted circle) rather than the actual stimulated screen location (gray circle). F, fovea; LH & RH, left and right hemispheres.

prior to the saccade, but most neurons remap in conjunction with, or even after, movement onset. The similarity in time course between these two phenomena supports the idea that they are related. Ross et al. (2001) suggested that perceptual mislocalization arises as a consequence of the temporal variability of remapping. Specifically, Ross et al. (2001) built on the observation that remapping in parietal cortex can begin at a range of latencies, starting prior to the saccade for some neurons, and as late as 100 ms after the saccade for others. They argued that the range of remapping latencies would produce an ambiguous percept around the time of the saccade. These observations indicate that the temporal features of remapping are consistent with the time course of altered spatial perception.

The second property of remapping that may contribute to mislocalization is spatial, rather than temporal (Nakamura and Colby, 2002; Kusunoki and Goldberg, 2003). Some neurons respond only to the location where the RF will be after the eyes have moved (the *new* RF). Other neurons, however, continue to respond at the original RF location (the *old* RF) after the eyes have moved. Around the time of the saccade, some neurons show dual responsiveness: they can respond to stimuli at both the old and new RF locations. Perisaccadic stimulus mislocalization may be related to this dual responsiveness.

One question not addressed by these single-unit studies is whether the cortical locus of activity corresponds to the perceived, rather than the actual, location of the stimulus. This problem is illustrated in Figure 59. Stimulus location is accurately perceived during fixation (Figure 59A). When the eyes are fixated on FP1, a stimulus located 8° from the fovea in the right visual field activates the appropriate portion of the retinotopic map in the left hemisphere (Figure 59B). As described above, the perceived location of stimuli presented prior to the eye movement is shifted in the direction of the saccade (Figure 59C). Where in the brain are perisaccadic stimuli represented? If cortical activation corresponds to the perceived, rather than the actual, stimulus location, activation should be shifted away from the veridical stimulus location (8°) and toward the misperceived stimulus location, that is closer to the fovea (Figure 59D). Such a cortical correlate of mislocalization would provide compelling evidence for a relationship between remapping and perception.

5.4 REMAPPING AND BEHAVIOR

Accurate oculomotor behavior depends on a representation of space that is updated in conjunction with saccades. The most compelling evidence in support of this link comes from behavioral studies. Both monkeys and humans are able to perform the double-step saccade task, and accurate behavior of the double-step task requires an updated representation of space (Mays and Sparks, 1980; Hallett and Lightstone, 1976). The physiological phenomenon of remapping provides a neural mechanism for spatial updating (Goldberg et al., 1990; Duhamel et al., 1992a). The existence of responses to the stimulus trace indicates that

visual representations are updated between sets of neurons in conjunction with voluntary eye movements. These responses in the single-step task are not directly related to behavior, however, because the information provided by the stimulus trace is not used to guide a subsequent saccade. Thus, absent from this line of reasoning is evidence linking physiology and behavior, remapping activity and task performance. In this section, we review single-unit and functional imaging data suggesting that the strength of activity in parietal and prefrontal cortex may correlate with spatially accurate behavior on a trial-by-trial basis.

The clearest example of a correlation between physiological activity and behavior comes from studies in monkey area MT (Newsome et al., 1990; Shadlen et al., 1996). Area MT is an intermediate-level visual area, primarily involved in the perception of motion direction. The strength of activity in single MT neurons strongly predicts the animal's judgment regarding the direction of a coherent motion stimulus embedded in noise. This finding, along with lesion and microstimulation studies, firmly established a link between physiological activity and a particular sensory percept.

In the case of motion detection in MT, accurate performance of the task requires a perceptual decision. This example therefore doesn't directly support a relationship between neuronal activity and the generation of spatially accurate motor output. Ideally, we would like to observe a correlation between activity in parietal cortex and performance on tasks that explicitly depend on updated spatial information, such as the double-step task. Neurophysiological recordings from area LIP during double-step performance suggest that such a correlation exists. Berman et al. (*in preparation*) found only a small correlation between single-cell activity and performance on the double-step task. At the population level, however, neuronal activity corresponded closely to task performance.

Do fMRI measurements reflect neuronal activity with sufficient fidelity to detect a correlation between physiological responses and behavior? The BOLD response reflects synaptic activity pooled over a relatively large cortical zone. Thus, one would naturally assume a relatively weak relationship between BOLD activity and behavior. Several studies, however, have demonstrated a surprisingly strong correlation between the two. In an illuminating study, Ress et al. (2000) discovered that fluctuations in the fMRI signal can be used to predict task performance on a trial by trial basis. Ress et al. (2000) tested subjects on a

pattern-detection task. Task difficultly was titrated by embedding the pattern in noise, thereby ensuring that subjects made a substantial number of detection errors, either missing the pattern when it was present (misses), or reporting it when it was not present (false alarms). The critical finding in this study is that activity in primary visual cortex predicted correct task performance. Activity was larger when subjects made a correct judgment regarding the presence or absence of the pattern. This functional imaging study bears some similarities with the single-unit studies in MT, described above. In both experiments, brain activity predicted perceptual performance. In neither example was neural activity found to be related to motor behavior.

Two recent functional imaging studies have provided further evidence that the level of brain activation can be used to predict behavior. Pessoa et al. (2002b) tested subjects on a delayed match to sample task. Subjects had to hold a complex visual pattern in memory during a delay period. At the end of the trial, a probe stimulus appeared, and subjects were required to indicate whether or not a probe stimulus matched the sample stimulus. The central finding in this study is that activity in parietal and prefrontal cortex during the delay period predicted accurate task performance. This study indicates that the magnitude of cortical activity reflects the fidelity of information stored over a delay period. However, again, this study does not provide a direct link between physiological responses and the generation of spatially-accurate motor output because the required response was simply a button press.

Curtis et al. (2004) tested subjects on a modified version of the delayed saccade task. The central finding of this study is that delay-period activity in both FEF and parietal cortex correlates with saccadic accuracy. Multiple processes likely contributed to delay-period activity in this task, including spatial memory and motor planning. However, this task may have involved remapping as well because subjects made saccades to previously flashed stimuli. Thus, this study suggests a link between physiological measures of remapping and oculomotor behavior. Future studies will be needed to test this link using tasks, such as the double-step task, that more effectively measure remapping activity in isolation from other cognitive and motor processes. We established that areas in human cortex, like areas in the non-human primate, are able to construct sophisticated visual representations that take our eye movements into account.

5.5 CONCLUSIONS

We move our eyes more often than our hearts beat. While saccades are perhaps less critical to our survival, both occur effortlessly and both typically proceed unnoticed. Saccades present a significant computational challenge to the visual system. The issue of how our brains construct a stable representation from chaotic retinal input has been a central question in vision science for over 150 years. We have learned a great deal about neural mechanisms of perceptual stability in just the last 20 years. Most of this work has been conducted in monkeys, one neuron at a time. The recent advent of fMRI and its widespread adoption in the last 10 years has opened up the possibility of investigating neural mechanisms for perceptual stability in humans. Our studies contribute to this effort by establishing that remapping occurs in humans and that it is amenable to investigation with fMRI. Functional imaging is still a nascent field, and new methods for analyzing fMRI data are very much needed to improve the sensitivity of statistical inference. Our Bayesian model of the BOLD response provides a flexible framework for achieving this goal. We have identified several regions that exhibit remapping, focusing on sensory areas in the posterior lobes of the brain.

BIBLIOGRAPHY

Aguirre, G. K. and D'Esposito, M. (2000). Experimental design for brain fmri. In Moonen, C. T. W. and Bandettini, P. A., editors, *In Functional MRI*, pages 369–381. Springer, Berlin. 2.3.2

Aguirre, G. K., Zarahn, E., and D'Esposito, M. (1998a). A critique of the use of the kolmogorov-smirnov (ks) statistic for the analysis of BOLD fMRI data. *Magn Reson Med*, 39(3):500–505. 3.2, 4.4.7.1

Aguirre, G. K., Zarahn, E., and D'Esposito, M. (1998b). The variability of human, BOLD hemodynamic responses. *Neuroimage*, 8(4):360–369. 3.2.3

Anagnostou, E., Kleiser, R., and Skrandies, W. (2000). Electrophysiological correlates of human intrasaccadic processing. *Exp Brain Res*, 130(2):177–187. 1.6.5

Andersen, R. A. (1997). Multimodal integration for the representation of space in the posterior parietal cortex. *Philos Trans R Soc Lond B Biol Sci*, 352(1360):1421–1428. 1.6

Andersen, R. A., Asanuma, C., Essick, G., and Siegel, R. M. (1990a). Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J Comp Neurol*, 296(1):65–113. 1.4, 1.5

Andersen, R. A., Asanuma, C., Essick, G., and Siegel, R. M. (1990b). Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J Comp Neurol*, 296(1):65–113. 1.5, 1.5

