Title
The role of attention in cortical and subcortical spatial mapping

Author
Fischer, Jason Thomas

Publication Date
2012

Peer reviewed|Thesis/dissertation

The role of attention in cortical and subcortical spatial mapping

By

Jason Thomas Fischer

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Psychology

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Whitney, Chair
Professor Jack Gallant
Professor Dennis Levi
Professor Michael Silver

Fall 2012

The role of attention in cortical and subcortical spatial mapping

by

Jason Thomas Fischer

Abstract

The role of attention in cortical and subcortical spatial mapping

by

Jason Thomas Fischer

Doctor of Philosophy in Psychology

University of California, Berkeley

Professor David Whitney, Chair

Perceiving objects' positions is one of the fundamental purposes of vision and is crucial to day-to-day life. Much work has been devoted to finding and characterizing spatial maps in the brain, but we still know strikingly little about how these maps support perceptual localization. The experiments in this thesis employ a combination of functional brain imaging and visual psychophysics, drawing on a body of behavioral literature on spatial perception and visual attention, to ask how position coding in the brain adapts to support the demands of the task at hand.

The experiments in Chapter 2 ask how and where *perceived* object positions are represented in the brain. Perceived position depends on many factors such as attention, frame of reference, adaptation, and motion. At what stage of visual processing is this information integrated into the brain's retinotopic maps? By measuring variability in perceived position alongside variability in the multivariate pattern of neural responses in a host of visual areas, we identify a percept-centered reference frame in high-level object-, face-, scene-, and motion-selective regions.

The experiments in Chapter 3 ask how the visual system achieves improved spatial resolution with focused attention. Attention is known to boost the amplitude of neural responses, but it might also sharpen position tuning at the neural population level, making the activations produced by adjacent objects more distinct within the cortex. Employing a combined imaging and modeling approach, we find that attention narrows the spatial spread of the fMRI BOLD response in early visual cortical areas, including V1. This narrowed population position tuning is an efficient means of achieving visual resolution improvements.

The experiments in Chapter 4 ask how the brain suppresses distracting visual input in order to limit its negative impact on performance. During a selective attention task, we measure the BOLD response in the pulvinar nucleus of the thalamus, an area whose damage appears to lead to distractor filtering deficits. We find that while attended items are represented with high spatial and featural precision in the pulvinar, ignored items are gated out, leaving no discernable trace in the pattern of response in the

pulvinar. These results suggest that the spatial maps in the pulvinar may serve as an interface in which distracting visual input is filtered out.

Collectively, the findings presented here speak to a remarkable flexibility in the way the brain represents spatial information. The brain's visual maps adapt on a moment-by-moment basis to support the demands of the task at hand.

Table of contents

page

Acknowledgments **iii**

Chapter 1: Introduction **1**

1.1 The representation of visual space in the brain 1
1.1.1 Position coding in high-level visual areas 2
1.1.2 Non-retinotopic coordinate frames 3
1.1.3 Perceptual localization – and mislocalization 4
1.2 Attentional modulation of the position code 5
1.3 Filtering out distracting visual input 6
1.4 Overarching goals 6

Chapter 2: The emergence of perceived position in the visual system **8**

2.1 Introduction 8
2.2 Methods 9
2.2.1 Functional localizers: stimuli and analysis 9
2.2.2 Main experiment stimuli 11
2.2.3 Experimental Design & Task 14
2.2.4 FMRI data acquisition and preprocessing 14
2.2.5 Data Analysis 15
2.2.6 Eye tracking data collection & analysis 19
2.3 Results 19
2.4 Discussion 32
2.5 Supplemental Materials 36

Chapter 3: Attention modulates position tuning of population responses in V1 **48**

3.1 Introduction 48
3.2 Results 50
3.3 Discussion 59
3.4 Supplemental Experimental Procedures 60
3.4.1 Stimuli 60
3.4.2 Attention task 60
3.4.3 fMRI data collection and analysis 61
3.4.4 ROI definition 62
3.4.5 Correlation analysis 63
3.4.6 Slice-based analysis 64

3.4.7 Modeling 64

Chapter 4: Attention gates visual coding in the pulvinar nucleus of the thalamus **76**

4.1 Introduction 76
4.2 Results 77
4.2.1 Experiment 1 – Position discrimination 77
4.2.2 Experiment 2 – Orientation discrimination 83
4.3 Discussion 86
4.4 Methods 88
4.4.1 Subjects 88
4.4.2 Experimental design 88
4.4.3 fMRI acquisition 89
4.4.4 Data analysis 89
4.4.5 ROI definition 90
4.5 Supplementary Methods 91
4.5.1 Experimental design & analysis 91
4.5.2 Position discrimination analysis 91
4.5.3 Orientation classification analysis 92

Chapter 5: Concluding Discussion **98**

5.1 Implications of flexible spatial mapping in the brain 98
5.2 What is the mechanism behind flexible spatial coding in the brain? 99
5.3 Population-level vs. single-unit neural measurements 99
5.4 Bridging the gap between fMRI and single-unit findings 100
5.5 Conclusions 102

References **103**

Acknowledgments:

I first want to thank David Whitney, my graduate mentor, for his immense investment in me both personally and professionally. He taught me how to conduct science, and led by example, approaching research with insatiable curiosity and humility. I relished the endless creativity in David's lab, which I will strive to carry forward in my research career. Dave, I am truly grateful for your guidance, confidence in my ability, passion for science, and friendship.

Thank you to my wife Laura, who I love more than anything. You are my companion through all of life's challenges and triumphs, and there is no way I would have finished my PhD without your support. No matter where we end up, we're always in this together. To my mom and dad, Wendy and Tom Fischer, thank you so much for your love and support. You've put enormous energy into ensuring that I'd have the opportunity to follow my dreams in life. The inspiring time that we spend together means more to me than I can possibly say. Thank you to my brother Mike – you are brilliant and I like to think that each of our creative pursuits will always reflect some of the other's influence. I'm looking forward to the time when we can set up shop next door to each other. I am always thinking of you all.

Thank you to Jack Gallant, Dennis Levi, and Michael Silver for generously investing their time on my dissertation committee. Their constructive and honest advice has helped shape how I conduct and think about research. I thank Petr Janata and Emilio Ferrer, who served on my qualifying committee at UC Davis and gave valuable feedback on my projects as I was getting started. I also thank Ron Mangun for his support during my time at UC Davis, and Bruno Olshausen for our inspiring conversations. I am grateful to Kevin Iga, an outstanding mentor and teacher during my undergraduate study at Pepperdine University who pushed me to understand mathematical concepts at their core.

To Bill Prinzmetal, thank you so much for our adventures in sailing and scientific pondering; I hope we'll share many more good times sailing on the bay. Thank you to Erica Whitney for saving the day on numerous occasions with invaluable edits to my manuscripts. To Matt Gryniewski-Peirson, thank you for sparking my interest in the human mind; here's to many more zany thought experiments. Thank you to David Bressler for helping me get started in the Whitney Lab, and providing the data that formed the basis for my first projects. Thanks to Jenny Shankey, who put in many long hours assisting with psychophysical data collection. I also want to thank the colleagues and friends with whom I formed some of my fondest memories from the past six years, and from whom I've learned a lot: Jesse Bengson, Summer Sheremata, David Dermidjian, Mayur Mudigonda, Britta Hummel, Allie Yamanashi Leib, and Stephen Keimel.

Most of all, I want to thank my labmates in the Whitney Lab, who are like family: Paul Bulakowski, Wes Chaney, Jennifer Corbett, Faraz Farzin, Jason Haberman, Steve Haroz, Tom Harp, David Horton, Kami Koldewyn, Anna Kosovicheva, Alina Liberman, Liz Louie, Gerrit Maus, Brent Parsons, Mita Puri, Niki Spotswood, Tim Sweeny, Santani Teng, Ben Wolfe, and Nikki Wurnitsch. I cherish the time we've spent together and what I've learned from you. I look forward to continuing to outdo our prior adventures at every opportunity. ST, it has been an honor and a pleasure to start graduate school together and progress on this quest side by side. I maintain that buttercream is objectively the best frosting, but I'm open to a dialogue about the merits of other varieties.

Chapter 1

Introduction

Our visual experience is fundamentally spatial, and accurate localization of the objects in the visual environment is among the most crucial products of the brain's myriad computations. To interact with the world, we must determine the locations of the objects around us, orient ourselves within the spatial environment, and execute motor actions with high spatial precision. In many cases, even a tiny localization error can mean the difference between success and failure in the task at hand. Moreover, maintaining precise neural representations of visual space is important not just in service of object localization per se, but also as a precursor to other core functions of the visual system. Binding elementary visual features into meaningful objects relies on spatial proximity (Treisman and Gelade 1980; Ashby, Prinzmetal et al. 1996), and crowding, a breakdown of object recognition in the face of visual clutter, may result from a loss of precise position information (Wolford 1975; Popple and Levi 2005; Levi 2008; Greenwood, Bex et al. 2009; Dakin, Cass et al. 2010). The experiments in this thesis focus on what kind of information is stored in the spatial maps in the brain, and whether and how topographic mapping in the brain can adapt to suit a person's present goals.

1.1 The representation of visual space in the brain

From the moment that light reaches the eye, the spatial organization in the retinal image is maintained and augmented over the course of visual processing, first by the optics of the eye, and then by successive topographic maps in the brain that persist through the multitude of transformations that visual information undergoes. Spatial maps are the medium in which visual computations take place. Retinotopic maps, which preserve the coordinate system of the retinae within their topographic layout, are found in many brain areas. Sixteen or more distinct retinotopic maps have been identified in the human cortex (Wandell, Dumoulin et al. 2007), in addition to subcortical retinotopic maps (Kastner, O'Connor et al. 2004; Schneider, Richter et al. 2004; Fischer and Whitney 2009; Schneider 2011), and maps of motor actions such as eye movements that reside in retinotopic coordinates (DuBois and Cohen 2000; Hagler Jr and Sereno 2006; Hagler, Riecke et al. 2007). So fundamental to the brain's organization is retinotopy, that it likely shapes the physical development of the occipital cortex, influencing the folding pattern of the cortical sheet within the skull (Rajimehr and Tootell 2009). The reproduction of retinal space across many brain areas means that there is redundancy in position mapping in the brain: even as visual information is transformed and abstracted, it maintains at least a coarse spatial order. As described next, this holds true even within higher-level ventral visual areas that were classically thought to possess little or no position information.

1.1.1 Position coding in high-level visual areas

One computational challenge for the visual system is to achieve invariant object recognition – the ability to assign a constant identity to a given object under variable viewing conditions (DiCarlo and Cox 2007). The same object may produce an infinite number of retinal images, depending on its position and orientation relative to the viewer, lighting conditions, visual noise, and many other factors. A perfectly invariant object detector would signal an object in the same way regardless of changes in, for example, orientation, translation, and size that might result from movement of the object or the observer. Much work in theoretical neuroscience has been invested into formulating biologically plausible systems that would achieve invariant object recognition in the brain (Olshausen, Anderson et al. 1993; Wallis and Rolls 1997; Ullman and Soloviev 1999; Yuille and Kersten 2006; Serre, Wolf et al. 2007), and many models would predict that high-level brain areas that code for objects should show little or no position selectivity. The same population of cells within the face-selective Fusiform Face Area (Kanwisher, McDermott et al. 1997), for example, would be active for a given face stimulus regardless of where that face appeared in the visual field. Hand in hand with this is the idea, first put forward by Ungerleider and Mishkin (Ungerleider and Mishkin 1982; Mishkin, Ungerleider et al. 1983), that form and location information are segregated in the brain into separate processing streams. This hypothesis and the related idea of separate visual streams for perception and action have found some support in lesion and neurophysiological studies (Goodale and Milner 1992).

More recent recordings from both monkeys (Op De Beeck and Vogels 2000) (DiCarlo and Maunsell 2003) and humans (Levy, Hasson et al. 2001; MacEvoy and Epstein 2007; Schwarzlose, Swisher et al. 2008) have overturned the notion that ventral object areas lack position information. In humans, position-selective responses have been found in the Fusiform Face Area (FFA(Levy, Hasson et al. 2001; Hemond, Kanwisher et al. 2007; Schwarzlose, Swisher et al. 2008; Kravitz, Kriegeskorte et al. 2010), Occipital Face Area (OFA(Hemond, Kanwisher et al. 2007; Schwarzlose, Swisher et al. 2008), Parahippocampal Place Area (PPA; (Levy, Hasson et al. 2001; MacEvoy and Epstein 2007; Schwarzlose, Swisher et al. 2008), and object-selective lateral occipital gyrus (LO; (Levy, Hasson et al. 2001; Larsson and Heeger 2006; Hemond, Kanwisher et al. 2007; McKyton and Zohary 2007; Schwarzlose, Swisher et al. 2008; Kravitz, Kriegeskorte et al. 2010), among others. In some cases the position information in high-level areas is apparently coarser than in early areas (for example, the topography in face- and scene-selective regions may be organized along a central-peripheral dimension rather than containing full retinotopic maps (Levy, Hasson et al. 2001)), but there is some form of reliable position information within these areas nonetheless.

Recent psychophysical evidence also argues against absolute translational invariance in object perception. Afraz et al. (Afraz, Pashkam et al. 2010) found that there is substantial heterogeneity in perceived facial gender and age for faces presented and different locations in the visual field. An identical face may look different when viewed at different locations, and physically different faces may appear identical if

they fall in regions of the visual field with different biases. The pattern of biases that Afraz et al. found differed across subjects, but was stable across multiple testing sessions within a subject. Thus, while we may experience the illusion of perfectly position-invariant object perception in day-to-day life, this is not the case when our object perception is subjected to close scrutiny.

1.1.2 Non-retinotopic coordinate frames

An ongoing debate centers on what kinds of transformations in spatial mapping occur in the visual system and specifically, whether spatial maps in non-retinotopic coordinates are explicitly computed and maintained in the brain. Given that we interact with objects in a three-dimensional environment, in which an object's physical location can remain constant even as its retinotopic location changes (or vice versa) due to head, body, and eye movements, an intuitively appealing idea is that the brain constructs a reference frame that is grounded in the coordinates of the physical environment that we interact with (allocentric space), rather than the coordinates of retinal input. Such a coordinate system must rely not only on visual input from the retina, but also on feedback from the oculomotor muscles, vestibular system, and other indicators of head and body orientation in order to update the positions of objects relative to the observer. Despite the theoretical appeal of an explicit spatial framework that is always in flux, continuously transformed by our every head, body, and eye movement, maintaining such a map would be very computationally demanding and perhaps unnecessary (Cavanagh, Hunt et al. 2010). Thus, there is ongoing debate surrounding the question of the existence of explicit allocentric maps in the brain.

There are mixed results as to whether we behave in a manner consistent with the maintenance of high fidelity allocentric maps in the brain. A number of studies have shown rapid and precise integration of visual information across saccades in world-centered coordinates, including integration of low-level pictorial information (Jonides, Irwin et al. 1982), form information (Hayhoe, Lachter et al. 1991; Melcher 2005), and motion information (Melcher and Morrone 2003; Ezzati, Golzar et al. 2008). The precision and speed of integration within allocentric coordinates would seem to indicate that world-centered object locations aren't simply being computed on the fly when needed, but rather must be maintained explicitly somewhere in the brain, ready for use. Along the same lines, saccadic adaptation reportedly occurs in world-centered, rather than retina-centered coordinates (Zimmermann, Burr et al. 2011). At the same time, other studies using similar methodology have failed to replicate the results on allocentric visual integration (O'Regan and Levy-Schoen 1983; Pashkam, Shim et al. 2007; Afraz and Cavanagh 2009). Additionally, recent evidence has shown that spatial attention (Golomb, Chun et al. 2008; Golomb, Nguyen-Phuc et al. 2010; Golomb, Pulido et al. 2010) and spatial memory (Golomb and Kanwisher 2012) is more precise when probed in retinotopic than spatiotopic (allocentric) coordinates.

Brain imaging results are similarly mixed. Recording from human motion-selective area MT+, d'Avossa et al. (d'Avossa, Tosetti et al. 2007) reported that the BOLD response in MT+ is modulated by gaze position such that responses to visual

motion occur in spatiotopic, not retinotopic coordinates. Similar results were also reported for object-selective area LO (McKyton and Zohary 2007). Other studies, looking more broadly across the visual cortex, have failed to find evidence of explicit allocentric coding in any visual area, including LO and MT+ (Gardner, Merriam et al. 2008; Golomb and Kanwisher 2010; Golomb and Kanwisher 2011). Given the similar measures used to evaluate spatiotopy vs. retinotopy among these studies, there is no obvious explanation for the contradicting results. The source of the discrepancy may lie in the degree to which the data was filtered for quality (Gardner, Merriam et al. 2008) or the way in which attention is controlled (Crespi, Biagi et al. 2011). The case for explicit spatiotopic representations is somewhat more consistent within multisensory regions in the parietal cortex, where single-unit studies in monkeys (Andersen, Essick et al. 1985; Galletti, Battaglini et al. 1993; Galletti, Battaglini et al. 1995; Andersen, Snyder et al. 1997; Duhamel, Bremmer et al. 1997) and human imaging studies (Bremmer, Schlack et al. 2001; Sereno and Huang 2006) have found agreement.

While this debate is ongoing, it misses a key point about human perception: no coordinate system based on physical geometry can account for the nature of localization errors that human observers make. As outlined below, moment-to-moment fluctuations in perceived position depend on a host of factors that are not accounted for by either retinotopic or allocentric coordinates.

1.1.3 Perceptual localization – and mislocalization

Despite the fundamental importance of position perception to everyday life, we are strikingly prone to localization errors. A host of factors such as visual motion (Ramachandran and Anstis 1990; De Valois and De Valois 1991; Whitney and Cavanagh 2000; Nijhawan 2002; Whitney 2002), adaptation (Whitaker, McGraw et al. 1995; Whitaker, McGraw et al. 1997; Snowden 1998), frame of reference (Roelofs 1935; Bridgeman, Peery et al. 1997), eye movements (Cai, Pouget et al. 1997; Ross, Morrone et al. 1997; Ross, Morrone et al. 2001), object size (McGraw, Roach et al. 2012), and many others can lead to a discrepancy between the perceived and physical (retinal) location of an object.

Among the most pronounced and well-known perceptual mislocalizations result from visual motion. In the flash lag illusion (MacKay 1961; Nijhawan 2002), a moving object appears to be shifted ahead along its motion trajectory, relative to an object briefly flashed at the same location as the moving object (i.e., the flash “lags behind” the moving object). Even remote motion can bias position perception: when a brief flash is presented along with nearby (but not overlapping) motion, the flash appears to shifted in the direction of the accompanying motion (Whitney and Cavanagh 2000). Why do we experience these distortions in perceived position? Are they the outcome of computational errors or bottlenecks that we’re simply not able to overcome? On the contrary, there is evidence that such mislocalizations can actually be adaptive for interacting with the physical environment (Whitney 2002; Whitney, Murakami et al. 2010), for example by updating the perceived position of a moving object to match its

true physical location, thereby compensating for the change in position that has occurred during neural processing delays (Nijhawan 1994).

Given that visual motion modulates perceived position, one might expect to find matching motion-induced distortions in retinotopic maps within the brain. Indeed, visual motion both within an object (Whitney, Goltz et al. 2003; Whitney and Bressler 2007) and remote from an object (Maus, Fischer et al. 2009) has been shown to alter the retinotopic mapping of position both in early visual cortex and in motion-selective MT+. Similar findings have been reported at the level of single units, with receptive field locations shifting in the presence of visual motion (Berry, Brivanlou et al. 1999; Fu, Shen et al. 2004; Sundberg, Fallah et al. 2006; Schwartz, Taylor et al. 2007).

Motion is just one of the many above-mentioned factors that can dramatically influence the perceived locations of objects. The experiments in Chapter 2 of this thesis more directly probe the encoding of perceived position in the visual system by asking subjects to perform a challenging localization task during fMRI scanning. Rather than explicitly manipulating perceived position, we capitalize on subjects' perceptual errors in order to ask which visual areas display flexible spatial coding that matches perceived position on a moment-by-moment basis.

1.2 Attentional modulation of the position code

Spatial attention can also modulate perceived position. Suzuki and Cavanagh (Suzuki and Cavanagh 1997) first reported the attentional repulsion effect: when attention is directed to a location (either voluntarily or involuntarily), the perceived locations of nearby objects are repelled away from the attended location. The attentional repulsion effect persists in the presence of visual landmarks, as demonstrated by the fact that it can cause an apparent reversal in the direction of misalignment of two slightly offset bars (vernier discrimination task; (Suzuki and Cavanagh 1997)). Spatially-directed attention can also modulate the perceived eccentricity of a peripheral target (Fortenbaugh and Robertson 2011), and even modulate perceived position in an object-specific fashion, compressing the apparent distances between targets while expanding the perceived distances between distractors in dynamic, multi-object displays (Liverence and Scholl 2011). Thus, perceived spatial relations are intimately tied to the locus and distribution of attention.

Another way in which attention can modulate spatial perception is by optimizing visual resolution for the task at hand. Exogenously cued (involuntary) spatial attention enhances visual acuity at the cued location (Eriksen and Collins 1969; Eriksen and Rohrbaugh 1970; Nakayama and Mackeben 1989). Interestingly, with involuntary attention, this boost in visual acuity can actually come at the cost of performance, when the task at hand calls for more, rather than less integration across space (Yeshurun and Carrasco 1998). However, when observers voluntarily direct attention to a location, the effects on resolution are more flexible, either increasing or decreasing resolution to match the scale of the most relevant information (Yeshurun, Montagna et al. 2008; Carrasco and Yeshurun 2009). By what mechanism does visual attention modulate

spatial resolution? Known effects of spatial attention, including signal gain (Buracas and Boynton 2007) and increased reliability of responses (Bressler and Silver 2010) are in principle sufficient to achieve sharpened resolution – the improved signal-to-noise ratio they afford makes adjacent points more reliably distinguishable from a signal detection standpoint. However, a more efficient way to improve visual resolution would be to additionally narrow the distribution of the neural response in the cortex corresponding to a given object or location (essentially, a narrowed point-spread function in the neural population response). The efficiency of such a mechanism comes from the fact that adjacent stimuli produce less overlapping distributions of response, making it easier to distinguish adjacent objects from a signal detection standpoint. Whether such a mechanism operates in the human visual cortex was unknown prior to the work presented here, and this is the question addressed by the experiments in Chapter 3.

1.3 Filtering out distracting visual input

The brain contains topographic maps of attended locations in the parietal cortex (Silver, Ress et al. 2005; Swisher, Halko et al. 2007), and these maps may be a source of attention signals that feed back to visual areas (Lauritzen, D'Esposito et al. 2009; Silver and Kastner 2009). What about *ignored* locations? At a given moment, only a fraction of the visual input to the brain is crucial for the task at hand, while other input can be distracting. Alongside highlighting important visual input, how does the attentional system suppress distracting input? A clue to answering this question comes from patients with focal stroke-related damage to the pulvinar nucleus of the thalamus. These patients exhibit difficulty with visual tasks in the presence of nearby distracting items, while their performance is relatively unaffected in the absence of distraction (Ward, Danziger et al. 2002; Snow, Allen et al. 2009). In accord with this, when Desimone et al. (Desimone, Wessinger et al. 1990) induced reversible lesions to the pulvinar in monkeys, they found similar deficits – visual discrimination was selectively impaired in the presence of salient distractors. These studies suggest that the pulvinar is important for distractor suppression and that studying the topographic maps in the pulvinar may be key to understanding how the brain filters out distracting visual input. The experiments in Chapter 4 characterize visual encoding in the healthy human pulvinar during challenging selective attention tasks. Using multivariate pattern analysis to measure position and orientation information in the fMRI BOLD response, we characterize the precision with which attended targets and ignored distractors are encoded in the pulvinar when the two are simultaneously present and compete for attentional resources.

1.4 Overarching goals

Collectively, the research I present here aims to probe how flexible and adaptive spatial coding in the brain can be. If spatial maps are the medium in which visual computations take place in the brain, how much can the maps themselves change in service of the current task demands? We test for moment-by-moment fluctuations in spatial mapping that match perceived position, sharpened position tuning in primary visual cortex in service of improved visual resolution at attended locations, and gating of

distracting information at ignored locations within the visual maps of the pulvinar. To anticipate, in all three cases we found a remarkable flexibility in how the brain encodes spatial information.

Chapter 2

The emergence of perceived position in the visual system

2.1 Introduction

While retinotopy has been extensively studied throughout the visual cortex, and is considered one of the fundamental principles by which visual areas are organized, there are many circumstances in which the perceived position of an object differs markedly from its retinal position. Object motion (Ramachandran and Anstis 1990; De Valois and De Valois 1991; Whitney 2002), eye movements (Cai, Pouget et al. 1997; Ross, Morrone et al. 1997; Ross, Morrone et al. 2001), attention shifts (Suzuki and Cavanagh 1997; Kerzel 2003), changes in frame of reference (Roelofs 1935; Bridgeman, Peery et al. 1997), and adaptation (Whitaker, McGraw et al. 1997) are among the many factors that can lead to disparate physical and perceived position information. These examples speak to the fact that reading off an object's retinal coordinates is only a single step in the complex task of object localization; perceived and physical object position are often dissociated. Given this, we might well expect that some visual areas encode position in percept-based, rather than retina-based coordinates.

Where might percept-centered position coding exist in the visual system? Retinotopy is well-characterized in striate and early extrastriate visual cortex, but the nature of spatial coding in higher-level object-, scene-, and motion-processing areas is still unclear. Activity in the fusiform face area (FFA) and parahippocampal place area (PPA) exhibits relatively weak position selectivity (Hemond, Kanwisher et al. 2007; MacEvoy and Epstein 2007; Schwarzlose, Swisher et al. 2008), but these areas show biases in response amplitude for centrally versus peripherally presented stimuli (Levy, Hasson et al. 2001). Retinotopic maps have been found in human lateral occipital (LO) cortex (Larsson and Heeger 2006), but there is also evidence that the predominant organization in LO is head- or body-centered rather than retina-centered (McKyton and Zohary 2007). Similarly, while position coding in the motion-selective middle temporal (MT) region has previously been reported as retinotopic (Huk, Dougherty et al. 2002), a recent study that manipulated gaze direction relative to a motion stimulus found that responses in MT were more consistent with a spatiotopic reference frame ((d'Avossa, Tosetti et al. 2007); c.f. (Gardner, Merriam et al. 2008)). These conflicting results may be reconciled if the aforementioned areas are integrating retinal and extraretinal sources of information in order to construct a representation of perceived position – perceived position may or may not match retinal position on a given trial or for a given experimental paradigm.

In the current study, we tested the hypothesis that higher-level visual areas preferentially represent perceived, rather than physical object position. We measured the relative precision of perceived versus physical position coding in five functionally localized higher-level visual areas: lateral occipital cortex (LO), posterior fusiform gyrus (pFs), the fusiform face area (FFA), the parahippocampal place area (PPA), and the middle temporal complex (MT+). Utilizing subjects' mislocalizations to dissociate physical and perceived stimulus position, we found that changes in the patterns of activity in each of these higher-level areas were more tightly coupled with changes in perceived position than in physical position. Our results reveal the existence of a percept-centered coordinate frame for position coding in higher-level visual areas.

2.2 Methods

Eight subjects with normal or corrected-to-normal vision participated in this study (six participated exclusively in the main experiment, one participated in both the main experiment and the eye tracking control, and one participated exclusively in the eye tracking control). Each subject provided informed consent prior to participation, and all scanning procedures were approved by the University of California, Davis Institutional Review Board.

2.2.1 Functional localizers: stimuli and analysis

In separate functional runs for each subject, we localized visual areas MT+, FFA, PPA, LO, and pFs, as well as V1, V2, V3, V3a, VP, and V4. To localize and demarcate areas V1 through V4, we used flickering wedge stimuli (“bowtie” patterns (Sereno, Dale et al. 1995)). Bowties consisted of counterphase flickering (7.5 Hz) radial sine wave patterns of 11.79 deg. radius, subtending an arc of 8.16 deg. There were three conditions in these bowtie runs; in two conditions, the bowties were centered on the vertical or horizontal meridians, and the third condition was a fixation baseline in which only the fixation point was present. Conditions were randomized in 36 ten-second blocks. In all conditions, subjects were instructed to fixate while performing a counting task at the fixation point. Small (0.98 deg.) radial and circular gratings appeared around the fixation point at random times during each ten-second block, and one pattern was always presented more often than the other. At the end of each block a white annulus (0.98 degrees diameter) appeared around the fixation point, prompting subjects to make a response indicating which pattern had occurred most often.

To define the boundaries of areas V1 through V4, we traced the mirror reversals in the cortical representations of the horizontal and vertical meridians, as identified by the horizontal and vertical bowtie stimuli (Sereno, Dale et al. 1995). To do this, we constructed a GLM contrast between the horizontal and vertical bowtie stimuli. This contrast yielded a striated map of activity across the early visual areas, revealing their retinotopic organization. We separately overlaid each subject’s contrast map on his or her inflated brain, and traced the boundaries of areas V1 through V4 by following the horizontal and vertical meridians.

To define regions of interest (ROIs) for areas LO and pFs, we conducted separate localizer scans using methods similar to those used by Grill-Spector and colleagues (Grill-Spector, Kushnir et al. 1998). Stimuli in these scans consisted of intact and scrambled objects (Supp. Fig. 2.1a). Subjects performed a one-back matching task, indicating whether or not the current object matched the previously presented object, while maintaining fixation throughout each run. Intact objects were gray-scaled and centered within 4.1 x 5.7 deg. rectangles. Scrambled objects were created by dividing each of the object images into a grid of 875 squares and shuffling the locations of these squares within the rectangle. Each run consisted of ten thirty-second stimulation blocks (five with intact objects and five with scrambled objects), interleaved with five twenty-second fixation periods. Within each stimulation block, 40 stimuli were presented (1.33 Hz). Each subject participated in 2 runs, except for FF, who participated in one run.

To localize areas PPA and FFA, we used a one-back matching task and block design similar to those that we used to localize LO and pFs (Kanwisher, McDermott et al. 1997; Epstein, Harris et al. 1999). Within each run, face and house stimuli (Supp. Fig. 2.1b & c) were presented in ten alternating thirty-second blocks, interleaved with five twenty-second fixation periods; within each block, 40 stimuli were presented. As with intact objects, the face and house stimuli were gray-scaled and centered within 4.1 x 5.7 deg. rectangles. Each subject participated in two runs.

ROIs for LO, pFs, PPA, and FFA were defined by conducting GLM analyses on data collected from the localizer runs. For all subjects, each area was separately defined as the region with the strongest contrast in activations when the subject viewed intact objects versus scrambled objects (LO and pFs), houses versus faces (PPA), or faces versus houses (FFA). To select ROIs for six of our subjects, we set the threshold at $t > 4.9$, $p < 0.05$, Bonferroni corrected. One subject had a weak BOLD response, so the threshold was reduced to $t > 3.6$, $p < 0.05$. The inclusion or exclusion of this subject's data did not significantly influence any results in this study. Several subregions of the LOC have been identified; however, these have not been clearly established and universally agreed upon (Grill-Spector, Kushnir et al. 1999; Larsson and Heeger 2006; McKyton and Zohary 2007). The Talairach coordinates we found corresponded most with studies that identified these subregions as LO and pFs; therefore, we chose to retain this nomenclature in our study (Grill-Spector, Kushnir et al. 1999; Avidan, Hasson et al. 2002; Altmann, Deubelius et al. 2004).

Each subject also participated in separate runs to functionally localize area hMT+ (the human homologue of monkey areas MT and MST, commonly referred to as MT+). The MT+ localizer runs consisted of three conditions: Gabors with inward motion (drifting toward fixation), Gabors with outward motion (drifting away from fixation), and a fixation baseline condition in which only the fixation point was present. The Gabors were situated as in the main experiment: one Gabor was presented in each visual quadrant at an eccentricity of 9 degrees from fixation. The Gabors had a spatial frequency of 0.38 degrees, and drifted inward or outward at 2.5 Hz for the duration of each block. The three stimulus conditions were randomized in 18 ten-second blocks, and during these blocks, subjects performed a task at the fixation point identical to the one described for

the bowtie localizer above. Subjects participated in a minimum of five MT+ localizer runs.

MT+ regions of interest were functionally defined for each subject by contrasting moving versus baseline responses in a GLM applied to the localizer runs. The threshold for inclusion in the ROI was $t > \pm 6.3$, $p < 0.001$, Bonferroni corrected. Two subjects had weak responses to all stimuli, so the threshold for inclusion was dropped to $t > \pm 3$, $p < 0.003$. The inclusion or exclusion of these subjects did not significantly change the results. The Talairach coordinates for the functionally defined MT+ ROIs were consistent with those estimated in previous studies (Watson, Myers et al. 1993; Tootell, Reppas et al. 1995; Dumoulin, Bittar et al. 2000; Kourtzi and Kanwisher 2000; Dukelow, DeSouza et al. 2001). See Supplemental Table 2.1 for the averaged Talairach coordinates of LO, pFs, PPA, FFA, and MT+ across all seven subjects (Talairach and Tournoux 1988).

2.2.2 Main experiment stimuli

Stimuli consisted of four flickering Gabor patterns (sinusoidal luminance modulations within Gaussian contrast envelopes; Fig. 2.1a). The Gaussian contrast envelope of the Gabors was defined as $L(x,y) = A \cdot e^{\frac{-r^2}{(\sigma M)^2}}$, where *A* is the peak contrast amplitude, *r* is the distance of *(x,y)* from the center of the Gaussian, *σ* is the standard deviation, and *M* is the maximum radius. Gabors had a spatial frequency of 0.38 cycles/deg, and flickered in counterphase at 7.5 Hz; the phase of each Gabor was independently randomized on each trial. One Gabor was positioned in each visual quadrant, with the peak contrast (87% Michelson) always located 9.04 deg. from a central fixation point. We used Gabors rather than the optimal stimulus for each ROI to avoid generating activity specific to one particular area. In addition, the physical characteristics of the Gabor stimuli (e.g., position, contrast, and spatial frequency) are easily quantified and can be independently manipulated; the features of objects, faces, or houses (e.g., surface properties, contours, and asymmetries) are difficult to control.

In each of five experimental conditions, the centroids of all four Gabors were set to one of five possible eccentricities. Gabor centroids were manipulated by applying a skew to the Gaussian contrast envelopes (Whitaker, McGraw et al. 1996). A central condition had no skew applied to the contrast envelope; its size (standard deviation) was 1.66 degrees and its centroid was at 9.039 degrees eccentricity. In two more foveal conditions, the contrast envelopes were skewed by 0.19 and 0.38 degrees toward fixation, and in two more eccentric conditions the contrast envelopes were skewed by 0.19 and 0.38 degrees away from fixation. The Gabors skewed in this way had centroids at 8.430, 8.735, 9.039, 9.343, and 9.647 degrees from fixation, based on the equation $(\sigma_2 - \sigma_1)\sqrt{\frac{2}{\pi}}$ (Whitaker, McGraw et al. 1996; Whitaker and McGraw 1998). Eccentricities were manipulated in this way rather than by shifting the peak contrast because skewing the contrast envelope is a better method of isolating the perceptual

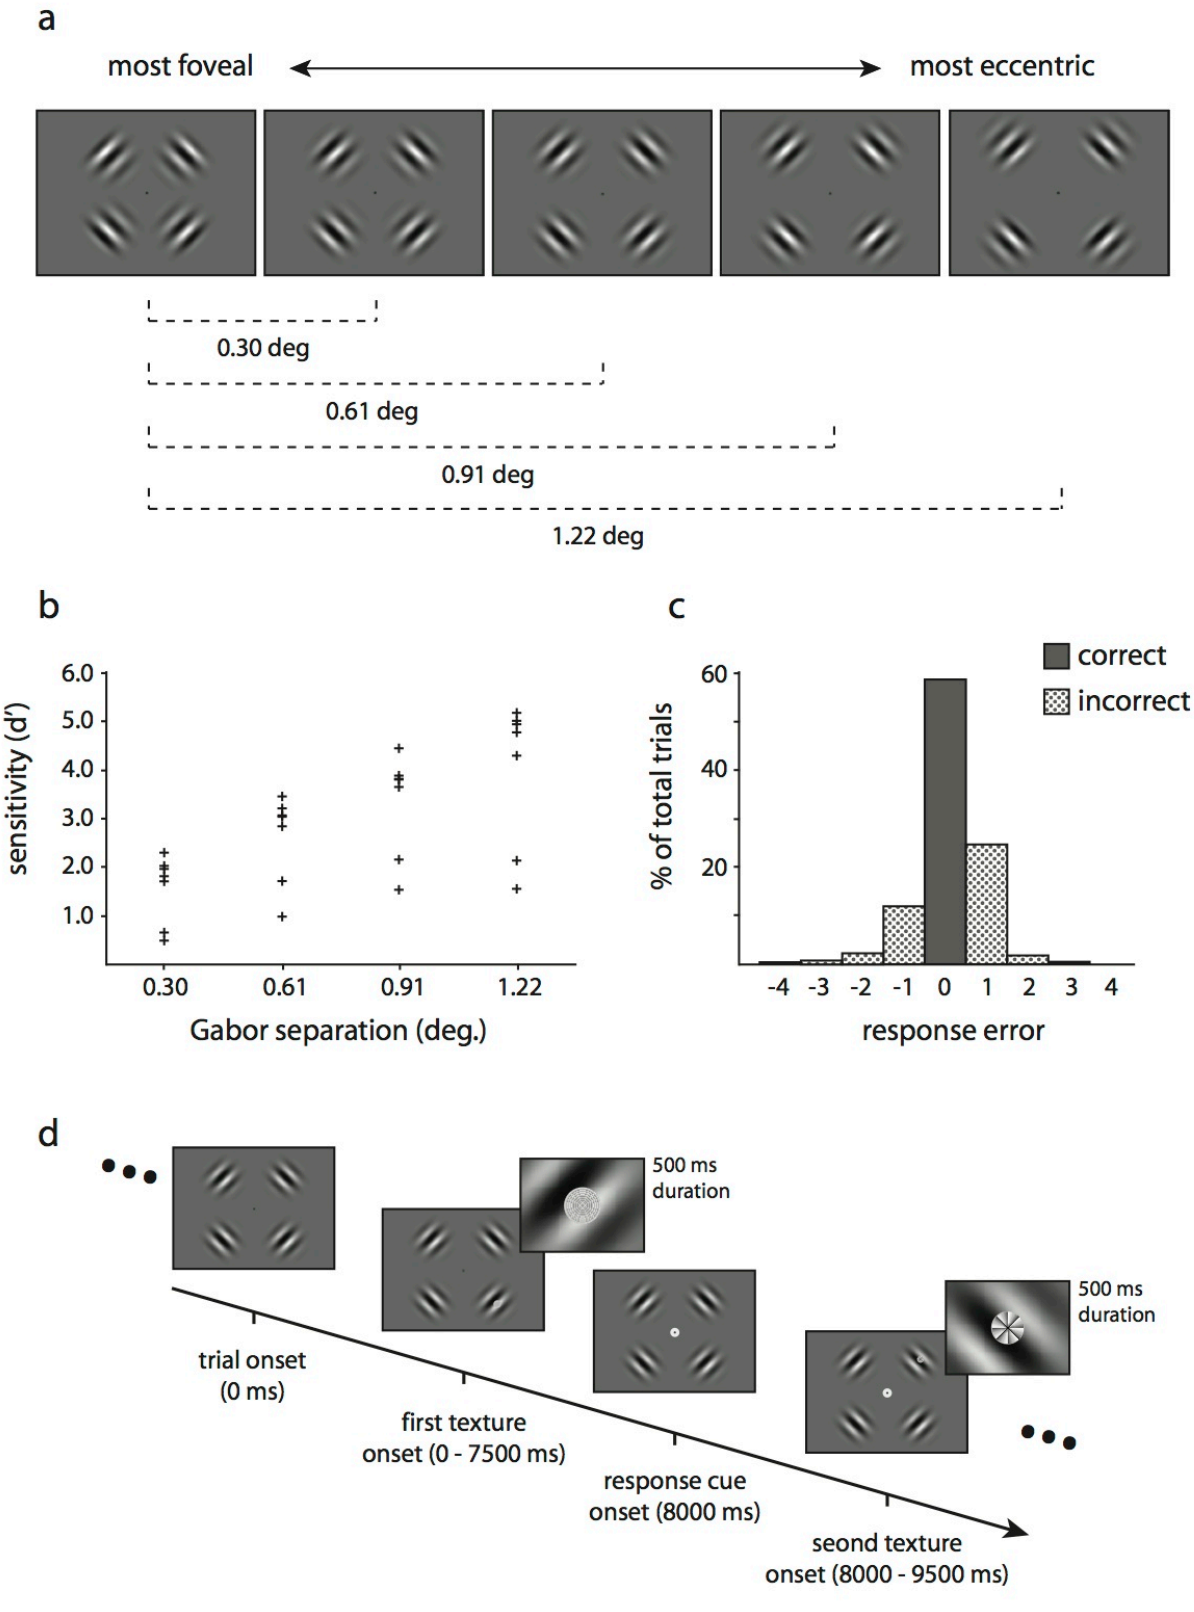
a
most foveal
most eccentric
0.30 deg
0.61 deg
0.91 deg
1.22 deg
b
sensitivity (d')
6.0
5.0
4.0
3.0
2.0
1.0
0.30
0.61
0.91
1.22
Gabor separation (deg.)
c
correct
incorrect
% of total trials
60
40
20
-4 -3 -2 -1 0 1 2 3 4
response error
d
500 ms
duration
500 ms
duration
trial onset
(0 ms)
first texture
onset (0 - 7500 ms)
response cue
onset (8000 ms)
seond texture
onset (8000 - 9500 ms)

Figure 2.1. Main experimental stimuli and behavioral results. (**a**) On each trial, Gabor stimuli appeared at one of five possible eccentricity conditions, ranging from 8.43 to 9.65 degrees from a central fixation point. The range of eccentricities pictured is exaggerated for visualization purposes. The Gabors flickered in counterphase at 7.5 Hz, and were present for the duration of each 10 second trial. A white annulus around the fixation point during the last 2 seconds of each trial cued subjects to respond, indicating which of the five conditions was present. In a sixth baseline condition, only the fixation point was present for 10 seconds. The behavioral and correlation analyses refer to the separations between the five stimulus conditions – these are indicated below the stimuli (in deg. visual angle). (**b**) Subjects' sensitivity (in d' units, see Macmillan and Creelman (MacMillan and Creelman 2004)) is plotted against stimulus separation. The positive trend in this plot indicates that the five stimulus positions sampled the dynamic range of subjects' discrimination sensitivity. The overall hit rate was 58%. (**c**) Histogram of the response errors for all trials in the main experiment, binned by the distance between the correct and actual responses. For example, if, on a given trial, the Gabors were presented in position 4 (9.34 degrees from fixation), but the subject responded with position 3, that trial is binned in response error (-1). The trials in solid gray at response error 0 are the trials in which subjects responded correctly; we exclusively used the *missed* trials (dotted bars; 42% of total trials), in which perceived and physical position were dissociated, in the analysis of the main experiment. (**d**) An example trial. The Gabor stimuli were present for the duration of each 10 s trial. At a random time during the first eight seconds of each trial, a small texture (either a radial or concentric grating) appeared for 500 ms at 9.04 deg. from fixation in a random quadrant. During the last two seconds of the trial, a second texture appeared at the same eccentricity in a random quadrant; the type of texture matched the first on 50% of the trials. At the end of the trial, subjects gave two responses, indicating the positions of the Gabors (5AFC) and whether or not the two presented textures were the same (2AFC).

mechanisms that perform centroid analysis (Whitaker, McGraw et al. 1996), and skewing the Gabor envelope dissociates peak contrast from centroid information. In a previous study, we showed that skewing a Gabor's Gaussian envelope versus shifting its peak contrast are both equally valid means of altering the retinotopic representation of the pattern (Whitney and Bressler 2007). A sixth condition consisted of a fixation baseline in which only the fixation point, a 0.39 deg. diameter bulls eye, was present.

In a control analysis, we presented faces instead of Gabors (Supp. Fig. 2.5). The stimuli consisted of four faces, situated symmetrically about the fixation point in the four visual quadrants, just as with the Gabor stimuli. The faces were drawn from the PICS database (University of Stirling Psychology Department; http://pics.psych.stir.ac.uk/). Each face was enveloped in a Gaussian contrast window to define its centroid and standardize its size. The Gaussian contrast profiles had a standard deviation of 1.66 degrees; to yield five position conditions, the envelope was skewed by –0.38, -0.19, 0, 0.19, or 0.38 degrees and applied to the faces, yielding centroids at 8.430, 8.735, 9.039, 9.343, and 9.647 degrees from fixation, just as with the Gabor stimuli. The identity of the four faces updated at 7.5 Hz, and on each update, the next identity was randomly drawn from a pool of 20 possible identities. In all other respects, the experimental parameters for the face position discrimination experiment were the same as those for the main experiment.

2.2.3 Experimental Design & Task

In each functional imaging run, the six stimulus conditions were randomly interleaved in 36 ten-second blocks (each condition was presented six times; runs were 360s in length). A blocked design was used to maximize signal to noise ratio in the resulting maps of BOLD response. Each subject participated in five functional runs.

Subjects maintained fixation at a central point throughout the entire experiment. On each trial, subjects attended to the locations of the surrounding Gabor stimuli and judged the position (centroid eccentricity) of the Gabors. This task was a five alternative-forced-choice (5AFC) classification task (MacMillan and Creelman 2004), with one Gabor position assigned to each of five buttons on a button box (method of single stimuli (Volkmann 1932; McKee, Silverman et al. 1986)). At the end of each 10 second trial, a white annulus (0.98 deg. diameter) appeared around the fixation point, prompting subjects to respond by pressing the button associated with the position in which the Gabors appeared.

To prevent subjects from making the position judgments on the stimuli quickly and then attending elsewhere for the remainder of each trial, subjects performed a secondary sustained attention task at the locations of the Gabors. During the first eight seconds of each ten-second block, at a randomly chosen time, a patterned circle was briefly superimposed on one of the four Gabors (chosen at random) for 500 ms. The patterned circle (either a circular or radial pattern) was always presented at an eccentricity of 9.04 degrees. A second patterned circle was presented again in the same manner during the last two seconds of each ten-second block, and subjects responded to indicate whether the first patterned circle presented matched the second circle (patterns matched with a probability of 50%). Figure 2.1d depicts an example trial. In this way, we ensured that subjects maintained attention at the locations of the surrounding stimuli for the entire 10 seconds of each trial. In addition to the functional runs described above, additional runs were interleaved in which subjects performed a task at the fixation point and did not perform the 5AFC position discrimination task. These runs were part of a separate analysis and are not included in the present results.

We also ran a subset of the subjects from the main experiment in an event-related experiment in which stimulation intervals were only 2 s long. We found very similar results to those from the main experiment, but the signal strength in the BOLD response was substantially reduced. Since the power of our pattern analysis depends on signal-to-noise ratio in the BOLD response (Fischer and Whitney 2009), our main experiment's blocked design was aimed at achieving an optimal compromise between the psychological and analytical demands of the study.

2.2.4 FMRI data acquisition and preprocessing

Imaging was conducted at the UC Davis Imaging Research Center on a 3-Tesla Siemens TRIO scanner. Each subject's head was placed in a Siemens eight-channel phased-array head coil, and padding was placed on the side and forehead of the

subject to restrict movement. Using a Digital Projection Mercury 5000HD projector, stimuli were back-projected at 75 Hz onto a semi-transparent screen from outside the bore. Subjects were able to see the screen and stimuli via a mirror angled at 45 degrees, located 9.5 cm directly above their eyes. Functional images were collected with a gradient-recalled echo EPI sequence. Whole-brain anatomical images were acquired with a high resolution (1 mm^3) Turbo Spin Echo scan. The acquisition parameters were: TR = 2000 ms, TE = 26 ms, FA = 90 deg, FOV = 22 x 22 cm^2, voxel size = 1.528 x 1.528 x 2.5 mm^3, 20 slices per volume. The imaging volume was centered on the calcarine sulcus, covering the occipital lobe.

All preprocessing and general linear model (GLM) analyses were conducted using Brain Voyager QX (Brain Innovation B.V., Maastricht, The Netherlands). Preprocessing included linear trend removal and 3D motion correction on a run-by-run basis. Prior to all GLM analyses, corrections for serial correlations (removal of first-order autocorrelations) were applied. The images from each functional run were individually aligned to the subject's respective high-resolution anatomical image, reducing the effects of head movement between runs. The anatomical images were then transformed into Talairach space.

2.2.5 Data Analysis

In the main analysis, we discarded the "hit" trials and analyzed only the "miss" trials in order to focus on the subset of the data in which there was a difference between the physical and perceived locations of the stimuli. To measure position selectivity within each ROI, we first generated a map of the BOLD response corresponding to each stimulus eccentricity condition. We did this separately for each subject's five functional runs in a GLM analysis, using the six stimulus conditions (five Gabor eccentricities, plus a baseline condition) as predictors. We separately contrasted each of the five Gabor eccentricities against the baseline condition to produce five statistical maps (t values, unthresholded) of the BOLD activity unique to the five stimulus eccentricities. The subsequent correlation analysis measured position discrimination by tracking systematic changes in the patterns of activity in these maps.

To quantify the precision of position discrimination, we generated a position discrimination plot for each ROI. To do this, we first computed the correlations between all possible pairings of activity maps within each ROI. Given any two of a subject's five maps of BOLD response (corresponding to the five stimulus eccentricities), we computed the correlation between the maps by pairing the t values from the two on a voxel-by-voxel basis, and computing a Pearson r for the resulting set of pairs. Figure 2.2 illustrates this process for two of the ten correlations: correlating the activity from adjacent stimuli generally yielded large r values (Fig. 2.2a), whereas correlating the activity from more distant stimuli yielded smaller r values (Fig. 2.2b). We converted the 10 resulting r values to Fisher z scores so that they could be linearly compared with each other. This process yielded ten z scores for each ROI; to create a position discrimination plot, we plotted each of the z scores against the spatial separation (in degrees visual angle) of the two stimulus eccentricities from which it was produced (Fig.

2.2c). Adjacent conditions had centroids separated by .304 degrees; more distant conditions were separated by multiples of this value (0.609, 0.912, and 1.216 degrees). In order to avoid a loss of information due to imperfect registration between runs, we performed this process separately for each run; thus, each subject's individual position discrimination plot for a given ROI contained a total of 50 points (Fig. 2.2c depicts data for a single run). For group analyses, we fit a regression to all subjects' data taken together. As we wished to combine data across several subjects in order to make inferences about the population at large, we used a random effects analysis to account for between-subject variability (Holmes and Friston 1998). Our regression model took the form $z_{ijkl} = \beta_0 + \tau_i + \beta_1 x_{jk} + \varepsilon_{ijkl}$, where *i* indexed the subjects, the pair (*j*, *k*) indexed the ten stimulus pairings, and l indexed the run number; τ_i accounted for baseline differences between subjects. The goodness of fit of the group regression (r) for each ROI provided an index of the precision of position coding there – the stronger the inverse relationship between the BOLD response correlations and their corresponding stimulus separations, the more precise the position discrimination(Bressler, Spotswood et al. 2007; Fischer and Whitney 2009; Fischer and Whitney 2009). We expressed the precision of position coding for each plot as the Fisher z-transformed r value for that plot. While the slope of a position discrimination plot also reflects the precision of position coding in the corresponding ROI, r serves as a better indicator of coding precision because it captures the *significance* of slope. That is, r reflects the steepness of the linear fit and the compactness of points around the regression line, both of which are important indicators of the precision of position coding. While the addition of an outlier to a plot might increase the slope of the linear fit, the corresponding r value will generally decrease, reflecting a greater scatter in the points around the regression line. The latter is desirable, as in our case, greater scatter means less systematic encoding of stimulus position. We took the Fisher z-transform of each r value to linearize the scale of the precision estimates, so that they could be directly compared. Table 2.1 gives the position discrimination fit estimates for each ROI.

Area V1 - subject 2, run 1

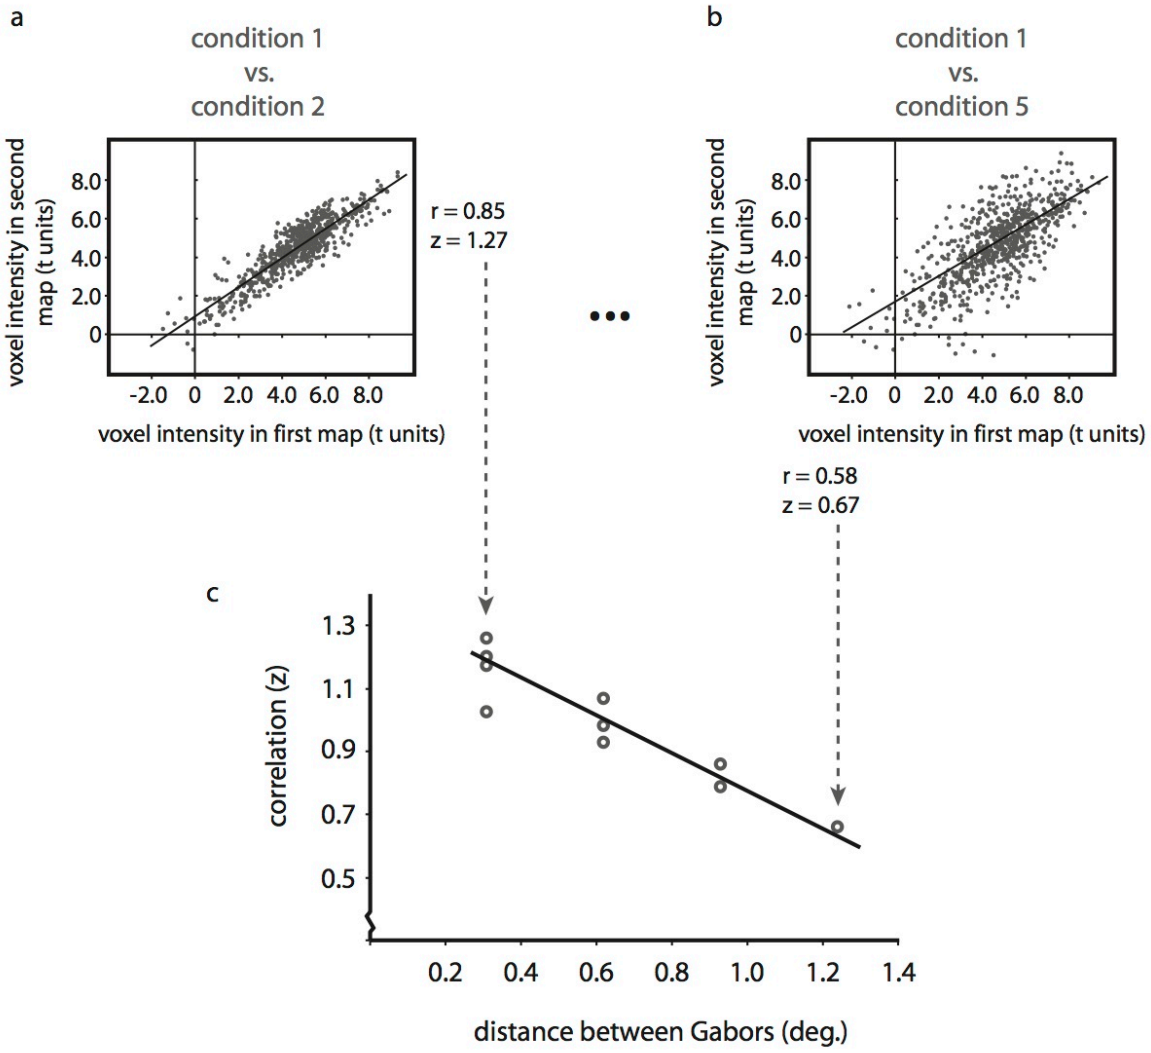

Figure 2.2. Construction of a position discrimination plot. To measure position selectivity within an ROI, we created a position discrimination plot, in which the similarity between the patterns of activity produced by any two of the five stimulus conditions is plotted against the spatial separation between the stimuli presented in those conditions. We used the correlation (Pearson's r) between two patterns of BOLD response as the measure of their similarity. To compute the correlation between a given pair of activity maps, we plotted the intensity values (t units) from one map against those from the other map on a voxel-by-voxel basis, and fit a linear regression to the plot. We transformed the resulting r value to a Fisher z to allow for a linear comparison among multiple correlations measured in this way. Two such correlations computed within V1 for an example subject are shown in panels (**a**) and (**b**). The plot in panel (**a**) corresponds to two adjacent stimulus conditions, while the plot in panel (**b**) corresponds to the two furthest separated stimulus eccentricities. Note that the correlation in panel (**a**) is substantially stronger than that in panel (**b**). The fact that retinotopically proximal stimuli produce similar patterns of BOLD response, while more distant stimuli produce less similar patterns of BOLD, is an indication that activity in the ROI is position-selective. To evaluate the precision of position selectivity in the BOLD response, we plotted each of the ten correlations against the distance between the stimuli that produced it , and fit a linear regression to the resulting plot (**c**). Panel **c** shows the correlations from a single run; we computed the correlations on a run-by-run basis, so a full position discrimination plot for one subject had fifty total points (see *Methods*). A significant negative trend in this position discrimination plot indicates that activity in the ROI is sensitive to the parametric position manipulation (Bressler, Spotswood et al. 2007; Fischer and Whitney 2009; Fischer and Whitney 2009).

We measured the precision of *perceived* position coding in the same manner, substituting BOLD response maps corresponding to subjects' responses for the original maps that reflected physical stimulus position (see Supp. Fig. 2.2). We created these perceived position maps by using the subjective reports provided by each subject as predictors in the GLM. For example, if on a given trial a subject indicated that the Gabors were at the most eccentric position, then that trial was coded as condition 5 in the GLM, regardless of what was actually presented during that trial. We used only "missed" trials in the analysis of both physical and perceived position coding, so each trial was coded differently in the two analyses. We performed the correlation analysis a second time, using each subject's five perceived position maps (corresponding to the five possible responses) to create perceived position discrimination plots for each region of interest. Panel a of Figure 2.3 shows the perceived and physical position discrimination plots for area LO. The two plots differ with respect to the goodness of their linear fits, and hence with respect to the precision of coding that they reveal for physical and perceived position. Table 2.1 gives the precision of perceived position discrimination for each ROI (group data). To compare the precision of physical and perceived position coding within each ROI, we performed a Z-test: $Z_{diff} = \frac{Z_{perceived} - Z_{physical}}{\sqrt{\frac{1}{N_1 - 3} + \frac{1}{N_2 - 3}}}$, where N_1 and N_2 are the numbers of points in the perceived and physical position discrimination plots, respectively. Note that in Figure 2.3a, there are fewer total points in the plot for perceived position than in the plot for physical position. It was occasionally not possible to create a map for a particular perceived position when no corresponding trials were available (for example, on a given run a subject may never have made an erroneous "5" response). The Z-test for significance

accommodates this difference in the number of data points (Snedecor and Cochran 1980). In subsequent control analyses, we also conducted the position discrimination analysis using all trials instead of only the missed trials. In this case, the plots for physical and perceived position had exactly the same number of points, and the results showed the same effects as in the main experiment (see Fig. 2.5 and Supp. Fig. 2.4).

To test for significant heterogeneity in the nature of position coding across the visual areas in the dorsal and ventral streams, we performed a chi-square test on the ($Z_{perceived}$ – $Z_{physical}$) scores for each of the two collections of visual areas. The test is given by $X^2 = \sum((N_i - 3)(zi - \bar{z})^2)$, where each z_i is a ($Z_{perceived}$ – $Z_{physical}$) score, $\bar{z}$ is the weighted mean of all such scores, and N_i is the number of points on each position discrimination plot. Subsequently, we tested for a trend in the nature of position coding across visual areas by computing a Spearman rank correlation coefficient between the areas' ordering in the visual processing hierarchy and their bias for physical or perceived position coding, given by ($Z_{perceived}$ – $Z_{physical}$). We ordered the early visual areas according to the order in which they are encountered when moving anteriorly across the cortex from V1 (allowing for duplicate ranks). As a guide, we used Figure 1 from Tootell et al (Tootell, Tsao et al. 2003) Thus, the early visual areas were ranked as follows: V1 – 1, V2 – 2, V3 – 3, VP – 3, V3a – 4, V4 – 4. All of the higher level visual areas were assigned the same rank of 5. We computed ρ and its corresponding p value in SPSS 17.0.

2.2.6 Eye tracking data collection & analysis

For two control subjects, we monitored eye position during scanning. Eye position was monitored at 60 HZ using an ASL Eye-Trac 6 series long-range eye tracker. Eye position data was collected using EyeTrac6000 software, and subsequently analyzed in Matlab 7.1. To determine whether the subjects' eye movements were correlated with either the stimuli or their responses, we aligned the eye position data with the stimulus presentation epochs for each run (Figure 2.6b). We then populated a separate list of sampled eye positions for each of the five physical positions, and each of the five perceived conditions. All of the eye positions in the physical condition lists were also present in the perceived condition lists, but many were assigned to a different condition number because the subject misperceived the stimulus position. For each list, we computed the mean eye position in the x direction (purple data points in Figure 2.6c) and in the y direction (green data points), as well as the mean squared distance from the fixation point as a measure of variability. For each subject, separately for the physical and perceived stimulus positions, we performed three one-way ANOVAs, testing for significant differences in x position, y position, and variability.

2.3 Results

In separate imaging runs, we functionally localized visual areas LO, pFs, FFA, PPA, and MT+ using standard techniques, including the presentation of objects, buildings, faces, and moving stimuli (Supp. Fig. 2.1; see *Methods*) (Sereno, Dale et al.

1995; Kanwisher, McDermott et al. 1997; Grill-Spector, Kushnir et al. 1998; Epstein, Harris et al. 1999). The Talairach coordinates of our regions of interest were in good agreement with those reported in previous studies (Supp. Table 2.1) (Watson, Myers et al. 1993; Gauthier, Tarr et al. 1999; Grill-Spector, Kushnir et al. 1999; Dumoulin, Bittar et al. 2000; Kourtzi and Kanwisher 2000; Avidan, Hasson et al. 2002; Epstein, Graham et al. 2003; Altmann, Deubelius et al. 2004; Grill-Spector, Knouf et al. 2004; Yi, Kelley et al. 2006). In the main experiment, our goal was to measure the precision of stimulus position information in the pattern of BOLD response in each region of interest (ROI). During scanning, we presented flickering Gabor patches at five possible eccentricities, ranging from 8.43 to 9.65 degrees from fixation (Fig. 2.1a). In each 10 second block, the stimuli were presented at one of the five eccentricities and subjects reported their apparent position in a five alternative forced choice (5AFC) response. Seven subjects participated in the main experiment.

Figure 2.1b shows behavioral performance: subjects' ability to discriminate between two different conditions (discrimination sensitivity, d') is plotted as a function of the separation between the two eccentricities (see Macmillan and Creelman (2004) for details on calculating pairwise sensitivity). A one-way ANOVA across the five conditions at the smallest (.30 deg.) separation showed no performance bias for any particular stimulus eccentricity ($F_{4,30} = 0.75$; $p = 0.57$). The positive trend in d' for increasing stimulus separation indicates that when subjects made errors, they were most likely to mistake the presented condition for an adjacent one. This tendency is also evident in Figure 2.1c, which shows a histogram of subjects' errors. Trials are binned according to the difference between the subject's response and the correct response; the trials in bin 0 are correct trials, and trials in positive bins are those in which the subject perceived the stimulus as more eccentrically positioned than it was presented. Subjects performed well at determining the positions of the Gabors (~58% correct; chance is 20%), but the task was sufficiently difficult to elicit a substantial number of errors. In the subsequent data analysis, we focused exclusively on these "missed" trials, in which subjects mislocalized the positions of the stimuli, in order to dissociate physical and perceived position.

Subjects also performed a secondary task to ensure that they attended at the locations of the Gabors for the entirety of each trial. During the first eight seconds of each trial, a small pattern appeared at the centroid of one of the Gabors. At a randomly chosen time during the final two seconds of each trial, a second pattern appeared in one of the Gabors; the second pattern matched the first on half of the trials. At the end of each trial, in addition to reporting the positions of the Gabors, subjects indicated whether the two patterns matched (Fig. 2.1d shows an example trial). Performance on this secondary task was 75.9% correct (chance was 50%). Subjects' responses on the secondary task were not correlated with their responses on the primary position discrimination task ($F_{4,1040} = 0.53$; $p = 0.72$), nor were they correlated with the physical positions of the stimuli ($F_{4,1040} = 0.25$; $p = 0.91$). Since responses on the primary and secondary tasks were not correlated, subjects did not use the patterns to judge the positions of the Gabor stimuli. The below ceiling performance on this task indicates that

it was demanding enough to hold subjects' attention at the location of the Gabors for the duration of each trial.

Within each ROI, we measured the selectivity for stimulus position by tracking the change in the spatial pattern of the BOLD response corresponding to the incremental shifts in stimulus position (see *Methods*). We first discarded the "hit" trials and kept only the "miss" trials for the main analysis; by doing so, we focused on the subset of the data in which there was a difference between the physical and perceived locations of the stimuli (see Fig. 2.5 and Supp. Fig. 2.4 for the results of the same analysis when all trials are included). For each subject, we produced five maps of BOLD response corresponding to the five stimulus positions by separately contrasting each condition against a fixation baseline in a general linear model. We then tested for a trend in the spatial pattern of the BOLD response by measuring the similarity between pairs of maps. For any given pairing, we computed the correlation between the two maps within the current ROI (Fig. 2.2a & b show example correlations computed for conditions separated by .30 deg. and 1.22 deg, respectively). Collecting the correlations from all ten possible pairings of the five maps, we plotted each correlation against the spatial separation between the stimuli that produced it to construct a position discrimination plot (Fig. 2.2c). An ROI that codes stimulus position will have a strong negative trend in its position discrimination plot, since stimuli that are closer together in space produce more highly overlapping patterns of BOLD response. In an ROI that does not contain position information, on the other hand, there is no reason to expect any trend in the correlations on the position discrimination plot. Thus, we used R, the goodness of fit of a linear regression applied to the data on the position discrimination plot, as an index of the precision of position coding in each ROI. The goodness of fit measure captures how tightly clustered the correlations are at each separation, as well as the slope of the data, which indicates how dramatically the pattern of BOLD changed with incremental changes in position (see *Methods*). We took the Fisher z-transform of each r value to linearize the scale of the precision estimates, so that they could be directly compared. We have previously shown that this pattern analysis technique is able to measure position selectivity in the BOLD response on a submillimeter scale (Bressler, Spotswood et al. 2007; Whitney and Bressler 2007; Fischer and Whitney 2009; Fischer and Whitney 2009).

To test the precision with which each ROI codes *perceived*, rather than physical stimulus position, we constructed position discrimination plots in the same way, this time using maps of BOLD response corresponding to the five possible responses from the position discrimination task. These maps, corresponding to the positions in which subjects perceived the stimuli on a trial-by-trial basis, were produced using the exact same trials as in the analysis of physical position coding; the only difference was in how the trials were coded in the GLM analysis (see Supp. Fig. 2.2). The resulting maps reflected activity corresponding to perceiving the stimuli in the five different locations, regardless of variations in the actual locations of the stimuli. Because we used only missed trials, every trial was coded differently when coded by physical position vs. the subject's response. In this way, we were able to dissociate physical and perceived

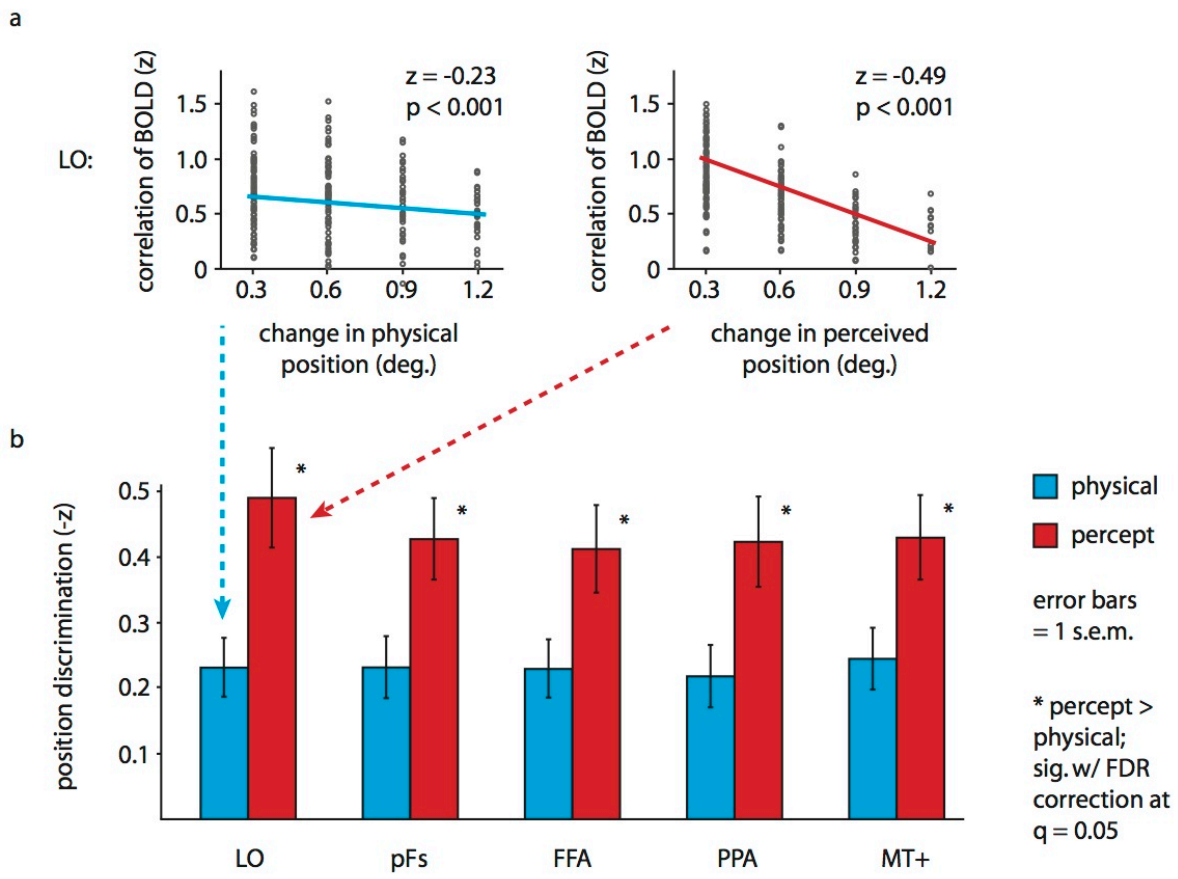

Figure 2.3. Comparison of physical and perceived position discrimination in higher-level visual areas. (**a**) For each ROI, we constructed a position discrimination plot, as outlined in Figure 2.2, separately for physical position (panel **a**, left) and perceived position (panel **a**, right). These plots are based on the exact same trials; the only difference between the two is whether the trials were coded according to their physical position, or according to subjects' responses, in the GLM analyses. The plots in panel **a** show data from all subjects; to perform a group-level analysis, we fit a linear regression to all subjects' data taken together, and included a random effect of subject in the regression model to account for between-subject variance (see *Methods*). The goodness of fit of the linear regression captures how tightly changes in BOLD were coupled with changes in physical or perceived stimulus position, and hence serves as an index of the precision of position coding. (**b**) The precision of physical (blue) and perceived (red) position coding are plotted side by side for each of the five higher-level visual areas. Y-axis units are in –Z, so a taller bar indicates a more strongly negative correlation on the position discrimination plot. Error bars indicate +/- 1 s.e.m. Coding of perceived position was significantly more precise than coding of physical position in every higher-level visual area we tested (see Table 2.1 for statistics).

position, and separately measure the amount of information about each in the BOLD response.

Figure 2.3 shows the results of the position discrimination analysis for area LO (Fig. 2.3a) and the other four higher-level areas we tested (Fig. 2.3b). The position

discrimination plots in Fig. 2.3a show all subjects' data plotted together; to measure a group-level effect, we fitted a regression to all subjects' data, and included a random effect of subject in the regression model to account for between-subject variability (Holmes and Friston 1998)see *Methods*). Precision of physical and perceived position coding is plotted as -z, reflecting the fact that a strongly negative linear fit to a position discrimination plot reflects precise position coding. Every area we tested showed significant discrimination of physical stimulus position ($p < 0.0001$ for all areas). Strikingly, however, the coding of perceived position was significantly more precise than coding of physical position in every area (red bars in Fig. 2.3; see Table 2.1 for statistics). Consistent with our hypothesis, comparing the relative precision of physical versus perceived position coding within each area revealed a preferential representation of perceived position in these higher-level visual areas. Supplemental Figure 2.3 shows physical and perceived position discrimination in the five higher-level areas from Figure 2.3, broken down by individual subject. With few exceptions, individual subjects showed effects in the same direction as in the group-level analysis. The fact that subjects showed a consistent pattern of results, with higher-level visual areas more precisely reflecting perceived position than physical position, speaks to the existence of an underlying percept-centered coding scheme. Each subject made a unique set of errors, and any alternative coding scheme tied solely to the physical positions of the stimuli would have been washed out in our analysis, in which a given condition could be, at various times, reassigned to any of the other four conditions.

In the main analysis, we used only unthresholded BOLD response maps because non-significant voxels can still carry precise stimulus information when analyzed as a multivariate pattern (Norman, Polyn et al. 2006). We also tried the same analysis using thresholded BOLD maps and found very similar results to those in the main experiment (see Supp. Fig. 2.4).

The smallest and largest eccentricities define the endpoints of the continuum of possible responses, so we wondered whether a dramatic difference between those two conditions in the 'percept' maps might be driving the higher precision of perceived position coding, relative to physical position coding. We performed the position discrimination analysis again after removing the correlations at the 1.22 deg. separation, and we still found stronger discrimination of perceived position than physical position in every higher-level visual area (the least significant difference was in the FFA: $Z = 2.06$, $p = 0.039$). We also measured the similarity of the physical and perceived BOLD response maps within each of the five position conditions to determine whether the most dramatic changes resulting from recoding were isolated to any particular condition(s). Correlating the physical and percept maps with each other in each ROI, we found no difference across the five conditions for any ROI (most significant was FFA, $F_{4,170} = 0.96$, $p = 0.43$). Thus, systematic differences between the patterns of BOLD corresponding to physical and perceived position are present across the whole continuum of stimulus positions.

	Physical Position discrimination (-Z)	Perceived Position discrimination (-Z)	Z test (percept > physical)	p
LO	0.231	0.487	-2.90	0.0038*
pFs	0.229	0.426	-2.23	0.0259*
FFA	0.227	0.409	-2.06	0.0396*
PPA	0.217	0.424	-2.34	0.0195*
MT+	0.241	0.428	-2.12	0.0339*

*significant with FDR correction for multiple comparisons at q = 0.05

Table 2.1. The precision of physical and perceived position coding in each higher-level ROI. The precision estimates were computed by fitting a linear regression to the position discrimination plot (grouped data) of each region of interest. An r value significantly different from zero indicates significant position selectivity within the ROI, and a more negative r implies more precise position coding (Fischer and Whitney 2009). We applied a Fisher z-transform to each precision estimate (r value) for the sake of linear comparison, and we present the precision estimates in negative z units so that larger values indicate more precise position selectivity. We performed a within-area comparison of perceived and physical position coding for each ROI using a Z-test (see *Methods*). Encoding of perceived position was more precise than that of physical position in every higher-level area we tested. We corrected for multiple comparisons by controlling the false discovery rate (FDR) to 0.05 (Benjamini and Hochberg 1995).

In a follow-up analysis, we measured physical and perceived position discrimination in early visual areas V1, V2, V3, V3a, VP, and V4, which we functionally defined using a standard retinotopic localizer (Supp. Fig. 2.1; see *Methods*) (Sereno, Dale et al. 1995). The collected data for all visual areas, divided into the dorsal and ventral visual streams, are presented in Figure 2.4. The within-area difference scores -($Z_{percept} - Z_{physical}$), plotted below the raw discrimination data, are the most informative indices of the nature of position coding in each visual area because they reveal the direction and strength of the bias within an area for encoding physical versus perceived position discrimination. It is important to note that due to potential inhomogeneities in signal quality across the brain, a comparison of position discrimination scores is only informative within, and not between, visual areas (see *Discussion*).

The data in Figure 2.4 reveal a trend in the nature of position representation: while the higher-level visual areas we tested showed a strong preference for coding perceived stimulus position, that effect is diminished or reversed in earlier areas. A Chi-square test revealed significant heterogeneity in coding preference across both the dorsal and ventral stream areas (X^2_{dorsal} = 20.92, p = 0.0003; $X^2_{ventral}$ = 35.29, p < 0.0001). To test for a systematic progression in the nature of position coding across areas, we computed a non-parametric correlation (Spearman rho) between the rank-ordered visual areas and their coding preference -($Z_{percept} - Z_{physical}$) scores. We

assigned all higher-level areas the same rank (5), and we ranked the earlier visual areas according to the order they are encountered moving anteriorly from V1 (for reference, we used the inflated human cortex image from Figure 1 in Tootell et al. (Tootell, Tsao et al. 2003)): V1 – 1; V2 – 2; V3 & VP – 3; V3a & V4 – 4. The correlation was highly significant ($\rho = 0.80$; $p = 0.003$). To evaluate this a priori ranking relative to all possible rankings of the visual areas, we computed a 25,000 sample bootstrapped distribution of rank correlations, with a randomly drawn ranking of areas for each sample. The correlation of 0.80 obtained with the a priori ranking was larger in absolute value than 99.5% of the bootstrapped samples, indicating that our ranking based on functional anatomy is a good match to the independently measured position discrimination estimates for each area. The strong correlation between an area's location in the visual processing hierarchy and its position coding bias reinforces the idea that the nature of position coding evolves as information progresses through the visual processing hierarchy, becoming relatively more strongly tied to perception in higher-level areas. This is not to say that early visual areas contain no information about perceived position, but that percept-centered information does not yet dominate position representations at the earliest stages of visual processing. Our results provide a clear answer to our underlying question regarding the nature of position coding in higher-level areas: the location in which an object is perceived matters more than the location in which it was presented.