Anderson, T. J., Jenkins, I. H., Brooks, D. J., Hawken, M. B., Frackowiak, R. S., and Kennard, C. (1994). Cortical control of saccades and fixation in man. a pet study. *Brain*, 117 (Pt 5):1073–1084. 1.6.2

Arrington, C. M., Carr, T. H., Mayer, A. R., and Rao, S. M. (2000). Neural mechanisms of visual attention: object-based selection of a region in space. *J Cogn Neurosci*, 12 Suppl 2:106–117. 1.6.3

Asanuma, C., Andersen, R. A., and Cowan, W. M. (1985). The thalamic relations of the caudal inferior parietal lobule and the lateral prefrontal cortex in monkeys: divergent cortical projections from cell clusters in the medial pulvinar nucleus. *J Comp Neurol*, 241(3):357–381. 1.5

Avidan, G., Hasson, U., Hendler, T., Zohary, E., and Malach, R. (2002). Analysis of the neuronal selectivity underlying low fMRI signals. *Curr Biol*, 12(12):964–972. 3.2.1

Awater, H. and Lappe, M. (2004). Perception of visual space at the time of pro- and anti-saccades. *J Neurophysiol*, 91(6):2457–2464. 1.2

Baizer, J. S. and Bender, D. B. (1989). Comparison of saccadic eye movements in humans and macaques to single-step and double-step target movements. *Vision Res*, 29(4):485–495. 2.2

Baizer, J. S., Desimone, R., and Ungerleider, L. G. (1993). Comparison of subcortical connections of inferior temporal and posterior parietal cortex in monkeys. *Vis Neurosci*, 10(1):59–72. 1.5

Baizer, J. S., Ungerleider, L. G., and Desimone, R. (1991). Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J Neurosci*, 11(1):168–190. 1.5

Baker, J. T., Patel, G. H., Corbetta, M., and Snyder, L. H. (2005). Distribution of activity across the monkey cerebral cortical surface, thalamus and midbrain during rapid, visually guided saccades. *Cereb Cortex*. 1.6.2

Bandettini, P. A. and Cox, R. W. (2000). Event-related fMRI contrast when using constant interstimulus interval: theory and experiment. *Magn Reson Med*, 43(4):540–548. 2.3.2

Bandettini, P. A., Jesmanowicz, A., Wong, E. C., and Hyde, J. S. (1993). Processing strategies for time-course data sets in functional MRI of the human brain. *Magn Reson Med*, 30(2):161–173. 3.2

Barash, S., Bracewell, R. M., Fogassi, L., Gnadt, J. W., and Andersen, R. A. (1991a). Saccade-related activity in the lateral intraparietal area. i. temporal properties; comparison with area 7a. *J Neurophysiol*, 66(3):1095–1108. 1.5, 1.6.2

Barash, S., Bracewell, R. M., Fogassi, L., Gnadt, J. W., and Andersen, R. A. (1991b). Saccade-related activity in the lateral intraparietal area. ii. spatial properties. *J Neurophysiol*, 66(3):1109–1124. 1.5, 1.6, 1.6.2, 2.5

Beauchamp, M. S., Petit, L., Ellmore, T. M., Ingeholm, J., and Haxby, J. V. (2001). A parametric fMRI study of overt and covert shifts of visuospatial attention. *Neuroimage*, 14(2):310–321. 1.6.2

Behrmann, M., Geng, J. J., and Shomstein, S. (2004). Parietal cortex and attention. *Curr Opin Neurobiol*, 14(2):212–217. 1.6.3

Bellebaum, C., Hoffmann, K.-P., and Daum, I. (2005). Post-saccadic updating of visual space in the posterior parietal cortex in humans. *Behav Brain Res*, 163(2):194–203. 1.6.5

Belmonte, M. and Yurgelun-Todd, D. (2001). Permutation testing made practical for functional magnetic resonance image analysis. *IEEE Trans Med Imaging*, 20(3):243–248. 3.2

Ben Hamed, S., Duhamel, J. R., Bremmer, F., and Graf, W. (2001). Representation of the visual field in the lateral intraparietal area of macaque monkeys: a quantitative receptive field analysis. *Exp Brain Res*, 140(2):127–144. 2.5

Bender, D. B. (1981). Retinotopic organization of macaque pulvinar. *J Neurophysiol*, 46(3):672–693. 1.5

Benevento, L. A. and Standage, G. P. (1983). The organization of projections of the retinorecipient and nonretinorecipient nuclei of the pretectal complex and layers of the superior colliculus to the lateral pulvinar and medial pulvinar in the macaque monkey. *J Comp Neurol*, 217(3):307–336. 1.5

Berman, R. A. and Colby, C. L. (2002). Spatial working memory in human extrastriate cortex. *Physiol Behav*, 77(4-5):621–627. 2.5.1, 5.2

Berman, R. A., Colby, C. L., Genovese, C. R., Voyvodic, J. T., Luna, B., Thulborn, K. R., and Sweeney, J. A. (1999). Cortical networks subserving pursuit and saccadic eye movements in humans: an fMRI study. *Hum Brain Mapp*, 8(4):209–225. 1.6.2, 3.2.2, 5.2

Berman, R. A., Heiser, L. M., Saunders, R. C., and Colby, C. L. (2005). Dynamic circuitry for updating spatial representations: I. behavioral evidence for interhemispheric transfer in the split-brain macaque. *J Neurophysiol*. 5.2

Birn, R. M. and Bandettini, P. A. (2005). The effect of stimulus duty cycle and "off" duration on BOLD response linearity. *Neuroimage*, 27(1):70–82. 4.4.6.3

Birn, R. M., Cox, R. W., and Bandettini, P. A. (2002). Detection versus estimation in event-related fMRI: choosing the optimal stimulus timing. *Neuroimage*, 15(1):252–264. 2.3.2

Birn, R. M., Saad, Z. S., and Bandettini, P. A. (2001). Spatial heterogeneity of the nonlinear dynamics in the fMRI BOLD response. *Neuroimage*, 14(4):817–826. 4.4.6.3

Bisley, J. W., Krishna, B. S., and Goldberg, M. E. (2004). A rapid and precise on-response in posterior parietal cortex. *J Neurosci*, 24(8):1833–1838. 1.5, 1.6.1

Blatt, G. J., Andersen, R. A., and Stoner, G. R. (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area lip) in the macaque. *J Comp Neurol*, 299(4):421–445. 1.5

Boynton, G. M., Demb, J. B., Glover, G. H., and Heeger, D. J. (1999). Neuronal basis of contrast discrimination. *Vision Res*, 39(2):257–269. 4.5.4

Boynton, G. M., Engel, S. A., Glover, G. H., and Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci*, 16(13):4207–4221. 2.4.2, 3.2.3, 19, 3.4.1, 4.4.1.1, 4.4.6.3

Brefczynski, J. A. and DeYoe, E. A. (1999). A physiological correlate of the 'spotlight' of visual attention. *Nat Neurosci*, 2(4):370–374. 4.2

Bremmer, F., Schlack, A., Shah, N. J., Zafiris, O., Kubischik, M., Hoffmann, K., Zilles, K., and Fink, G. R. (2001). Polymodal motion processing in posterior parietal and premotor cortex: a human fMRI study strongly implies equivalencies between humans and monkeys. *Neuron*, 29(1):287–296. 5.2

Brewer, A. A., Liu, J., Wade, A. R., and Wandell, B. A. (2005). Visual field maps and stimulus selectivity in human ventral occipital cortex. *Nat Neurosci*, 8(8):1102–1109. 4.3.6, 5.2

Brewer, A. A., Press, W. A., Logothetis, N. K., and Wandell, B. A. (2002). Visual areas in macaque cortex measured using functional magnetic resonance imaging. *J Neurosci*, 22(23):10416–10426. 5.2

Bridgeman, B. and Stark, L. (1991). Ocular proprioception and efference copy in registering visual direction. *Vision Res*, 31(11):1903–1913. 1.1

Bridgeman, B., van der Heijden, A. H. L., and Velichkovsky, B. M. (1994). A theory of visual stability across saccadic eye movements. *Behavioral and Brain Sciences*, 17:247–292. 1.1

Brindley, G. S. and Merton, P. A. (1960). The absence of position sense in the human eye. *J Physiol*, 153:127–130. 1.1

Bruce, C. J., Goldberg, M. E., Bushnell, M. C., and Stanton, G. B. (1985). Primate frontal eye fields. ii. physiological and anatomical correlates of electrically evoked eye movements. *J Neurophysiol*, 54(3):714–734. 1.5, 3.2.2

Bruno, R. M. and Simons, D. J. (2002). Feedforward mechanisms of excitatory and inhibitory cortical receptive fields. *J Neurosci*, 22(24):10966–10975. 3.2.1

Bullmore, E. T., Suckling, J., Overmeyer, S., Rabe-Hesketh, S., Taylor, E., and Brammer, M. J. (1999). Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural mr images of the brain. *IEEE Trans Med Imaging*, 18(1):32–42. 3.2

Buracas, G. T. and Boynton, G. M. (2002). Efficient design of event-related fMRI experiments using m-sequences. *Neuroimage*, 16(3 Pt 1):801–813. 3.5.2, 4.3.3