If subjects' errors had been random or not directly related to perceived position, could recoding the condition labels according to those (uninformative) responses ever yield a spurious improvement in the precision of position coding measured by the correlation analysis? To evaluate how uniquely predictive subjects' actual responses were relative to other possible ways in which the trials could have been relabeled, we performed a bootstrapping analysis in which we repeated the correlation analysis for each subject and each ROI 1000 times. On each iteration, we added "errors" to the physical trials labels; the errors were randomly drawn from the distribution of subjects' actual errors, such that over all iterations, the distribution of simulated errors matched the distribution of actual errors that subjects made (Fig. 2.1c). For each ROI, we ran the position discrimination analysis with the 1000 sets of simulated trial labels, and obtained a distribution of discrimination values (Fig. 2.5). Since the errors for each set of trials labels were randomly drawn, the bootstrapped discrimination values reflect the precision of position discrimination we would expect to measure if subjects' responses reflected random errors. In Figure 2.5, the bootstrapped discrimination values are shown in gray histograms, and the discrimination measured with the actual physical and percept trial labelings are indicated with blue and red bars respectively (these values differ slightly from those measured in the main experiment; here, we used all trials in the analysis rather than only missed trials to allow for a direct comparison of all iterations). In each of the higher-level visual areas, the discrimination measured with the actual percept labels falls in the extreme upper tail of the distribution of discrimination values for all possible trial relabelings (least significant was pFs; $-Z = 3.41$, $p < 0.001$). The degree to which subjects' responses outperform other possible labelings is striking given that both the percept labelings and the bootstrapped labelings had, on average, 58% of their labels in common with the physical coding; the only difference is the

a dorsal areas

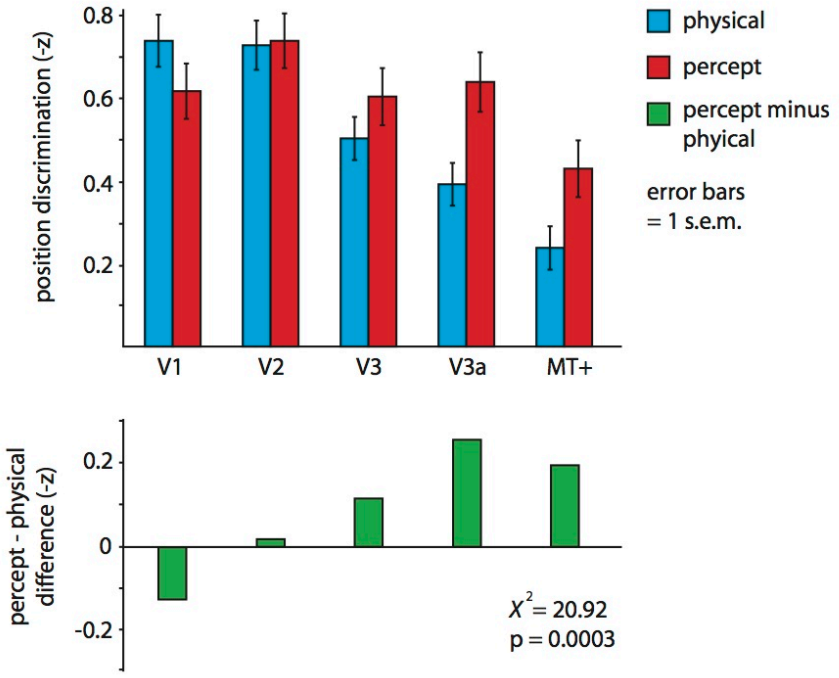

b ventral areas

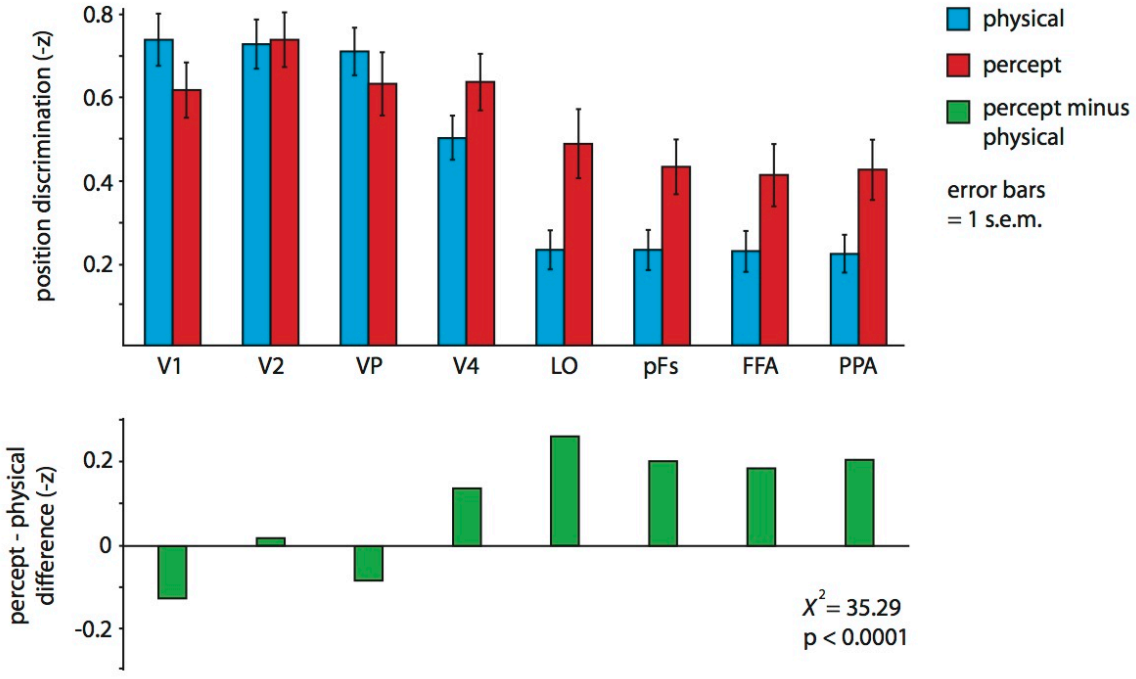

Figure 2.4. Comparison of position discrimination in higher-level visual with that in early visual areas. (**a**) Precision of physical (blue) and perceived (red) position discrimination is shown for the dorsal visual areas in ascending order. Coding preference is captured in the green bars: they show the within-area comparison of physical versus perceived position discrimination for each visual area, obtained by subtracting the precision of physical coding from that of percept coding $-(Z_{percept} - Z_{physical})$. Positive values in the green bars indicate a preferential coding of perceived position, whereas negative values indicate a preferential coding of retinal position. While activity in MT+ reflects perceived position more precisely than physical position, in earlier areas this bias is diminished or reversed. A Chi-square test revealed that the nature of position coding differed significantly among these dorsal areas (X^2_{dorsal} = 20.92, p = 0.0003). Panel (**b**) shows a comparison of perceived and physical position discrimination in the ventral visual stream. Here, too, there was significant variability between areas in the bias for representing physical or perceived position ($X^2_{ventral}$ = 35.29, p < 0.0001). To test for a systematic progression in the nature of position coding across areas, we ranked all eleven areas according to their locations in the visual processing stream (see *Methods*), and computed a Spearman rho rank correlation between the visual area ordering and position coding bias $-(Z_{percept} - Z_{physical})$. The correlation was highly significant (ρ = 0.80; p = 0.003). To test this a priori ordering against all other possible rakings of visual areas, we computed a 25,000 sample bootstrapped distribution of rank correlations, with a randomly drawn ranking of areas for each sample. The correlation of 0.80 obtained with the a priori ranking was larger in absolute value than 99.5% of the bootstrapped samples, indicating that our ranking based on functional anatomy is a good match to the independently measured position discrimination estimates for each area. The strong correlation between an area's position in the visual processing stream and its bias for representing physical or perceived position reinforces the idea that the nature of position coding evolves as information progresses through the visual processing hierarchy, becoming relatively more strongly tied to perception in higher-level areas.

precise arrangement of the errors. Thus, it is indeed critical to use the precise responses that subjects made in order to obtain the dramatic improvement in the precision of position coding that we measured in higher-level areas.

The bootstrapping analysis also allowed us to further disentangle the measurements of percept-and retina-centered encoding. In an area dominated by retinotopic coding such as V1, one might ask whether there is also some degree of encoding of perceived position. In the main experiment, we found significant discrimination of the perceived stimulus positions in V1, but because subjects' responses were highly correlated with the physical stimulus positions (Fig. 2.1c), that measurement could have been carried entirely by the underlying retinotopic coding. However, the histograms in Figure 2.5 provide a means of testing for a unique contribution of percept-related information after taking into account the correlation between subjects' responses and the physical stimulus positions. If there was no unique percept information in V1, then the discrimination of subjects' responses (the red bar) should fall near the center of the bootstrapped distribution, indicating that labeling the trials according to subjects' responses is no better than using a random perturbation of the physical trial labels. In fact, while a percept-based labeling scheme is substantially worse than a physical labeling scheme at predicting the pattern of BOLD in V1, it still falls 2.38 standard deviations (p = 0.017) above what would be expected by chance if there were no percept-specific information in V1. Thus, there is some information about

perceived position encoded in V1, even after accounting for the correlation between subjects' responses and the physical stimulus positions (the bootstrapped labels were, on average, as strongly correlated with the physical labels as subjects' responses were). Nonetheless, the relative precision of retinotopic coding versus percept-centered coding is stronger in V1 than in any other area we measured.

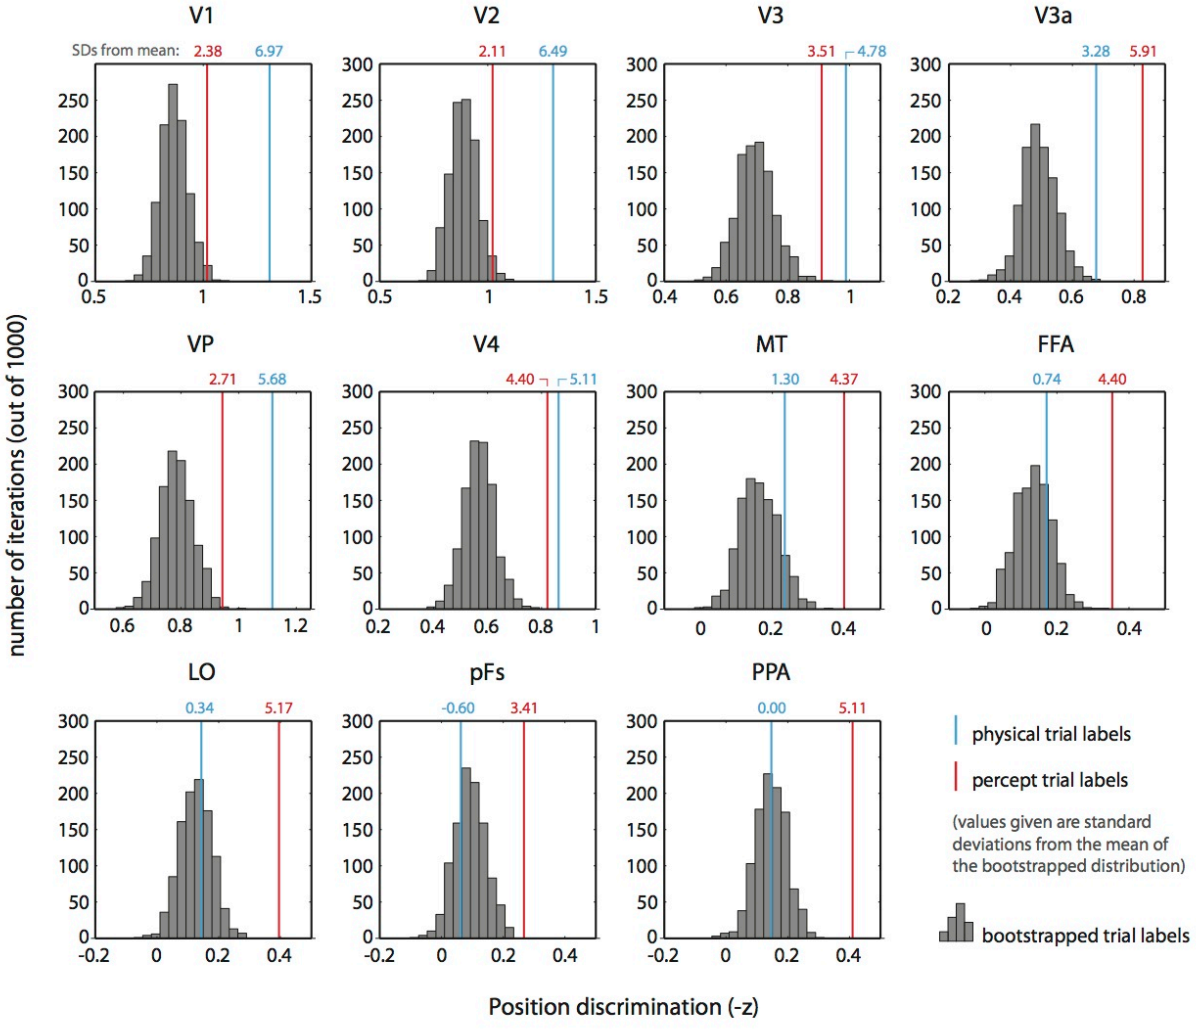

Figure 2.5. Bootstrapped position discrimination analysis. In order to evaluate how uniquely predictive subjects' responses were of changes in the pattern of BOLD response, relative to other possible codings of the GLM predictors, we performed the position discrimination analysis on a bootstrapped sample of 1000 sets of trial labels per subject. To generate each bootstrapped sample, we started with the physical stimulus positions and added "errors" sampled from the distribution of errors that subjects actually made during the experiment (Fig. 2.1C). Thus, over the course of all 1000 iterations, the average errors in the GLM design matrices matched the distribution in Figure 2.1C. On each iteration, we obtained a group-wise position discrimination score (-Z) for each ROI; the collected position discrimination scores are plotted in the gray histograms. Of particular interest was whether subjects' responses performed better at predicting the pattern of BOLD in higher-level areas than did the random errors in the bootstrapped trial labels, which were, on average, equally well-correlated with the physical stimulus positions as subjects' responses were. In fact, in every higher-level area, the percept encoding was an extreme outlier from the bootstrapped distribution (least significant was pFs; $-Z = 3.41$, $p < 0.001$), indicating that the precise errors that subjects made were uniquely predictive of changes in the pattern of BOLD in higher-level areas. Likewise, the fact that the physical trial labels were extreme outliers from the bootstrapped distributions in lower-level areas indicated that even a slight perturbation of the precise retinotopic labeling of the trials resulted in a dramatic decrease in the ability to predict changes in the pattern of BOLD response.

A potential concern is that subjects' eye movements might have been correlated with either the physical or perceived stimulus positions, displacing the stimuli on the retina in a systematic fashion. In fact, the main experiment was designed specifically to factor out the impact of eye movements: since the same trials were used to compare the relative precision of physical and perceived position coding in every visual area, any effects of eye movements would have had the same impact in every ROI, and could not have produced effects in opposite directions, as we found in higher-level visual areas versus early areas (Fig. 2.4). Nonetheless, to test for a possible correlation between eye movements and subjects' responses, we conducted a control experiment in which we tracked two subjects' eye positions during scanning (Fig. 2.6). The position discrimination data from the control subjects were consistent with the results of the main experiment (Fig. 2.6a; see caption for stats). Figure 2.6b shows a sample eye trace (sampled at 60 Hz) for one run from subject 1a, with the stimulus conditions for that run indicated in shades of blue. The correlation between eye position and the stimulus conditions for this run was $r = 0.031$, $p = 0.86$. The largest correlation for any run was $r = 0.048$, $p = 0.78$. Figure 2.6c shows the mean *x* position (purple) and *y* position (green) of gaze during each condition for both subjects. We performed one-way ANOVAs for each of x position, y position, and variability in gaze position, separately for the eye measurements grouped by physical position and perceived position. In neither subject was gaze position or variability related to the physical or perceived stimulus positions (subj. 1a: most significant test was $F_{y\_pos \times percept} = 1.00$, $p = 0.41$; subj. 2a: most significant test was $F_{var \times physical} = 1.57$, $p = 0.18$), indicating that subjects did not make eye movements of different magnitudes or directions in the different conditions. Mean position and variability are easily visualized in Figure 2.6d, which show scatterplots of the recorded gaze positions for each condition for subject 1a. Together, these data show that eye movements cannot account for our results.

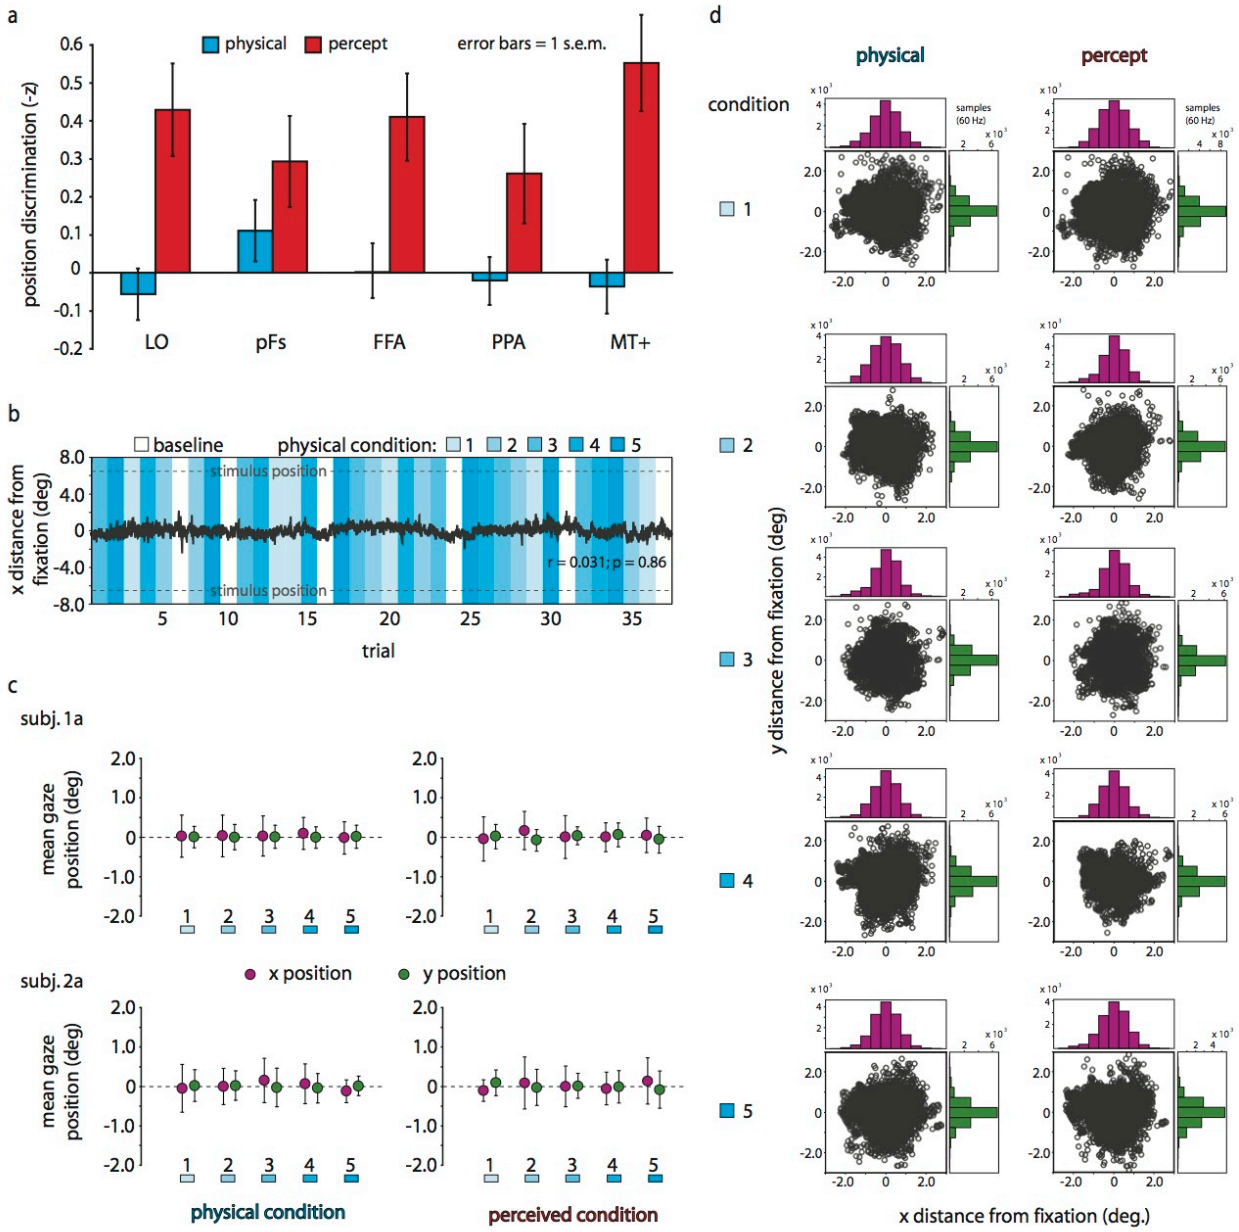
a
physical
percept
error bars = 1 s.e.m.
position discrimination (-z)
LO
pFs
FFA
PPA
MT+
b
baseline
physical condition: 1 2 3 4 5
stimulus position
stimulus position
x distance from fixation (deg)
r = 0.031; p = 0.86
trial
c
subj. 1a
subj. 2a
mean gaze position (deg)
x position
y position
physical condition
perceived condition
d
physical
percept
condition
samples (60 Hz)
y distance from fixation (deg)
x distance from fixation (deg.)

Figure 2.6. Control analysis with eye tracking during scanning. (**a**) Position discrimination data for two additional subjects scanned on the same task as in the main experiment, with eye tracking. Error bars represent +/- 1 s.e.m. The data are consistent with the results of the main experiment; the position discrimination estimates from these subjects fell within the range of estimates acquired for subjects in the main experiment in each visual area (most significant differences were MT+: $z_{physical}$ = 1.81, p = 0.07 and PPA: $z_{percept}$ = 0.87, p = 0.39, for physical and percept discrimination scores, respectively). Discrimination of physical position hovered near zero for these subjects; this pattern is not atypical compared with the individual subject data from the main experiment, although there were also subjects that showed substantial discrimination of physical position (see Supp. Fig. 2.3). (**b**) Sample eye trace for one run from subject 1a. X position of gaze (sampled at 60 Hz) is plotted for the duration of the run, and the presentation of the five position conditions are indicated behind the trace in shades of blue. The correlation between eye position and the stimulus conditions for this run was r = 0.031, p = 0.86. The largest correlation for any run was r = 0.048, p = 0.78. The mean stimulus position, indicated with gray dashed lines, was at +/- 6.4 degrees from fixation. (**c**) Mean x position (purple) and y position (green) are shown for each of the five physical positions (left plots) and perceived positions (right plots) for the two control subjects. Eye position was not correlated with the stimulus position, or the responses made, for either subject (see *Results* for statistics). There was also no correlation between physical or perceived position and variability of eye position (see *Results* for statistics). (**d**) Collections of the recorded gaze positions corresponding to each of the physical and perceived conditions for subject 1a. The fixation point fell at position (0,0) in the center of each plot; the scatterplots show the collected position measurements sampled during the presentation of each condition (physical plots on the left), as well as the collected eye positions corresponding to each reported condition (percept plots on the right). The adjoining histograms show the x and y distributions of recorded eye positions composing each scatterplot.

Perhaps the faithful representation of retinal position within a visual area depends on the presentation of its preferred stimulus. To explore whether the preferential coding of perceived position in higher-level visual areas is affected by stimulus type, we conducted an additional control experiment in which we presented faces instead of Gabors (Supp. Fig. 2.5a). The faces were enveloped in a Gaussian contrast profiles to define their centroids, and they were updated in identity at 7.5 Hz. In all other respects, the stimuli and task were identical to those in the main experiment (see *Methods*). Two subjects from the main experiment participated in this control experiment, so that we had data for both Gabor stimuli and face stimuli for those subjects. We compared the precision of physical and perceived position coding within the FFA, for which faces are an optimal stimulus (Kanwisher, McDermott et al. 1997). Supplemental Figure 2.5b shows the position discrimination scores measured independently using Gabor stimuli and face stimuli. Of critical interest is whether the perceived versus physical bias in the FFA, given by $-(Z_{percept} - Z_{physical})$ (the green bars in Supp. Fig. 2.5b), differed depending on stimulus type. A Z-test on the two bias estimates revealed no significant difference (Z = 0.10, p = 0.92); in fact, the bias estimates using the two different stimulus types were remarkably similar. Notably, the overall selectivity for both perceived and physical position was higher in the FFA for faces than Gabors (this difference was borderline significant: t = 11.90, p = 0.053). This suggests that while optimizing the stimulus for a given visual area might improve the overall precision of position discrimination there, changes in stimulus type do not change the underlying nature of position coding.

A final concern is that higher-level visual areas may possess a representation of the digits used to make the 5AFC position discrimination response. Since subjects made two independent responses on each trial (a 5 AFC position discrimination on the right hand and a same/different texture discrimination on the left hand), we were able to test for evidence of digit coding or response planning in each ROI, independently of the main position discrimination task. We applied the same correlation analysis as in the main experiment after recoding the design matrix according to subjects' responses on the secondary (pattern matching) task (Supp. Fig. 2.6). No ROI was able to discriminate the digit used to make the secondary response (Supp. Fig. 2.6a; most significant was MT+; $-Z = 0.19$, $p = 0.24$). More importantly, across visual areas, variation in the precision of digit coding was not correlated with the degree of bias for encoding physical or perceived position (Supp. Fig. 2.6b; $r = 0.13$, $p = 0.68$). Thus, we can be confident that the percept-centered position coding that we measured in higher-level visual areas is not simply due to encoding of response planning or tactile input.

2.4 Discussion

Our results reveal a precise representation of perceived object position in every higher-level visual area we tested. While these areas do carry some information about a stimulus's retinotopic position, each carries significantly more information about the location in which the stimulus is perceived. The remarkable precision of *perceived* position coding that we found in areas pFs, FFA, and PPA, in which position information has previously been regarded as coarse at best (Grill-Spector and Malach 2001; Hemond, Kanwisher et al. 2007; MacEvoy and Epstein 2007), suggests that it is critical to take into account a subject's perceptual experience – not just the stimulus properties – when measuring stimulus coding in higher-level visual areas.

It is important to have a concrete interpretation of what it means to reassign the GLM predictors according to subjects' responses, as we did in the current analysis. Because subjects reported the apparent positions of the stimuli, there were effectively two distinct stimulus dimensions that we could define as predictors in the GLM (see Supp. Fig. 2.2). The first was physical object position (given by the five possible stimulus eccentricities), and the second was perceived object position (given by the five possible responses). The resulting two sets of BOLD maps revealed the variations in neural activity due to changes in physical stimulus position, and the variations in activity due to changes in perceived position, respectively. Since we included only the missed trials in our primary analysis, each trial had a different value on the perceived and physical dimensions, which provided the maximum possible independence between the two dimensions of interest. Similar techniques investigating "miss" trials are not uncommon in, for example, memory research (Eldridge, Knowlton et al. 2000; Henson, Hornberger et al. 2005). Subsequent analyses showed that including all trials (not just missed trials) yields the same evolution of position coding from retina-centered to percept-centered (Fig. 2.5).

An assumption of our analysis is that subjects' incorrect responses actually reflect the perceived stimulus position on a trial-by-trial basis, rather than simply

reflecting errors in motor execution or random guesses (lapses). While lapses undoubtedly resulted in some missed trials, if missed trials reflected only noise, then modeling that noise would have driven the position discrimination slopes toward zero, rather than producing the dramatic improvements that we found in higher-level areas. Indeed, the bootstrapping analysis presented in Figure 2.5 showed that in every higher-level visual area we tested, subjects' actual responses were far better at predicting changes in the BOLD than random responses that were equally well correlated with the physical positions of the stimuli. Moreover, within each subject we used the same set of five 'percept' activity maps to test position discrimination in every ROI. Any influence of errors and guessing would be uniform across all of the areas we tested and could not account for the opposite effects that we found in control (early) areas versus higher-level areas.

How do we know that some other, non-retinotopic coordinate frame doesn't better account for the results presented here? Could it be that object, face, and place regions code position information in saccade-centered, head-centered, world-centered, or some other physically-defined coordinate frame? We can rule out any single physically defined coordinate frame because all of the precision improvements that we found in our perceived position analysis were driven by missed trials, in which subjects misperceived the stimulus location. Despite this, there was improved systematic coding in every high-level ROI we tested—objects perceived to be located in the same position, regardless of where they were physically located, produced the most similar patterns of activity. As noted above, perceived position coding must incorporate a combination of many coordinate frames.

While the activity we measured in primary visual cortex (V1) was strongly tied to the retinal positions of the stimuli, we found that there was also some unique information about perceived position represented there (Figure 2.5). This is consistent with the results of a recent study by Murray et al. (Murray, Boyaci et al. 2006), which investigated the influence of perceived object size on representations in V1. The authors presented a stimulus of constant retinal size, but they manipulated its perceived size by changing the apparent depth of the object with contextual cues. They found that when the object was perceived as being larger, it activated a larger region of cortex in V1; the pattern of activation was similar to that observed for a physical increase in stimulus size. It is not yet possible to say whether the information about perceived size in V1 shown by Murray et al. (Murray, Boyaci et al. 2006) and the information about perceived position in V1 found in the present study reflect feedback input from higher-level areas that possess strong percept-centered coding, or whether some preliminary percept-related information is computed in V1 itself. While these findings and others (Tong 2003) show that even V1 cannot be said to be strictly retinotopic, in our present results, activity in V1 encoded retinal position much more precisely than perceived position (Fig. 2.5), and there was a steady accumulation of percept-centered information at higher stages in the visual processing hierarchy.

Consistent with the general principle that the receptive fields of individual neurons increase in size at successive stages of visual processing (Desimone and

Duncan 1995), we found steadily decreasing precision of physical position coding as we ascended the visual processing hierarchy. Could the superior coding of perceived stimulus position that we measured in higher-level areas simply be a result of coarser retinotopic maps? We would actually expect just the opposite: the capacity for coding perceived position should also be degraded by increasing RF size, unless there are additional sources of information present that are contributing to perceived position coding. Our results show that the precision of retinotopic coding does not reflect an absolute limitation on the capacity of an area for carrying position information. Higher-level areas can clearly support more precise position discrimination than is revealed by measuring retinotopic coding, when measured along the perceived position dimension, which incorporates additional information sources. In this sense, measuring physical position coding (e.g., retinotopy) only taps into a portion of the potential capacity for position information in higher-level areas. This fact is most apparent in Figure 2.5: while there is very little unique retinotopic information in higher-level visual areas, there is substantial information about perceived position.

Higher-level visual areas are, in general, more strongly modulated by attention (Maunsell 2004). Could the pattern of results that we found be due to attentional influences that manifest at higher levels? As highlighted in the introduction, attention is indeed one of the factors that can displace an object's perceived position from its retinal position (Suzuki and Cavanagh 1997), so we would expect attentional influences to contribute to the construction of a percept-centered framework. However, because we used the exact same trials to measure perceived versus physical position coding in our experiment, the influence of attention on the BOLD was the same for each analysis. As with the host of factors that contribute to perceived position, attention could only contribute to a difference between perceived and physical coding in our analysis by virtue of modulating the BOLD in a manner that is correlated with perceived, rather than physical, stimulus position. The combined influence of such factors is exactly what we aimed to measure in higher-level visual areas.

The position discrimination analysis is designed to characterize the nature of position coding *within* a given ROI. For a single ROI, factors such as number of voxels, signal strength, scanner noise, and motion artifacts are held constant across the physical and percept analyses, so the comparison of physical and perceived position coding reflects only the difference in how the GLM predictors were coded. Across ROIs, those factors confound direct precision comparisons. For this reason, we focused on the *relative* precisions of perceived and physical coding within each visual area in order to show differences in the underlying nature of position coding across the visual processing hierarchy. Similarly, on an absolute scale, V1 showed more precise coding of perceived position than did any higher-level area. However, the bootstrapped distribution in Figure 2.5 shows that we would expect strong discrimination of perceived position in V1 based solely on the combination of precise retinotopy and the correlation between subjects' responses and the physical stimulus positions (Fig. 2.1). By considering the relative precision of physical vs. percept coding ($-(Z_{percept} - Z_{physical})$), it becomes clear that coding in V1 is strongly biased toward retinal position, whereas coding in higher-level areas is dominated by percept information.

The cross-correlation analysis that we employed in this study is a simple yet powerful way to measure stimulus discrimination within a region of interest. Recently, a number of studies have demonstrated the promise of multivariate pattern analyses in detecting stimulus discrimination in the BOLD response (Haxby, Gobbini et al. 2001; Carlson, Schrater et al. 2003; Haynes and Rees 2005; Kamitani and Tong 2005; Bressler, Spotswood et al. 2007). While the sensitivity of all of these approaches derives from their multivariate nature, our correlation analysis is particularly well-suited to studying position discrimination because it allows us to track systematic changes in the BOLD response corresponding to incremental, parametric manipulations of the stimulus location. When the correlation analysis yields a significant position discrimination fit within an ROI, it implies not only that different stimulus positions produced detectably different patterns of activity in that ROI, but also that the similarity of the activity patterns was dependent on the similarity of the underlying stimuli. The analysis is sensitive to information encoded in complex coordinate frames, but it requires *topography*, not just different responses to different stimuli.

The perceived location of an object depends on many factors. For example, eye movements, gaze direction, scene and object motion, visual reference frames, and attention all influence perceived position (Matin, Picoult et al. 1982; Ross, Morrone et al. 1997; Suzuki and Cavanagh 1997; Nijhawan 2002; Schlag and Schlag-Rey 2002; Whitney 2002; Dassonville, Bridgeman et al. 2004). Thus, percept-centered mapping in the visual cortex reflects the accumulation of a copious array of retinal and extraretinal information. In this study, we measured the aggregate information about perceived position in each ROI, but we cannot yet say whether such percept information is supported by a single, highly complex percept-based topographic map, or the coexistence of many maps registered to each other but supported by distinct subpopulations of neurons. Complex spatial representations that reflect multiple coordinate frames have been found in parietal and motor cortex (Andersen, Snyder et al. 1997; Graziano and Gross 1998; Graziano 2006). Our results, demonstrating that position coding in higher-level visual areas adapts to reflect an object's perceived position on a trial-by-trial basis, show that a similarly nuanced picture of multiplexed spatial maps is necessary to understand how object position is computed and encoded in the human visual system.

2.5 Supplemental Materials

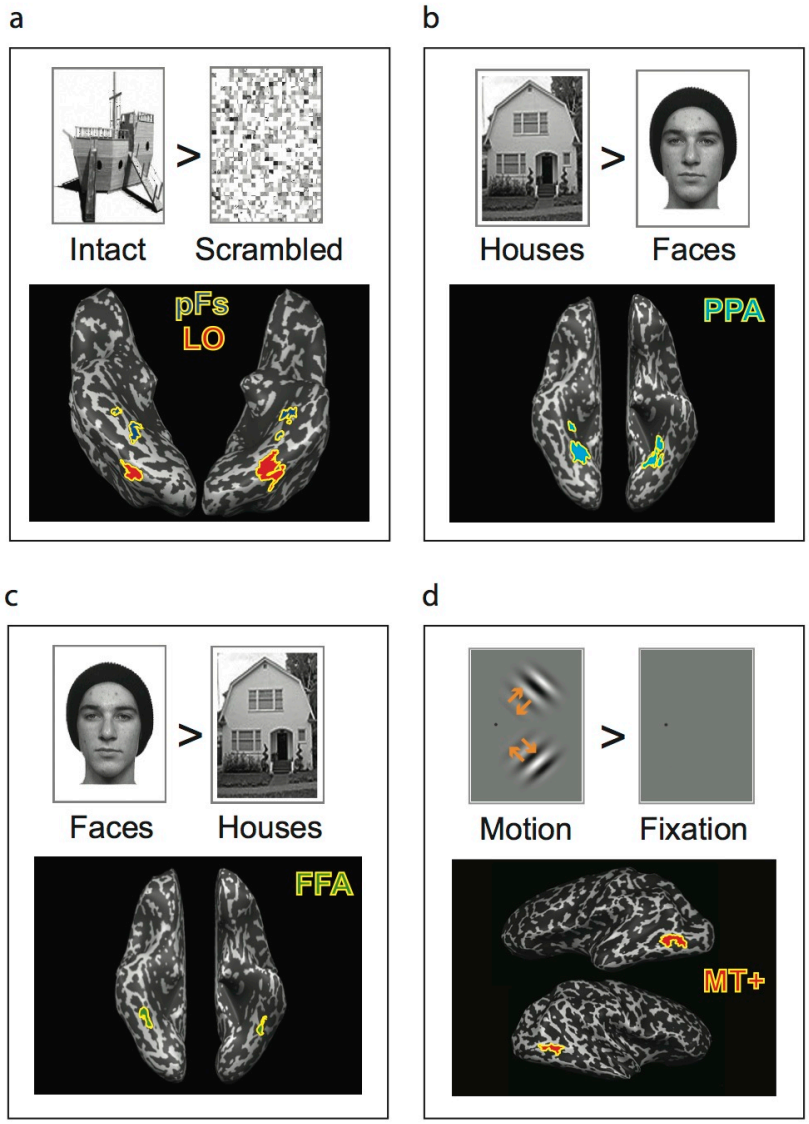

Supplemental Figure 2.1. Localizer stimuli for higher-level visual areas. Localizers for these regions were similar to those used in past studies (Sereno, Dale et al. 1995; Kanwisher, McDermott et al. 1997; Grill-Spector, Kushnir et al. 1998; Epstein, Harris et al. 1999). (**a**) LOC localizer stimuli consisted of gray-scaled intact and scrambled objects superimposed on the fixation point. LO (pictured below the stimuli in red) and pFs (pictured in blue) were defined as the regions most responsive to intact objects versus scrambled (threshold for inclusion was $t = 6.0$, $p < 0.01$). PPA and FFA localizer stimuli consisted of gray-scaled faces and houses. PPA was defined as the region most responsive to houses versus faces (**b**), and FFA was defined as the region showing the strongest response to faces in contrast to houses (panel **c**; $t = 6.0$, $p < 0.01$). (**d**) MT+ was localized by contrasting moving Gabors against a baseline condition in which only the fixation point was present ($t = 6.3$, $p < 0.001$). In each case, a representative ROI from one subject is shown on inflated views of the brain.

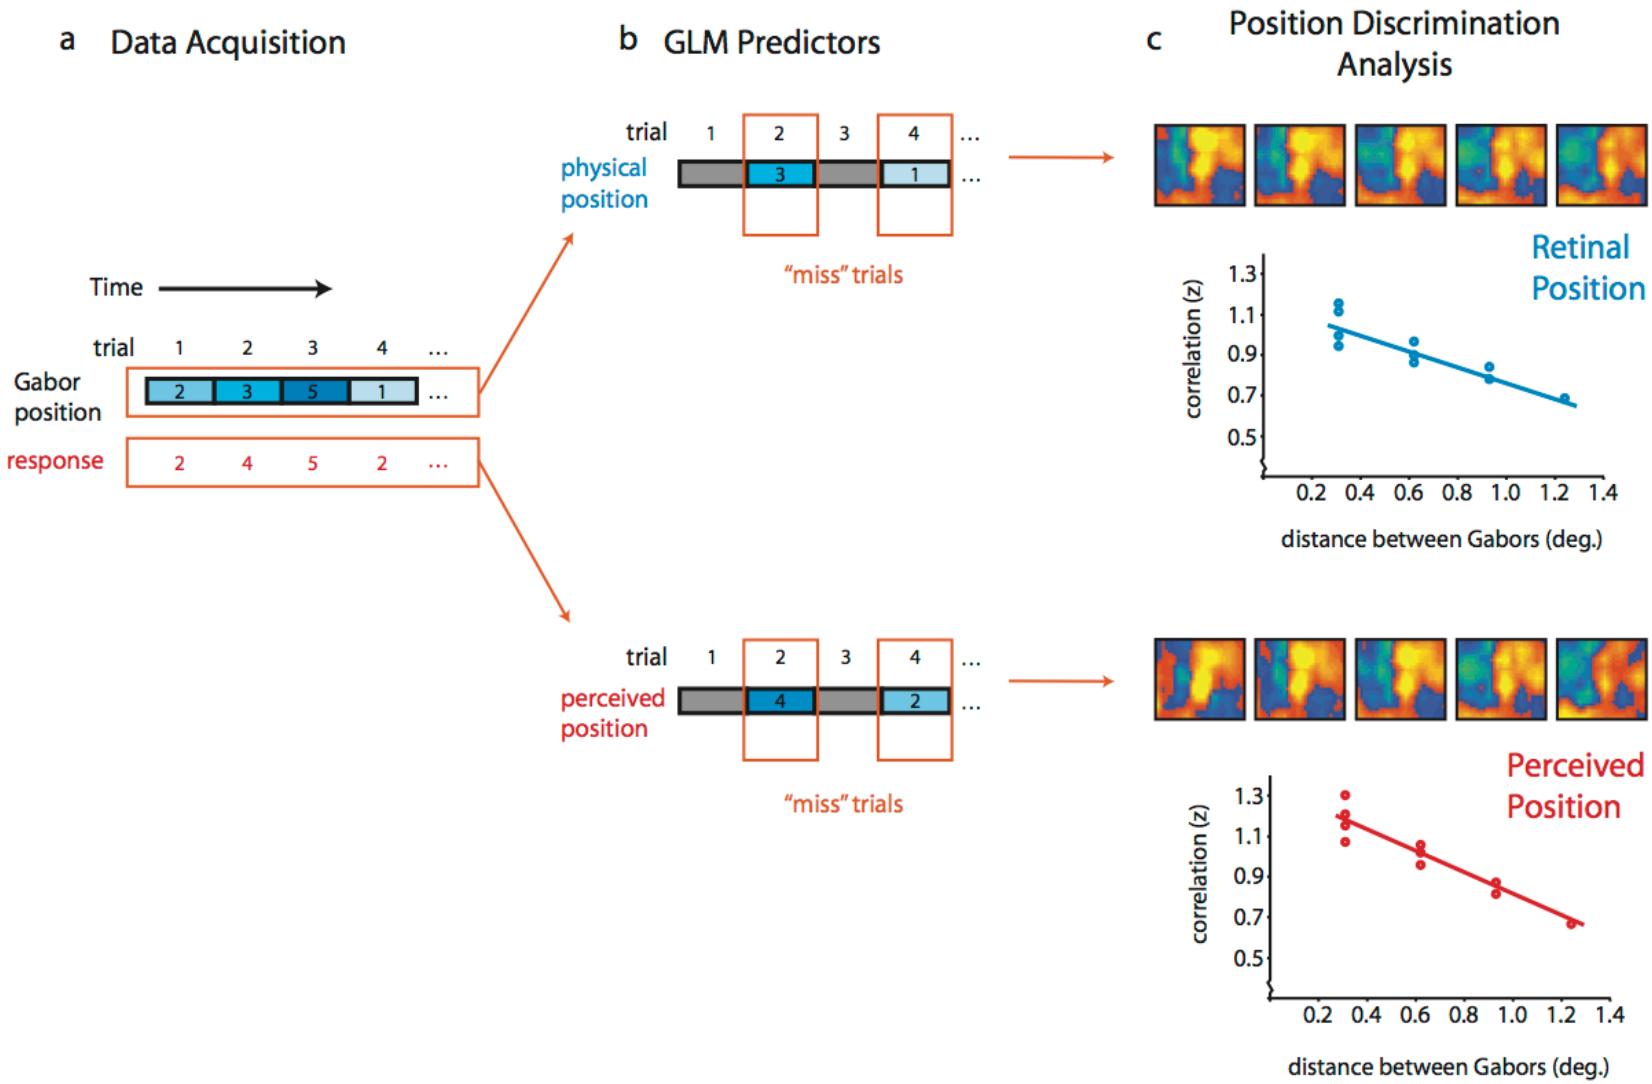
a Data Acquisition
b GLM Predictors
c Position Discrimination Analysis
Time
trial 1 2 3 4 …
Gabor position 2 3 5 1 …
response 2 4 5 2 …
trial 1 2 3 4 …
physical position 3 1 …
"miss" trials
trial 1 2 3 4 …
perceived position 4 2 …
"miss" trials
Retinal Position
correlation (z)
1.3
1.1
0.9
0.7
0.5
0.2 0.4 0.6 0.8 1.0 1.2 1.4
distance between Gabors (deg.)
Perceived Position
correlation (z)
1.3
1.1
0.9
0.7
0.5
0.2 0.4 0.6 0.8 1.0 1.2 1.4
distance between Gabors (deg.)