Burock, M. A. and Dale, A. M. (2000). Estimation and detection of event-related fMRI signals with temporally correlated noise: a statistically efficient and unbiased approach. *Hum Brain Mapp*, 11(4):249–260. 3.2.3

Bushnell, M. C., Goldberg, M. E., and Robinson, D. L. (1981). Behavioral enhancement of visual responses in monkey cerebral cortex. i. modulation in posterior parietal cortex related to selective visual attention. *J Neurophysiol*, 46(4):755–772. 1.6.3

Buxton, R. B., Uludag, K., Dubowitz, D. J., and Liu, T. T. (2004). Modeling the hemodynamic response to brain activation. *Neuroimage*, 23 Suppl 1:220–233. 4.4.6.3

Cavada, C. and Goldman-Rakic, P. S. (1989). Posterior parietal cortex in rhesus monkey: Ii. evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal lobe. *J Comp Neurol*, 287(4):422–445. 1.5

Clower, D. M., West, R. A., Lynch, J. C., and Strick, P. L. (2001). The inferior parietal lobule is the target of output from the superior colliculus, hippocampus, and cerebellum. *J Neurosci*, 21(16):6283–6291. 1.5

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., and Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, 386(6625):604–608. 3.2.1

Cohen, M. S. (1997). Parametric analysis of fMRI data using linear systems methods. *Neuroimage*, 6(2):93–103. 3.2.3, 22

Cohen, M. S. and Bookheimer, S. Y. (1994). Localization of brain function using magnetic resonance imaging. *Trends Neurosci*, 17(7):268–277. 3.2

Colby, C. (1996). A neurophysiological distinction between attention and intention. In McClelland and Inui, T., editors, *Attention and performance XVI: Information integration in perception and communication*, pages 157–177. MIT Press. 1.6.3, 2.5.1

Colby, C. L. and Duhamel, J. R. (1991). Heterogeneity of extrastriate visual areas and multiple parietal areas in the macaque monkey. *Neuropsychologia*, 29(6):517–537. 5.2

Colby, C. L., Duhamel, J. R., and Goldberg, M. E. (1995). Oculocentric spatial representation in parietal cortex. *Cereb Cortex*, 5(5):470–481. 1.4, 5.3

Colby, C. L., Duhamel, J. R., and Goldberg, M. E. (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J Neurophysiol*, 76(5):2841–2852. 1.5, 1.6, 1.6.1, 1.6.2, 1.6.3, 2.5, 4.2

Colby, C. L., Gattass, R., Olson, C. R., and Gross, C. G. (1988). Topographical organization of cortical afferents to extrastriate visual area po in the macaque: a dual tracer study. *J Comp Neurol*, 269(3):392–413. 1.4

Colby, C. L. and Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annu Rev Neurosci*, 22:319–349. 1.4, 1.5, 2.2

Connor, C. E., Gallant, J. L., Preddie, D. C., and Van Essen, D. C. (1996). Responses in area V4 depend on the spatial relationship between stimulus and attention. *J Neurophysiol*, 75(3):1306–1308. 4.5.6

Connor, C. E., Preddie, D. C., Gallant, J. L., and Van Essen, D. C. (1997). Spatial attention effects in macaque area V4. *J Neurosci*, 17(9):3201–3214. 4.5.6

Corbetta, M. (1993). Positron emission tomography as a tool to study human vision and attention. *Proc Natl Acad Sci U S A*, 90(23):10901–10903. 1.6.3

Corbetta, M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Ollinger, J. M., Drury, H. A., Linenweber, M. R., Petersen, S. E., Raichle, M. E., Van Essen, D. C., and Shulman, G. L. (1998). A common network of functional areas for attention and eye movements. *Neuron*, 21(4):761–773. 5.2

Corbetta, M., Burton, H., Sinclair, R. J., Conturo, T. E., Akbudak, E., and McDonald, J. W. (2002). Functional reorganization and stability of somatosensory-motor cortical topography in a tetraplegic subject with late recovery. *Proc Natl Acad Sci U S A*, 99(26):17066–17071. 1.6.3, 5.2

Corbetta, M., Kincade, J. M., Ollinger, J. M., McAvoy, M. P., and Shulman, G. L. (2000). Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nat Neurosci*, 3(3):292–297. 1.6.3

Courtney, S. M., Petit, L., Maisog, J. M., Ungerleider, L. G., and Haxby, J. V. (1998). An area specialized for spatial working memory in human frontal cortex. *Science*, 279(5355):1347–1351. 2.5.1

Cox, R. W. and Hyde, J. S. (1997). Software tools for analysis and visualization of fMRI data. *NMR Biomed*, 10(4-5):171–178. 2.3.4.2

Culham, J. C., Cavanagh, P., and Kanwisher, N. G. (2001). Attention response functions: characterizing brain areas using fMRI activation during parametric variations of attentional load. *Neuron*, 32(4):737–745. 5.2

Culham, J. C. and Kanwisher, N. G. (2001). Neuroimaging of cognitive functions in human parietal cortex. *Curr Opin Neurobiol*, 11(2):157–163. 1.6.3

Curtis, C. and D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends Cogn Sci*, 7(9):415–423. 1.6.4

Curtis, C. E., Rao, V. Y., and D'Esposito, M. (2004). Maintenance of spatial and motor codes during oculomotor delayed response tasks. *J Neurosci*, 24(16):3944–3952. 1.6.4, 3.2.1, 5.2, 5.4

Curtis, C. E., Sun, F. T., Miller, L. M., and D'Esposito, M. (2005). Coherence between fMRI time-series distinguishes two spatial working memory networks. *Neuroimage*, 26(1):177–183. 3.2.1

Dale, A. M. and Buckner, R. L. (1997). Selective averaging of rapidly presented inficual trials using fMRI. *Human Brain Mapping*, 5:329–340. 3.2, 3.2.3, 4.4.1.1

Dale, A. M., Fischl, B., and Sereno, M. I. (1999). Cortical surface-based analysis. i. segmentation and surface reconstruction. *Neuroimage*, 9(2):179–194. 2.3.4.2, 3.2.3, 4.3.5.1

Darby, D. G., Nobre, A. C., Thangaraj, V., Edelman, R., Mesulam, M. M., and Warach, S. (1996). Cortical activation in the human brain during lateral saccades using epistar functional magnetic resonance imaging. *Neuroimage*, 3(1):53–62. 1.6.2

DeSimone, K., Weiner, K., Schneider, K. A., and Kastner, S. (2005). Topographic maps in human frontal and parietal cortex: a high-resolution fMRI study. *Soc Neurosci Abstr*, 508.18. 5.2

Desimone, R. (1995). Neuropsychology. is dopamine a missing link? *Nature*, 376(6541):549–550. 4.5.6

DeSouza, J. F. X., Dukelow, S. P., and Vilis, T. (2002). Eye position signals modulate early dorsal and ventral visual areas. *Cereb Cortex*, 12(9):991–997. 4.2

DeYoe, E. A., Carman, G. J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., Miller, D., and Neitz, J. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. *Proc Natl Acad Sci U S A*, 93(6):2382–2386. 1.6.1, 2.2, 4.3.6, 4.4.3

Donaldson, I. M. (2000). The functions of the proprioceptors of the eye muscles. *Philos Trans R Soc Lond B Biol Sci*, 355(1404):1685–1754. 1.1

Dorris, M. C. and Glimcher, P. W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron*, 44(2):365–378. 1.6

Dougherty, R. F., Koch, V. M., Brewer, A. A., Fischer, B., Modersitzki, J., and Wandell, B. A. (2003). Visual field representations and locations of visual areas V1/2/3 in human visual cortex. *J Vis*, 3(10):586–598. 4.3.6, 4.4.3

Downar, J., Crawley, A. P., Mikulis, D. J., and Davis, K. D. (2000). A multimodal cortical network for the detection of changes in the sensory environment. *Nat Neurosci*, 3(3):277–283. 1.6.3

Downar, J., Crawley, A. P., Mikulis, D. J., and Davis, K. D. (2001). The effect of task relevance on the cortical response to changes in visual and auditory stimuli: an event-related fMRI study. *Neuroimage*, 14(6):1256–1267. 1.6.3

Downar, J., Crawley, A. P., Mikulis, D. J., and Davis, K. D. (2002). A cortical network sensitive to stimulus salience in a neutral behavioral context across multiple sensory modalities. *J Neurophysiol*, 87(1):615–620. 1.6.3

Duann, J.-R., Jung, T.-P., Kuo, W.-J., Yeh, T.-C., Makeig, S., Hsieh, J.-C., and Sejnowski, T. J. (2002). Single-trial variability in event-related BOLD signals. *Neuroimage*, 15(4):823–835. 2.3.4.3

Duhamel, J. R., Colby, C. L., and Goldberg, M. E. (1992a). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255(5040):90–92. 1.4, 6, 1.4, 1.5, 1.6.5, 2.2, 2.5.1, 2.5.2, 4.2, 4.4.3, 4.4.7, 4.5.1, 4.5.3, 5.4