Supplemental Figure 2.2. Parallel analyses of physical and perceived position discrimination. During scanning, subjects reported the perceived positions of the stimuli on a trial-by-trial basis (shown in red in panel **a**). The data were subsequently analyzed with two general linear models (GLMs), one using the actual stimulus presentations as predictors (panel **b**, top), and the other using subjects' responses as predictors (panel **b**, bottom). We restricted our analysis to the "miss" trials in which subjects misreported the positions of the stimuli, so that every trial was coded differently for the physical and percept GLM analyses. Each GLM generated five maps of BOLD response. In the case of the physical position discrimination analysis, these five maps reflected the activity unique to each stimulus position (panel **c**, top), and in the perceived position analysis, the five maps reflected the activity unique to each of the five possible responses (panel **c**, bottom). These two sets of activity maps were used to create position discrimination plots for retinal and perceived position, respectively (panel **c**; see *Methods*). The visualized maps of BOLD activity here are examples from one ROI, and the differences between the maps are easily visible; however, these differences are not always apparent, hence the value of using a multivariate pattern analysis. Additionally, the multivariate correlation analysis allows for a straightforward quantification of the precision of position discrimination in each ROI, rather than qualitatively tracking the activations.

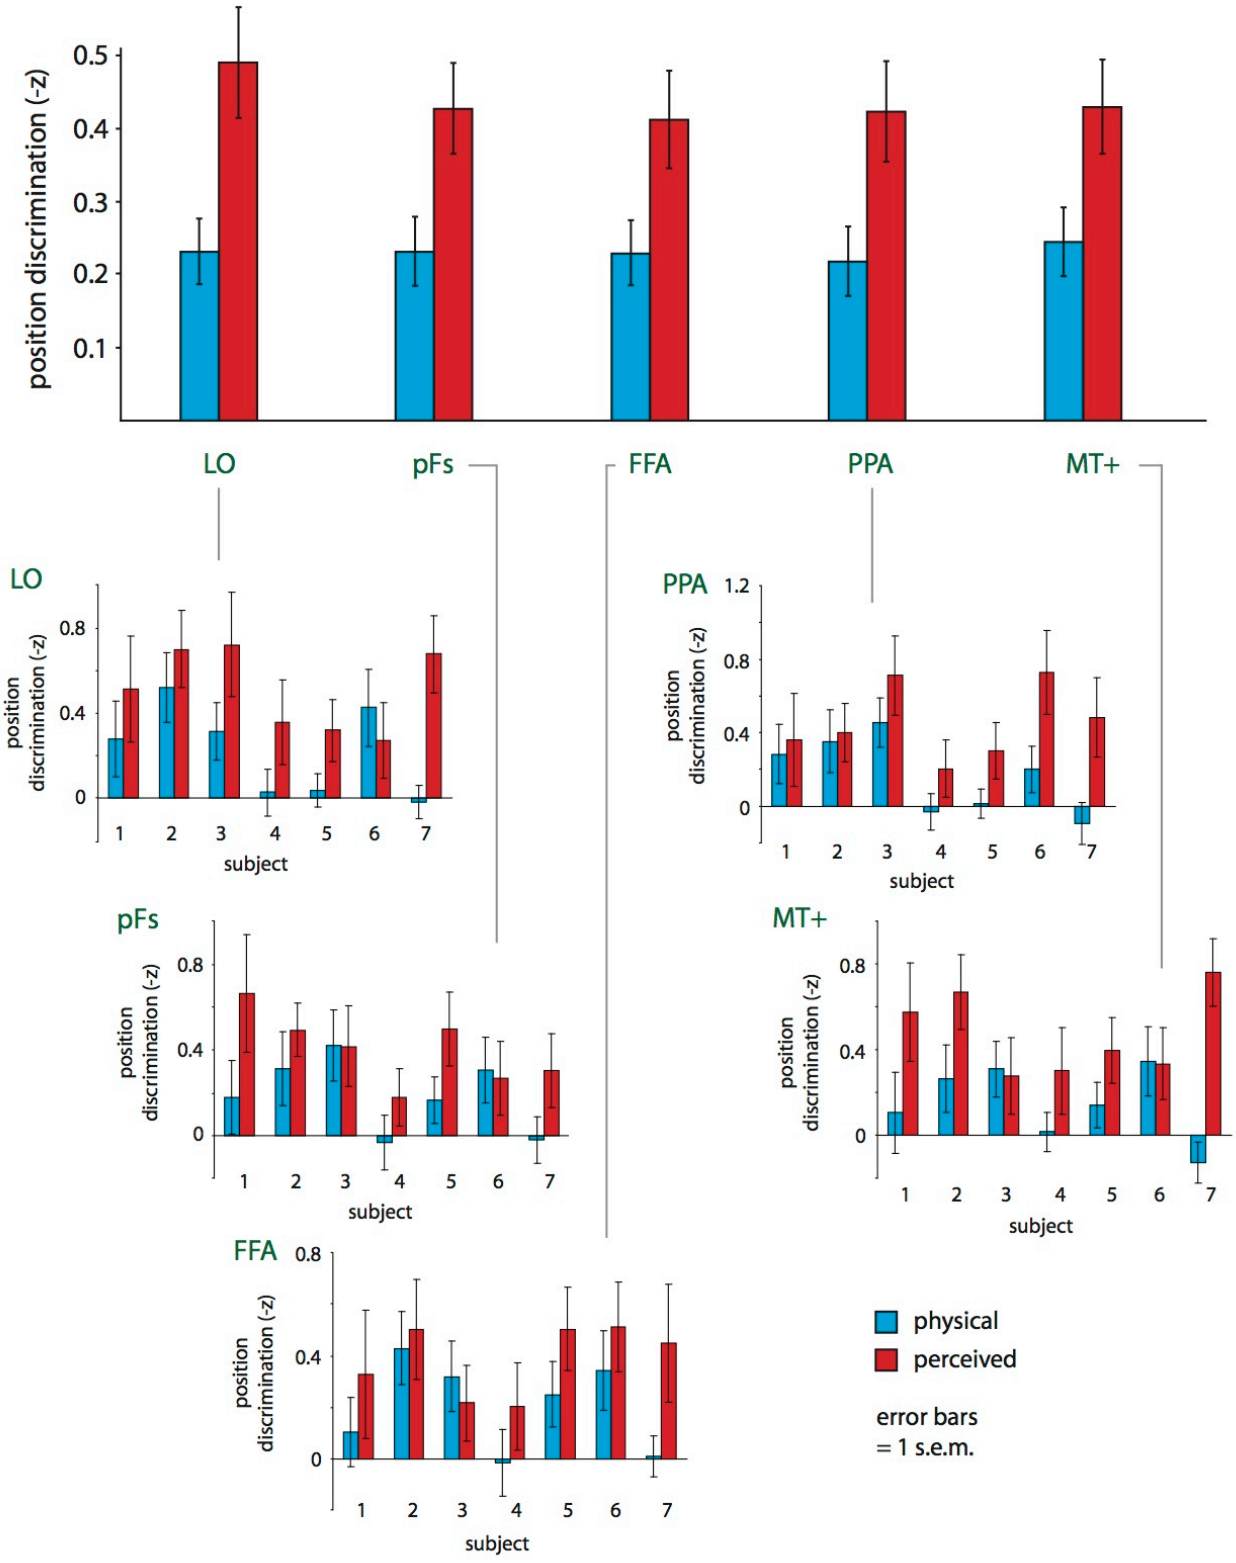
position discrimination (-z)
0.5
0.4
0.3
0.2
0.1
LO
pFs
FFA
PPA
MT+
position discrimination (-z)
0.8
0.4
0
1 2 3 4 5 6 7
subject
1.2
physical
perceived
error bars
= 1 s.e.m.

Supplemental Figure 2.3. Individual subject breakdown of position discrimination in higher-level visual areas. Position discrimination plots were computed separately for each visual area in each subject; individual subjects' physical and perceived position discrimination measurements are collected by visual area to show the data that composes each group-level precision estimate. Discrimination of physical position is shown in blue, and discrimination of perceived position is shown in red; error bars are +/- 1 s.e.m. With few exceptions, individual subjects showed effects in the same direction as the grouped data. Two of the subjects (2 & 5) were markedly less sensitive at discriminating among the stimulus positions than the other subjects (see Fig. 2.1b). Excluding their data from the analysis did not significantly change the group-level physical or perceived position discrimination measurement in any visual area (most significant change was in pFs, Z_{diff} = 1.59, p = 0.11).

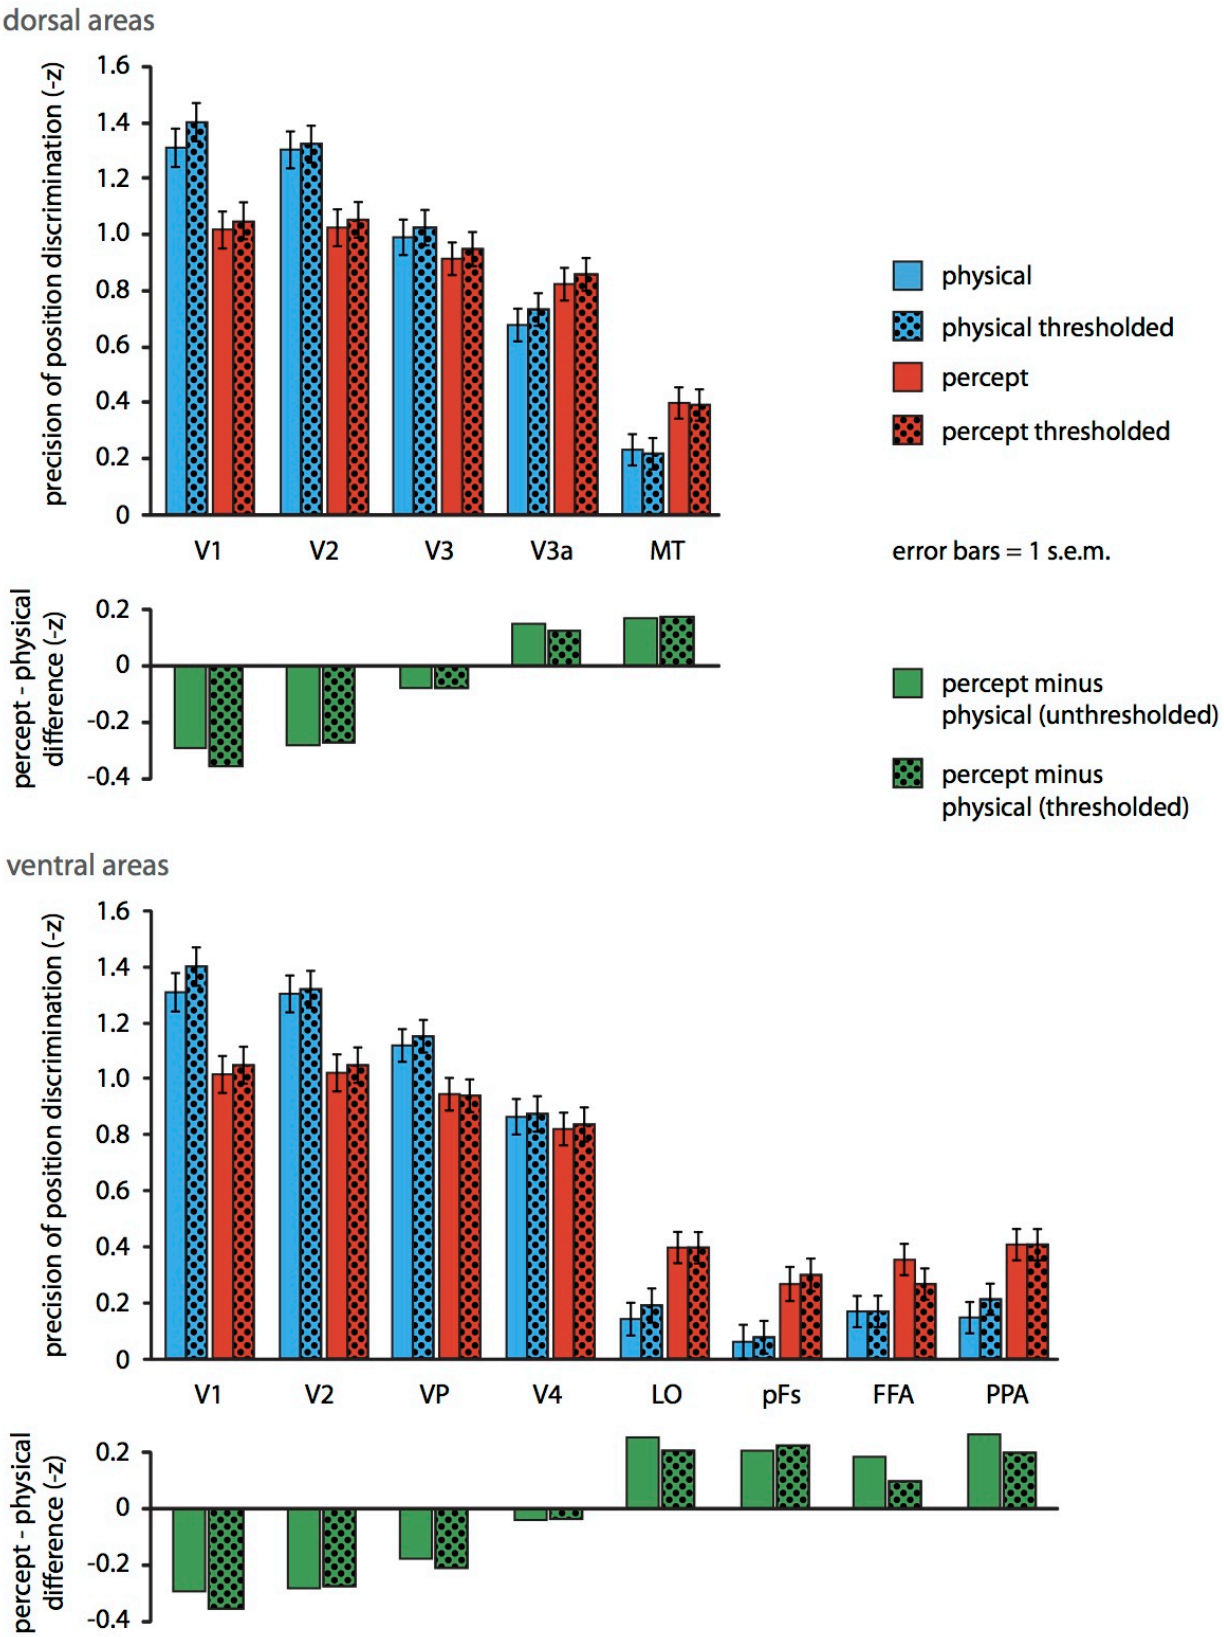
dorsal areas
precision of position discrimination (-z)
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0
V1
V2
V3
V3a
MT
physical
physical thresholded
percept
percept thresholded
error bars = 1 s.e.m.
percept - physical difference (-z)
0.2
0
-0.2
-0.4
percept minus physical (unthresholded)
percept minus physical (thresholded)
ventral areas
precision of position discrimination (-z)
V1
V2
VP
V4
LO
pFs
FFA
PPA
percept - physical difference (-z)

Supplemental Figure 2.4. Comparison of the position discrimination analysis using thresholded vs. unthresholded BOLD response maps. We reran the correlation analysis after thresholding the BOLD maps at $p < 0.05$, which in this case yielded similar but slightly more stringent thresholds than a false discovery rate (FDR) correction. Discrimination scores for the thresholded physical and percept maps are shown in dotted blue and dotted red, respectively, and discrimination scores for the unthresholded maps are shown in solid blue and red (here we used all trials, rather than only missed trials, so these discrimination scores differ slightly from the ones presented in the main experiment). The difference between perceived and physical position discrimination, which serves as an index of a bias in the nature of coding, is plotted in dotted green for the thresholded data and solid green for the unthresholded data. There was no significant difference between the precision of coding as measured with thresholded vs. unthresholded maps for any of the areas we tested (the most significant difference was in V1 for the physical thresholded vs. unthresholded maps: $Z = 1.22$, $p = 0.22$). However, in some higher-level areas thresholding decreased the size of the measured percept-physical difference, and in FFA the difference was no longer significant ($Z = 1.24$, $p = 0.22$). Thresholding yielded an overall decrease in the heterogeneity of the measured coding biases across the ventral areas ($X^2_{unthresholded} = 126.6$, $X^2_{thresholded} = 120.2$; both $p < 0.0001$), but an increase in ventral areas ($X^2_{unthresholded} = 62.0$, $X^2_{thresholded} = 68.3$; both $p < 0.0001$), where raw BOLD activations were generally more robust. Thus, at least in higher-level areas, the correlation analysis appears to be more powerful when using unthresholded BOLD response maps, consistent with the idea that groups of voxels that are not significantly activated in a univariate analysis can nonetheless carry precise stimulus information (Norman, Polyn et al. 2006).

a

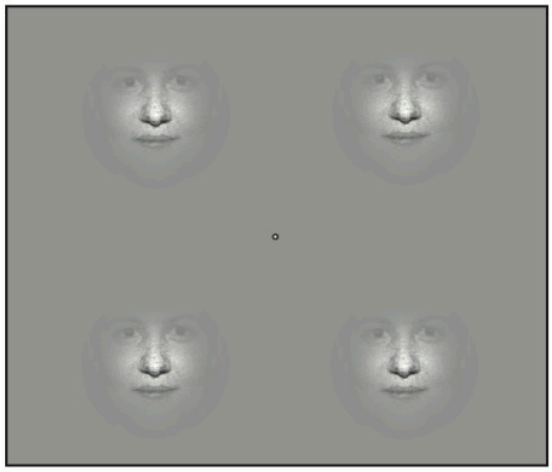

b

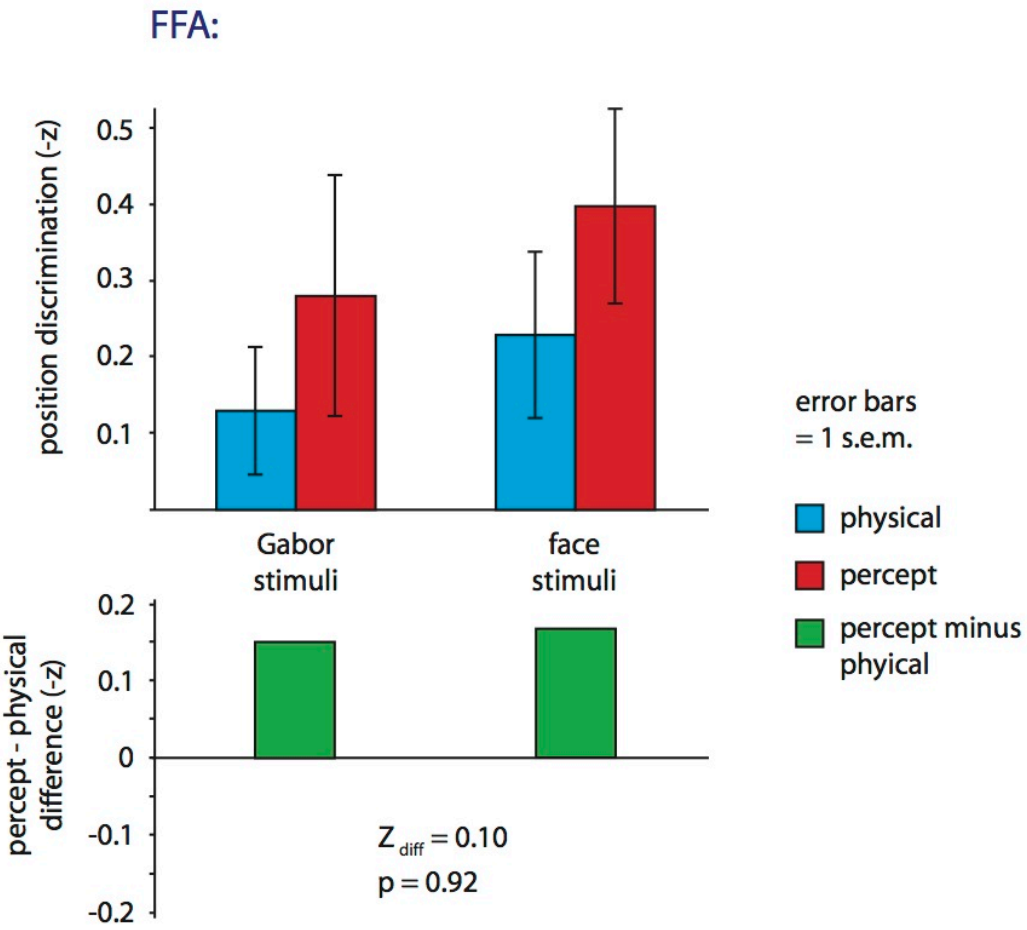
FFA:
position discrimination (-z)
0.5
0.4
0.3
0.2
0.1
Gabor stimuli
face stimuli
error bars = 1 s.e.m.
physical
percept
percept minus phyical
percept - physical difference (-z)
0.2
0.1
0
-0.1
-0.2
Z diff = 0.10
p = 0.92

Supplemental Figure 2.5. Control experiment with face stimuli. To test whether percept-centered position coding is influenced by stimulus preference, we performed a control experiment in which we presented faces instead of Gabors (**a**). The faces were positioned just as the Gabors in the main experiment were, with one in each visual quadrant and at five possible eccentricities. Each face was enveloped in a 1.66 degree standard deviation Gaussian contrast profile to define its centroid, and updated in identity at 7.5 Hz (example faces here have slightly larger contrast envelopes for printing purposes). In all other respects, the stimuli and task were the same as in the main experiment (see *Methods*). (**b**) Comparison of position discrimination measurements made using face stimuli and Gabor stimuli. Data for the two subjects who participated in this control experiment are plotted together. Of interest was whether the nature of position coding that we measured within an area depended on the stimuli used to probe it. To this end, we tested for a difference in the percept vs. physical bias ($-(z_{percept} - z_{physical})$; green bars), measured independently using face stimuli and Gabor stimuli, within the fusiform face area (FFA), for which faces are the optimal stimulus. A Z-test on the two bias estimates revealed no significant difference ($Z = 0.10$, $p = 0.92$). The agreement between our two measurements of the nature of position coding in the FFA, made using optimal and non-optimal stimuli for evoking responses there, suggests that the nature of position coding in a visual area does not depend on the type of stimulus it is representing.

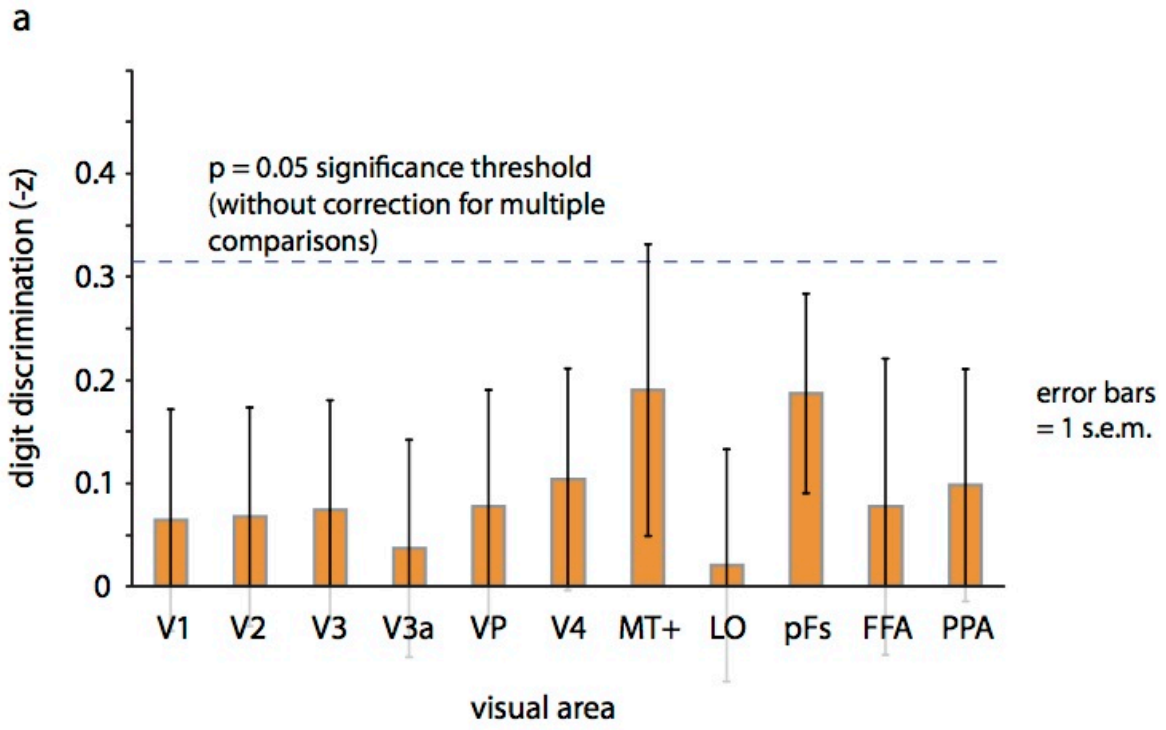
a
digit discrimination (-z)
0.4
0.3
0.2
0.1
0
p = 0.05 significance threshold
(without correction for multiple
comparisons)
error bars
= 1 s.e.m.
V1
V2
V3
V3a
VP
V4
MT+
LO
pFs
FFA
PPA
visual area

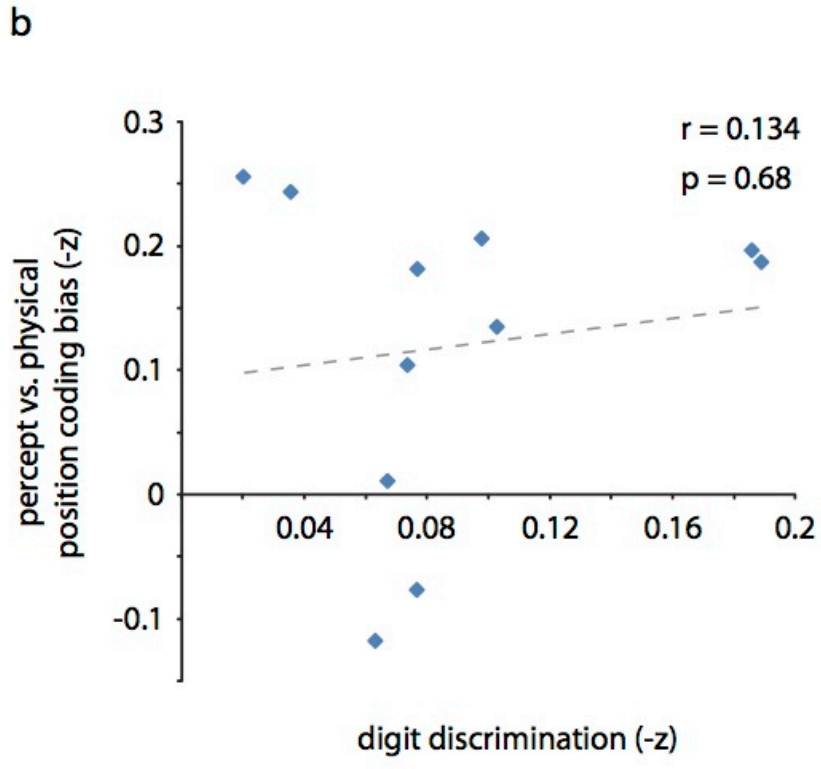
b
percept vs. physical
position coding bias (-z)
0.3
0.2
0.1
0
-0.1
r = 0.134
p = 0.68
0.04
0.08
0.12
0.16
0.2
digit discrimination (-z)

Supplemental Figure 2.6. Investigation of response digit discrimination in visual cortex. To evaluate the possibility that the visual areas we tested also encode information about response planning or tactile input, we applied the same correlation analysis as in the main experiment after recoding the design matrix according to subjects' responses on the secondary (pattern matching) task. Subjects' responses were randomly divided into two groups, such that there were five total predictors: baseline, index.1, index.2, middle.1, and middle.2. Thus, the output of the analysis is a test of the hypothesis that the patterns of activity from trials in which the subject responded with the same digit are better correlated than the patterns of activity from trials in which the subject responded with different digits. The precision of digit coding (-Z) is shown for each area in panel **a**; discrimination was not significant in any ROI (most significant was MT+; -Z = 0.19, p = 0.24). Panel **b** shows the index of bias for physical or percept coding $-(z_{percept} - z_{physical})$ as a function of the precision of digit discrimination (each ROI is a point). Variation in the precision of digit coding was not correlated with the degree of bias for encoding physical or perceived position (r = 0.13, p = 0.68). These results show that the encoding of response planning or tactile input does not explain the percept-dominated coding that we measured in higher-level visual areas.

	Left			Right		
	x	y	z	x	y	z
This study						
LO	-40.5 ± 4.5	-68 ± 3.5	-6.7 ± 5.4	39.8 ± 4.8	-72.8 ± 7.7	-6.7 ± 3.9
pFs	-33.7 ± 2.3	-39 ± 6.2	-13.5 ± 2.9	33.5 ± 5.5	-46.7 ± 9.4	-13.2 ± 2.6
FFA	-37.3 ± 2.9	-45.1 ± 8.8	12.7 ± 4.7	35.9 ± 4.3	-49.7 ± 13	-13.1 ± 3.3
PPA	-24.9 ± 4.5	-44.3 ±8	-6 ± 2	25 ± 3	-40.7 ± 4.3	-6 ± 2
MT+	-43.4 ± 5.0	-68.4 ± 4.0	4.3 ± 6.2	39.5 ± 2.8	-67.3 ± 3.7	5.3 ± 5.6
Previous Studies						
Grill-Spector et al. (1999)						
LO	-36 ± 7	-71 ± 7	-13 ± 5	37 ± 5	-69 ± 7	-10 ± 4
pFs	-38 ± 5	-50 ± 6	-17 ± 5	33 ± 4	-47 ± 6	-14 ± 4
Avidan et al. (2002)						
Experiment 1						
LO	-42 ± 5	-69 ± 8	-6 ± 5	42 ± 4	-60 ± 6	-7 ± 5
pFs	-36 ± 4	-48 ± 10	-19 ± 3	36 ± 5	-45 ± 10	-17 ± 5
Experiment 2						
LO	-44 ± 2	-66 ± 4	-8 ± 5	43 ± 5	-63 ± 5	4 ± 5
pFs	-34 ± 4	-53 ± 11	-19 ± 4	35 ± 2	-51 ± 3	-17 ± 2
Altmann et al. (2004)						
LO	-38.1 ± 9.1	-72.2 ± 7.3	-6.7 ± 5.0	37.9 ± 1.8	-69.4 ± 4.3	-8.1 ± 3.7
pFs	-35.2 ± 6.4	-48.9 ± 9.9	-15.4 ± 2.6	31.5 ± 3.6	-46.8 ± 5.0	-15.8 ± 2.1
Epstein et al. (2003)						
Experiment 1						
PPA	-26	-48	-9	28	-46	-10
Experiment 2						
PPA	-28	-47	-10	29	-43	-11
Yi et al. (2006)						
PPA	-23	-38	-9	20	-38	-8
Gauthier et al. (1999)						
FFA	-40	-46	-12	41	-55	-10
Grill-Spector et al. (2004)						
FFA	-37 ± 4	-42 ± 7	-16 ± 5	39 ± 3	-40 ± 7	-16 ± 5
Watson et al. (1993)						
MT+	–41 ± 5.6	–69 ± 6.0	2 ± 5.3	41 ± 3.7	–67 ± 4.7	2 ± 3.2
Dumoulin et al. (2000)						
MT+	–47 ± 3.8	76 ± 4.9	2 ± 2.7	44 ± 3.3	–67 ± 3.1	0 ± 5.1
Kourtzi & Kanwisher (2000)						
MT+	–45.8 ± 6.3	–67.4 ± 7.1	0.2 ± 6.9	42.9 ± 5.5	–72.5 ± 4.7	4.4 ± 4.3

Supplemental Table 2.1. Talairach coordinates for higher-level visual areas, averaged across seven subjects, ± standard deviation in mm. Our coordinates were consistent with those of previous studies, confirming our ROI selections (Watson, Myers et al. 1993; Gauthier, Tarr et al. 1999; Grill-Spector, Kushnir et al. 1999; Dumoulin, Bittar et al. 2000; Kourtzi and Kanwisher 2000; Avidan, Hasson et al. 2002; Epstein, Graham et al. 2003; Altmann, Deubelius et al. 2004; Grill-Spector, Knouf et al. 2004; Yi, Kelley et al. 2006).

Chapter 3

Attention modulates position tuning of population responses in V1

3.1 Introduction

In the experiments in Chapter 2, attentional demands were held constant – subjects always attended at the locations of the peripheral stimuli in order to perform the localization task. However, attention itself is a factor that can modulate both perceived position and visual resolution (Suzuki and Cavanagh 1997; Yeshurun and Carrasco 1998). The experiments in Chapter 3 ask whether changes in the locus of visual attention produce changes in the spatial mapping in early visual cortex.

Psychophysical research has shown that spatially-directed attention can improve the resolution of visual perception, even on basic visual tasks (Balz and Hock 1997; Yeshurun and Carrasco 1998; Yeshurun and Carrasco 1999; Yeshurun, Montagna et al. 2008). When attentional allocation is stimulus-driven (transient), this improvement in spatial resolution happens even when it comes at a cost to performance. Yeshurun and Carrasco (1998) used a texture segregation task, in which performance suffers when visual resolution is too fine, to show that when transient attention was cued to targets near the fovea, it actually impaired performance by sharpening spatial resolution. However, the same group showed that sustained attention, which is deliberately directed to the stimulus position rather than automatically drawn to the stimulus, improves performance at all eccentricities on the same task by optimally adjusting the scale of visual resolution (Yeshurun, Montagna et al. 2008). These studies demonstrated that while both transient and sustained attention can afford substantial improvements in visual resolution, sustained attention uniquely extends and optimizes the dynamic range of spatial resolution, allowing humans to perceive either fine-scale detail or coarse-scale texture and surface properties.

What is the mechanism of sustained attention that mediates this ability to modulate the spatial resolution of position coding and perception? One contributor is response gain. When attention is directed to a location in the visual scene, cortical cells in the corresponding retinotopic positions fire more strongly in response to their preferred stimuli (Haenny and Schiller 1988; Spitzer, Desimone et al. 1988; Motter 1993; McAdams and Maunsell 1999), and even increase their baseline activity in the absence of visual stimulation (Luck, Chelazzi et al. 1997). This multiplicative signal gain occurs with greater magnitude at later stages in the visual processing stream (Treue 2001; Maunsell 2004). Spatially-directed attention produces a robust signal gain in the fMRI BOLD response as well, both with (Tootell, Hadjikhani et al. 1998; Brefczynski and

DeYoe 1999; Gandhi, Heeger et al. 1999; Somers, Dale et al. 1999), and without (Kastner, Pinsk et al. 1999) visual stimulation.

In addition to boosting the amplitude of neuronal responses, attention might also increase a cell's stimulus specificity by narrowing its tuning function around the optimal stimulus for a given attribute. This question has been investigated in the featural domain (for example, color, orientation, and spatial frequency), as well as in the spatial domain. In both cases, there are conflicting reports on whether attention narrows neural tuning properties. The earliest studies investigating attentional tuning in the featural domain reported narrowing of orientation and color tuning in single units in V4 (Haenny and Schiller 1988; Spitzer, Desimone et al. 1988). These studies used task difficulty and reward expectation to manipulate attention, whereas a subsequent study by McAdams and Maunsell (McAdams and Maunsell 1999), in which attention was directed either inside or outside the receptive field of the recorded neuron, found that the effect of attention on orientation tuning in V4 was purely multiplicative.

At the neural population level, Murray and Wojciulik (2004) reported that attention sharpens object orientation selectivity in human lateral occipital complex (LOC), as measured in the fMRI BOLD response. While this might also be due to changes in the adaptability of neurons or gain in a more selective subset of cells (Boynton 2004), this result hints at the possibility that tuning narrowing manifests itself only at the population level, perhaps via a mechanism such as feature similarity-dependent gain (that is, a multiplicative gain which depends on the similarity of the attended feature and the cell's preferred feature; (Treue and Martinez Trujillo 1999). Consistent with this, while Treue and Martinez-Trujillo (1999) did not find a change in the direction selectivity of individual MT cells with attention, they found that feature-based attention increased direction selectivity at the population level (Martinez-Trujillo and Treue 2004).

In the spatial domain, there is general agreement that attention can increase position selectivity of cells in later visual areas; the debate now revolves around how early these changes in position selectivity occur in the visual processing stream, and how such changes in position tuning come about. In the earliest related study, recording from V4, Moran and Desimone (1985) showed that when an attended and an unattended stimulus both fall within a cell's receptive field, the neuronal response is primarily determined by the attended stimulus; responses to the unattended stimulus are attenuated. Similarly, Luck et al. (1997) showed that in V2 and V4, when preferred and non-preferred stimuli are simultaneously present within a cell's receptive field, the cell's responses are larger when attention is directed to the preferred stimulus. In these experiments, the recorded receptive fields effectively shrunk around the attended stimuli, resulting in increased position selectivity. However, because such experiments have not been feasible in V1 due to prohibitively small receptive fields, the possibility remains that attention affects only signal gain in V1, which then feeds forward in a fashion that gives rise to position selectivity changes downstream (Maunsell and McAdams 2000). It is this unresolved issue that we address in the present study.