Duhamel, J. R., Goldberg, M. E., Fitzgibbon, E. J., Sirigu, A., and Grafman, J. (1992b). Saccadic dysmetria in a patient with a right frontoparietal lesion. the importance of corollary discharge for accurate spatial behaviour. *Brain*, 115 (Pt 5):1387–1402. 1.3, 2.2, 4.5.1, 4.5.2

Duong, T. Q., Kim, D. S., Ugurbil, K., and Kim, S. G. (2000). Spatiotemporal dynamics of the BOLD fMRI signals: toward mapping submillimeter cortical columns using the early negative response. *Magn Reson Med*, 44(2):231–242. 3.2, 3.2.3

Eddy, W. F., Fitzgerald, M., and Noll, D. C. (1996). Improved image registration by using fourier interpolation. *Magn Reson Med*, 36(6):923–931. 2.3.4.1, 4.3.5.2

Engel, S. A., Glover, G. H., and Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex*, 7(2):181–192. 1.6.1, 2.2, 4.3.6, 4.4.3

Engel, S. A., Rumelhart, D. E., Wandell, B. A., Lee, A. T., Glover, G. H., Chichilnisky, E. J., and Shadlen, M. N. (1994). fMRI of human visual cortex. *Nature*, 369(6481):525–525. 1.6.1, 4.3.6, 4.4.3

Esposito, F., Formisano, E., Seifritz, E., Goebel, R., Morrone, R., Tedeschi, G., and Di Salle, F. (2002). Spatial independent component analysis of functional MRI time-series: to what extent do results depend on the algorithm used? *Hum Brain Mapp*, 16(3):146–157. 3.2.1

Fischl, B., Sereno, M. I., Tootell, R. B., and Dale, A. M. (1999). High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum Brain Mapp*, 8(4):272–284. 2.3.4.2, 4.3.5.1

Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., and Noll, D. C. (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. *Magn Reson Med*, 33(5):636–647. 3.2.2

Formisano, E. and Goebel, R. (2003). Tracking cognitive processes with functional MRI mental chronometry. *Curr Opin Neurobiol*, 13(2):174–181. 3.2.1

Formisano, E., Linden, D. E. J., Di Salle, F., Trojano, L., Esposito, F., Sack, A. T., Grossi, D., Zanella, F. E., and Goebel, R. (2002). Tracking the mind's image in the brain i: time-resolved fMRI during visuospatial mental imagery. *Neuron*, 35(1):185–194. 3.2.1

Fox, P. T., Fox, J. M., Raichle, M. E., and Burde, R. M. (1985). The role of cerebral cortex in the generation of voluntary saccades: a positron emission tomographic study. *J Neurophysiol*, 54(2):348–369. 1.6.2

Fox, P. T., Miezin, F. M., Allman, J. M., Van Essen, D. C., and Raichle, M. E. (1987). Retinotopic organization of human visual cortex mapped with positron-emission tomography. *J Neurosci*, 7(3):913–922. 1.6.1

Fox, P. T., Mintun, M. A., Raichle, M. E., Miezin, F. M., Allman, J. M., and Van Essen, D. C. (1986). Mapping human visual cortex with positron emission tomography. *Nature*, 323(6091):806–809. 1.6.1

Friston, K. J., Glaser, D. E., Henson, R. N. A., Kiebel, S., Phillips, C., and Ashburner, J. (2002a). Classical and bayesian inference in neuroimaging: applications. *Neuroimage*, 16(2):484–512. 3.3

Friston, K. J., Holmes, A. P., Poline, J. B., Grasby, P. J., Williams, S. C., Frackowiak, R. S., and Turner, R. (1995). Analysis of fMRI time-series revisited. *Neuroimage*, 2(1):45–53. 3.2

Friston, K. J., Josephs, O., Rees, G., and Turner, R. (1998). Nonlinear event-related responses in fMRI. *Magn Reson Med*, 39(1):41–52. 4.4.6.3

Friston, K. J., Mechelli, A., Turner, R., and Price, C. J. (2000a). Nonlinear responses in fMRI: the balloon model, volterra kernels, and other hemodynamics. *Neuroimage*, 12(4):466–477. 4.4.6.3

Friston, K. J., Mechelli, A., Turner, R., and Price, C. J. (2000b). Nonlinear responses in fMRI: the balloon model, volterra kernels, and other hemodynamics. *Neuroimage*, 12(4):466–477. 4.4.6.3

Friston, K. J. and Penny, W. (2003). Posterior probability maps and spms. *Neuroimage*, 19(3):1240–1249. 3.3, 3.5.3

Friston, K. J., Penny, W., Phillips, C., Kiebel, S., Hinton, G., and Ashburner, J. (2002b). Classical and bayesian inference in neuroimaging: theory. *Neuroimage*, 16(2):465–483. 3.3

Gandhi, S. P., Heeger, D. J., and Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proc Natl Acad Sci U S A*, 96(6):3314–3319. 4.2, 4.5.5

Gardner, J. L., Sun, P., Waggoner, R. A., Ueno, K., Tanaka, K., and Cheng, K. (2005). Contrast adaptation and representation in human early visual cortex. *Neuron*, 47(4):607–620. 3.2.1

Gattass, R., Gross, C. G., and Sandell, J. H. (1981). Visual topography of V2 in the macaque. *J Comp Neurol*, 201(4):519–539. 4.4.3, 4.4.5.1

Gattass, R., Sousa, A. P., and Gross, C. G. (1988). Visuotopic organization and extent of V3 and V4 of the macaque. *J Neurosci*, 8(6):1831–1845. 4.4.3, 4.4.5.1

Gauthier, G. M., Nommay, D., and Vercher, J. L. (1990). Ocular muscle proprioception and visual localization of targets in man. *Brain*, 113 (Pt 6):1857–1871. 1.1

Genovese, C. (1998a). Functional magnetic resonance imaging and spatio-temporal inference. In Bernardo, J., Berger, J., Dawid, A., and Smith, A., editors, *Bayesian Statistics*, volume 6. Oxford University Press. 3.4.1, 3.4.1, 21

Genovese, C. (1998b). Statistical inference in functional magnetic resonance imaging. Technical Report 674, CMU Statistics. 3.4.1

Genovese, C. R. (2000). A bayesian time-course model for functional magnetic resonance imaging data (with discussion). *Journal of the American Statistical Association*, 95(451):691–719. 3.2.1, 3.3, 3.4.1, 3.4.2, 3.6

Genovese, C. R., Lazar, N. A., and Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, 15(4):870–878. 2.3.4.2, 3.5.3

Gitelman, D. R. (2002). Ilab: a program for postexperimental eye movement analysis. *Behav Res Methods Instrum Comput*, 34(4):605–612. 2.3.5, 4.3.4

Glover, G. H. (1999). Deconvolution of impulse response in event-related BOLD fMRI. *Neuroimage*, 9(4):416–429. 3.2, 3.2.3, 3.5.1

Gnadt, J. W. and Andersen, R. A. (1988). Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res*, 70(1):216–220. 1.6.4

Goldberg, M. E., Bisley, J., Powell, K. D., Gottlieb, J., and Kusunoki, M. (2002). The role of the lateral intraparietal area of the monkey in the generation of saccades and visuospatial attention. *Ann N Y Acad Sci*, 956:205–215. 1.6

Goldberg, M. E. and Bushnell, M. C. (1981). Behavioral enhancement of visual responses in monkey cerebral cortex. ii. modulation in frontal eye fields specifically related to saccades. *J Neurophysiol*, 46(4):773–787. 1.5

Goldberg, M. E., Colby, C. L., and Duhamel, J. R. (1990). Representation of visuomotor space in the parietal lobe of the monkey. *Cold Spring Harb Symp Quant Biol*, 55:729–739. 1.4, 1.5, 2.2, 5.2, 5.4

Gottlieb, J. P., Kusunoki, M., and Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, 391(6666):481–484. 1.6.3

Grefkes, C., Weiss, P. H., Zilles, K., and Fink, G. R. (2002). Crossmodal processing of object features in human anterior intraparietal cortex: an fMRI study implies equivalencies between humans and monkeys. *Neuron*, 35(1):173–184. 5.2

Grüsser, O. J. (1994). Early concepts on efference copy and reafference. *Behavioral and Brain Sciences*, 17:262–264. 2

Grunewald, A., Linden, J. F., and Andersen, R. A. (1999). Responses to auditory stimuli in macaque lateral intraparietal area. i. effects of training. *J Neurophysiol*, 82(1):330–342. 2.5

Gössl, C., Fahrmeir, L., and Auer, D. P. (2001). Bayesian modeling of the hemodynamic response function in BOLD fMRI. *Neuroimage*, 14(1 Pt 1):140–148. 3.3

Haenny, P. E., Maunsell, J. H., and Schiller, P. H. (1988). State dependent activity in monkey visual cortex. ii. retinal and extraretinal factors in V4. *Exp Brain Res*, 69(2):245–259. 4.5.6