Related to the question of narrowed position tuning is the contrast gain model of attention (Reynolds and Desimone 1999; Reynolds, Pasternak et al. 2000). According to the contrast gain model, a neural response at the attended location behaves as if the stimulus contrast has been increased, which amounts to a leftward shift in the cell's contrast response function. This hypothesis is supported by the fact that transient covert attention improves observer contrast thresholds (Cameron, Tai et al. 2002), and increases perceived stimulus contrast (Carrasco, Ling et al. 2004). Importantly, increasing stimulus contrast has also been shown to shrink the spatial summation areas of cells in V1 corresponding to the attended location (Sceniak, Ringach et al. 1999). There is currently an active debate as to whether the contrast gain model or a multiplicative response gain model better explains the effects of attention on neural responses (Williford and Maunsell 2006); of the two, only the contrast gain model predicts that attention will narrow the position tuning properties of cells in V1.

A recent study by Roberts et al. (2007) took aim at a question similar to ours by measuring attention-related changes in the length tuning of V1 cells. They found that attention reduced preferred stimulus length when the attended stimulus was near the fovea (~2-3 degrees eccentricity), but *increased* preferred length at more peripheral locations. Their results are apparently at odds with the fact that attention can dramatically improve visual resolution at peripheral locations (Yeshurun and Carrasco 1998; Yeshurun, Montagna et al. 2008). Given this discrepancy, it is critical to study whether attention improves population position selectivity in V1, because spatial perception ultimately depends on population coding (Pouget, Dayan et al. 2000; Purushothaman and Bradley 2005).

The goal of the present study was to test whether sustained attention narrows population tuning for position in the earliest visual areas (V1 through V4), and to relate increased position selectivity to improved precision of retinotopic encoding in these areas. To address these questions, we used a multivariate pattern analysis in which we cross-correlated the patterns of BOLD activity produced by attended and unattended stimuli to measure how the precision of retinotopic coding in V1 through V4 improves with attention.

3.2 Results

To test whether spatially-directed attention narrows population position tuning, we parametrically varied the position of a stimulus in an fMRI experiment and measured the degree of overlap in the neural activity patterns produced by adjacent stimuli. The stimuli were four flickering (7.5 Hz) Gabor patches, one in each visual quadrant. In five conditions, the Gabors were positioned at five different eccentricities ranging from 8.43 to 9.65 degrees from fixation (Fig. 3.1A). These five conditions, and a sixth fixation baseline condition, were presented in randomly ordered 10 second blocks during each 6 minute scanning run. In separate runs, subjects were instructed to attend to the Gabors to determine their eccentricity (*attended* trials) or to ignore the Gabors and attend to fixation for a counting task (*unattended* trials; Fig. 3.1B). During *attended* trials, subjects reported the Gabors' eccentricity in a five alternative forced choice (5AFC) task (Fig.

3.1C) while maintaining attention at the Gabors' positions to detect two small patterns presented randomly during each block. Subjects reported whether the two patterns were the same or different at the end of each block (group performance was 86.6 % correct +/- 3% sd). During *unattended* trials, subjects counted the occurrences of two different textures that appeared around the fixation point and reported which appeared most often during the block (group performance was 76.4 % correct +/- 4% sd). Besides the task instructions, the stimuli were identical for *attended* and *unattended* trials.

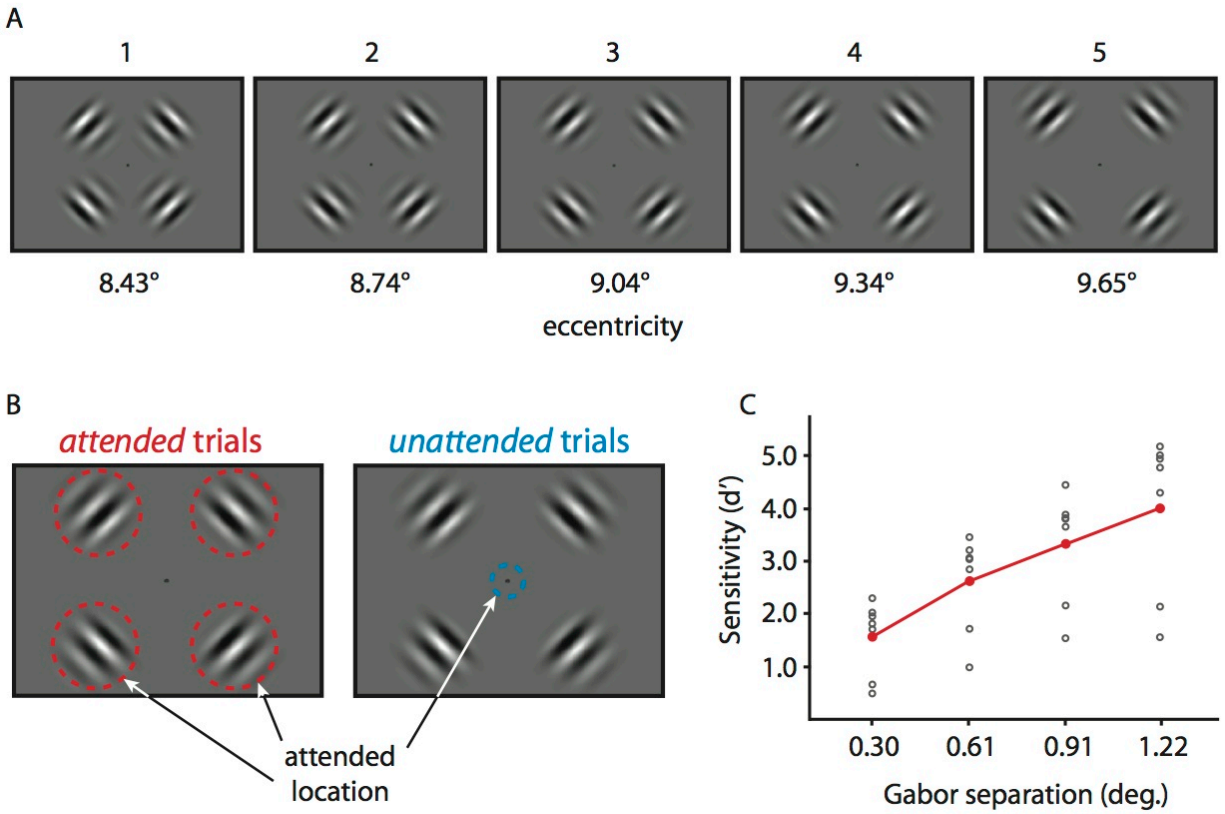

Figure 3.1: Experimental stimuli and psychophysical performance. A) On each trial, four Gabor patches (0.38 cycles/deg, with 1.66 deg. sd Gaussian envelopes) were presented at 8.430, 8.735, 9.039, 9.343, or 9.647 degrees from fixation. The Gabors flickered in counterphase at 7.5 Hz for the duration of each 10s block. **B)** In separate trials, subjects either attended to the fixation point to perform a counting task, or attended to the surrounding Gabors to judge their eccentricity and perform a texture matching task (see Supplemental Experimental Procedures for details). The same stimuli were presented on both types of trials – the only difference between the two was the locus of attention. Keeping the loci of attention in *attended* and *unattended* trials distinct, along with preventing subjects from achieving ceiling performance on the eccentricity judgment task, motivated the peripheral presentation of the Gabor stimuli. **C)** On *attended* trials, subjects made a 5AFC discrimination on the eccentricity of the Gabors. Sensitivity (d') is plotted against the spatial separation between Gabor centroids (individual subject data is plotted in gray, averaged group data is plotted in red). The positive slope indicates that while subjects often mistook the presented eccentricity for an adjacent one, they rarely mistook it for an eccentricity that was three or four increments away. Thus, our stimuli sampled in the dynamic range of subjects' psychophysical discrimination.

Psychophysical studies indicate that spatially-directed attention can improve visual resolution to facilitate the performance of fine-scaled tasks (Yeshurun, Montagna et al. 2008). To complement these behavioral results, we tested for improved spatial coding in early visual areas with attention, using a cross-correlational pattern analysis (Bressler, Spotswood et al. 2007; Fischer and Whitney 2007). For each subject, we generated five maps of BOLD response by contrasting the five stimulus positions with the baseline. Within independently localized ROIs for areas V1- V4, we cross-correlated the five BOLD maps, obtaining a correlation for each possible pairing of maps. We then plotted these correlations as a function of the physical separation between the stimuli that yielded each pair of BOLD maps, producing a position discrimination plot (Fig. 3.2; see Supplemental Experimental Procedures for details). A negative slope on the position discrimination plot indicates that stimuli produce more distinct loci of activation in the ROI as they move further apart in the visual field. Figure 3.2 shows data from all subjects, plotted separately for *attended* and *unattended* runs. In every visual area, attention significantly improved the precision of spatial coding, yielding a steeper slope on the position discrimination plot ($Z = 6.07$ for V1, $Z = 5.54$ for V2, $Z = 5.36$ for V3, $Z = 6.92$ for V4; all $p < 0.001$). Thus, consistent with psychophysical results (Balz and Hock 1997; Yeshurun and Carrasco 1998; Yeshurun and Carrasco 1999; Yeshurun, Montagna et al. 2008), attention improved the precision of retinotopic encoding in the cortex by representing fine position differences with more substantial changes in the pattern of BOLD response.

The improved precision of position coding with attention could result from signal gain alone, which improves the SNR in the BOLD response, or from a combination of signal gain and narrowed population position tuning, which would produce more spatially distinct activity regions for adjacent stimuli. To resolve this, we examined correlations at the largest (1.22 deg.) separation on the position discrimination plots. While the SNR improvement afforded by a BOLD gain would increase the measured correlations, position tuning narrowing would work conversely, decreasing the correlations by reducing the overlap in the regions of BOLD response stimuli produced. This expected pattern holds across the different types of gain that have been distinguished in the attention literature (see Supplemental Figure 3.1); here, *signal gain* refers to a multiplicative vertical scaling of the BOLD response profile. A decrease in the measured correlations with attention could only be due to the influence of narrowed position tuning, and such a decrease would be most prominent at the largest stimulus separation, where the corresponding regions of activity are already most distinct.

Figure 3.3A compares the correlations, at 1.22 deg. separation, for the *attended* and *unattended* trials in areas V1 and V2. In every subject, attending to the Gabor patches significantly reduced the correlation between their resulting patterns of activity (paired t-tests; $t_{(6)} = 5.04$, $p < 0.001$ for V1; $t_{(6)} = 5.69$, $p < 0.001$ for V2). Because signal gain cannot be responsible for these reduced correlations, these data show that attention reduced the spatial spread of the BOLD response around the attended stimuli.

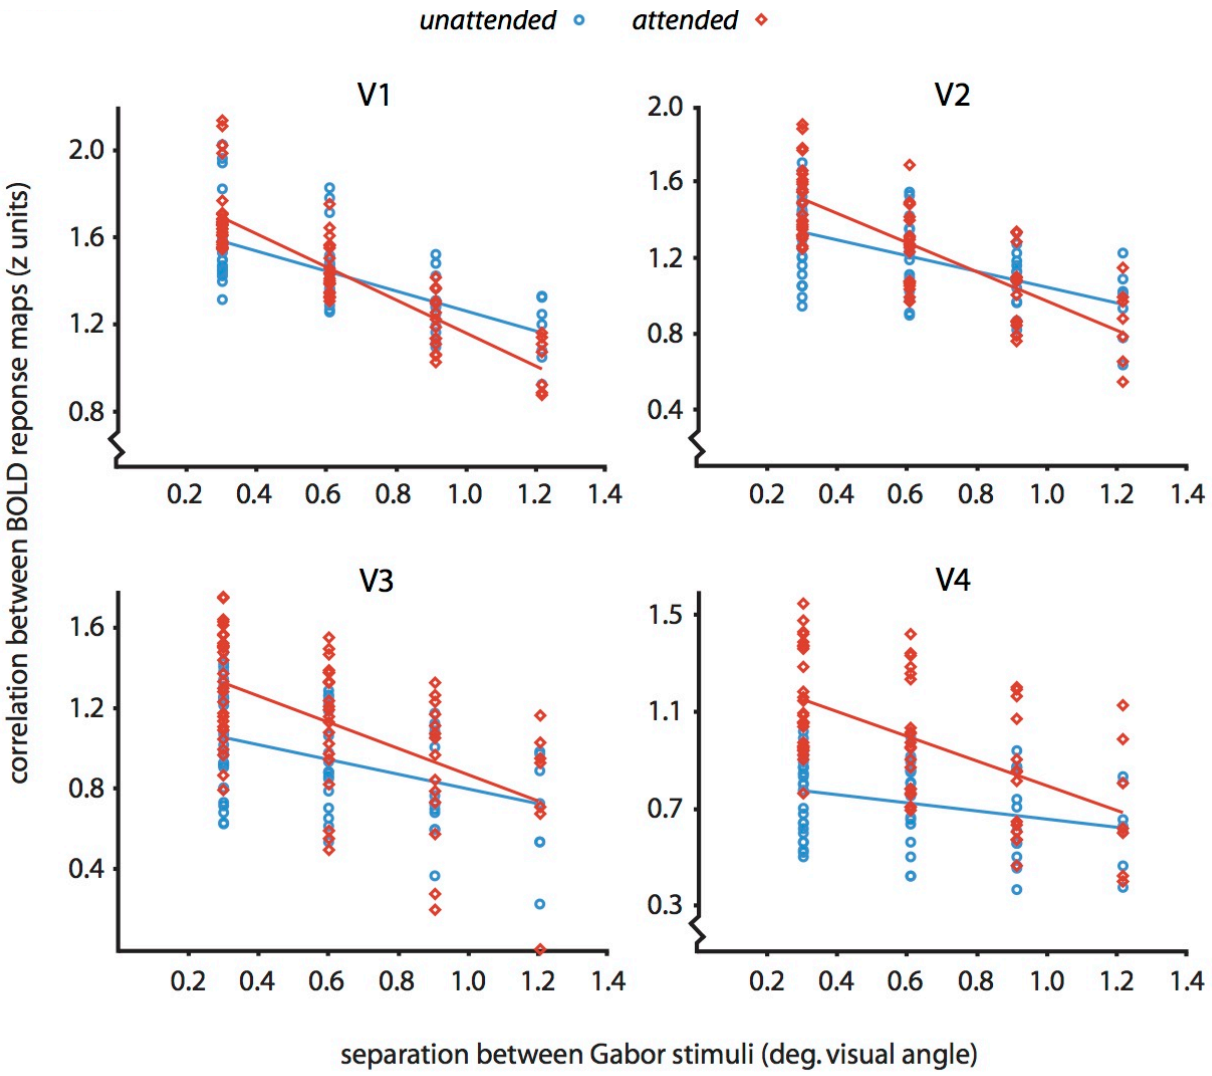

Figure 3.2: Position discrimination plots for V1 through V4. Each plot shows the relationship between the retinal proximity of a pair of stimulus conditions (x axis) and the correlation between their resulting patterns of BOLD response (y axis). A negative linear trend in the correlations reflects the fact that stimuli presented in close proximity to each other produced highly overlapping regions of BOLD response within the ROI being tested, whereas stimuli presented at more distant locations produced more distinct patterns of activation. We analyzed the unthresholded BOLD response within each ROI in order to include position information carried by nonsignificant and negative voxels (Norman, Polyn et al. 2006; Bressler, Spotswood et al. 2007). We averaged each subject's correlations across runs to produce ten Fisher z scores per subject, and plotted all subjects' points together to produce group position discrimination plots for areas V1 through V4. Attended trials are shown in red, and unattended trials are shown in blue. We fit separate linear regressions to the attended and unattended data to quantify the precision of position coding; each regression included a random effects variable to account for plotting multiple subjects together (see Supplemental Experimental Procedures). The linear fit for the attended data was significantly steeper than that for the unattended data in each visual area (the least significant difference was for area V3, $Z = 5.36$, $p < 0.001$), indicating that attention improved the precision of retinotopic encoding in each of V1 through V4.

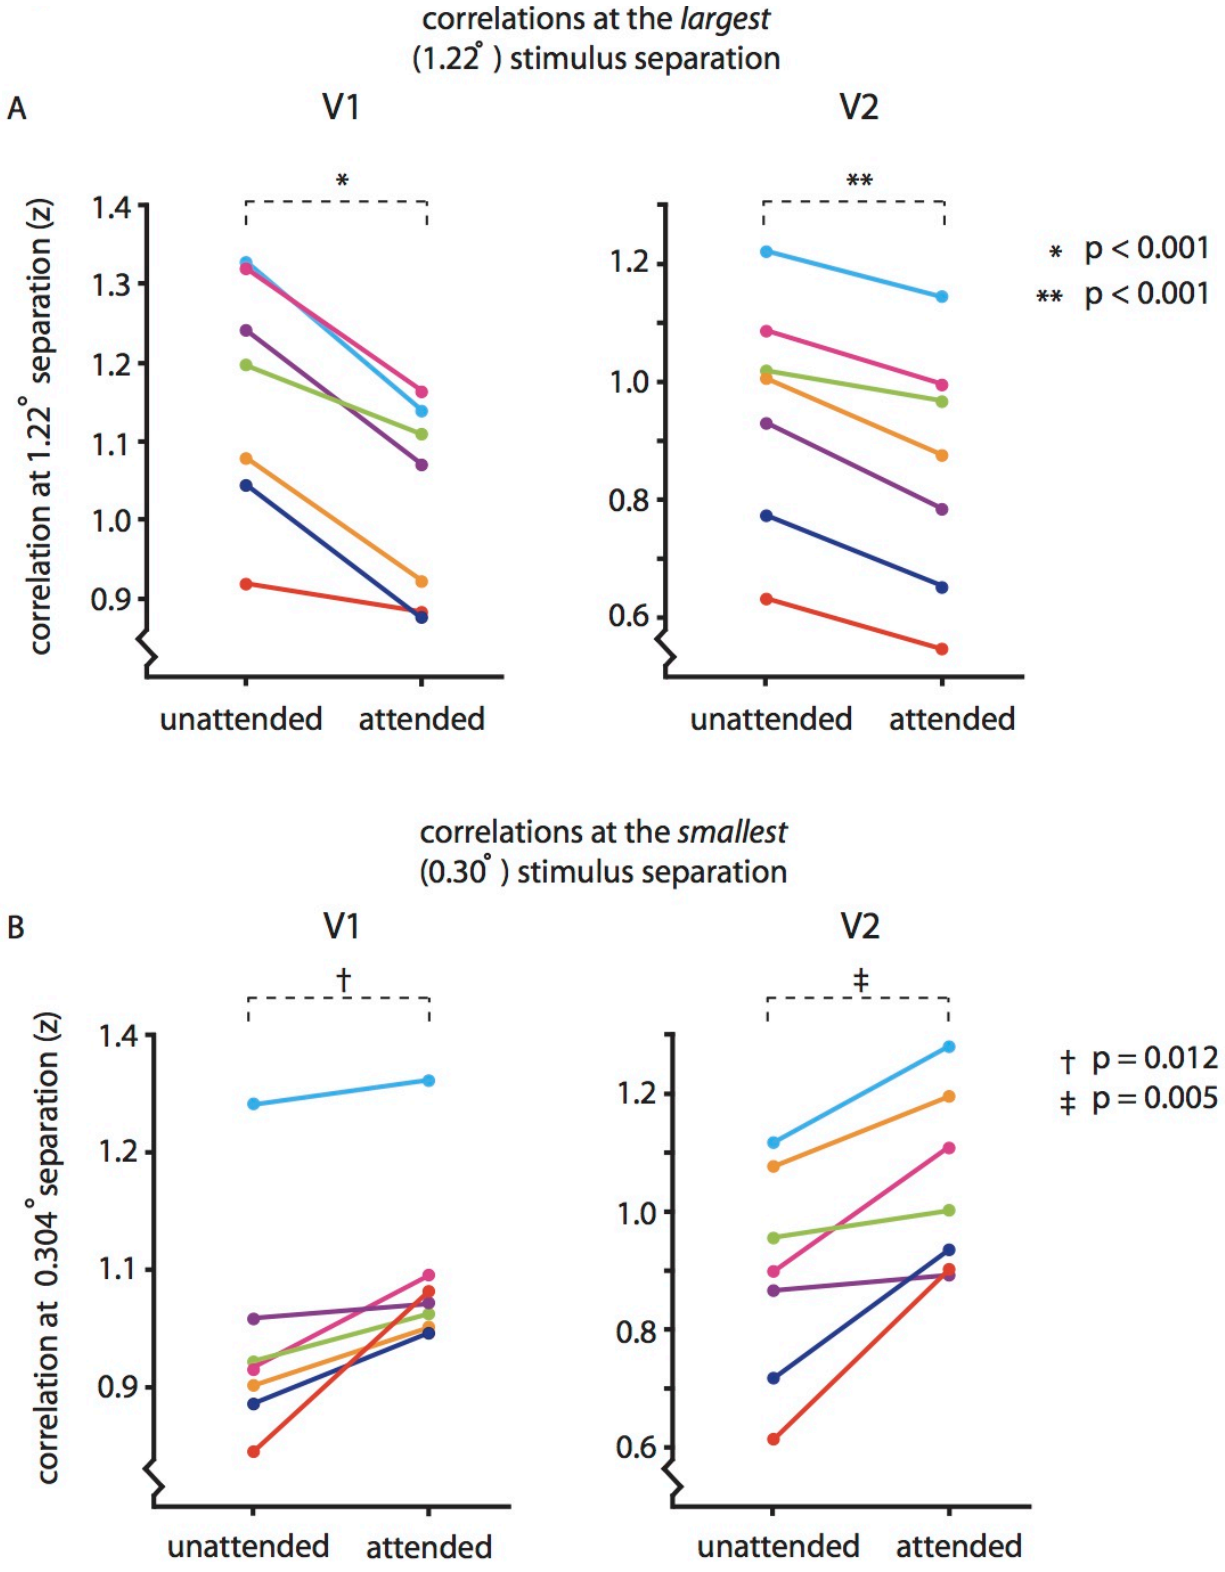
correlations at the *largest*
(1.22°) stimulus separation
A
V1
V2
*
**
correlation at 1.22° separation (z)
1.4
1.3
1.2
1.1
1.0
0.9
0.8
0.6
unattended
attended
* p < 0.001
** p < 0.001
correlations at the *smallest*
(0.30°) stimulus separation
B
V1
V2
†
‡
correlation at 0.304° separation (z)
† p = 0.012
‡ p = 0.005

Figure 3.3: Attention reduced the overlap between adjacent patterns of activity, while increasing the peak BOLD response. A) To evaluate whether attention narrowed position tuning in the BOLD response, we examined the correlation between the patterns of activity produced by the most foveal and most eccentric stimuli. Each subject's z scores for the 1.22 degree stimulus separation are shown separately for attended and unattended conditions (each color is one subject). In both V1 and V2, the correlations were significantly lower for the attended stimuli, indicating a reduction in the spatial spread of the BOLD response with attention (paired t test; $t_{(6)} = 5.04$, $p < 0.001$ for V1; $t_{(6)} = 5.69$, $p < 0.001$ for V2). **B)** To test for an increased peak BOLD response with attention, we examined the correlations at the smallest stimulus separation (.304 degrees). As in panel **A**, subjects' z scores are plotted separately for attended and unattended conditions in V1 and V2. For every subject, the correlation increased with attention in both V1 and V2, indicating a robust signal gain with attention. These increases were significant in both areas (paired t-tests; $t_{(6)} = 3.57$, $p = 0.012$ for V1; $t_{(6)} = 4.32$, $p = 0.005$ for V2).

We made the same comparison for the stimuli at the two most foveal positions, separated by only 0.304 deg. If attention increased the peak amplitude of the BOLD response in addition to narrowing its spatial spread, we would expect increased correlations with attention for very closely spaced stimuli due to an increase in the signal, relative to noise (in the extreme case of stimuli presented in the same position, the correlation between the two patterns of response would track SNR independently of position tuning narrowing). In fact, at this smallest stimulus separation, the correlations for all subjects increased significantly with attention (Fig. 3.3B; paired t-tests; $t_{(6)} = 3.57$, $p = 0.012$ for V1; $t_{(6)} = 4.32$, $p = 0.005$ for V2). We independently verified that attention increased the peak BOLD response by comparing the maximum responses in V1 and V2 in the attended and unattended conditions (paired t-tests; $t_{(6)} = 3.23$, $p = 0.018$ for V1; $t_{(6)} = 3.90$, $p = 0.008$ for V2). This confirms that correlations can simultaneously reveal both signal gain and position tuning narrowing in the BOLD response.

The decrease in correlation at the largest stimulus separation found in V1 and V2 was not found in V3 or V4. However, the regression lines for the *attended* and *unattended* data intersect well above the horizontal axis in the plots for both areas (at (1.27, 0.70) for V3, and (1.40, 0.59) for V4), suggesting that measuring at larger stimulus separations would likely reveal the same decrease in correlation with attention as in V1 and V2. Coarser retinotopy in V3 and V4 compared with V1 and V2, and a stronger gain effect of attention at later stages in the cortical hierarchy (Treue 2001; Maunsell 2004), likely underlie the need to sample at larger stimulus separations in V3 and V4.

To better understand how attention-related tuning and gain are reflected in the position discrimination plots, we simulated the correlation analysis with a modeled BOLD response profile (Fig. 3.4A). Figure 3.4B shows a green baseline curve reflecting correlation as a function of separation between patterns of BOLD, along with curves simulating a 10% signal gain (red) and a 10% narrowing of population position tuning (blue). Both the red and blue curves are steeper than the baseline curve, indicating more precise position coding; importantly, signal gain only increased correlations, whereas narrowed position tuning only decreased correlations. This pattern reflects the

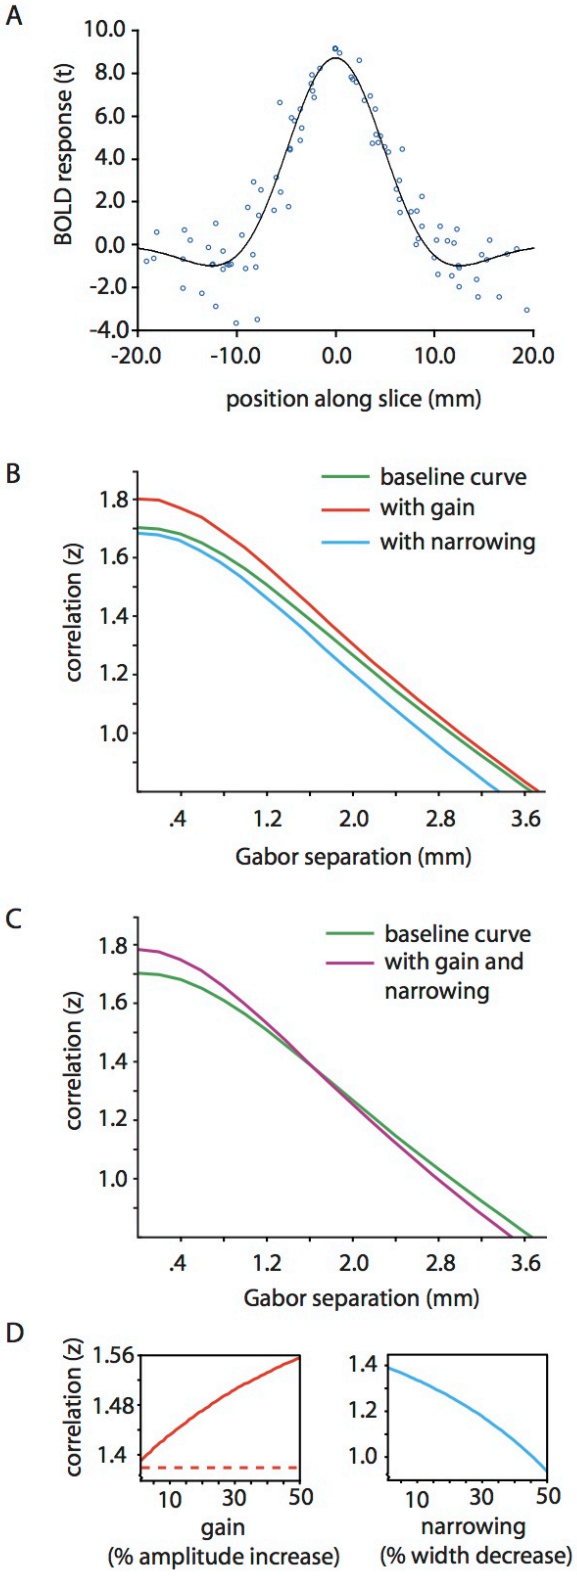
A
10.0
8.0
6.0
4.0
2.0
0.0
-2.0
-4.0
BOLD response (t)
-20.0
-10.0
0.0
10.0
20.0
position along slice (mm)
B
baseline curve
with gain
with narrowing
1.8
1.6
1.4
1.2
1.0
correlation (z)
.4
1.2
2.0
2.8
3.6
Gabor separation (mm)
C
baseline curve
with gain and
narrowing
1.8
1.6
1.4
1.2
1.0
correlation (z)
.4
1.2
2.0
2.8
3.6
Gabor separation (mm)
D
correlation (z)
1.56
1.48
1.4
10
30
50
gain
(% amplitude increase)
1.4
1.2
1.0
10
30
50
narrowing
(% width decrease)

Figure 3.4: Modeling the effects of signal gain and position tuning narrowing. A) We obtained a BOLD response curve for use in subsequent simulations by measuring the BOLD signal along slices on each subject's inflated cortex (see Supplemental Experimental Procedures). Data from all subjects were plotted together and normalized by peak amplitude to generate a prototypical BOLD response profile for use in modeling. These data were fit with a difference-of-Gaussians function (black curve), and 200 points were sampled evenly along this curve for use in subsequent simulations. **B)** Using the resulting BOLD response curve, we produced a simulated position discrimination plot (like those in Fig. 3.2), sampling on a continuous range of stimulus separations (green curve). To produce this plot, in a Monte Carlo simulation, we added realistic noise (SNR = 15.33; see Supplemental Experimental Procedures) to two copies of the prototype BOLD response curve, and varied the offset between them, computing the correlation at each separation. To model the effects of gain and narrowing on the position discrimination plot, we repeated the same Monte Carlo simulation after adding a 10% amplitude gain (red curve) and, separately, a 10% width narrowing (blue curve) to the BOLD response profiles. The application of signal gain increased the correlation between the simulated BOLD response patterns at every separation along the abscissa, while narrowing the BOLD response profiles decreased the correlation at every separation; both gain and tuning narrowing made the slope of the curve steeper. **C**) Simultaneously applying both gain and narrowing to the simulated BOLD response profiles produced the purple curve. This curve crosses the baseline (green) curve in the same fashion that the attended and unattended regression lines cross in the position discrimination data presented in Figure 3.2. **D)** To model the effects across a broad range of signal gain strengths and levels of position tuning narrowing, we performed a similar Monte Carlo simulation as above, this time holding the separation between the two simulated BOLD response curves constant, and instead varying the gain level (left plot in red) or the curve width (right plot in blue) over a wide range of values. We separately tested the effects of response gain, activity gain, and additive gain (see Supplemental Figure 3.1): response gain and activity gain yielded the same monotonically increasing curve (solid red), while additive gain left the correlations unchanged at all levels (dashed red). Narrowing the curve widths always decreased their correlation (blue curve).

fact that signal gain improves position coding by increasing SNR, while narrowed position tuning improves position coding by making the BOLD response patterns more distinct. The purple curve in Figure 3.4C, the result of simultaneous gain and tuning narrowing, matches the characteristic crossover pattern in the position discrimination plots for areas V1 and V2, revealing evidence for both attentional mechanisms at work.

Could signal gain alone ever *decrease* the correlation between two BOLD response profiles? In an additional simulation, we fixed the separation between two modeled BOLD response curves and measured the correlation between the curves over a range of signal gain levels. Vertically scaling the BOLD response by a positive factor (response gain or activity gain) always increased the correlation between the two curves (solid red curve in Fig. 3.4D), whereas rigidly shifting the curves upward (additive gain) had no effect on the correlations (dashed red line in Fig. 3.4D). There was no instance in which signal gain of any type reduced the correlation between the patterns of modeled BOLD response (see Supplemental Figure 3.2 for a proof that both response and activity gain always increase a preexisting positive correlation between two response patterns). In a complementary simulation, we adjusted the BOLD response curve widths over a broad range; narrowing the curves always reduced the correlation between the two profiles (blue curve in Fig. 3.4D).

One concern is that subjects may have attended to a stimulus feature of the Gabors (e.g., orientation) rather than its location. To ensure the position tuning narrowing we measured was due to spatially-directed attention, in a control experiment we presented broadband noise patches rather than Gabors (Supplemental Fig. 3.3A). The results for two subjects (Supplemental Fig. 3.3C) were consistent with the main experiment, revealing significant position tuning narrowing with attention (a decrease in the correlations at the largest separation in attended vs. unattended trials; $t = 2.92$, $p = 0.02$ in V1, $t = 3.27$, $p = 0.01$ in V2). Because these noise stimuli contained no dominant spatial frequency or orientation to which subjects could attend (Supplemental Fig. 3.3B), we are confident that the position tuning narrowing observed is due to spatially-directed attention.

To complement the main analysis, we directly examined the BOLD response along slices through the peak activations in V1 in each subject (Supplemental Figure 3.4). We fit a difference of Gaussian curve to each measured BOLD response profile and compared the width parameters of the positive lobes for attended and unattended conditions to test for narrowing of the BOLD response with attention. A 3-way (position x attention x subject) ANOVA revealed a significant main effect of attention on curve width ($F = 10.24$, $p = 0.002$); attended curves were significantly narrower than unattended curves. We also found that attention significantly increased the curve amplitude parameters ($F = 12.57$, $p < 0.001$), consistent with the simultaneous signal gain and decreased spread of the BOLD response found using the correlational approach.

This analysis verified that attention acted over the entire BOLD response profile, rather than just a restricted portion. If attentional effects were localized to a subregion of the stimulus, one or both of its edges along the slices should be free of attentional modulation. We tested for an attentional effect at the negative “dips” at the edges of the BOLD response profiles, and found that attention significantly *reduced* the responses at both edges ($F = 8.28$, $p = 0.006$ at the foveal edge; $F = 11.13$, $p = 0.001$ at the eccentric edge). Thus, there was significant attentional modulation over the entire extent of the BOLD stimulus representation. This also shows that the narrowed position tuning is not simply due to a positive signal gain that drops off as a function of distance from the center of the spotlight.

We also tested for a difference in the curve phase parameters between the attended and unattended conditions. Attention had no effect on the peak positions of the BOLD response profiles ($F = 0.19$, $p = 0.67$), and there was no significant attention x stimulus position interaction, which might have indicated shifts in different directions for different stimulus positions ($F = 0.003$, $p \approx 1.0$). Attention narrowed the BOLD response spread without significantly shifting its distribution. This slice-based approach allowed us to directly visualize the BOLD response’s shape along a single dimension, complementing and supporting the more powerful multivariate analysis that took into account the changes in the BOLD response pattern in all dimensions simultaneously.

3.3 Discussion

Our results show that spatially-directed attention improves retinotopic coding precision in V1 through V4 by boosting signal amplitude *and* by narrowing position tuning at the neural population level. This is consistent with an emerging pattern of results in the feature tuning domain: attention narrows the tuning of population responses (Martinez-Trujillo and Treue 2004; Murray and Wojciulik 2004), even when tuning for the same features shows little or no change at the single-unit level (McAdams and Maunsell 1999; Treue and Martinez Trujillo 1999). This makes sense, because psychophysical resolution ultimately relies on population coding (Pouget, Dayan et al. 2000; Purushothaman and Bradley 2005).

What mechanism underlies the narrowing effect that spatial attention has on population position tuning? Is it at the single unit level? Several studies show that a V4 neuron's response when two stimuli were positioned inside its receptive field selectively reflected the attended stimulus, suggestive of shrinking receptive fields with attention (Moran and Desimone 1985) (Treue and Maunsell 1996; Luck, Chelazzi et al. 1997). However, this has only been established in extrastriate areas, where narrowed position tuning could result from a spatially specific gain in V1, even if position tuning remained unchanged there (Maunsell and McAdams 2000). Whether attention can narrow the position tuning of single cells in V1 remains unresolved, but narrowed single-unit position tuning could shrink the spread of the population response.

Still, spatial attention need not narrow position tuning at the single-unit level to yield narrowed population position tuning—in a coarse coding framework, the most precise population codes are achieved by larger, highly-overlapping receptive fields (Eurich and Schwegler 1997). As such, increasing the *number* of active units at the attended location while decreasing them in the surrounding region would decrease the spread of the BOLD response we measured while providing a higher resolution population code at the attended location by increasing RF overlap there. Such an attentional mechanism, which has heterogeneous effects across the extent of the attentional "spotlight," could narrow population position tuning while leaving single-unit tuning unchanged.

From a perception standpoint, an attentional mechanism that modulates visual resolution by controlling both signal gain and population position tuning is far more efficient than one using signal gain alone. For many attentionally demanding tasks, performance benefits from the greatest resolution the visual system can achieve; some tasks, such as texture perception (Yeshurun and Carrasco 1998), are better performed at lower resolutions. The attentional system's challenge, then, is not simply to achieve the maximum possible resolution but the broadest possible dynamic range of resolutions. The ability to adjust population position tuning affords the attentional system dynamic control of visual resolution over a broad and continuous range, even if signal gain is held constant or manipulated independently to subserve additional functions such as coordinate frame transformations (Salinas and Abbott 1997). Thus, a dual gain

and tuning mechanism effectively maximizes the flexibility and dynamic range of visual resolution.

3.4 Supplemental Experimental Procedures

3.4.1 Stimuli

The stimuli in the main experiment consisted of four Gabors (sine wave grating within a Gaussian contrast envelope), one presented in each quadrant of the visual field around a central fixation point. The Gabors appeared at five possible eccentricities: 8.430, 8.735, 9.039, 9.343, and 9.647 degrees from fixation. The positions of the Gabors were manipulated by applying a skew to the Gaussian contrast envelopes (Whitaker, McGraw et al. 1996). A central condition had no skew applied to the contrast envelope; its size (standard deviation) was 1.66 degrees and its centroid was at 9.039 degrees eccentricity. In two more foveal conditions the contrast envelopes were skewed by 0.19 and 0.38 degrees toward fixation, and in two more eccentric conditions the contrast envelopes were skewed by 0.19 and 0.38 degrees away from fixation. Eccentricities were manipulated in this way rather than by shifting the peak contrast because skewing the contrast envelope is a better method of isolating the mechanisms that perform centroid analysis (Whitaker, McGraw et al. 1996). A sixth condition consisted of a fixation baseline in which only the fixation point was present.

In all conditions the spatial frequency of the Gabors was 0.38 cycles/degree, peak contrast was 87% (Michelson), and each Gabor was flickered in counterphase at 7.5 Hz. The phase of each Gabor was randomized on every trial. During scanning, the six conditions were randomly interleaved in 36 ten-second blocks per run. Subjects maintained fixation at a central point (0.39 deg. diameter) throughout the entire experiment.

In a control experiment, we presented windowed noise rather than Gabors. We first created random noise with a Michelson contrast of 87%, with each light or dark element in the noise occupying a 4 x 4 block of pixels. We then rotated the image to a random orientation (to avoid peaks in the power spectrum at any given orientation) and added skewed Gaussian contrast envelopes just as with the Gabors in the main experiment. During presentation, we updated the noise and its orientation at 10Hz, so that subjects saw dynamic random noise rather than flickering Gabors. In all respects besides this difference in stimuli, the experimental design was identical for the control experiment and the main experiment.