Hagler, D. J. and Sereno, M. I. (2003). Retinotopic maps in frontal cortex. *Neuroimage*, 19:S420. 5.2

Hallett, P. E. and Lightstone, A. D. (1976). Saccadic eye movements to flashed targets. *Vision Res*, 16(1):107–114. 1.2, 2.2, 5.4

Hansen, K. A., David, S. V., and Gallant, J. L. (2004). Parametric reverse correlation reveals spatial linearity of retinotopic human V1 BOLD response. *Neuroimage*, 23(1):233–241. 3.2, 3.2.3, 4.3.3

Hardy, S. G. and Lynch, J. C. (1992). The spatial distribution of pulvinar neurons that project to two subregions of the inferior parietal lobule in the macaque. *Cereb Cortex*, 2(3):217–230. 1.5

Harting, J. K., Huerta, M. F., Frankfurter, A. J., Strominger, N. L., and Royce, G. J. (1980). Ascending pathways from the monkey superior colliculus: an autoradiographic analysis. *J Comp Neurol*, 192(4):853–882. 1.5

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., and Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science*, 303(5664):1634–1640. 3.2.3

Haynes, J.-D., Deichmann, R., and Rees, G. (2005). Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. *Nature*. 5.2

Heeger, D. J., Huk, A. C., Geisler, W. S., and Albrecht, D. G. (2000). Spikes versus BOLD: what does neuroimaging tell us about neuronal activity? *Nat Neurosci*, 3(7):631–633. 4.5.4

Heeger, D. J. and Ress, D. (2002). What does fMRI tell us about neuronal activity? *Nat Rev Neurosci*, 3(2):142–151. 4.5.5

Heide, W., Binkofski, F., Seitz, R. J., Posse, S., Nitschke, M. F., Freund, H. J., and Kompf, D. (2001). Activation of frontoparietal cortices during memorized triple-step sequences of saccadic eye movements: an fMRI study. *Eur J Neurosci*, 13(6):1177–1189. 2.2, 2.5.1, 4.5.1

Heide, W., Blankenburg, M., Zimmermann, E., and Kompf, D. (1995). Cortical control of double-step saccades: implications for spatial orientation. *Ann Neurol*, 38(5):739–748. 1.3, 1.5, 1.5, 2.2, 4.5.1, 4.5.2

Heiser, L. M., Berman, R. A., Saunders, R. C., and Colby, C. L. (2005). Dynamic circuitry for updating spatial representations: Ii. physiological evidence for interhemispheric transfer in area lip of the split-brain macaque. *J Neurophysiol*. 5.2

Hikosaka, O., Miyauchi, S., Takeichi, H., and Shimojo, S. (1996). Multimodal spatial attention visualized by motion illusion. In Inui, T. and McClelland, J. L., editors, *Attention and Performance XVI: Information Integration in Perception and Communication*, pages 237–262. The MIT Press, Cambridge, MA. 2.5.1

Hinrichs, H., Scholz, M., Tempelmann, C., Woldorff, M. G., Dale, A. M., and Heinze, H. J. (2000). Deconvolution of event-related fMRI responses in fast-rate experimental designs: tracking amplitude variations. *J Cogn Neurosci*, 12 Suppl 2:76–89. 3.2.3

Honda, H. (1989). Perceptual localization of visual stimuli flashed during saccades. *Percept Psychophys*, 45(2):162–174. 1.2, 4.5.5

Honda, H. (1991). The time courses of visual mislocalization and of extraretinal eye position signals at the time of vertical saccades. *Vision Res*, 31(11):1915–1921. 4.5.5, 5.3

Honda, H. (1999). Modification of saccade-contingent visual mislocalization by the presence of a visual frame of reference. *Vision Res*, 39(1):51–57. 4.5.5

Hopfinger, J. B., Buonocore, M. H., and Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nat Neurosci*, 3(3):284–291. 1.6.3

Huettel, S. A. and McCarthy, G. (2001). The effects of single-trial averaging upon the spatial extent of fMRI activation. *Neuroreport*, 12(11):2411–2416. 4.4.6.3

Jonides, J. and Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. *Percept Psychophys*, 43(4):346–354. 1.6.3

Jouve, B., Rosenstiehl, P., and Imbert, M. (1998). A mathematical approach to the connectivity between the cortical visual areas of the macaque monkey. *Cereb Cortex*, 8(1):28–39. 1.5

Kalatsky, V. A. and Stryker, M. P. (2003). New paradigm for optical imaging: temporally encoded maps of intrinsic signal. *Neuron*, 38(4):529–545. 4.3.6

Kanwisher, N. and Wojciulik, E. (2000). Visual attention: insights from brain imaging. *Nat Rev Neurosci*, 1(2):91–100. 1.6.3

Kara, P. and Reid, R. C. (2003). Efficacy of retinal spikes in driving cortical responses. *J Neurosci*, 23(24):8547–8557. 3.2.1

Kastner, S., De Weerd, P., Desimone, R., and Ungerleider, L. G. (1998). Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science*, 282(5386):108–111. 4.5.5

Kastner, S., De Weerd, P., Pinsk, M. A., Elizondo, M. I., Desimone, R., and Ungerleider, L. G. (2001). Modulation of sensory suppression: implications for receptive field sizes in the human visual cortex. *J Neurophysiol*, 86(3):1398–1411. 4.4.3, 4.4.5.1

Kastner, S., Nothdurft, H. C., and Pigarev, I. N. (1999a). Neuronal responses to orientation and motion contrast in cat striate cortex. *Vis Neurosci*, 16(3):587–600. 1.6.3, 4.2

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., and Ungerleider, L. G. (1999b). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, 22(4):751–761. 4.5.5

Kellman, P., Gelderen, P. v., de Zwart, J. A., and Duyn, J. H. (2003). Method for functional MRI mapping of nonlinear response. *Neuroimage*, 19(1):190–199. 4.3.3

Khayat, P. S., Spekreijse, H., and Roelfsema, P. R. (2004). Correlates of transsaccadic integration in the primary visual cortex of the monkey. *Proc Natl Acad Sci U S A*, 101(34):12712–12717. 4.5.7

Kim, S.-G. (2003). Progress in understanding functional imaging signals. *Proc Natl Acad Sci U S A*, 100(7):3550–3552. 4.5.4

Kleiser, R. and Skrandies, W. (2000). Neural correlates of reafference: evoked brain activity during motion perception and saccadic eye movements. *Exp Brain Res*, 133(3):312–320. 1.6.5

Kowler, E., Anderson, E., Dosher, B., and Blaser, E. (1995). The role of attention in the programming of saccades. *Vision Res*, 35(13):1897–1916. 1.1, 2.5.1

Koyama, M., Hasegawa, I., Osada, T., Adachi, Y., Nakahara, K., and Miyashita, Y. (2004). Functional magnetic resonance imaging of macaque monkeys performing visually guided saccade tasks: comparison of cortical eye fields with humans. *Neuron*, 41(5):795–807. 1.6.2, 5.2

Krekelberg, B., Kubischik, M., Hoffmann, K.-P., and Bremmer, F. (2003). Neural correlates of visual localization and perisaccadic mislocalization. *Neuron*, 37(3):537–545. 4.5.5

Kusunoki, M. and Goldberg, M. E. (2003). The time course of perisaccadic receptive field shifts in the lateral intraparietal area of the monkey. *J Neurophysiol*, 89(3):1519–1527. 4.5.1, 5.3, 5.3

Lange, N. (2000). Statistical procedures for functional mri. In Moonen, C. T. W. and Bandettini, P. A., editors, *In Functional MRI*, pages 301–332. Springer, Berlin. 3.2

Lange, N. and Zeger, S. L. (1997). Nonlinear fourier time series analysis for human brain mapping by functional magnetic resonance imaging. *Applied Statistics*, 46:1–29. 3.2.4

Lappe, M., Awater, H., and Krekelberg, B. (2000). Postsaccadic visual references generate presaccadic compression of space. *Nature*, 403(6772):892–895. 1.2

Larsson, J. and Heeger, J. D. (2005). Human visual cortex between dorsal V3 and V5/MT+ contains two complete maps of the contralateral visual hemifield. *Soc Neurosci Abstr*, 508.15. 4.3.6, 5.2

Lee, S.-H., Blake, R., and Heeger, D. J. (2005). Traveling waves of activity in primary visual cortex during binocular rivalry. *Nat Neurosci*, 8(1):22–23. 19, 3.2.4, 4.5.5

Lewis, R. F. and Zee, D. S. (1993). Abnormal spatial localization with trigeminal-oculomotor synkinesis. evidence for a proprioceptive effect. *Brain*, 116 (Pt 5):1105–1118. 1.1

Li, C. S. and Andersen, R. A. (2001). Inactivation of macaque lateral intraparietal area delays initiation of the second saccade predominantly from contralesional eye positions in a double-saccade task. *Exp Brain Res*, 137(1):45–57. 1.3, 1.5