3.4.2 Attention task

We presented stimuli in a 10-second blocked design in order to test the influence of sustained attention on the distribution of BOLD response in the cortex. This blocked presentation also allowed for improved SNR in the measured BOLD response versus an event-related design. In half of the functional runs, referred to here as *attended* runs, we ensured that subjects attended to the location of the Gabors for the entire duration of

the trial using a sustained attention task (Bressler, Spotswood et al. 2007). In the remaining functional runs, which we refer to as *unattended* runs, subjects' attention was held at the fixation point, and they ignored the Gabors in the periphery.

The same task-related stimuli were present in attended and unattended trials. At a randomly chosen time during the first 8 seconds of each ten second block, a small texture pattern (either circular or radial grating, chosen randomly; 1.09 degrees diameter) appeared for 500 ms superimposed on one of the 4 Gabors (chosen randomly), at an eccentricity of 9.04 degrees (center to center) from fixation. The flashed texture appeared during every condition, including the fixation baseline. During the last two seconds of each ten second block a second flashed texture appeared at the center of a randomly chosen Gabor for 500 ms, and a white annulus (0.98 degrees diameter) was presented continuously around the fixation point, indicating that subjects should make a response. The second textured flash matched the first one with a probability of 50%.

In attended runs, we instructed subjects to maintain fixation at all times and make two judgments. First, subjects discriminated the eccentricity of the Gabors (the degree of skew in their envelopes) in a 5 alternative-forced-choice task. Subjects also needed to continuously monitor the 4 Gabors for the texture flashes. At the end of each trial, when the white annulus cue was presented around the fixation point, subjects indicated both the eccentricity of the Gabors and whether or not the two presented textures matched. This task ensured that subjects attended to the surrounding Gabors for the entire duration of each trial.

Additionally, a small grating annulus (either circular or radial grating, chosen randomly; 0.98 degrees diameter) appeared for 500 ms surrounding the fixation point at least 7 times but no more than 11 times per trial, and one texture always appeared once more than the other. These flashes were uncorrelated with the textured flash superimposed on the Gabors and were therefore uninformative. During attended runs, we instructed subjects to ignore the flashes at the fixation point, and to attend entirely at the positions of the Gabors. During runs in which the Gabors were unattended, subjects ignored the eccentricity of the Gabors and the appearance of the textures, and instead counted how many times each grating annulus appeared around the fixation point. When the white annulus cue was presented at the end of each trial, subjects made a 2AFC judgment of which grating had appeared more often. Thus, during attended and unattended runs (referring to whether or not the Gabor stimuli were attended to), demanding tasks held subjects' attention at the location of the Gabors and at the fixation point, respectively. Attended and unattended runs were identical in terms of the stimuli presented; the only difference between the two was in the judgments that subjects were instructed to make.

3.4.3 fMRI data collection and analysis

Seven subjects participated in the main experiment, and two subjects participated in windowed noise control experiment. Scanning protocols were approved

by the University of California, Davis, Human Subject Review Board. Imaging was conducted on a 3-Tesla Siemens TRIO scanner located at the UC Davis Imaging Research Center. Each participant's head was rested in a Siemens eight-channel phased-array head coil. Stimuli were back-projected with a Digital Projection Mercury 5000HD projector (75 Hz) onto a semi-transparent screen from outside the bore. A mirror angled at 45 deg, located 9.5 cm directly above the subject, provided a reflected view of the stimuli. Functional images were acquired with a gradient-recalled echo EPI sequence. Whole-brain structural images were collected with a high resolution (1 mm^3) Turbo Spin Echo scan that was used to align functional images. The acquisition parameters were: TR = 2000 ms, TE = 26 ms, FA = 90 deg, FOV = 22 x 22 cm^2, voxel size = 1.528 x 1.528 x 2.5 mm^3, 20 slices per volume. The imaging volume was parallel to and centered on the calcarine sulcus, covering the occipital lobe.

All fMRI data processing, including linear trend removal, 3D motion correction, and GLM analyses, was conducted with Brain Voyager QX (Brain Innovation B.V., Maastricht, The Netherlands). The images were not spatially smoothed. A correction for serial correlations (removal of first-order autocorrelation) was used prior to all GLM analyses. Each subject's high resolution anatomical image was transformed to Talaraich coordinates (Talairach and Tournoux 1988), and the data for each functional run was individually aligned to the subject's Talairach-transformed anatomical image. Performing individual alignments for each functional run, within each subject, mitigated any effect of subject movement between functional runs.

To generate an individual map of BOLD response for each eccentricity condition, we fit a GLM to the data with five predictors (corresponding to the five Gabor positions). Each predictor consisted of a boxcar function representing the blocked presentations of the condition, convolved with a canonical HRF. We separately contrasted each of these five eccentricity conditions against a sixth predictor (the fixation baseline) to generate a unique map of BOLD response for each stimulus condition. We performed this process separately for all subjects, creating distinct maps for attended and unattended runs.

3.4.4 ROI definition

We defined the regions of interest (ROIs) for areas V1 through V4 by tracing reversals in each subject's retinotopic mapping in the cortex (Sereno, Dale et al. 1995). In localizer runs, which were presented separately from the main experimental runs, we presented flickering bowtie stimuli to identify the borders of visual areas V1 through V4. The bowties consisted of radial sine wave patterns of 11.79 deg. radius, subtending an arc of 8.16 deg. The bowties flickered in counterphase at 7.5 Hz. There were three conditions in these retinotopy runs; in two conditions, the bowties were centered on the vertical or horizontal meridians, and the third condition was a fixation baseline in which only the fixation point was present. Conditions were randomized in 36 ten-second blocks. We performed general linear model (GLM) analyses on the data in the retinotopy runs to define visual areas V1 through V4. Bowties along the vertical meridians were contrasted with those on the horizontal meridians, which yielded a striated map of activity across early visual areas; we defined the boundaries of each

visual area V1 through V4 using the mirror reversals in the representations of the horizontal and vertical meridians, as in Sereno et al. (Sereno, Dale et al. 1995).

3.4.5 Correlation analysis

To obtain estimates of the precision of retinotopic coding within the V1 through V4 regions of interest, we produced a position discrimination plot for each ROI. To produce each position discrimination plot, we began with the five maps of BOLD response, corresponding to the five different stimulus eccentricities. Given any two maps, we found the correlation between them within the ROI by plotting t values from one map against the corresponding t values from the other map, on a voxel-by-voxel basis. The resulting plot had as many points as there were voxels in the ROI. We computed a Pearson r value for every such voxel plot, and converted the r values to Fisher z scores so that they could be linearly compared with each other. We produced a voxel plot and a corresponding correlation for every pairing of activity maps, resulting in ten z scores. To create a final position discrimination plot, we plotted each of the ten z scores against the spatial separation (in deg. visual angle) of the two stimuli that produced that particular correlation (smaller separations typically implied larger z scores, and vice versa; (Bressler, Spotswood et al. 2007; Fischer and Whitney 2007)). Adjacent conditions had centroids separated by .304 degrees; more distant conditions were separated by multiples of this value (.609, .912, and 1.216 degrees). We plotted subjects together (Fig. 3.2), and fit a regression line to the position discrimination plot in order to provide an index of the precision of retinotopic coding within the specific ROI that was used to create the plot. The linear regression included an additional random effects term to account for plotting multiple subjects together. This model took the form $z_{ijk} = \beta_0 + \tau_i + \beta_1 x_{jk} + \varepsilon_{ijk}$, where *i* indexed the subjects and the pair (*j*, *k*) indexed the ten stimulus pairings; τ_i provided for rigid vertical shifts between subjects that did not affect the slope. The more negative the slope of the regression line, the better the position discrimination. We produced position discrimination plots for attended and unattended conditions separately in order to allow a comparison of the precision of spatial coding with and without attention. To test whether this difference in slopes was significant, we computed Z statistic on the difference of r values between the attention and fixation plots, given by $Z = \frac{r_1' - r_2'}{\sqrt{\frac{1}{N_1 - 3} + \frac{1}{N_2 - 3}}}$, where r_1' and r_2' are the Fisher-transformed Pearson r values from the attended and unattended plots, and N_1 and N_2 are the number of points in each plot respectively.

This correlational process is based on a simple concept: in a cortical area which encodes object position, the activity resulting from two spatially adjacent stimuli will be highly correlated, while the activity resulting from more distant stimuli will be less strongly correlated. Thus a z score plot constructed as described above will have a negative slope if BOLD response strongly discriminates stimulus position, and a slope near zero if position information is not encoded in the voxels within the ROI for which the plot was constructed. The slope of such a plot indexes the precision of position

coding because it indicates the minimum distance between stimuli required to produce significantly different BOLD response patterns (Fischer and Whitney 2007).

3.4.6 Slice-based analysis

To directly examine the profile of the BOLD response, we measured activity in V1 along a slice across each subject's inflated cortex. We created inflated meshes from subjects' high-resolution anatomical scans, and then mapped a GLM contrast of the most eccentric stimulus condition versus the most foveal condition onto the inflated meshes to identify the retinotopic location of BOLD response in V1 and to define the direction of retinotopic progression for each subject (Supplemental Fig. 3.4A). We then drew a slice on each subject's right hemisphere surface map, through the region of peak activation in V1 and along the calcarine sulcus, parallel to the direction of retinotopic progression of the stimuli. We measured the BOLD response along each slice separately for attended and unattended conditions, and plotted BOLD amplitude for each as a function of cortical distance (millimeters in Talairach space) along the slices.

To characterize the shape of the BOLD response, we fit a difference-of-Gaussians (DoG) function to each subject's slice data (Supplemental Fig. 3.4B shows four subjects). The DoG function had the form $y = a_1 e^{-\frac{(x-b_1)^2}{2c_1^2}} + a_2 e^{-\frac{(x-b_2)^2}{2c_2^2}}$, where a_1 was constrained to be positive and a_2 was constrained to be negative, so that the first and second halves of the equation represent the positive and negative components of the curve, respectively. We were then able to compare the height, width, and peak positions of the positive lobes of the attended and unattended curves by comparing the a_1, c_1, and b_1 parameters respectively. We performed such comparisons using three-way ANOVAs (stimulus position x attention condition x subject), carried out in Matlab 7.1 (The Mathworks, Natick, Massachusetts). When comparing activity at the endpoints of the BOLD response profile, we located the negative lobes of a BOLD slice by taking the smallest responses in the foveal and eccentric halves of the response profile. These provided additional "parameter estimates" for the amplitude of the negative dips – we used this approach, rather than using the amplitude parameter for the negative component of the DoG curve, in order to measure the foveal and eccentric dips separately. After obtaining these measurements of the amplitude of the negative dips, we compared them for attended and unattended conditions using three-way ANOVAs, just as with the DoG fit parameters.

3.4.7 Modeling

To model the affects of gain and tuning on the slope of a position discrimination plot, we first created a prototypical BOLD response profile using subjects' BOLD response profiles from the slice-based analysis. We aligned the response curves for the unattended conditions from all subjects peak-to-peak, normalizing by peak amplitude, and plotting them together (Figure 3.4A). Then we fit a difference-of-Gaussians (DoG)

function to this combined data (black curve in Figure 3.4A), and used a discretely sampled version of this curve with 200 points in subsequent simulations.

To generate the green curve in Figure 3.4B, representing the position discrimination slope for unattended conditions, we used two copies of the empirically-defined DoG curve to simulate two different Gabor positions in a Monte Carlo simulation. In each iteration, random noise was added to each curve independently, to produce an SNR of 15.33. We estimated this SNR from subjects' BOLD response maps using the correlations at the smallest separation for each subject. SNR is related to correlation by $\frac{S}{N} = \frac{\rho}{1-\rho}$ in the case where two copies of the same signal are corrupted by independent noise sources of the same energy. Since the correlations at the smallest separation reflect very similar but non-identical signals, our SNR estimate was necessarily conservative, and indicative of a lower bound. This estimate reflects the SNR of the BOLD response maps (resulting from GLM contrasts), and not the raw fMR timecourse. We computed SNR for the unattended runs for each subject and then averaged across subjects to yield this value for modeling. In each iteration of the Monte Carlo simulation, the two curves were offset from each other by a peak-to-peak separation of δ in the range of [0.0, 4.0] mm, corresponding to the x axis in Fig. 3.4B & C. A correlation was then computed between the two curves to produce a Fisher z. At each separation δ, 5000 iterations were completed in this way, with new random noise added in each iteration. We averaged the resulting z scores at each separation, and these average z scores were plotted against the separation δ at which they were computed to produce the green curve in Figure 3.4B & C. This curve simulates the slope of a position discrimination plot for unattended stimuli (Fig. 3.2), modeled on a continuous scale.

To generate the remaining curves in Figure 3.4B & C, we ran the same Monte Carlo simulation, after scaling the DoG curves. For the gain plot (red), we increased the amplitude of the DoG curves by 10% to simulate a multiplicative amplitude gain, and all other parameters were kept the same. For the tuning plot (blue), we decreased the width of the DoG by 10% to simulate position tuning narrowing, and all other parameters were kept the same. For the combined gain and tuning plot (purple), both a 10% amplitude gain and a 10% narrowing were applied.

To test how the correlation between patterns of activity is affected by other levels of gain and tuning, we used the same pair of DoGs as described above, now fixed at a realistic separation of 1.51 mm, which was established by taking all subjects' average peak-to-peak separation of the BOLD response profiles produced by the most eccentric and most foveal conditions (corresponding to the largest separation on the position discrimination plot). We added noise to each curve as before, and carried out Monte Carlo simulations by varying the amount of signal gain or width narrowing applied to the modeled curves (Figure 3.4D). In the case of signal gain, we modeled the outcome of multiplicatively scaling the response curves around zero (corresponding to response gain, in which baseline activity is not modulated by attention), multiplicatively scaling the response curves after adding back in a baseline constant (corresponding to activity

gain, in which baseline activity is also scaled by attention), and simply adding a baseline constant (corresponding to additive gain). We modeled the last two over a range of baseline constants (from 1% - 1000% of the amplitude of the modeled BOLD response profile), but the value of the baseline constant had no impact on the resulting correlations; activity gain produced an identical curve to that of response gain (solid red curve in Fig. 3.4D), and additive gain had no impact on the correlations at any level of gain (dashed red line in Fig. 3.4D). We performed all modeling in Matlab 7.1.

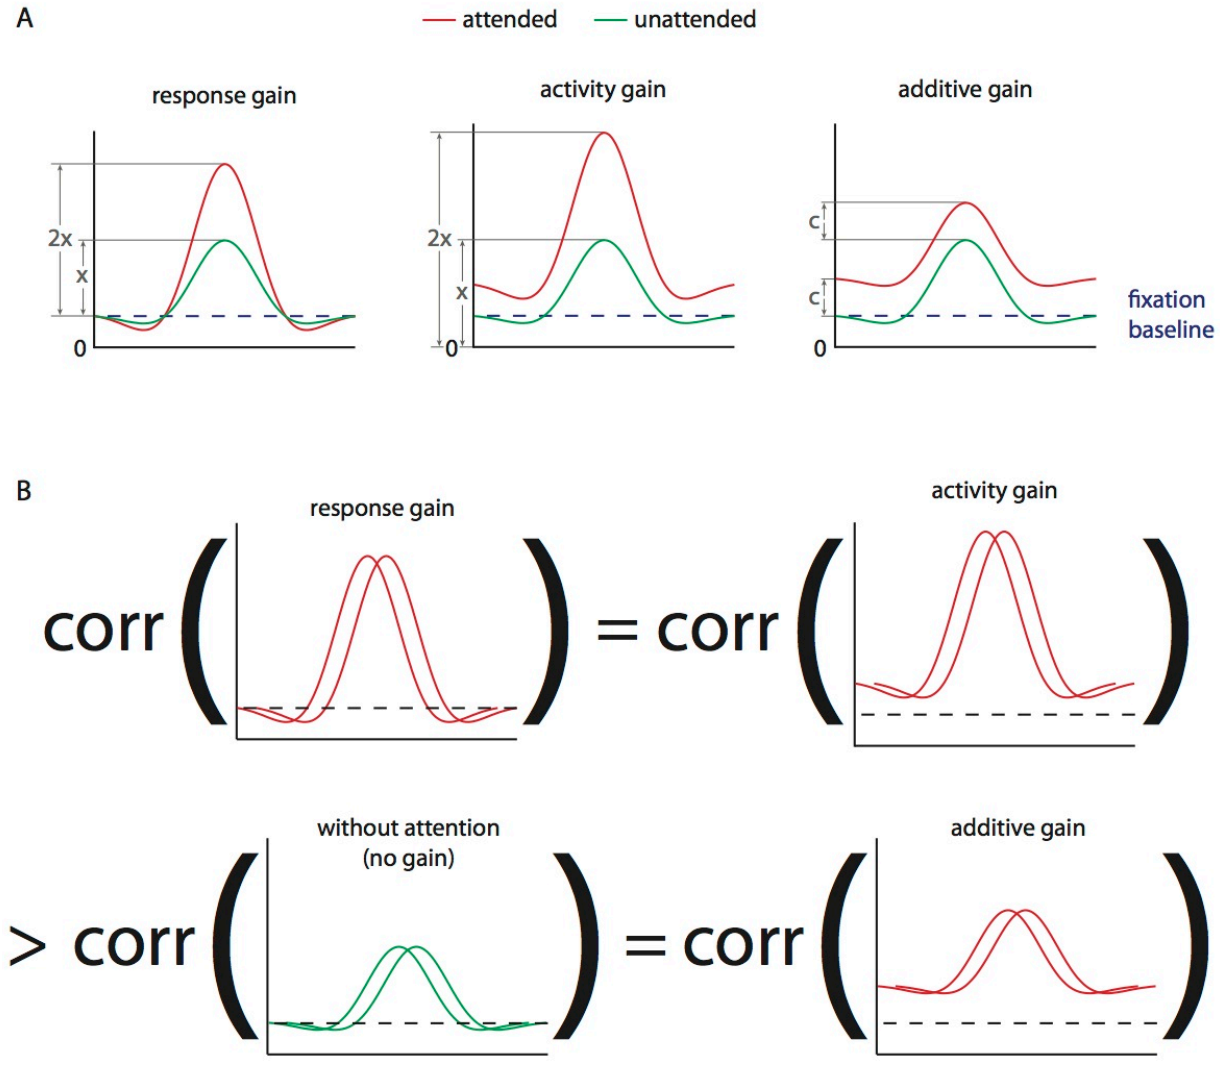

Supplemental Figure 3.1: The effect of response gain, activity gain, and additive gain on the correlation between BOLD patterns. In attention research, distinctions have been drawn among *response gain,* in which attention has a multiplicative effect on above-baseline activity (stimulus-evoked responses), *activity gain*, in which all activity, including baseline, is multiplicatively increased by attention, and *additive gain*, in which the activity level is increased by a constant, rather than scaled. These three scenarios are shown in panel **A)**. While activity gain yields larger overall responses than response gain, the shapes of the BOLD response curves generated by response gain and activity gain are identical; activity gain yields a vertically-shifted version of the response gain curve. Similarly, additive gain yields a vertically-shifted version of the unattended (green) curve. **B)** Response gain and activity gain produce an equivalent increase in the correlation between two curves, because they yield identically-shaped attended (red) curves. For the same reason, additive gain leaves the correlation between two curves unchanged, because it merely shifts the curves vertically. In this study, we use the term *signal gain* to refer to a vertical scaling of the BOLD response that could result from either response gain or activity gain. Signal gain always increases the correlation between two BOLD response curves, assuming the presence of noise, and positive initial correlation between the unattended curves (see Supplemental Figure 3.2 for a proof).

Let X and Y be random variables with variances σ^2_x and σ^2_y, respectively.

Say that X and Y are corrupted by random noise ε_1 and ε_2, both with mean 0 and variance greater than zero, such that ε_1 and ε_2 are independent of X, Y, and each other. Then

$$\rho(X+\varepsilon_1,Y+\varepsilon_2)=\frac{Cov(X+\varepsilon_1,Y+\varepsilon_2)}{\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2}}=\frac{Cov(X,Y)+Cov(X,\varepsilon_2)+Cov(\varepsilon_1,Y)+Cov(\varepsilon_1,\varepsilon_2)}{\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2}}$$

$=\frac{Cov(X,Y)}{\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2}}$ by the independence of X and ε_2, Y and ε_1, and ε_1 and ε_2.

Now, looking at the denominator,

$$\sigma_{X+\varepsilon_1}\sigma_{Y+\varepsilon_2}=\sqrt{\sigma^2{}_{X+\varepsilon_1}}\sqrt{\sigma^2{}_{Y+\varepsilon_2}}=\sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}};$$

Thus, $\rho(X+\varepsilon_1,Y+\varepsilon_2)=\frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}}$.

Now, introduce a multiplicative gain of magnitude *a* to X and *b* to Y, where *a, b* > 1.

The new correlation is given by

$$\rho(aX+\varepsilon_1,bY+\varepsilon_2)=\frac{Cov(aX,bY)}{\sqrt{\sigma^2{}_{aX}+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_{bY}+\sigma^2{}_{\varepsilon_2}}}=\frac{abCov(X,Y)}{\sqrt{a^2\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{b^2\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}}=$$

$$\frac{abCov(X,Y)}{a\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}}b\sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}}}=\frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}}\sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}}}.$$

Now, since $\sigma^2{}_{\varepsilon_1},\sigma^2{}_{\varepsilon_2}>0$ and $a^2, b^2>1$,

$$\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}}<\sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}} \text{ and } \sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}}<\sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}.$$

Thus,

$$\rho(aX+\varepsilon_1,bY+\varepsilon_2)=\frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\frac{\sigma^2{}_{\varepsilon_1}}{a^2}}\sqrt{\sigma^2{}_Y+\frac{\sigma^2{}_{\varepsilon_2}}{b^2}}}>\frac{Cov(X,Y)}{\sqrt{\sigma^2{}_X+\sigma^2{}_{\varepsilon_1}}\sqrt{\sigma^2{}_Y+\sigma^2{}_{\varepsilon_2}}}=\rho(X+\varepsilon_1,Y+\varepsilon_2)$$

when X and Y covary positively. Therefore, a multiplicative gain applied to one or both signals in the presence of uncorrelated noise increases the correlation between those signals, given that they covary positively.

Supplemental Figure 3.2: Algebraic proof that multiplicative gain increases the correlation between two positively correlated signals. We assume that two signals of arbitrary shape are positively correlated, and are in the presence of nonzero noise. Applying a multiplicative scaling by a factor greater than one to one or both signals (as is the case with response gain or activity gain) increases the correlation between the two signals by increasing the signal strength while the noise level remains constant. Importantly, if the original signals are negatively correlated, then a multiplicative scaling of one or both signals will increase the strength of the correlation between them, yielding a smaller (more negative) correlation. In this study, all correlations that we measured were positive.

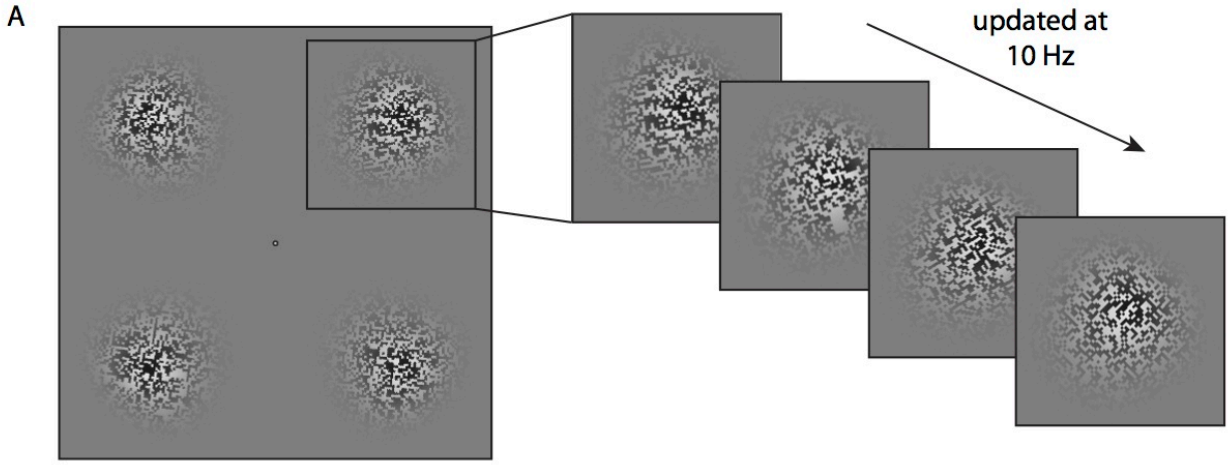
A
updated at
10 Hz

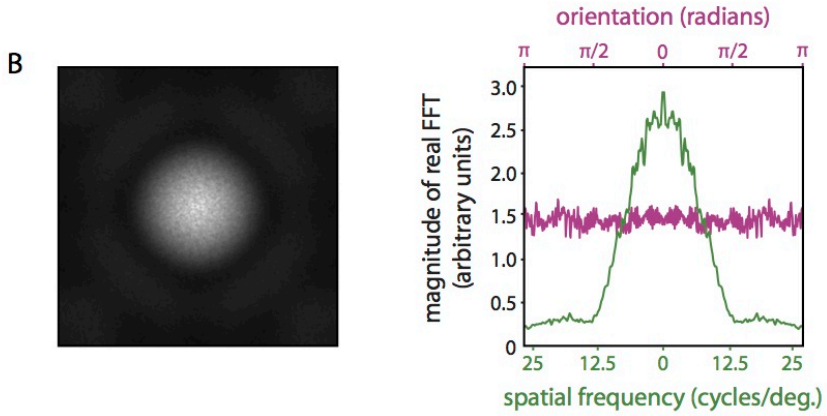
B
orientation (radians)
π
π/2
0
π/2
π
magnitude of real FFT
(arbitrary units)
3.0
2.5
2.0
1.5
1.0
0.5
0
25
12.5
0
12.5
25
spatial frequency (cycles/deg.)

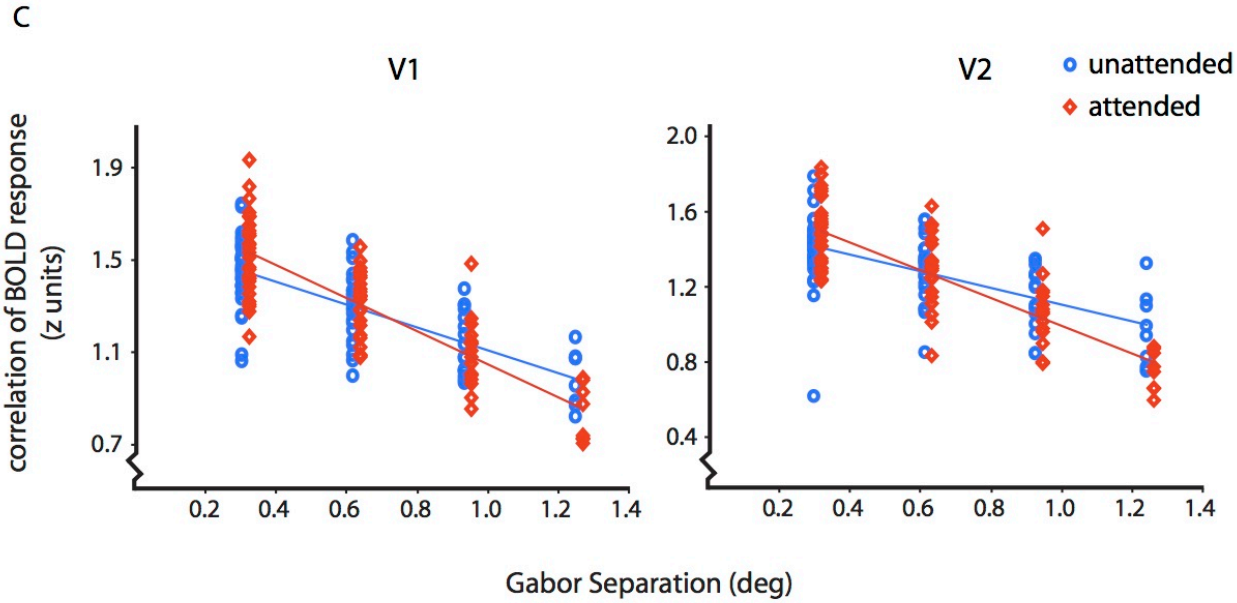
C
V1
V2
unattended
attended
correlation of BOLD response
(z units)
1.9
1.5
1.1
0.7
2.0
1.6
1.2
0.8
0.4
0.2
0.4
0.6
0.8
1.0
1.2
1.4
Gabor Separation (deg)

Supplemental Figure 3.3: Control experiment to rule out feature-based attention. A) To ensure that the effects we measured in the main experiment were due to spatially-directed, rather than feature-based attention, we performed a control experiment with two subjects in which we presented Gaussian-windowed noise instead of Gabor patches. Dark and light patches in the noise were 4 pixels by 4 pixels, allowing us to update the orientation of the grid on which the noise was created rapidly (10 Hz). The result was a stimulus that was broadband in spatial frequency and orientation. **B)** On the left is the 2d Fourier transform of a noise stimulus corresponding to the middle eccentricity, and on the right is a visualization of the power across the full range of spatial frequencies at orientation 0 (green), and orientations at 10 cycles/dig (purple; the DC component is omitted for plotting purposes). There are no dramatic spikes in the power spectrum, and power over spatial frequencies drops smoothly in a roughly 1/f fashion. In all other respects besides the stimulus, the control experiment was identical to the main experiment. **C)** We plotted both subjects together as in the main experiment (Figure 3.2); here, individual runs are plotted separately, and the x positions of the attended and unattended data are offset slightly for visualization purposes. Both V1 and V2 exhibit the crossover between attended and unattended correlations that we found in the main experiment, and there was a significant drop in the correlations with attention at the largest separation ($t = 2.92$, $p = 0.02$ in V1; $t = 3.27$, $p = 0.01$ in V2) indicating reduced spread of the BOLD response with attention. The increase in correlations at the largest separation corresponding to signal gain was not quite significant ($t = 1.87$, $p = 0.07$ in V1; $t = 1.65$, $p = 0.11$ in V2), but by comparing direct measurements of the largest responses in attended and unattended conditions, we verified that attention increased the peak responses in both ROIs (Subj. DH: $t = 4.03$, $p < 0.001$ in V1, $t = 9.42$, $p < 0.0001$ in V2; Subj. NS: $t = 3.87$, $p = 0.001$ in V1, $t = 6.31$, $p < 0.0001$ in V2). Because these noise stimuli contained no dominant spatial frequency or orientation, we can be confident that spatially-directed attention narrows population position tuning.

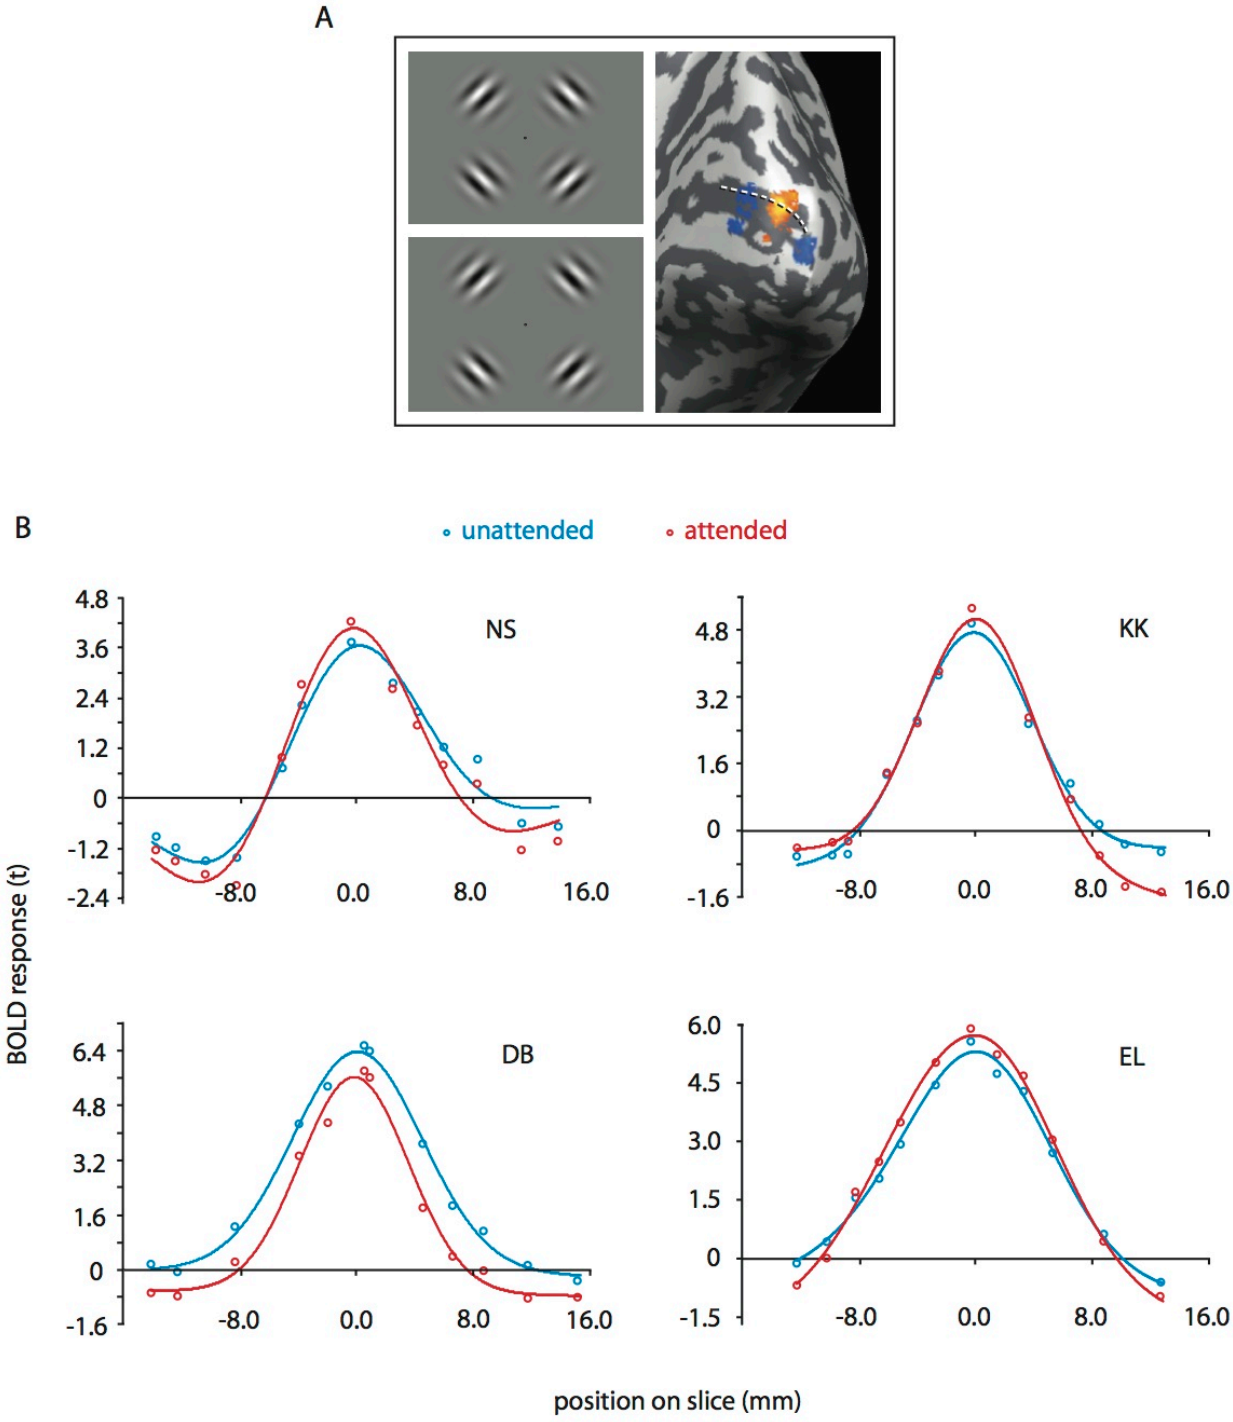
A
B
unattended
attended
NS
KK
DB
EL
BOLD response (t)
position on slice (mm)
4.8
3.6
2.4
1.2
0
-1.2
-2.4
3.2
1.6
-1.6
6.4
4.8
3.2
1.6
0
-1.6
6.0
4.5
3.0
1.5
0
-1.5
-8.0
0.0
8.0
16.0

Supplemental Figure 3.4: Direct measurement of the shape of the BOLD response. A) We found the direction of retinotopic progression of the stimuli by contrasting the least and most eccentric stimulus positions in a general linear model. We overlaid each subject's resulting map on his or her inflated cortex, and drew a slice through the peak activation in V1, along the direction of the progression of the stimulus eccentricity manipulation. We measured the BOLD response for each of the five stimulus eccentricities, for attended and unattended trials, and fit difference of Gaussians curves to each set of measurements to characterize the shape of the BOLD response profiles. We compared the amplitude, width, and phase parameters of the curve fits for attended and unattended conditions, and found significant narrowing ($F = 10.24$, $p = 0.002$) and signal gain ($F = 12.57$, $p < 0.001$) with attention, but no change in the positions of the BOLD responses ($F = 0.19$, $p = 0.67$). We also tested for a decrease in the BOLD response with attention at the locations of the two negative lobes by taking the lowest responses at the foveal and eccentric edges of each subject's attended and unattended BOLD response curves. We found that both edges of the stimulus representation were significantly modulated by attention ($F = 8.28$, $p = 0.006$ at the foveal edge; $F = 11.13$, $p = 0.001$ at the eccentric edge). These results rule out a spatially asymmetric effect of attention on the BOLD response profiles, as well as showing that attention acted over the entire stimulus representation. **B)** We created a visualization of the BOLD response profiles for four subjects by aligning the five eccentricity conditions peak-to-peak and averaging across the five conditions. Though subtle, a decrease in the spread of the BOLD response is evident in the fact that the attended (red) curves have a smaller width than the unattended curves at the x-axis.