Linden, J. F., Grunewald, A., and Andersen, R. A. (1999). Responses to auditory stimuli in macaque lateral intraparietal area. ii. behavioral modulation. *J Neurophysiol*, 82(1):343–358. 2.5

Liu, T. T. (2004). Efficiency, power, and entropy in event-related fMRI with multiple trial types. part ii: design of experiments. *Neuroimage*, 21(1):401–413. 3.5.2, 4.3.3

Liu, T. T. and Frank, L. R. (2004). Efficiency, power, and entropy in event-related fMRI with multiple trial types. part i: theory. *Neuroimage*, 21(1):387–400. 3.5.2, 4.3.3

Liu, T. T., Frank, L. R., Wong, E. C., and Buxton, R. B. (2001). Detection power, estimation efficiency, and predictability in event-related fMRI. *Neuroimage*, 13(4):759–773. 2.3.2

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843):150–157. 4.5.4, 4.5.5

Logothetis, N. K. and Wandell, B. A. (2004). Interpreting the BOLD signal. *Annu Rev Physiol*, 66:735–769. 4.4.3, 4.5.4, 4.5.5

Luck, S. J., Chelazzi, L., Hillyard, S. A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J Neurophysiol*, 77(1):24–42. 4.5.5, 4.5.6

Luna, B., Thulborn, K. R., Strojwas, M. H., McCurtain, B. J., Berman, R. A., Genovese, C. R., and Sweeney, J. A. (1998). Dorsal cortical regions subserving visually guided saccades in humans: an fMRI study. *Cereb Cortex*, 8(1):40–47. 1.6.2, 5.2, 5.2

Lynch, J. C., Graybiel, A. M., and Lobeck, L. J. (1985). The differential projection of two cytoarchitectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. *J Comp Neurol*, 235(2):241–254. 1.5

Lynch, J. C., Hoover, J. E., and Strick, P. L. (1994). Input to the primate frontal eye field from the substantia nigra, superior colliculus, and dentate nucleus demonstrated by transneuronal transport. *Exp Brain Res*, 100(1):181–186. 1.5

Ma, L., Worsley, K., and Evans, A. (1999). Variability of spatial location of activation in fmri and pet cbf images. *NeuroImage*, 9:S178. 3.2.2

Mach, E. (1959/1914). *The analysis of sensations*. Open Court, Chicago.

Marois, R., Leung, H. C., and Gore, J. C. (2000). A stimulus-driven approach to object identity and location processing in the human brain. *Neuron*, 25(3):717–728. 1.6.3

Marrelec, G., Benali, H., Ciuciu, P., Pelegrini-Issac, M., and Poline, J.-B. (2003). Robust bayesian estimation of the hemodynamic response function in event-related BOLD fMRI using basic physiological information. *Hum Brain Mapp*, 19(1):1–17. 3.3

Marrelec, G., Ciuciu, P., Pelegrini-Issac, M., and Benali, H. (2004). Estimation of the hemodynamic response in event-related functional MRI: Bayesian networks as a framework for efficient bayesian modeling and inference. *IEEE Trans Med Imaging*, 23(8):959–967. 3.3

Matin, E., Clymer, A. B., and Matin, L. (1972). Metacontrast and saccadic suppression. *Science*, 178(57):179–182. 1.2

Mays, L. E. and Sparks, D. L. (1980). Saccades are spatially, not retinocentrically, coded. *Science*, 208(4448):1163–1165. 5.4

Mazzoni, P., Bracewell, R. M., Barash, S., and Andersen, R. A. (1996). Motor intention activity in the macaque's lateral intraparietal area. i. dissociation of motor plan from sensory memory. *J Neurophysiol*, 76(3):1439–1456. 2.5

McAdams, C. J. and Maunsell, J. H. (1999). Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron*, 23(4):765–773. 4.5.5, 4.5.6

McKeown, M. J., Jung, T. P., Makeig, S., Brown, G., Kindermann, S. S., Lee, T. W., and Sejnowski, T. J. (1998a). Spatially independent activity patterns in functional MRI data during the stroop color-naming task. *Proc Natl Acad Sci U S A*, 95(3):803–810. 3.2.1

McKeown, M. J., Makeig, S., Brown, G. G., Jung, T. P., Kindermann, S. S., Bell, A. J., and Sejnowski, T. J. (1998b). Analysis of fMRI data by blind separation into independent spatial components. *Hum Brain Mapp*, 6(3):160–188. 3.2.1

McMains, S. A. and Somers, D. C. (2004). Multiple spotlights of attentional selection in human visual cortex. *Neuron*, 42(4):677–686. 4.2

McNamee, R. L. and Eddy, W. F. (2001). Visual analysis of variance: a tool for quantitative assessment of fMRI data processing and analysis. *Magn Reson Med*, 46(6):1202–1208. 4.3.5.2

Medendorp, W. P., Goltz, H. C., Vilis, T., and Crawford, J. D. (2003). Gaze-centered updating of visual space in human parietal cortex. *J Neurosci*, 23(15):6209–6214. 4.5.1

Melamed, E. and Larsen, B. (1979). Cortical activation pattern during saccadic eye movements in humans: localization by focal cerebral blood flow increases. *Ann Neurol*, 5(1):79–88. 1.6.2

Merriam, E. P. and Colby, C. L. (2005). Active vision in parietal and extrastriate cortex. *Neuroscientist*, 11(5):484–493. 1.4

Merriam, E. P., Colby, C. L., Thulborn, K. R., Luna, B., Olson, C. R., and Sweeney, J. A. (2001). Stimulus-response incompatibility activates cortex proximate to three eye fields. *Neuroimage*, 13(5):794–800. 1.6.2

Merriam, E. P., Genovese, C. R., and Colby, C. L. (2003). Spatial updating in human parietal cortex. *Neuron*, 39(2):361–373. 3.4.2, 4.2, 4.5.1

Mohler, C. W., Goldberg, M. E., and Wurtz, R. H. (1973). Visual receptive fields of frontal eye field neurons. *Brain Res*, 61:385–389. 1.5

Moran, J. and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229(4715):782–784. 4.5.6

Morel, A. and Bullier, J. (1990). Anatomical segregation of two cortical visual pathways in the macaque monkey. *Vis Neurosci*, 4(6):555–578. 1.5

Morrone, M. C., Ross, J., and Burr, D. C. (1997). Apparent position of visual targets during real and simulated saccadic eye movements. *J Neurosci*, 17(20):7941–7953. 1.2

Murray, S. O. and Wojciulik, E. (2004). Attention increases neural selectivity in the human lateral occipital complex. *Nat Neurosci*, 7(1):70–74. 4.3.3

Nadell, D. and Heeger, D. (2003). Spatiotemporal summation in the bold response. *Soc Neurosci Abstr*, 69.5. 3.2.3

Nakamura, K. and Colby, C. L. (2000). Visual, saccade-related, and cognitive activation of single neurons in monkey extrastriate area V3a. *J Neurophysiol*, 84(2):677–692. 1.5, 4.4.5.2

Nakamura, K. and Colby, C. L. (2002). Updating of the visual representation in monkey striate and extrastriate cortex during saccades. *Proc Natl Acad Sci U S A*, 99(6):4026–4031. 1.5, 1.6.5, 2.5.2, 4.2, 4.4.7, 4.4.7.2, 4.5.1, 4.5.2, 57, 4.5.4, 4.5.5, 5.3, 5.3

Newsome, W. T., Britten, K. H., Salzman, C. D., and Movshon, J. A. (1990). Neuronal mechanisms of motion perception. *Cold Spring Harb Symp Quant Biol*, 55:697–705. 5.4

Nichols, T. E. and Holmes, A. P. (2002). Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp*, 15(1):1–25. 3.2

Nobre, A. C., Sebestyen, G. N., Gitelman, D. R., Mesulam, M. M., Frackowiak, R. S., and Frith, C. D. (1997). Functional localization of the system for visuospatial attention using positron emission tomography. *Brain*, 120 (Pt 3):515–533. 1.6.3

O'Connor, D. H., Fukui, M. M., Pinsk, M. A., and Kastner, S. (2002). Attention modulates responses in the human lateral geniculate nucleus. *Nat Neurosci*, 5(11):1203–1209. 5.2

O'Driscoll, G. A., Alpert, N. M., Matthysse, S. W., Levy, D. L., Rauch, S. L., and Holzman, P. S. (1995). Functional neuroanatomy of antisaccade eye movements investigated with positron emission tomography. *Proc Natl Acad Sci U S A*, 92(3):925–929. 1.6.2

Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis*, 10(4):437–442. 4.3.2

Penny, W., Kiebel, S., and Friston, K. (2003). Variational bayesian inference for fMRI time series. *Neuroimage*, 19(3):727–741. 3.3

Penny, W. D., Trujillo-Barreto, N. J., and Friston, K. J. (2005). Bayesian fMRI time series analysis with spatial priors. *Neuroimage*, 24(2):350–362. 3.3

Pessoa, L., Gutierrez, E., Bandettini, P., and Ungerleider, L. (2002a). Neural correlates of visual working memory: fMRI amplitude predicts task performance. *Neuron*, 35(5):975–987. 1.6.4

Pessoa, L., Kastner, S., and Ungerleider, L. G. (2002b). Attentional control of the processing of neural and emotional stimuli. *Brain Res Cogn Brain Res*, 15(1):31–45. 5.4

Pessoa, L., Kastner, S., and Ungerleider, L. G. (2003). Neuroimaging studies of attention: from modulation of sensory processing to top-down control. *J Neurosci*, 23(10):3990–3998. 1.6.3, 4.5.5

Petersen, S. E., Robinson, D. L., and Keys, W. (1985). Pulvinar nuclei of the behaving rhesus monkey: visual responses and their modulation. *J Neurophysiol*, 54(4):867–886. 1.5

Petersson, K. M., Nichols, T. E., Poline, J. B., and Holmes, A. P. (1999a). Statistical limitations in functional neuroimaging. i. non-inferential methods and statistical models. *Philos Trans R Soc Lond B Biol Sci*, 354(1387):1239–1260. 3.2

Petersson, K. M., Nichols, T. E., Poline, J. B., and Holmes, A. P. (1999b). Statistical limitations in functional neuroimaging. ii. signal detection and statistical inference. *Philos Trans R Soc Lond B Biol Sci*, 354(1387):1261–1281. 3.2

Petit, L. and Haxby, J. V. (1999). Functional anatomy of pursuit eye movements in humans as revealed by fMRI. *J Neurophysiol*, 82(1):463–471. 3.2.2, 5.2

Petit, L., Orssaud, C., Tzourio, N., Salamon, G., Mazoyer, B., and Berthoz, A. (1993). Pet study of voluntary saccadic eye movements in humans: basal ganglia-thalamocortical system and cingulate cortex involvement. *J Neurophysiol*, 69(4):1009–1017. 1.6.2

Petrides, M. and Pandya, D. N. (1984). Projections to the frontal cortex from the posterior parietal region in the rhesus monkey. *J Comp Neurol*, 228(1):105–116. 1.5, 1.5

Pigarev, I. N., Rizzolatti, G., and Scandolara, C. (1979). Neurons responding to visual stimuli in the frontal lobe of macaque monkeys. *Neurosci Lett*, 12(2-3):207–212. 1.5

Pouget, P., Emeric, E. E., Stuphorn, V., Reis, K., and Schall, J. D. (2005). Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. *J Neurophysiol*, 94(3):2086–2092. 1.5

Press, W. A., Brewer, A. A., Dougherty, R. F., Wade, A. R., and Wandell, B. A. (2001). Visual areas and spatial summation in human visual cortex. *Vision Res*, 41(10-11):1321–1332. 4.3.6, 5.2, 5.2

Price, C. J. and Friston, K. J. (1997). Cognitive conjunction: a new approach to brain activation experiments. *Neuroimage*, 5(4 Pt 1):261–270. 5.2

Purdon, P. L., Solo, V., Weisskoff, R. M., and Brown, E. N. (2001). Locally regularized spatiotemporal modeling and model comparison for functional MRI. *Neuroimage*, 14(4):912–923. 3.2.4

Quaia, C., Optican, L., and Goldberg, M. (1998). The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields. *Neural Netw*, 11(7-8):1229–1240. 1.5

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L. (2001). A default mode of brain function. *Proc Natl Acad Sci U S A*, 98(2):676–682. 3.2.1

Rees, G., Friston, K., and Koch, C. (2000). A direct quantitative relationship between the functional properties of human and macaque V5. *Nat Neurosci*, 3(7):716–723. 4.5.4

Reid, R. C., Victor, J. D., and Shapley, R. M. (1997). The use of m-sequences in the analysis of visual neurons: linear receptive field properties. *Vis Neurosci*, 14(6):1015–1027. 3.5.2, 4.3.3

Reppas, J. B., Usrey, W. M., and Reid, R. C. (2002). Saccadic eye movements modulate visual responses in the lateral geniculate nucleus. *Neuron*, 35(5):961–974. 4.5.7

Ress, D., Backus, B. T., and Heeger, D. J. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nat Neurosci*, 3(9):940–945. 4.2, 4.5.5, 5.4

Ress, D. and Heeger, D. J. (2003). Neuronal correlates of perception in early visual cortex. *Nat Neurosci*, 6(4):414–420. 4.5.4

Roitman, J. D. and Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci*, 22(21):9475–9489. 1.6

Rosano, C., Krisky, C. M., Welling, J. S., Eddy, W. F., Luna, B., Thulborn, K. R., and Sweeney, J. A. (2002). Pursuit and saccadic eye movement subregions in human frontal eye field: a high-resolution fMRI investigation. *Cereb Cortex*, 12(2):107–115. 3.2.2

Ross, J., Morrone, M. C., and Burr, D. C. (1997). Compression of visual space before saccades. *Nature*, 386(6625):598–601. 1.2, 4.5.5, 5.3

Ross, J., Morrone, M. C., Goldberg, M. E., and Burr, D. C. (2001). Changes in visual perception at the time of saccades. *Trends Neurosci*, 24(2):113–121. 1.1, 5.3

Saad, Z. S., Ropella, K. M., Cox, R. W., and DeYoe, E. A. (2001). Analysis and use of fMRI response delays. *Hum Brain Mapp*, 13(2):74–93. 2.3.2, 2.3.4.3

Schacter, D. L., Buckner, R. L., Koutstaal, W., Dale, A. M., and Rosen, B. R. (1997). Late onset of anterior prefrontal activity during true and false recognition: an event-related fMRI study. *Neuroimage*, 6(4):259–269. 3.2.3

Schall, J. D. (1991). Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: comparison with supplementary eye fields. *J Neurophysiol*, 66(2):559–579. 1.5

Schall, J. D., Morel, A., King, D. J., and Bullier, J. (1995). Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. *J Neurosci*, 15(6):4464–4487. 1.5, 1.5, 4.5.2

Schmolesky, M. T., Wang, Y., Hanes, D. P., Thompson, K. G., Leutgeb, S., Schall, J. D., and Leventhal, A. G. (1998). Signal timing across the macaque visual system. *J Neurophysiol*, 79(6):3272–3278. 1.5

Schneider, K. A. and Kastner, S. (2005). Visual responses of the human superior colliculus: a high-resolution functional magnetic resonance imaging study. *J Neurophysiol*, 94(4):2491–2503. 5.2

Schneider, K. A., Richter, M. C., and Kastner, S. (2004). Retinotopic organization and functional subdivisions of the human lateral geniculate nucleus: a high-resolution functional magnetic resonance imaging study. *J Neurosci*, 24(41):8975–8985. 5.2

Serences, J. T., Shomstein, S., Leber, A. B., Golay, X., Egeth, H. E., and Yantis, S. (2005). Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychol Sci*, 16(2):114–122. 1.6.3

Serences, J. T., Yantis, S., Culberson, A., and Awh, E. (2004). Preparatory activity in visual cortex indexes distractor suppression during covert spatial orienting. *J Neurophysiol*. 3.2.3, 3.5.1

Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., and Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, 268(5212):889–893. 1.6.1, 2.2, 2.3.2, 4.3.6, 4.4.3

Sereno, M. I., Pitzalis, S., and Martinez, A. (2001). Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science*, 294(5545):1350–1354. 1.6, 2.5, 2.5.1, 5.2

Sereno, M. I. and Tootell, R. B. H. (2005). From monkeys to humans: what do we now know about brain homologies? *Curr Opin Neurobiol*, 15(2):135–144. 1.6, 4.5.5

R Development Core Team (2005). R: *A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0. 2.3.4.3

Shadlen, M. N., Britten, K. H., Newsome, W. T., and Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci*, 16(4):1486–1510. 5.4

Sharma, J., Dragoi, V., Tenenbaum, J. B., Miller, E. K., and Sur, M. (2003). V1 neurons signal acquisition of an internal representation of stimulus location. *Science*, 300(5626):1758–1763. 4.5.7

Shulman, G. L., Ollinger, J. M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Petersen, S. E., and Corbetta, M. (1999). Areas involved in encoding and applying directional expectations to moving objects. *J Neurosci*, 19(21):9480–9496. 1.6.3

Silver, M. A., Ress, D., and Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *J Neurophysiol*, 94(2):1358–1371. 1.6.1, 1.6.3

Simon, O., Mangin, J. F., Cohen, L., Le Bihan, D., and Dehaene, S. (2002). Topographical layout of hand, eye, calculation, and language-related areas in the human parietal lobe. *Neuron*, 33(3):475–487. 5.2

Skrandies, W. and Anagnostou, E. (1999). Electroencephalographic cortical oscillations and saccadic eye movements in humans. *Neurosci Lett*, 261(1-2):57–60. 1.6.5

Skrandies, W. and Laschke, K. (1997). Topography of visually evoked brain activity during eye movements: lambda waves, saccadic suppression, and discrimination performance. *Int J Psychophysiol*, 27(1):15–27. 1.6.5