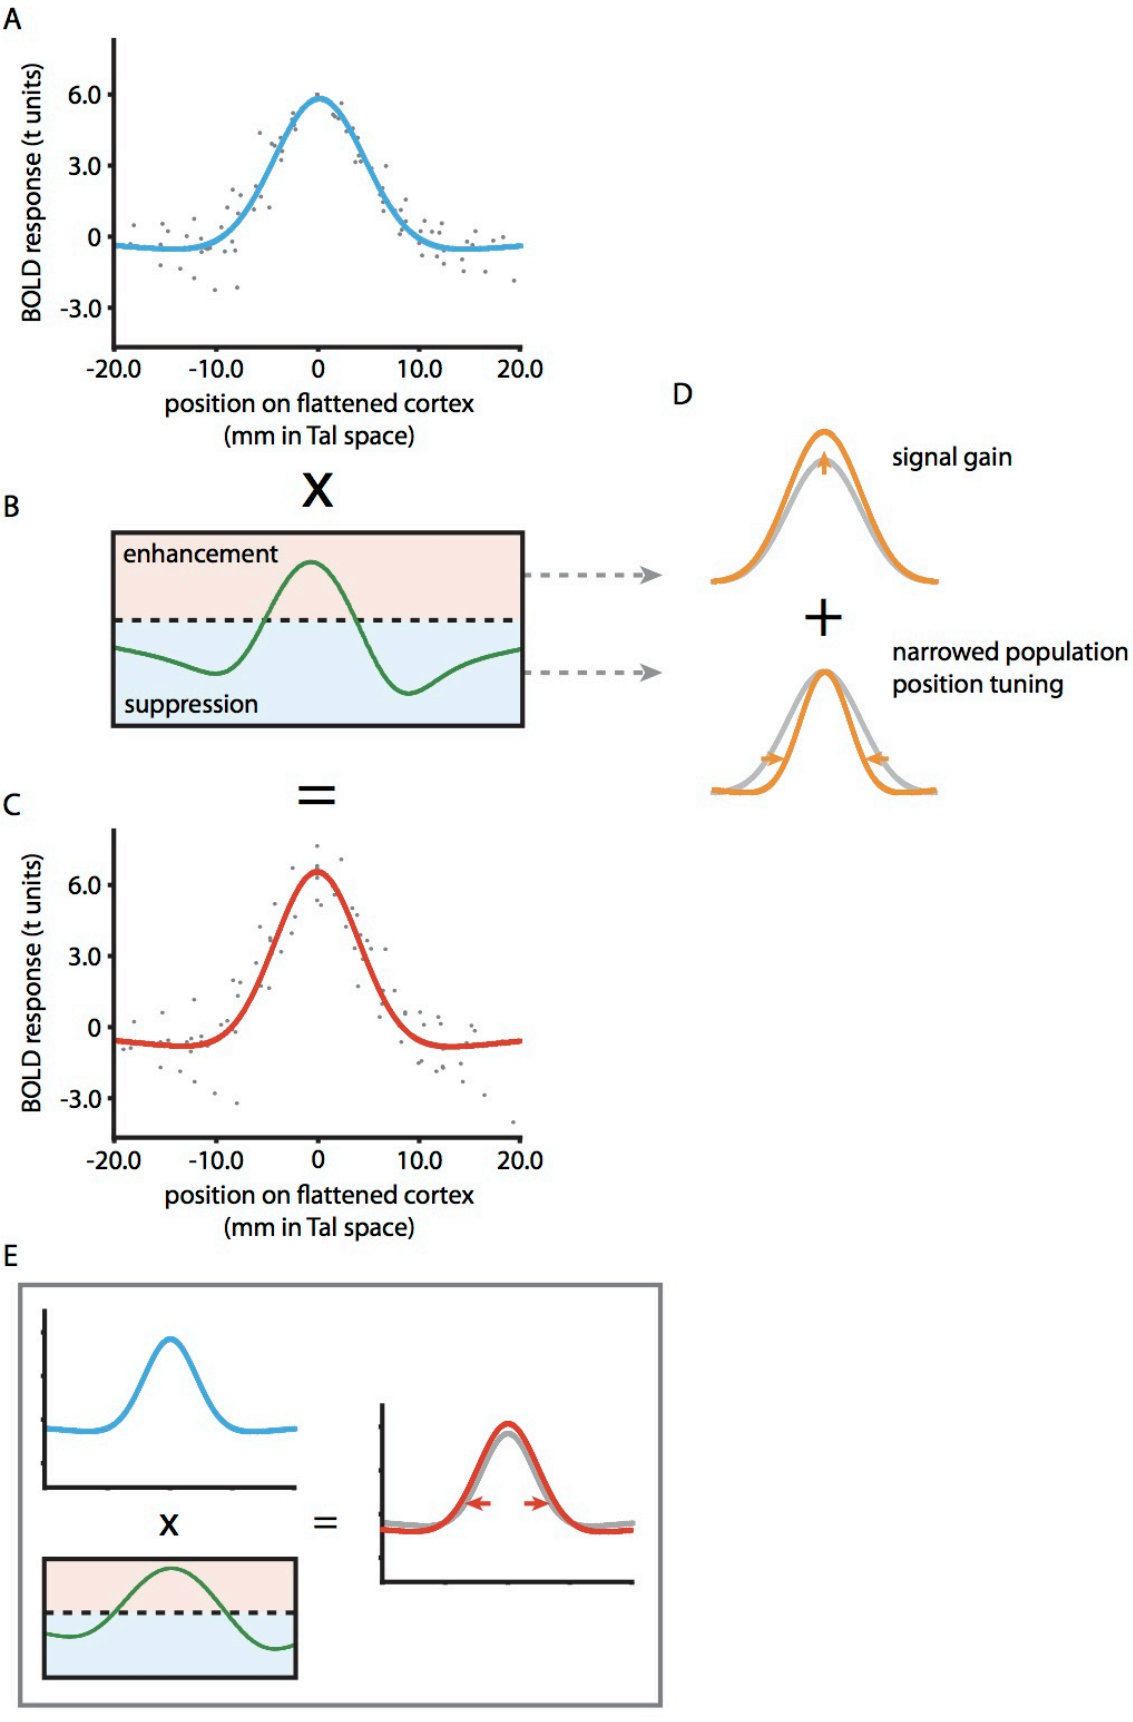
A
BOLD response (t units)
6.0
3.0
0
-3.0
-20.0
-10.0
0
10.0
20.0
position on flattened cortex
(mm in Tal space)
B
X
enhancement
suppression
C
=
BOLD response (t units)
6.0
3.0
0
-3.0
-20.0
-10.0
0
10.0
20.0
position on flattened cortex
(mm in Tal space)
D
signal gain
+
narrowed population
position tuning
E
X
=

Supplemental Figure 3.5: Visualizing signal gain and position tuning narrowing as a spatially heterogeneous gain function. A) - C) The simultaneous signal gain and narrowing produced by focused attention can be thought of as a spatially distributed field of enhancement and suppression. We generated group BOLD response profiles for unattended (blue; panel **A**) and attended (red; panel **C**) stimuli in order to visualize how attention scaled the BOLD response as a function of position along the slice. We plotted all subjects' slice data together (see Supp. Fig. 3.4), normalizing to the group mean peak BOLD response for the unattended data. The blue and red profiles are difference of Gaussians curve fits to the attended and unattended slice data, respectively. By deconvolving the attended and unattended curves, we recovered the attentional modulation kernel in panel **B**, showing the factor by which attention scales the BOLD response as a function of position along the slice. The dashed horizontal line is the line of unity; where the green curve intersects this line, attention left the BOLD response unchanged. Above the line of unity, attention scaled the BOLD by a factor greater than 1, and below the line of unity, attention scaled the BOLD by a factor less than 1. Because the regions of suppression encompass much of the positive lobe of the unattended BOLD response profile in **A)** above, scaling the blue curve by the attentional modulation kernel yields a narrowed red curve in panel **C** below. **D)** The attentional modulation kernel, which scales the BOLD independently at each location along the slice, can be parsimoniously summarized by two simple operations acting on the entire distribution: vertical and horizontal scaling of the population response. Each of these operations is directly related to visual resolution: vertical scaling affects the signal-to-noise ratio, and horizontal scaling affects the overlap between neighboring response distributions. **E)** The existence of suppression in unattended regions need not predict narrowed population position tuning at the location of the stimulus. As reported by Tootell et al. (Tootell, Hadjikhani et al. 1998), Somers et al. (Somers, Dale et al. 1999), Smith et al. (Smith, Singh et al. 2000), Müller & Kleinschmidt (Muller and Kleinschmidt 2004), and others, focused attention can yield a decrease in the BOLD response relative to baseline at locations remote from the attended region. These studies established that attention can have heterogeneous effects on the BOLD response which extend far beyond the attended location, but, critically, they did not resolve how attention alters the distribution of population responses corresponding to the stimulus itself. As illustrated here, the mere existence of such suppressive regions does not imply that attention will narrow the distribution of BOLD response at the location of the stimulus. In the possibility depicted in panel **E**, attention increases the peak of the BOLD response but also *broadens* the distribution of the BOLD, despite having a suppressive effect in the surrounding regions. In light of this possibility, previous studies cannot reach a conclusion about how attention affects population position tuning because they did not directly measure the distribution of the BOLD response at the location of the stimulus, as we have done here.

Chapter 4

Attention gates visual coding in the pulvinar nucleus of the thalamus

4.1 Introduction

Perhaps as important as what is represented in the brain's spatial maps is what's *not* represented there. The majority of the visual input to the brain at a given moment is unnecessary for the task at hand, and can even be detrimental to performance. Alongside examining how flexible spatial mapping in the brain may contribute to the privileged encoding of attended information (Chapter 3), I wanted to know how flexibility in position coding might be used to gate out distracting input. The experiments in Chapter 4 focus on the pulvinar nucleus of the thalamus, a prime candidate region for such a gating function.

The pulvinar nucleus of the thalamus is believed to have an important, integrative function in the mammalian brain, likely related to visual attention (Casanova 2004; Shipp 2004). While pinpointing the pulvinar's specific role in vision has proven notoriously difficult, one line of evidence suggests the pulvinar may be important in isolating behaviorally relevant objects from surrounding distractors (Desimone, Wessinger et al. 1990; LaBerge and Buchsbaum 1990; Laberge, Carter et al. 1992; Buchsbaum, Buchsbaum et al. 2006; Snow, Allen et al. 2009). When Desimone et al. (1990) deactivated the macaque pulvinar by muscimol injection, visual discrimination was impaired in the visual field contralateral to the injection site, but only when distractors were present; performance was normal in the absence of distractors. Recently (Snow, Allen et al. 2009), a similar pattern was reported in human subjects with pulvinar lesions due to stroke: patients had compromised visual discrimination in the contralesional visual field, but only when salient distractors were present. These results point toward an involvement of the pulvinar in resolving competition between visual targets and distractors (Robinson and Petersen 1992; Desimone and Duncan 1995). However, other studies based on lesions have failed to reinforce this notion. In a flanker interference task (Danziger, Ward et al. 2004) and a global/local interference task (Danziger, Ward et al. 2001), subjects with pulvinar lesions showed no impairment in suppressing the distracting dimensions.

These results reveal two missing pieces of the puzzle with regard to distractor filtering in the pulvinar. First, lesion studies still paint a mixed picture of whether intact pulvinar function is necessary for normal distractor suppression. This may be due partly to the wide variety of stimuli and tasks used across the studies. Second, and more importantly, the current evidence for the pulvinar's role in distractor suppression comes

almost exclusively from monkey and human lesion studies. To establish a role of the human pulvinar in distractor suppression, it is critical to show that the pulvinar performs such a function in healthy humans.

Here, we tested for distractor filtering in the human pulvinar by characterizing how the pulvinar encodes targets and distractors when they compete for attentional resources. We capitalize on a newly emerging understanding of the human pulvinar's organization (Kastner, O'Connor et al. 2004; Cotton and Smith 2007; Fischer and Whitney 2009; Schneider 2011); specifically, the presence of precise, lateralized spatial maps in both hemispheres. Using a multi-voxel pattern analysis, we measured the precision with which attended and ignored stimuli are coded in the pulvinar when both were present simultaneously within the same visual hemifield. Our results show that attention gates both spatial and featural information in the pulvinar: the positions and orientations of attended targets were coded with high precision in pulvinar responses, while there was no detectable encoding of distractor positions or orientations, even though the targets and distractors differed only in their behavioral relevance. These data support the hypothesis that the pulvinar is involved in filtering distracting visual information and highlighting behaviorally relevant targets.

4.2 Results

4.2.1 Experiment 1 – Position discrimination

Subjects viewed four visual stimuli at a time; two were attended and two were ignored. The stimuli were Gabor patches in the four visual quadrants (Fig. 4.1a); in each ten-second stimulation block, they appeared at one of five possible eccentricities ranging from 8.4 to 9.6 degrees (Fig. 4.1b). The top two Gabors were fixed to have the same eccentricity as each other, as were the bottom two, but the eccentricities for the upper and lower Gabors were random with respect to each other. On alternating runs, subjects attended to the Gabors in either the upper or the lower visual field (Fig. 4.1c), watching for slight contrast decrements that occurred an average of five times per block. Attention was manipulated between the upper and lower visual fields (rather than left versus right) so that an attended and an ignored stimulus were always present within the same visual field, and hence represented in the pulvinar in the same hemisphere (see Discussion). At the end of each block, the subject responded whether there were more contrast decrements on the left, right, or the same number on the left and right (3 AFC; sensitivity [d'] for the task was 1.59, indicating that the task was challenging but above threshold). The contrast decrements also occurred in the ignored stimuli, such that they were not just irrelevant but also distracting; subjects were motivated to ignore those stimuli as completely as possible. We tracked the eye positions of 3 of the 5 subjects during scanning (Supp. Fig. 4.1).

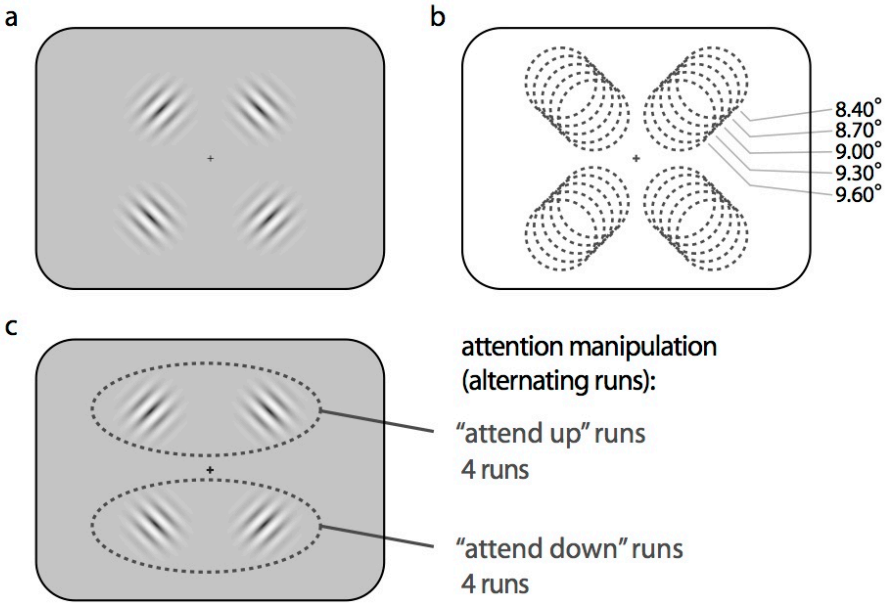

Figure 4.1: Stimuli and experimental design. a) The stimuli consisted of four Gabor patches in the four visual quadrants. **b)** The Gabors appeared at five possible eccentricities, ranging from 8.40 deg. to 9.60 deg. From trial to trial, the positions of the upper two Gabors varied independently of the positions of the lower Gabors. **c)** In alternating runs, subjects were instructed to attend to the Gabors in either the upper visual field or the lower visual field, to detect slight contrast decrements in the stimuli.

Independently manipulating the positions of the upper and lower Gabors allowed us to separately measure position coding at the attended and ignored locations within the same functional run. We extracted patterns of blood oxygenation level dependent (BOLD) responses corresponding to the attended and ignored stimuli from the same set of voxels by running two separate general linear models: one in which the five predictors were coded according to the positions of the attended stimuli, and one in which they were coded according to the positions of the ignored stimuli (Supp. Fig. 4.2). In each GLM, we separately contrasted the five predictors against the fixation baseline condition to obtain five maps of t-values. The result was a set of maps corresponding to the five attended positions and another set corresponding to the five ignored positions. Because the positions of the attended and ignored stimuli varied independently of each other, information about one of the dimensions was randomly distributed in the maps corresponding to the other dimension. Our subsequent pattern analysis tested for systematic variation in these activity maps as a function of stimulus position.

We used a cross-correlation approach (Haxby, Gobbini et al. 2001; Haushofer, Livingstone et al. 2008; Kriegeskorte, Mur et al. 2008; Fischer and Whitney 2009; Fischer, Spotswood et al. 2011) to test for position selectivity in the BOLD response in the pulvinar. The analysis (Fig. 4.2) tested whether stimuli presented closer together in space produced more similar patterns of activity. The center (position 3) map was used

as a baseline, and we compared the four other positions with it by cross-correlating the activity patterns (Fig. 4.2a). Cross-correlating all of the activity patterns for all stimulus positions against each other yielded similar results (Supp. Fig. 4.3), but using the central map as a baseline is more conservative because it avoids making assumptions about the linearity of the relationship between physical stimulus separation and correlation, as well as the independence of points in the plot (for more information, see Supplementary Methods). In the plots in Figure 4.2a, each point represents one voxel within the region of interest, its value from one BOLD map plotted on the x-axis (t units), and its value from a second BOLD map plotted on the y-axis. The correlation across voxels between a given pair of BOLD maps served as a measure of their similarity within that region of interest (ROI); we plotted the correlation between each pair as a function of the distance between the corresponding stimuli in visual space (Fig. 4.2b). A significant negative slope on this plot indicates precise discrimination of the stimulus positions—the multivariate pattern of response within the ROI in question changed systematically with changes in stimulus position. This analytical approach allowed us to separately measure the encoding of the attended and ignored stimuli within the pulvinar without making any assumption about the nature of the underlying topographic organization there (see Discussion). We performed this analysis in pulvinar ROIs defined based on the Talairach Atlas (Talairach and Tournoux 1988) but fit to each subject's individual anatomy (Fig. 4.2c; see Methods).

Figure 4.3a shows position discrimination for attended objects (red data) and ignored objects (blue data) in the pulvinar. Data from the five subjects are plotted together and fit with a regression model that included a random effect of subject(Holmes and Friston 1998) (see Methods). Both the left and right pulvinar showed significant coding of the attended stimulus positions, as indicated by significant negative slopes in the regression fits (Fisher $z = -0.49$, $p = 0.011$ in left pulvinar and $z = -0.67$, $p = 0.001$ in right pulvinar; $n = 20$ for each test; p values were assessed using a nonparametric bootstrap test described in the Methods). All subjects showed the same trend individually. However, there was no detectable encoding of ignored objects ($z = 0.11$ in left pulvinar and $z = 0.05$ in right pulvinar, both $p > 0.5$; $n = 20$ for each test; nonparametric bootstrap test). Figure 4.3b shows the same analysis performed in bilateral pulvinar ROIs; discrimination of attended positions was significantly better than discrimination of ignored positions ($p = 0.01$; $n = 20$ points in each plot; bootstrap test for a difference in model fits between attended and ignored). It is important to note that the data in the attended and ignored plots came from the same functional runs, analyzed within the same set of voxels, with the same stimuli present; no difference in signal-to-noise ratio, retinotopically-specific adaptation, or other aspects of signal quality could explain the difference in coding precision we found for attended and ignored objects in the pulvinar (see Supplementary Methods). These data demonstrate an essentially absolute gating of spatial information in the pulvinar by attention: only attended positions are encoded in the multi-voxel pattern of activity within the pulvinar.

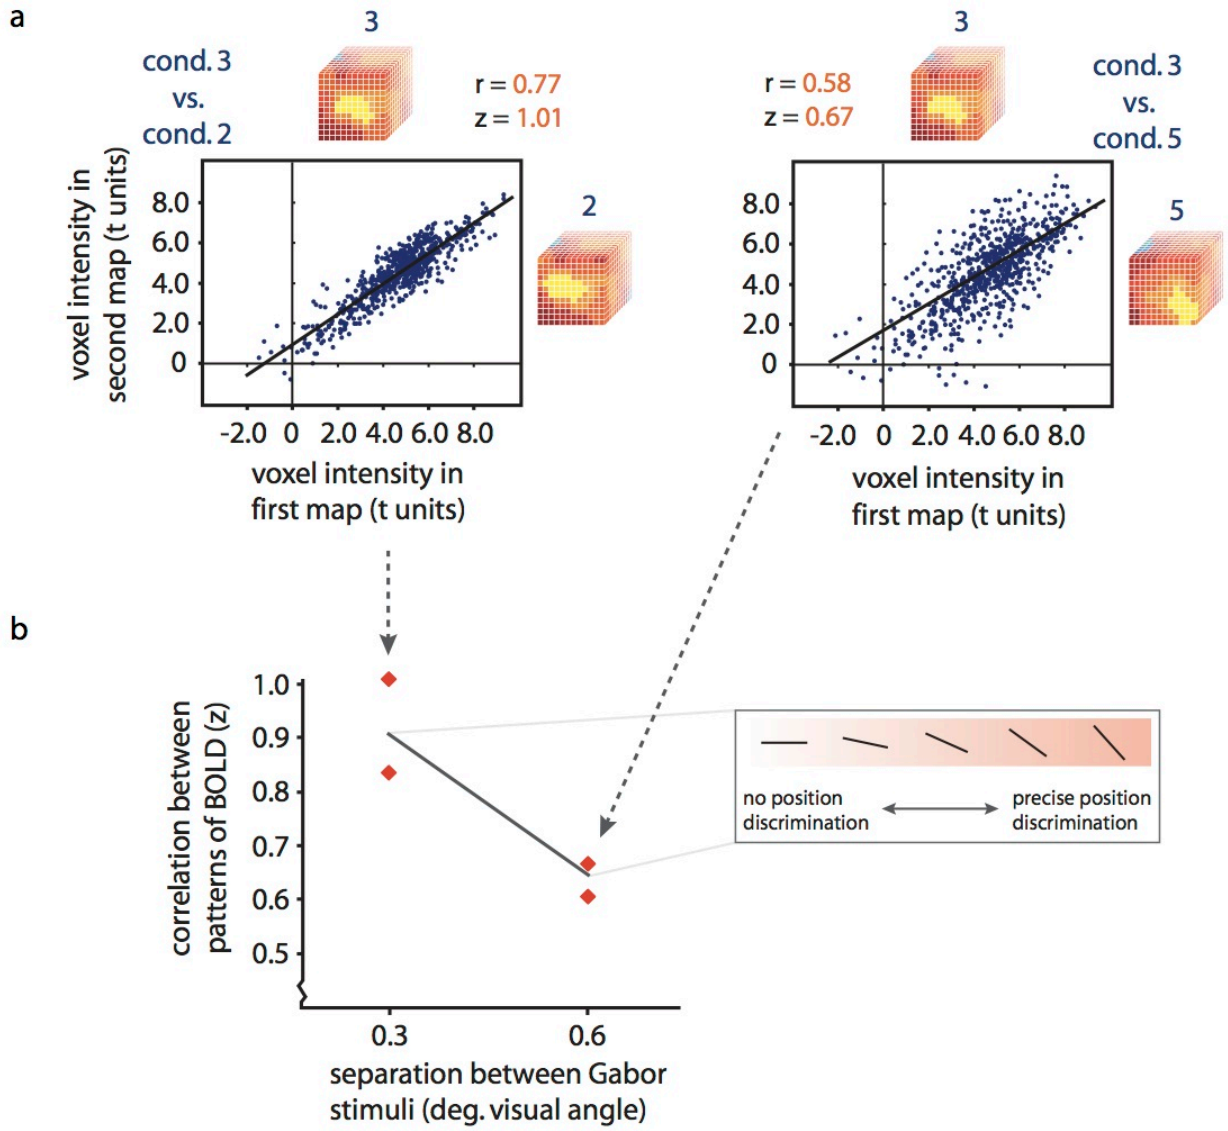
a
3
cond. 3
vs.
cond. 2
r = 0.77
z = 1.01
2
voxel intensity in
second map (t units)
8.0
6.0
4.0
2.0
0
-2.0 0 2.0 4.0 6.0 8.0
voxel intensity in
first map (t units)
3
r = 0.58
z = 0.67
cond. 3
vs.
cond. 5
5
8.0
6.0
4.0
2.0
0
-2.0 0 2.0 4.0 6.0 8.0
voxel intensity in
first map (t units)
b
correlation between
patterns of BOLD (z)
1.0
0.9
0.8
0.7
0.6
0.5
0.3
0.6
separation between Gabor
stimuli (deg. visual angle)
no position
discrimination
precise position
discrimination
c
Talairach
Adjusted
y = -24
R
L
y = -24
R
L

Figure 4.2: Correlation-based analysis of position selectivity. a) We computed the pairwise correlations between the patterns of activity produced by the five stimulus positions (the central condition vs. each of the other conditions). Each point on the plot represents one voxel, its value from the first BOLD map plotted against its value from the second BOLD map. We then tested for a trend in the correlations indicating position selectivity by plotting them as a function of the visual distance between the corresponding stimuli **(b)**. A negative trend in the plot of correlation values indicates that as the stimuli became more retinotopically distinct, they produced more distinct multivariate patterns of activity(Fischer and Whitney 2009). A steeper negative slope on the plot indicates more precise encoding of stimulus position. **c)** To generate regions of interest for the pulvinar, we used the coordinates for the pulvinar drawn from the Talairach Atlas(Talairach and Tournoux 1988) (shown in blue), and made minor adjustments to match the ROIs to each subject's individual anatomy (an example adjusted ROI is shown in green; see Methods).

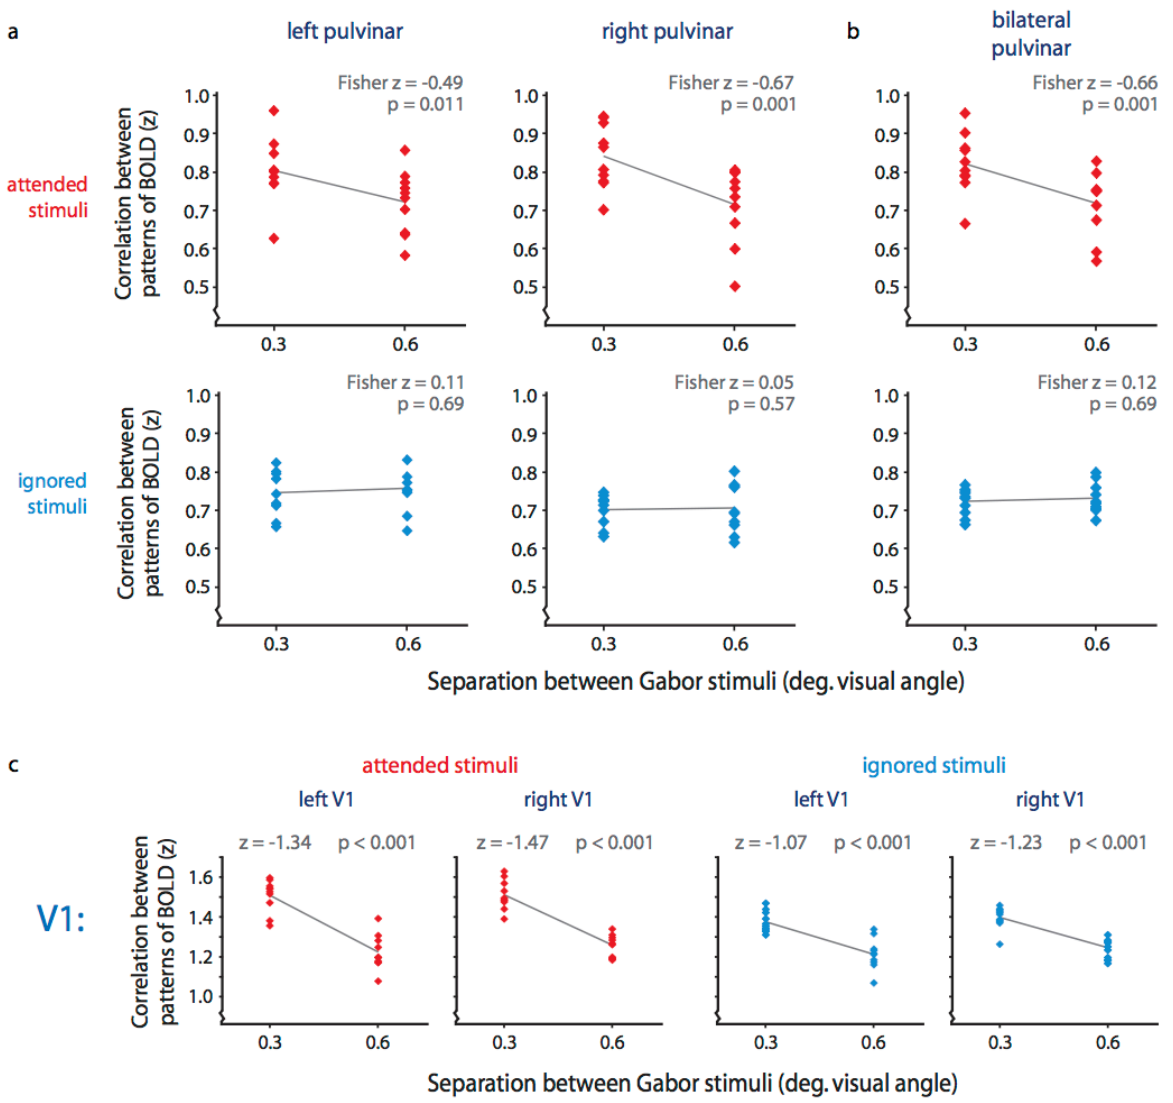

Figure 4.3: Attention gates visual encoding in the pulvinar. a) Correlation-based analysis for left and right pulvinar, using the BOLD maps for the attended positions (red data) and the ignored positions (blue data). Each subject contributes four points to each plot – two independently computed correlations at each separation. Because multiple points were plotted for each subject, we fit the data with a regression model that included a random effect of subject(Holmes and Friston 1998). Correlations are presented as Fisher z scores for the sake of linear comparison. In both hemispheres, there was significant discrimination of the attended positions (Fisher z = -0.49, p = 0.011 in left pulvinar; z = -0.67, p = 0.001 in right pulvinar; n = 20 for each test; nonparametric bootstrap test), but no discrimination of the ignored positions (z = 0.11, p = 0.69 in left pulvinar; z = 0.05, p = 0.57 in right pulvinar; n = 20 for each test; nonparametric bootstrap test). **b)** Position discrimination analyzed in bilateral pulvinar ROIs. Discrimination of attended positions was significantly better than discrimination of ignored positions (p = 0.01; n = 20 points in each plot; bootstrap test for a difference in model fits between attended and ignored). **c)** In primary visual cortex (V1), there was significant discrimination of both the attended and ignored stimulus positions, consistent with previous results(Fischer and Whitney 2009). Attentional modulation of position discrimination ($z_{attended}$ minus $z_{ignored}$) was significantly greater in the pulvinar than in V1 (p = 0.04; n = 20 points in each plot; bootstrap test for a greater attended-ignored difference in the pulvinar vs. V1).

If the attentional gating of spatial information we observed in the pulvinar is related to filtering out distractors, we would expect measurable encoding of the ignored stimulus positions at earlier stages of visual processing, including in the input signals to the pulvinar. We repeated the same analysis in seven independently localized visual areas: V1, V2, V3, V3a, VP, V4v, and MT+ (see Methods for localizer details). Figure 4.3c shows position discrimination plots, constructed in the same way as those shown for the pulvinar, for V1: V1 exhibited precise coding of both the attended and the ignored stimuli (attended: left – z = 1.34, right – z = 1.46; ignored: left – z = 1.07, right – z = 1.23; all $p < 0.001$; n = 20 for each test; nonparametric bootstrap test). In contrast to the pulvinar, each of the seven comparison visual areas showed significant discrimination of the ignored stimulus positions (nonparametric bootstrap test with n = 20 for each test; all $p < 0.02$; significant with FDR control for multiple comparisons (Benjamini and Hochberg 1995) with q = 0.05). Only the pulvinar displayed a complete gating of the ignored object positions, and attentional modulation of position coding was significantly stronger in the pulvinar than in the visual cortex ($p = 0.008$; n = 20 points in each plot; bootstrap test for a greater attended-ignored difference in the pulvinar vs. visual cortex). Of particular note is the significant encoding of ignored positions in MT+, which is thought to have a high density of driving input to the pulvinar (Grieve, Acuña et al. 2000).

4.2.2 Experiment 2 – Orientation discrimination

We sought to test whether attention gates types of stimulus information besides position coding in the pulvinar. We also sought to verify that attentional gating in the pulvinar happens for task-relevant stimulus features (in Experiment 1, subjects attended at the locations of the cued stimuli, but were not required to make a position-related judgment). Three subjects participated in this second experiment. The experimental design was similar to the first experiment, but now subjects were required to detect orientation changes in the cued Gabors. Each Gabor was oriented at either -45 deg. or +45 degrees relative to vertical, and the orientations in the upper and lower visual fields varied independently of each other from block to block (Fig. 4.4a). The left and right Gabors within the upper or lower visual field always had opposite orientations. An average of eight times per block, one of the four Gabors changed in orientation by 8 deg., for 200 ms At the beginning of each run, subjects were instructed to attend to the orientations of either the upper two or lower two Gabors, and respond any time they detected an orientation change in one of the cued Gabors but to ignore orientation changes in the other two. Sensitivity (d') for the orientation change detection task was 1.15; the task was challenging but above threshold.

We measured orientation coding within the pulvinar and visual cortex separately for the attended and ignored stimuli, which were simultaneously present at all times. To measure orientation coding, we used a support vector machine (SVM) pattern classifier to predict which orientation had been presented on each block (see Methods). We used SVM in the second experiment rather than the correlation analysis because the power of the latter is diminished when discriminating between just two stimulus categories rather than a multi-level parametric stimulus manipulation (in the analysis of the first

experiment, though we paired stimuli according to two possible separations, we utilized all five stimulus conditions); SVM, however, is well-suited to this case of binary classification (Mourão-Miranda, Bokde et al. 2005).

Classification results are shown for the pulvinar and early visual cortex in panels **b** and **c** of Figure 4.4, respectively (data are collapsed across the three subjects). In the pulvinar, we found robust encoding of the attended stimulus orientations (54.7% correct; 7.8 standard deviations above the mean of the bootstrapped chance distribution; $p < 0.001$) but no evidence of encoding of ignored orientations (49.9% correct; 0.34 standard deviations below bootstrapped chance; $p = 0.62$). Classification of the attended orientations was significantly better than classification of ignored orientations ($p < 0.001$; nonparametric bootstrap test for whether the attended-chance difference was larger than the ignored-chance difference). On the other hand, in the visual cortex we found significant classification of both attended and ignored stimulus orientations (56.3% correct for attended and 54.3% correct for ignored, 10.1 and 6.8 standard deviations above bootstrapped chance, respectively; both $p < 0.001$). Consistent with previous findings(Kamitani and Tong 2005; Serences, Saproo et al. 2009; Jehee, Brady et al. 2011)see Discussion), attended orientations were classified significantly better than ignored orientations in early visual areas ($p = 0.016$; bootstrap test for [attended-chance] > [ignored-chance]), but the modulation of orientation coding by attention was significantly stronger in the pulvinar than in early visual areas ($p = 0.012$; bootstrap test for [$\text{pulvinar}_{\text{attended}}$-$\text{pulvinar}_{\text{ignored}}$] > [early $\text{vis}_{\text{attended}}$-early $\text{vis}_{\text{ignored}}$]). These results demonstrate that attention gates not only spatial representations but also orientation information within the pulvinar. They also demonstrate attentional gating of a task-relevant feature within the pulvinar. Collectively, our results show a gating of ignored information within the pulvinar that generalizes across at least two stimulus features (position and orientation) and is distinct from the pattern of attentional modulation we observed in the visual cortex.

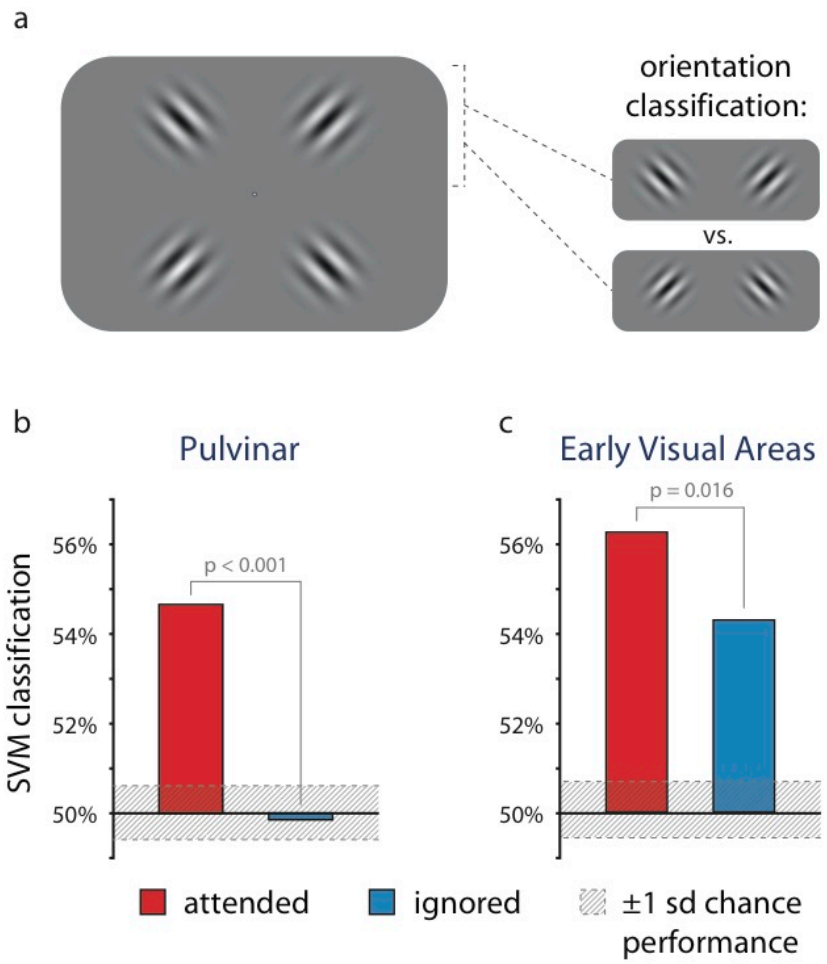

Figure 4.4: Attention gates orientation encoding in the pulvinar. a) Stimuli for the orientation decoding control experiment. Gabor patches were presented in the four visual quadrants; the Gabors could be oriented at -45 or +45 deg. with respect to vertical, and the orientations in the upper and lower visual fields varied independently (left and right Gabors within the upper or lower visual field were always oriented orthogonally). Subjects attended for brief orientation changes in the cued Gabors, ignoring the others. **b)** Classification of attended and ignored orientations within the pulvinar. ±1sd of chance classification, estimated by classifying with shuffled labels, is shown by the shaded region (see Supplementary Methods). There was reliable classification of attended orientations (red bar; 54.7% correct; 7.8 standard deviations above the mean of the bootstrapped chance distribution; $p < 0.001$), but chance classification of ignored orientations (blue bar; 49.9% correct; 0.34 standard deviations below bootstrapped chance; $p = 0.62$) in the pulvinar (data shown are collapsed across three subjects). Classification of attended orientations was significantly better than classification of ignored orientation ($p < 0.001$; nonparametric bootstrap test for whether the attended-chance difference was larger than the ignored-chance difference). **c)** In early visual areas (V1 through MT+), orientation classification was significantly above chance for both attended and ignored Gabors (56.3% correct for attended and 54.3% correct for ignored, 10.1 and 6.8 standard deviations above bootstrapped chance, respectively; both $p < 0.001$). Attended orientations were classified significantly better than ignored orientations in early visual areas ($p = 0.016$; bootstrap test for [attended-chance] > [ignored-chance]), but the modulation of orientation coding by attention was significantly stronger in the pulvinar than in early visual areas ($p = 0.012$; bootstrap test for [$\text{pulvinar}_{\text{attended}}$-$\text{pulvinar}_{\text{ignored}}$] > [$\text{early vis}_{\text{attended}}$-$\text{early vis}_{\text{ignored}}$]).

4.3 Discussion

Our results show that stimulus encoding in the human pulvinar is gated by attention: when behaviorally relevant visual stimuli compete with distractors for attentional resources within the same visual field, the pulvinar precisely represents attended but not ignored objects. Thus, activity in the pulvinar reflects an isolated representation of attended targets, which can serve to selectively confer further processing benefits on important stimuli.

Previous studies (Cotton and Smith 2007; Smith, Cotton et al. 2009; Schneider 2011) have reported significant fMRI responses in the pulvinar to unattended stimuli. For example, when Smith et al. (2009) presented subjects with optic flow stimuli and examined the BOLD response in the pulvinar when attention was directed either toward or away from the stimuli, they found significant pulvinar activation in both the attended and unattended cases, with a ~20% amplitude modulation with attention. How do our present results square with these previous findings? First, it is important to distinguish between the amplitude of the raw BOLD response and information content in fMRI patterns of activity, which can be dissociated (Haynes and Rees 2006; Jehee, Brady et al. 2011; Tong and Pratte 2012); our study takes the important step of decoding position and orientation information within the multivariate pattern of activity in the pulvinar, providing converging evidence for attentional gating from two information-based MVPA approaches. Equally important, though, is a key difference in experimental design between the above-mentioned studies and our study. The spatial attention manipulations in both Smith et al. (2009) and Schneider (2011) had attended and ignored stimuli on opposite sides of the horizontal meridian, and thus in different hemispheres in the pulvinar. Given the evidence for at least partially independent attentional resources in the left and right hemispheres (Luck, Hillyard et al. 1989; Alvarez and Cavanagh 2005), one possibility is that a left/right attentional manipulation induces less competition between attended and unattended stimuli than if both stimuli fell in the same visual field. Thus, a key aspect of our design is the upper/lower manipulation of attention, forcing targets and distractors to compete for attentional resources within the same visual field and same hemisphere.