Smith, A. M., Lewis, B. K., Ruttimann, U. E., Ye, F. Q., Sinnwell, T. M., Yang, Y., Duyn, J. H., and Frank, J. A. (1999). Investigation of low frequency drift in fMRI signal. *Neuroimage*, 9(5):526–533. 3.2.3

Smith, A. T., Singh, K. D., Williams, A. L., and Greenlee, M. W. (2001). Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cereb Cortex*, 11(12):1182–1190. 4.4.3

Smith, M., Putz, B., Auer, D., and Fahrmeir, L. (2003). Assessing brain activity through spatial bayesian variable selection. *Neuroimage*, 20(2):802–815. 3.3

Snyder, L. H., Batista, A. P., and Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. *Nature*, 386(6621):167–170. 1.6

Sommer, M. A. and Wurtz, R. H. (2001). Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. *J Neurophysiol*, 85(4):1673–1685. 1.5

Sommer, M. A. and Wurtz, R. H. (2002). A pathway in primate brain for internal monitoring of movements. *Science*, 296(5572):1480–1482. 1.5

Sommer, M. A. and Wurtz, R. H. (2004). What the brain stem tells the frontal cortex. ii. role of the sc-md-fef pathway in corollary discharge. *J Neurophysiol*, 91(3):1403–1423. 1.5

Sommer, M. A. and Wurtz, R. H. (2005). Shifting visual receptive fields in the frontal eye field: Do they use corollary discharge from the thalamus? *Soc Neurosci Abstr*. 1.5

Spitzer, H., Desimone, R., and Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science*, 240(4850):338–340. 4.5.6

Stanton, G. B., Bruce, C. J., and Goldberg, M. E. (1995). Topography of projections to posterior cortical areas from the macaque frontal eye fields. *J Comp Neurol*, 353(2):291–305. 1.5, 1.5, 4.5.2

Sun, F. T., Miller, L. M., and D'Esposito, M. (2005). Measuring temporal dynamics of functional networks using phase spectrum of fMRI data. *Neuroimage*, 28(1):227–237. 3.2.1

Sutter, E. E. (2001). Imaging visual function with the multifocal m-sequence technique. *Vision Res*, 41(10-11):1241–1255. 3.5.2, 4.3.3

Sweeney, J. A., Mintun, M. A., Kwee, S., Wiseman, M. B., Brown, D. L., Rosenberg, D. R., and Carl, J. R. (1996). Positron emission tomography study of voluntary saccadic eye movements and spatial working memory. *J Neurophysiol*, 75(1):454–468. 1.6.2, 2.5.1, 5.2, 5.2

Sylvester, R., Haynes, J.-D., and Rees, G. (2005). Saccades differentially modulate human lgn and V1 responses in the presence and absence of visual stimulation. *Curr Biol*, 15(1):37–41. 4.2, 4.5.7

Sylvester, R. and Rees, G. (2005). Extraretinal saccadic signals in human lgn and early retinotopic cortex. *Neuroimage*. 4.5.7

Sàenz, M., Buracas, G. T., and Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nat Neurosci*, 5(7):631–632. 4.4.3, 4.5.5

Thompson, K. G., Hanes, D. P., Bichot, N. P., and Schall, J. D. (1996). Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *J Neurophysiol*, 76(6):4040–4055. 1.5

Tolias, A. S., Moore, T., Smirnakis, S. M., Tehovnik, E. J., Siapas, A. G., and Schiller, P. H. (2001). Eye movements modulate visual receptive fields of V4 neurons. *Neuron*, 29(3):757–767. 4.4.5.2, 4.5.6, 4.5.7

Tootell, R. B. and Hadjikhani, N. (2001). Where is 'dorsal V4' in human visual cortex? retinotopic, topographic and functional evidence. *Cereb Cortex*, 11(4):298–311. 4.3.6, 5.2

Tootell, R. B., Hadjikhani, N. K., Vanduffel, W., Liu, A. K., Mendola, J. D., Sereno, M. I., and Dale, A. M. (1998). Functional analysis of primary visual cortex (V1) in humans. *Proc Natl Acad Sci U S A*, 95(3):811–817. 2.5, 4.2

Tootell, R. B., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., Sereno, M. I., and Dale, A. M. (1997a). Functional analysis of V3a and related areas in human visual cortex. *J Neurosci*, 17(18):7060–7078. 4.3.5.3, 5.2

Tootell, R. B., Mendola, J. D., Hadjikhani, N. K., Ledden, P. J., Liu, A. K., Reppas, J. B., Sereno, M. I., and Dale, A. M. (1997b). Functional analysis of V3a and related areas in human visual cortex. *J Neurosci*, 17(18):7060–7078. 4.3.6

Umeno, M. M. and Goldberg, M. E. (1997). Spatial processing in the monkey frontal eye field. i. predictive visual responses. *J Neurophysiol*, 78(3):1373–1383. 1.5, 1.6.5, 2.5.2, 4.2, 5.2

Umeno, M. M. and Goldberg, M. E. (2001). Spatial processing in the monkey frontal eye field. ii. memory responses. *J Neurophysiol*, 86(5):2344–2352. 1.5, 4.2, 5.2

Van Essen, D. C., Drury, H. A., Dickson, J., Harwell, J., Hanlon, D., and Anderson, C. H. (2001). An integrated software suite for surface-based analyses of cerebral cortex. *J Am Med Inform Assoc*, 8(5):443–459. 5.2

Vandenberghe, R., Duncan, J., Dupont, P., Ward, R., Poline, J. B., Bormans, G., Michiels, J., Mortelmans, L., and Orban, G. A. (1997). Attention to one or two features in left or right visual field: a positron emission tomography study. *J Neurosci*, 17(10):3739–3750. 1.6.3

Vazquez, A. L. and Noll, D. C. (1998). Nonlinear aspects of the BOLD response in functional MRI. *Neuroimage*, 7(2):108–118. 3.2.3, 4.4.6.3

Wager, T. D., Vazquez, A., Hernandez, L., and Noll, D. C. (2005). Accounting for nonlinear BOLD effects in fMRI: parameter estimates and a model for prediction in rapid event-related studies. *Neuroimage*, 25(1):206–218. 3.2.3, 4.4.6.3, 4.4.6.3, 4.5.3

Walker, M. F., Fitzgibbon, E. J., and Goldberg, M. E. (1995). Neurons in the monkey superior colliculus predict the visual result of impending saccadic eye movements. *J Neurophysiol*, 73(5):1988–2003. 1.5, 1.6.5, 4.2, 5.2

Wandell, B. A., Chial, S., and Backus, B. T. (2000). Visualization and measurement of the cortical surface. *J Cogn Neurosci*, 12(5):739–752. 4.3.5.1

Ward, B. D. (1998). *Deconvolution analysis of fMRI time series data, Documentation for the* AFNI *software package*. National Institutes of Health, Bethesda, MD. 2.3.4.2, 3.2.3, 3.5.1

Wardak, C., Olivier, E., and Duhamel, J.-R. (2004). A deficit in covert attention after parietal cortex inactivation in the monkey. *Neuron*, 42(3):501–508. 1.6.3

Wexler, M. (2005). Anticipating the three-dimensional consequences of eye movements. *Proc Natl Acad Sci U S A*, 102(4):1246–1251. 1

Wojciulik, E. and Kanwisher, N. (1999). The generality of parietal involvement in visual attention. *Neuron*, 23(4):747–764. 5.2

Woolrich, M. W., Behrens, T. E. J., Beckmann, C. F., Jenkinson, M., and Smith, S. M. (2004a). Multilevel linear modelling for fMRI group analysis using bayesian inference. *Neuroimage*, 21(4):1732–1747. 3.3

Woolrich, M. W., Jenkinson, M., Brady, J. M., and Smith, S. M. (2004b). Fully bayesian spatio-temporal modeling of fMRI data. *IEEE Trans Med Imaging*, 23(2):213–231. 3.3

Worsley, K., Marrett, S., Neelin, P., Vandal, A., Friston, K., , and Evans, A. (1996). A unified statistical approach for determining significant signals in images of cerebral activation. *Human Brain Mapping*, 4:58–73. 3.2.2

Worsley, K. J. and Friston, K. J. (1995). Analysis of fMRI time-series revisited–again. *Neuroimage*, 2(3):173–181. 3.2

Wunderlich, K., Schneider, K. A., and Kastner, S. (2005). Neural correlates of binocular rivalry in the human lateral geniculate nucleus. *Nat Neurosci*. 5.2

Wurtz, R. H. and Albano, J. E. (1980). Visual-motor function of the primate superior colliculus. *Annu Rev Neurosci*, 3:189–226. 1.5

Zeki, S., Watson, J. D., Lueck, C. J., Friston, K. J., Kennard, C., and Frackowiak, R. S. (1991). A direct demonstration of functional specialization in human visual cortex. *J Neurosci*, 11(3):641–649. 1.6.1