Previously (Fischer and Whitney 2009), we reported position sensitivity in the pulvinar under passive viewing conditions (attention was engaged at the fixation point by a task unrelated to the stimuli). Critically, the stimuli used to measure spatial discrimination were neither targets nor distractors—they were irrelevant to the fixation task that subjects performed. Other studies of retinotopy in the healthy human pulvinar have used undivided attention (Kastner, O'Connor et al. 2004; Cotton and Smith 2007) or featural attention (Schneider 2011), but our study is the first to test how the healthy human pulvinar responds to *competing* stimuli in the same visual field. We introduced competition between simultaneously presented stimuli, which resulted in the ignored objects being gated out from encoding in the pulvinar. Unlike putative attentional source regions such as the IPS areas (Silver, Ress et al. 2005), which show only weak responses to passive visual stimulation, the pulvinar precisely encodes passively viewed stimuli (Fischer and Whitney 2009) but not *distracting* ones.

We intentionally analyzed ROIs that encompassed the entire pulvinar. While the primate pulvinar has been classically divided into several subnuclei based on cytoarchitecture (Kaas and Lyon 2007), no structural or functional homology between the monkey and human pulvinars has been firmly established (Snow, Allen et al. 2009). Further, interactions between attentional signals and visual maps in the pulvinar likely involve more than a single subregion. While the lateral and inferior portions of the monkey pulvinar are the most visually responsive (Bender 1981; Casanova 2004), the pulvinar's connections with areas in the fronto-parietal attention network project from the medial portion (Hardy and Lynch 1992). We therefore considered the multivariate pattern of activity within the pulvinar as a whole (see Methods).

Our results dovetail with recent reports that visual awareness is correlated with variability in the strength of responses in the pulvinar (Wilke, Mueller et al. 2009; Padmala, Lim et al. 2010). For example, Padmala et al. (2010) used affective conditioning to increase the salience of otherwise neutral stimuli. Then, presenting conditioned and non-conditioned stimuli in a rapid serial visual presentation task during scanning, they found that stimuli detection was correlated with larger responses in the pulvinar; this effect was much larger for affectively conditioned stimuli. Our present results are complementary and demonstrate the flip side of the coin: we show that distracting stimuli are gated out from encoding in the pulvinar, while they demonstrated that the more salient an attended target is, the more robust a response it produces in the pulvinar, particularly if it has conditioned affective significance. Future work needs to study whether and how affective significance interacts with the attentional gating effect we report here.

Our orientation classification results complement and extend previous work on decoding orientation information from the human visual cortex (Kamitani and Tong 2005; Serences, Saproo et al. 2009; Jehee, Brady et al. 2011). Prior work has shown that attending to an orientation biases signals in the visual cortex to preferentially represent that orientation, yielding a boost in classification accuracy (Kamitani and Tong 2005; Serences, Saproo et al. 2009). This attention-related enhancement can occur selectively for task-relevant features (Jehee, Brady et al. 2011), and is carried at least in part by a boosted response within the subpopulation of voxels that is tuned to the attended orientation (Serences, Saproo et al. 2009). In Experiment 2, we similarly found that attention significantly improved the precision of orientation coding in early visual cortex (Fig. 4.4c). We extended the classification framework to measure orientation-selective responses in the human pulvinar for the first time, and found a significantly larger attentional modulation of orientation information in the pulvinar than in the visual cortex (Fig. 4.4b), such that there was no measurable information about the ignored orientations, even though the ignored and attended stimuli differed only in their task-relevance. The gating of orientation information in the pulvinar by attention sets the pulvinar apart from the pattern of results found in the visual cortex, and lends support to the hypothesis that the pulvinar plays an important role in distractor filtering.

Selective attention reduces the deleterious influences of distracting information on perception and behavior (Treisman 1969; Desimone and Duncan 1995; Mangun

1995; Carrasco 2011) and is therefore a critical mechanism for many visuomotor and cognitive functions. The pulvinar has long been suspected to be a neural interface where attention signals can gate out distractors, but this filtering function has never been observed in action in the pulvinar. Our results show that in the healthy human pulvinar, visual distractors are gated out from spatial coding, leaving a map that can serve to isolate behaviorally relevant objects and their features from among competing information.

4.4 Methods

4.4.1 Subjects

Seven healthy subjects, ranging in age from 22 to 31 years, participated in the study. All subjects provided written consent prior to participation, and all experimental procedures were approved by the UC Davis and UC Berkeley institutional review boards.

4.4.2 Experimental design

Before each scanning run, subjects were cued to attend to the Gabors in either the upper or lower visual field, and the subject attended at that location for the duration of the entire run. Attend-upper and attend-lower runs were interleaved so that no trends across the scanning session could differentially affect the two conditions. Within a run, the Gabor positions in the upper and lower visual fields were random with respect to each other, and the ordering of the conditions within the upper and lower visual fields was pseudo-randomized by shuffling within a group of the five position conditions plus a baseline condition, and then presenting six such shuffled groups in sequence. The Gabors (0.4 cycles/deg, peak contrast 87% (Michelson), 1.66 deg. sd envelope, 8 Hz counterphase flicker) appeared at five possible eccentricities, ranging from 8.40 deg. to 9.60 deg. in increments of 0.3 deg. Attend-upper and attend-lower runs were constructed in a completely identical fashion; the only difference between the two run types was the attended location. This upper/lower attention manipulation allowed the positions of the attended and ignored Gabors to vary independently of each other, and for both attended and ignored stimuli to be represented in each hemisphere. Stimuli were presented in 10 second blocks. We used a blocked design, rather than a rapid event-related design, to maximize the power in the position discrimination analysis (see Supplementary Methods).

In Experiment 2, we measured orientation discrimination rather than position discrimination. Subjects viewed four Gabors (0.4 cycles/deg, 87% peak Michelson contrast, 1.66 deg. sd envelope, 8 Hz counterphase flicker), one in each visual quadrant (Fig. 4.4a). The orientations of the Gabors in the upper and lower visual fields varied independently (Fig. 4.4a), and the positions of the Gabors were randomly jittered from block to block with a standard deviation of .47 deg. (mean 9.0 deg. eccentricity), independently in the upper and lower visual fields and independently of the orientations, in order to ensure that orientation decoding generalized across changes in position and

was not simply capitalizing on localized luminance differences between orientations. Successful performance on the orientation change detection task (see Results) required attending to both Gabors in the cued location (upper or lower visual field), while ignoring the Gabors in the non-cued field as completely as possible. As in Experiment 1, the stimuli in attend-upper and attend-lower runs were constructed in a completely identical fashion; the only difference was the verbal instruction at the beginning of the run.

4.4.3 fMRI acquisition

Scanning was conducted on a Siemens Trio 3T system at the UC Davis Imaging Research Center. Functional images were collected with a gradient-recalled echo EPI sequence using an eight-channel phased-array head coil. Whole-brain anatomical images were acquired with a high resolution (1 mm^3) Turbo Spin Echo scan. The acquisition parameters were: TR = 2000 ms, TE = 26 ms, FA = 90 deg, FOV = 22 x 22 cm^2, voxel size = 1.528 x 1.528 x 2.5 mm^3. The 24-slice imaging volume was centered on the thalamus and ran through the occipital pole. Runs were 360 seconds long (180 volumes). Using a Digital Projection Mercury 5000HD projector, stimuli were back-projected at 75 Hz onto a semi-transparent screen from outside the bore. Subjects were able to see the screen and stimuli via a mirror angled at 45 degrees, located 9.5 cm directly above their eyes. Each subject participated in eight main experimental runs and one additional retinotopic localizer run.

4.4.4 Data analysis

fMRI data preprocessing, ROI definition, and general linear models were performed in BrainVoyager QX v. 2.0; all other analysis was performed in Matlab R2008a. Data preprocessing consisted of slice scan time correction with cubic spline interpolation, 3D motion correction with trilinear interpolation, and high-pass filtering with a cutoff of 2 cycles/run. The data were not otherwise spatially smoothed. The data were normalized to Talairach space to allow for atlas-assisted definition of the pulvinar ROIs (see Fig. 4.2c).

Experiment 1 – Position discrimination: For each functional run, we constructed two general linear models, one with the five attended stimulus positions as five predictors, and the other with the five ignored stimulus positions as predictors. In addition, there was a sixth predictor for the baseline (fixation only) blocks. Supplementary Figure 4.2 depicts the GLM recoding procedure. For each GLM, we separately contrasted each of the five stimulus position predictors against the fixation baseline to generate five maps of BOLD response (t values) for use in the correlation analysis (Fig. 4.2).

In the group analysis of position discrimination, we used a random effects model to account for between-subject variability (Holmes and Friston 1998). The regression model took the form $z_{ijk} = \beta_0 + \beta x_{jk} + U_i + \varepsilon_{ijk}$, where *i* indexed the subjects, the pair (*j*, *k*) indexed the stimulus pairings, and U_i accounted for intercept differences between

subjects. This approach is comparable to a paired t-test and yields identical p values; we express the model in the general regression framework in order to make clear how the random effect of subject is incorporated into the model, and to visualize the slope of the data across stimulus separations. For significance testing, we generated bootstrapped confidence intervals for the group regression by sampling with replacement within each subject on 2000 iterations. See Supplementary Methods for more details.

Experiment 2 – Orientation discrimination: To measure the encoding of orientation information in the pulvinar and visual cortex, we used a support vector machine (SVM) classification analysis implemented in Matlab using the LIBSVM library (Lin and Chang 2011). Classification was run on block-by-block beta weights, rather than the raw BOLD timecourse, to improve SNR in the classification and account for hemodynamic lag. Within each subject and each region of interest, we concatenated the block-wise beta maps from all runs, and ran 1000 iterations of SVM classification with a linear kernel, each time training on a randomly selected 80% of blocks and testing on the remaining 20%. Overall classification accuracy was taken as the average classifier performance across the 1000 classification iterations. A label shuffling procedure confirmed that the training process did not "peek" at the test data (see Supplementary Methods for more details).

4.4.5 ROI definition

No functional localizer exists for the pulvinar, and portions of the pulvinar can be difficult or impossible to delineate based on anatomical scans alone. Because of this, we based our pulvinar ROI definitions on standardized coordinates drawn from the Talairach Atlas (Talairach and Tournoux 1988). To define individual pulvinar regions of interest for each subject, we began with standard pulvinar regions of interest drawn from the Talairach atlas (shown in blue in Fig. 4.2c), and made minor manual corrections based on each subject's anatomy. Corrections were only made where the atlas-based ROIs encroached on ventricle space or clearly disobeyed gray/white matter boundaries. We used the coronal plane running through the posterior commissure as a conservative anterior boundary to ensure that no gray matter outside the pulvinar was included the ROIs. The result was a left- and right- pulvinar ROI for each subject that fit the individual anatomy while staying close to the atlas coordinates (example ROIs shown in green in Fig. 4.2c).

To localize the early retinotopic visual areas in each subject, subjects participated in a separate scanning run consisting of standard localizer stimuli including flickering checkerboard wedge and drifting stimuli presented along the horizontal and vertical meridians (Sereno, Dale et al. 1995). A map that contrasted the horizontal and vertical meridians was overlaid on an inflated surface map for each subject, and the meridians were used to delineate the early retinotopic areas.

4.5 Supplementary Methods

4.5.1 Experimental design & analysis

In both experiments, we presented the stimuli in a blocked design. We recently explored the power of event-related vs. blocked designs for the correlation-based position discrimination analysis (Fischer and Whitney 2009), and found that while both yielded qualitatively similar results, the blocked design was better suited to generating high SNR maps for the correlation analysis.

An important aspect of the design for both experiments is the simultaneous presence of attended and ignored stimuli, coupled with the up/down attention manipulation. Not only does the simultaneous presentation of targets and distractors induce competition that may be critical to engaging the pulvinar (Desimone, Wessinger et al. 1990; LaBerge and Buchsbaum 1990; Laberge, Carter et al. 1992; Snow, Allen et al. 2009), it also allowed us to measure the precision of position coding for attended and ignored stimuli within the same set of voxels, and within the same functional runs. We conducted our analysis within regions of interest encompassing the entire pulvinar, and did not preselect the most active voxels as is sometimes done in classification analyses. Accordingly, the data that entered into the analyses for the attended and ignored locations was identical. Signal-to-noise ratio, vascular distribution, motion artifacts, and any other scanning artifacts were identical for the analyses of the attended and ignored stimuli, and cannot explain systematic differences between the precision of information that we measured in the two dimensions.

The up/down manipulation of attention (rather than left/right) was critical to the experimental design. Since the human pulvinar responds to stimulation in the contralateral visual field (Cotton and Smith 2007; Schneider 2011), by placing both a target and a distractor within the same visual hemifield, we were able to test for encoding of attended and ignored stimuli within the *same* hemisphere, in the *same* functional runs. A left/right attention manipulation would confound the attended location with visual field; hence any comparison of the encoding of attended and ignored stimuli within the same hemisphere would have had to be done across runs, for stimuli that were not presented simultaneously. Our design excludes the possibility that any differences across runs or hemispheres contributed to the attentional gating we measured in the pulvinar.

4.5.2 Position discrimination analysis

In the position discrimination analysis, we used the central (9 deg. eccentricity) stimulus condition as a baseline for several reasons. First, doing so allowed us to fit a linear model to the resulting correlation plots. Our previous modeling has shown that the correlation between two patterns of BOLD response falls off nonlinearly with increasing separation between the corresponding stimuli (Fischer and Whitney 2009). By using the central map as a baseline, we use each of the four non-baseline maps exactly once, and reduce the risk of a nonlinearity or floor effect in the correlations washing out our ability to measure position discrimination. Second, using a single baseline increases the

independence of correlations in the regression and reduces autocorrelations in the residuals (which would otherwise be a violation of the linear regression). Finally, and most importantly, the attentional modulation (gating) interaction (Fig. 4.3) still holds even when all position separations are analyzed ($p = 0.004$; bootstrap test for $z_{attended} > z_{ignored}$; $n = 50$ points for attended data and 50 points for ignored data). The results are therefore not contingent on the use of a baseline condition. Using the central condition as a baseline, however, is a more conservative approach for the reasons described above and is therefore adopted in the analysis in Fig. 4.3.

A potential influence on our position discrimination measurements is retinotopically-specific adaptation—perhaps responses to the central stimulus position were reduced due to repeated stimulation at that location. In a control analysis, to compare the amplitude of the raw BOLD response across the five stimulus positions, we conducted ANOVAs for the five attended and five ignored conditions separately. The input to the ANOVAs was the average BOLD response within the pulvinar (t values), for each of the five stimulus positions, across 8 runs, for two hemispheres per subject, totaling 400 measurements. Each ANOVA used a mixed-effects model with fixed effects of hemisphere and stimulus position, and a random effect of subject, as well as the three two-way interactions among these factors. We found no adaptation effects in the raw BOLD response in the pulvinar: the BOLD response did not differ significantly across the five stimulus positions within the maps for the attended or ignored stimuli (main effect of stimulus position for attended stimuli: $F_{4,395} = 1.24$, $p = 0.29$; for ignored stimuli: $F_{4,395} = 0.65$, $p = 0.63$; based on five stimulus positions and 400 total measurements; no significant interactions with hemisphere or subject).

To establish a bootstrapped null distribution against which to compare the strength of the correlations we measured in the pulvinar and V1 (Supp. Fig. 4.3), we performed a label shuffling procedure. On each of 1000 iterations, we randomly shuffled the trial labels in the GLM design matrix for each run, then recomputed the BOLD response map for each of the five stimulus positions. The shuffling included randomizing the fixation baseline labels, so as to remove all stimulus-related information from the GLM predictors. From the resulting five maps of BOLD response, we computed the pairwise correlations among the maps in the same way that we did for the position discrimination analysis in Experiment 1 (see Fig. 4.2), and recorded the mean correlation across all runs and all map pairs. Repeating this procedure 1000 times, we arrived at a distribution of correlation measurements that reflect a baseline level of correlation among response maps in our data set, independent of any stimulus-related information. This baseline level of correlation represents a floor for the correlations we might expect to find in the position discrimination analysis, and may result from factors such as partial voluming or local vascular structure.

4.5.3 Orientation classification analysis

The training stage of the classification analysis included a feature selection step that estimated an appropriate variance threshold to apply to the data before classifying (voxels below this threshold were excluded when classifying); the same variance

threshold was then applied when testing the performance of the model. This training process, including variance threshold estimation, was performed completely independently of the test data, and introduced no information about the test data as confirmed by the label shuffling procedure described below.

Theoretical chance for the 2-class classification in Experiment 2 is 50%. To determine an empirical chance distribution for significance testing, we repeated the classification procedure 1000 times, each time shuffling the label assignments of the blocks (after running a GLM to recover the block-by-block beta weights, but before performing the classification step). The resulting distributions (measured separately for each subject and each ROI) reflected the range of classifier performance to be expected by chance. To test the significance of the classification accuracy for an ROI, we compared the actual classification accuracy with this permuted bootstrapped chance distribution, computing the proportion of the chance distribution that was larger than the actual classification accuracy.

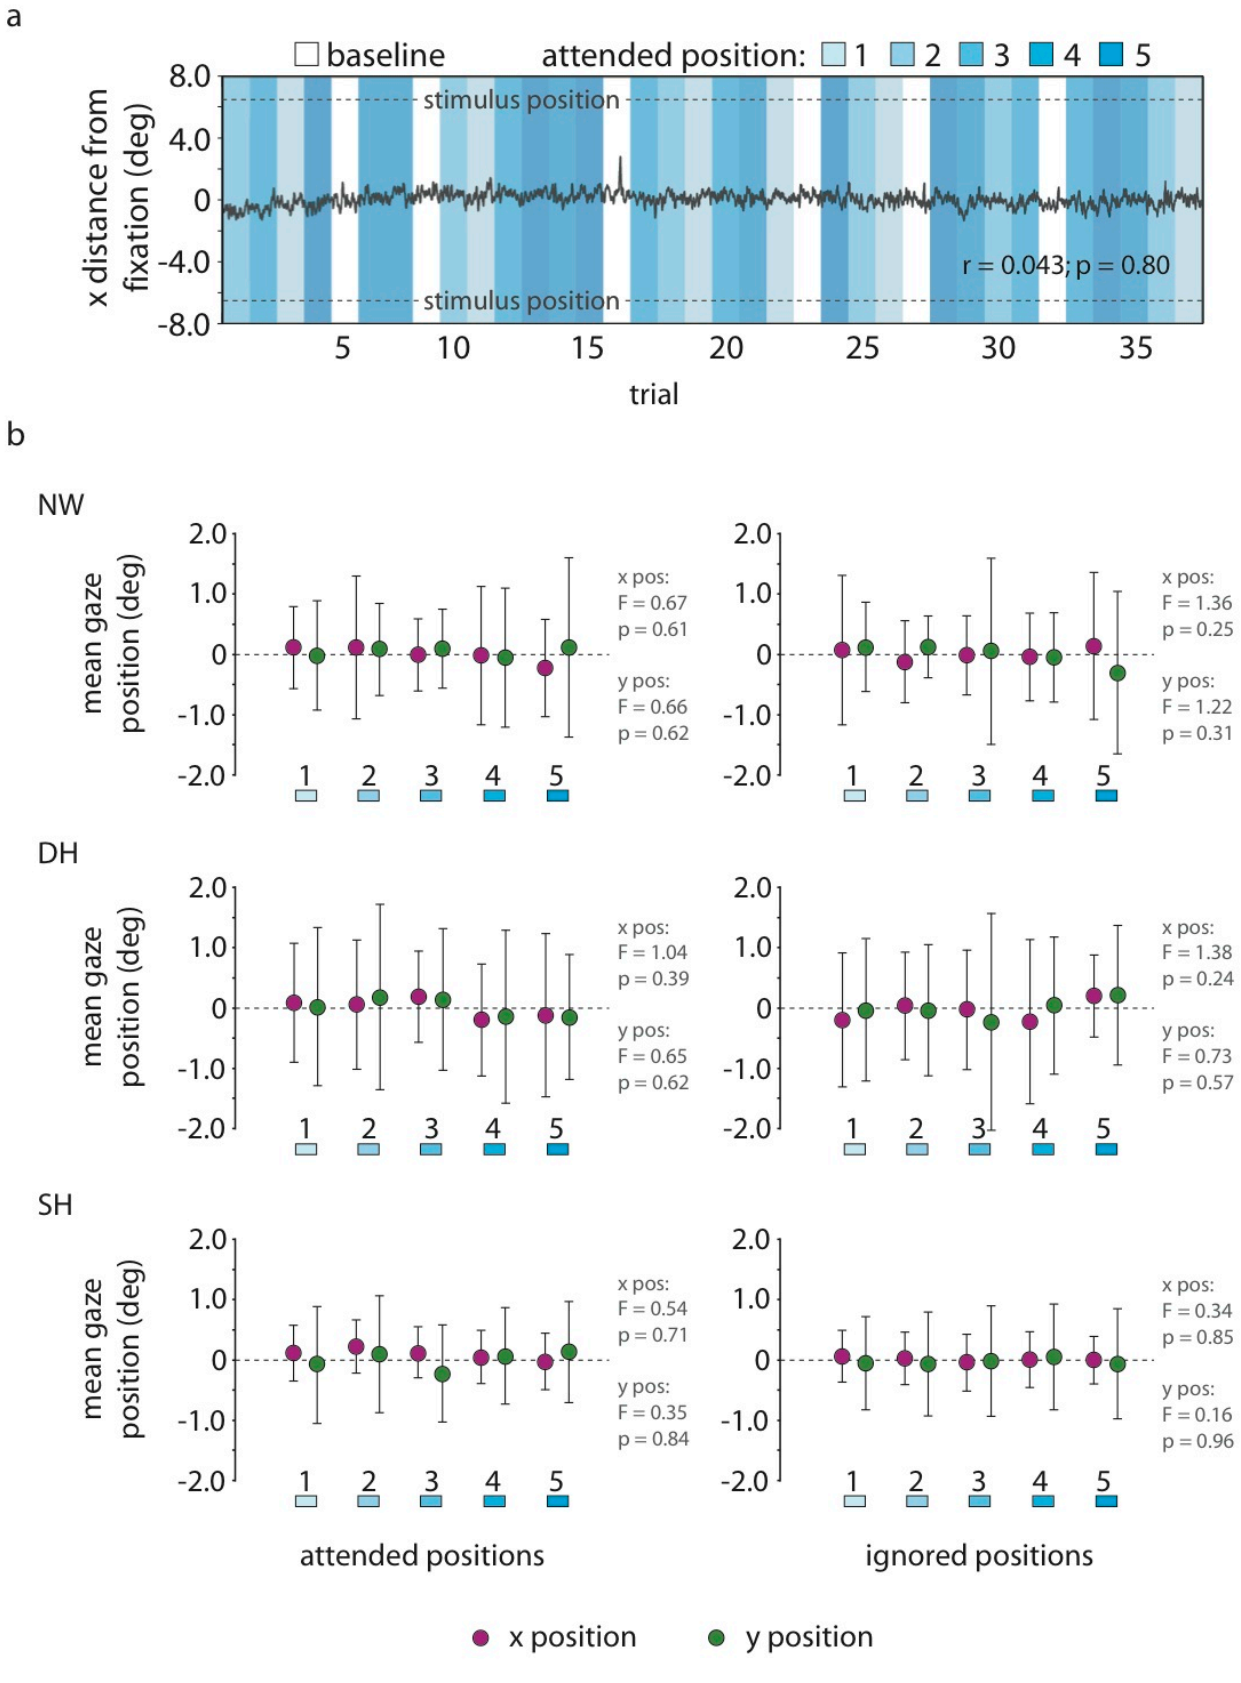
a
baseline
attended position: 1 2 3 4 5
stimulus position
stimulus position
x distance from fixation (deg)
8.0
4.0
0
-4.0
-8.0
5
10
15
20
25
30
35
trial
r = 0.043; p = 0.80
b
NW
DH
SH
mean gaze position (deg)
2.0
1.0
0
-1.0
-2.0
1
2
3
4
5
x pos:
F = 0.67
p = 0.61
y pos:
F = 0.66
p = 0.62
x pos:
F = 1.36
p = 0.25
y pos:
F = 1.22
p = 0.31
x pos:
F = 1.04
p = 0.39
y pos:
F = 0.65
p = 0.62
x pos:
F = 1.38
p = 0.24
y pos:
F = 0.73
p = 0.57
x pos:
F = 0.54
p = 0.71
y pos:
F = 0.35
p = 0.84
x pos:
F = 0.34
p = 0.85
y pos:
F = 0.16
p = 0.96
attended positions
ignored positions
x position
y position

Supplementary Figure 4.1: Analysis of eye position data. a) Example eye trace, showing the horizontal eye position across one run for subject DH. Shades of blue indicate the position of the attended stimuli on each trial. For this run, the correlation between mean horizontal eye position within each block and the stimulus eccentricity was r = 0.043 (p = 0.80; based on 36 blocks per run; df = 34). The correlation between eye position and stimulus condition was not significant for any run, for either the attended or ignored stimuli. **b)** Mean and standard deviation eye position is shown for each of the five attended positions and the five ignored positions separately for each subject. We ran four one-way ANOVAs within each subject, testing for a systematic difference in horizontal or vertical eye position across the five stimulus locations for either the attended or ignored stimuli. None of these tests showed significant differences in eye position across the five stimulus locations (most significant was in subject DH for the ignored stimulus locations x horizontal eye position: $F_{4,119995} = 1.38$, p = 0.24; all p values reported reflect tests for a main effect of eye position across the five stimulus locations, with 120,000 total eye position measurements). These results rule out the influence of systematic eye movements on our results.

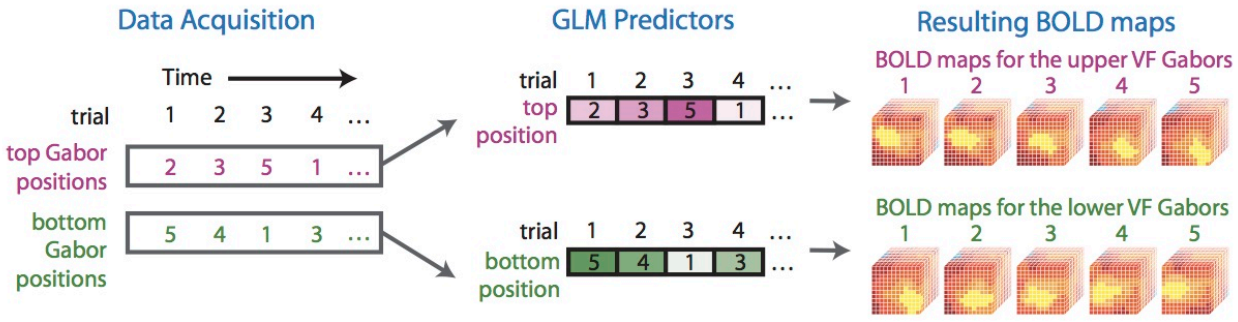

Supplementary Figure 4.2: Parallel analysis of the attended and ignored stimuli in the BOLD response. In order to separately isolate the patterns of activity corresponding to the attended stimuli and the ignored stimuli, we fit each functional run with two separate general linear models. One had predictors coded according to positions of the Gabors in the upper visual field (shown in purple), while the other had predictors coded according to positions of the Gabors in the lower visual field (shown in green). By separately contrasting each of the five stimulus positions against a fixation baseline in each GLM, we generated ten maps of BOLD response: five corresponding to positions of the stimuli in the upper visual field, and five corresponding to the positions of the stimuli in the lower visual field. This approach allowed us to tease apart information about the attended and ignored stimulus positions within the same set of voxels by using two separate and independent encoding models, each of which explained a unique portion of the pattern of activity in the Pulvinar ROI.

Pulvinar:

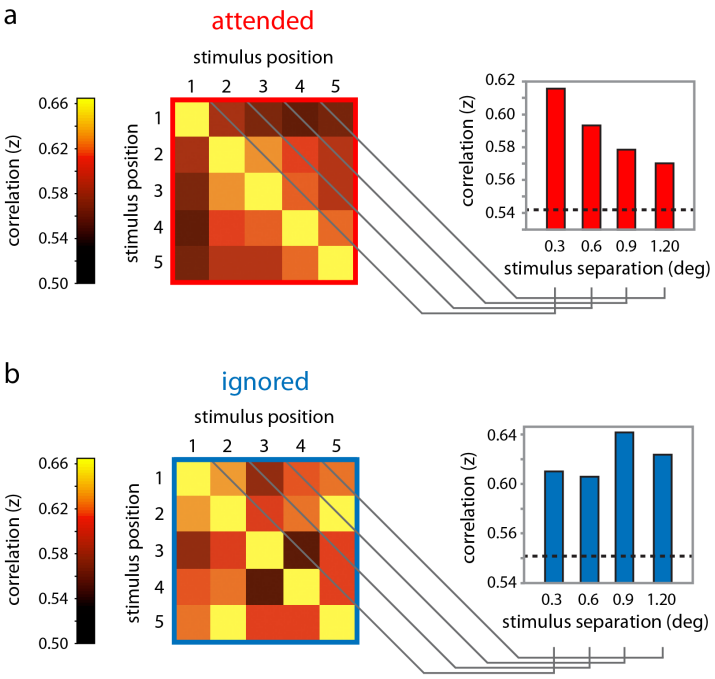
a
attended
stimulus position
1 2 3 4 5
correlation (z)
0.66
0.62
0.58
0.54
0.50
stimulus position
1
2
3
4
5
correlation (z)
0.62
0.60
0.58
0.56
0.54
0.3 0.6 0.9 1.20
stimulus separation (deg)
b
ignored
stimulus position
1 2 3 4 5
correlation (z)
0.66
0.62
0.58
0.54
0.50
stimulus position
1
2
3
4
5
correlation (z)
0.64
0.60
0.56
0.54
0.3 0.6 0.9 1.20
stimulus separation (deg)

V1:

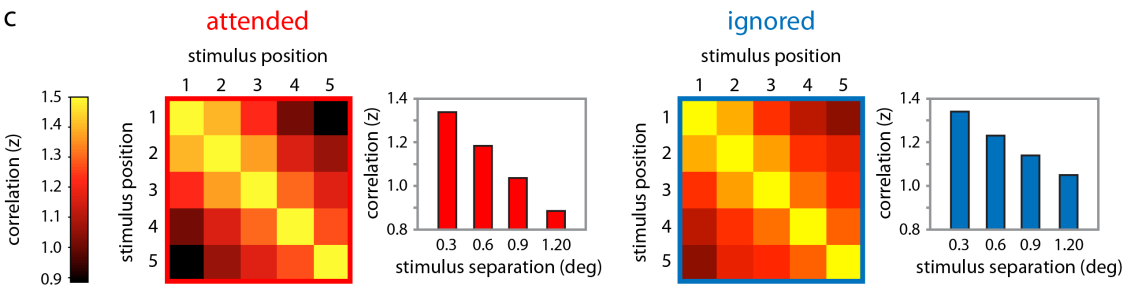
c
attended
stimulus position
1 2 3 4 5
correlation (z)
1.5
1.4
1.3
1.2
1.1
1.0
0.9
stimulus position
1
2
3
4
5
correlation (z)
1.4
1.2
1.0
0.8
0.3 0.6 0.9 1.20
stimulus separation (deg)
ignored
stimulus position
1 2 3 4 5
stimulus position
1
2
3
4
5
correlation (z)
1.4
1.2
1.0
0.8
0.3 0.6 0.9 1.20
stimulus separation (deg)

Supplementary Figure 4.3: Correlation matrices for the pulvinar and V1. a) Correlations among the maps for the five attended stimulus positions within the pulvinar, collapsed across hemispheres and subjects. Red bars show the mean correlation for each of the four stimulus separations, and the dashed line within the bar plot shows the mean of the bootstrapped null distribution (±1 sd of the distribution fall within the thickness of the dashed line; see Supplemental Methods). The null distribution reflects a baseline level of correlation among the five maps that cannot be attributed to stimulus-related activity, and might instead reflect factors such as partial voluming, vascular structure, etc. The correlation at each of the four stimulus separations was significantly greater than baseline (all $p < 0.001$, evaluated as the proportion of the bootstrapped null distribution that was greater than the actual correlation between attended maps). There was a significant negative trend in correlation with increasing stimulus separation (Fisher $z = -0.46$; $p < 0.001$; based on 50 points in a position discrimination plot constructed as in Fig. 4.3, but now including data at stimulus separations of 0.9 and 1.2 deg). Correlations here were measured on an individual run basis and position discrimination was calculated from the full cross correlation matrix; the resulting values are necessarily slightly different from those in Figure 4.3, but show the same pattern of results. **b)** Correlation matrix for the five maps of BOLD response corresponding to the ignored stimulus positions, measured within the pulvinar. The correlation at each of the four stimulus separations was significantly greater than the bootstrapped null distribution (all $p < 0.001$, evaluated as the proportion of the bootstrapped null distribution that was greater than the actual correlation between ignored maps), but the correlations did not show a systematic progression (slope) as a function of stimulus separation (Fisher $z = 0.18$; $p = 0.85$; nonparametric bootstrap test based on 50 points). **c)** Correlation matrices for the attended and ignored activity maps within bilateral V1. For both the attended and ignored stimulus positions, there was a systematic decrease in correlation between BOLD response maps with increasing stimulus separation (Fisher $z = -1.55$ for attended; $p < 0.001$; Fisher $z = -1.23$ for ignored; $p < 0.001$; nonparametric bootstrap tests based on 50 points per test).

Chapter 5

Concluding Discussion

The story that these studies collectively tell is that visuospatial mapping in the brain is not hard-wired to any physical coordinate frame, retinotopic or otherwise. Rather, spatial maps adapt according to our visual experience (Chapter 2), and our current goals (Chapters 3 and 4). Flexible spatial mapping is a means by which the brain adapts to the visual demands of the present moment. Here I discuss some implications of these findings, how they can be interpreted in the context of what we know from single-cell recordings, and the direction that I aim to take this research in the future.

5.1 Implications of flexible spatial mapping in the brain

The results presented here show that different regions of the brain can have position codes that do not match each other and inaccurately reflect the physical locations of objects in the world. While objects are represented at their retinal positions in primary visual cortex, those same objects are often represented at different locations in higher-level areas, corresponding to their perceived locations. Ignored objects are represented with almost the same fidelity as attended objects in V1, but they are not discernably represented in the position maps in the pulvinar. How does the brain reconcile its mismatched position codes? The answer to this question is presently unknown. Among the possibilities is that there exists a “master map” of object position in the brain – a map that simultaneously integrates all influences on perceived position, and serves as the ultimate reference for perceptual localization. This is a plausible scenario but perhaps unlikely, because the regions that we found to be the most strongly percept-centered in Chapter 2 are also the most specialized for particular stimulus classes and are therefore unlikely to subserve position perception in general. Another possibility, then, is that the brain draws on multiple maps when constructing percepts of objects’ locations, using position information from the visual area that is specialized for the type of stimulus being viewed (e.g., using position information from the FFA to evaluate the location of a face). It may seem that the large receptive fields in higher-level visual areas preclude the possibility of these areas supporting fine spatial judgments, but large receptive fields can nonetheless lead to high spatial precision when position information is pooled over a population of cells (Eurich and Schwegler 1997). The latter scenario would imply that different object classes may be localized with different degrees of precision, thereby providing a potential means of testing this hypothesis psychophysically (for example, by performing a vernier discrimination task with upright vs. inverted faces).

5.2 What is the mechanism behind flexible spatial coding in the brain?

What changes in neural responses allow the brain's spatial maps to adapt on a moment-by-moment basis? A prevalent theory is that spatial transformations in the brain are accomplished through the influence of gain fields (Andersen and Mountcastle 1983; Zipser and Andersen 1988). In the gain field account of spatial transformations, the tuning properties of individual neurons, including their receptive field locations, remain constant, but the amplitude of their responses is modulated by other factors such as gaze direction or head orientation. Gain fields transform spatial information at the population level by modulating the relative contribution of each unit to the overall collective output of a brain area. Thus, an influence of gain fields can give rise to shifted receptive fields in downstream regions that pool over the population output of an area that is influence by gain modulation. Though the origin and role of gain fields is still debated, modulation by gain fields has been observed in many regions throughout the primate brain, including primary visual cortex (Blohm and Crawford 2009). The flexible spatial mapping reported in this thesis may be accomplished through the influence of gain fields. For example, the encoding of perceived object position in high-level visual areas may be the product of the combined influence of many overlaid gain fields transmitting information about factors such as frame of reference, size, motion, and attention. While this is a plausible descriptive account of how flexible mapping in the brain may be achieved, it is still in some sense just a restatement of the fact that each of these factors plays a role in determining perceived position. A more difficult and as yet unanswered question is how this diverse and complex set of gain fields is computed to begin with.

5.3 Population-level vs. single-unit neural measurements

Functional MRI infers neural activity from measurements of blood flow at the scale of millimeters. As such, inferences based on fMRI are confined to the level of neural populations – the integrated activity of thousands of cells – rather than the activity of single neurons. For each of the main findings reported here, the same question could be asked at the scale of single-unit response properties, and the answer needn't necessarily be the same. For example, the narrowed position tuning in the population response of V1 reported in Chapter 3 could be a result of shrinking receptive fields in V1, or, equally plausible, shifting receptive fields, localized changes in the number of active units, localized changes in the response amplitude of active units, or any combination thereof. Changes in the neural population response reflect an integration over the changes in the responses of many single units, but this does *not* imply that population-level activity can be inferred or understood by examining responses in a small subset of single cells. For example, spike rate increases in the responses of single neurons at an attended location, once thought to underlie attention-related behavioral benefits, actually only account for a small fraction of the signal improvement at the population level – a much larger share of the improvement is due to increased independence of the signals among the neurons in the population (Cohen and Maunsell 2009). Perception ultimately depends on pooling the information contained within an entire population of cells (Pouget, Dayan et al. 2000;

Purushothaman and Bradley 2005), and with the advent of multivariate pattern analysis techniques for fMRI data (Haxby, Gobbini et al. 2001; Kamitani and Tong 2005; Kay, Naselaris et al. 2008; Tong and Pratte 2012), we now have access to some of the information content in the neural population response through examining the fMRI BOLD signal. This makes fMRI well-suited for evaluating brain responses alongside perception and behavior.

5.4 Bridging the gap between fMRI and single-unit findings

Moving forward, a primary goal of mine is to connect the findings presented here to what is known about the primate brain from neurophysiological recordings. The pulvinar nucleus in particular is an area where interesting findings are emerging at both the population and single-unit levels of description, with no immediately obvious way to connect the two. Based on its widespread and intricately organized connections with the cortex, it has been hypothesized that the pulvinar may serve as an interface for modulating the synchronization of firing between cortical areas (Guillery and Sherman 2002; Shipp 2003; Saalmann and Kastner 2009). This notion was recently confirmed with single-unit recordings from the pulvinar, V4, and the temporo-occipital area (TEO) in macaque monkeys: synchrony between the responses in V4 and TEO increased when attention was directed to the location of their receptive fields, and the pulvinar was found to be a driving source of this modulation in synchrony (Saalmann, Pinsk et al. 2012).

How do our results on the gating of ignored information from the spatial maps in the pulvinar, presented here in Chapter 4 (Fischer and Whitney 2012), connect with these findings on neural synchrony at the single-unit level? In an ongoing project, I am testing for a connection between spatial mapping in the pulvinar and cortico-cortical communication between visual areas. During an attention task, I am measuring the functional connectivity (signal coherence) between V1 and V2 on individual runs and asking whether this connectivity is stronger when position selectivity in the pulvinar is stronger. The data collected so far (Figure 5.1) indicate that signal coherence between V1 and V2 is strongly correlated with spatial coding in the pulvinar within attended locations but is weakly or even negatively correlated with spatial coding in the pulvinar at unattended locations. These results provide a means of connecting our findings from Chapter 4 with the latest results emerging from single-unit recordings in the pulvinar.

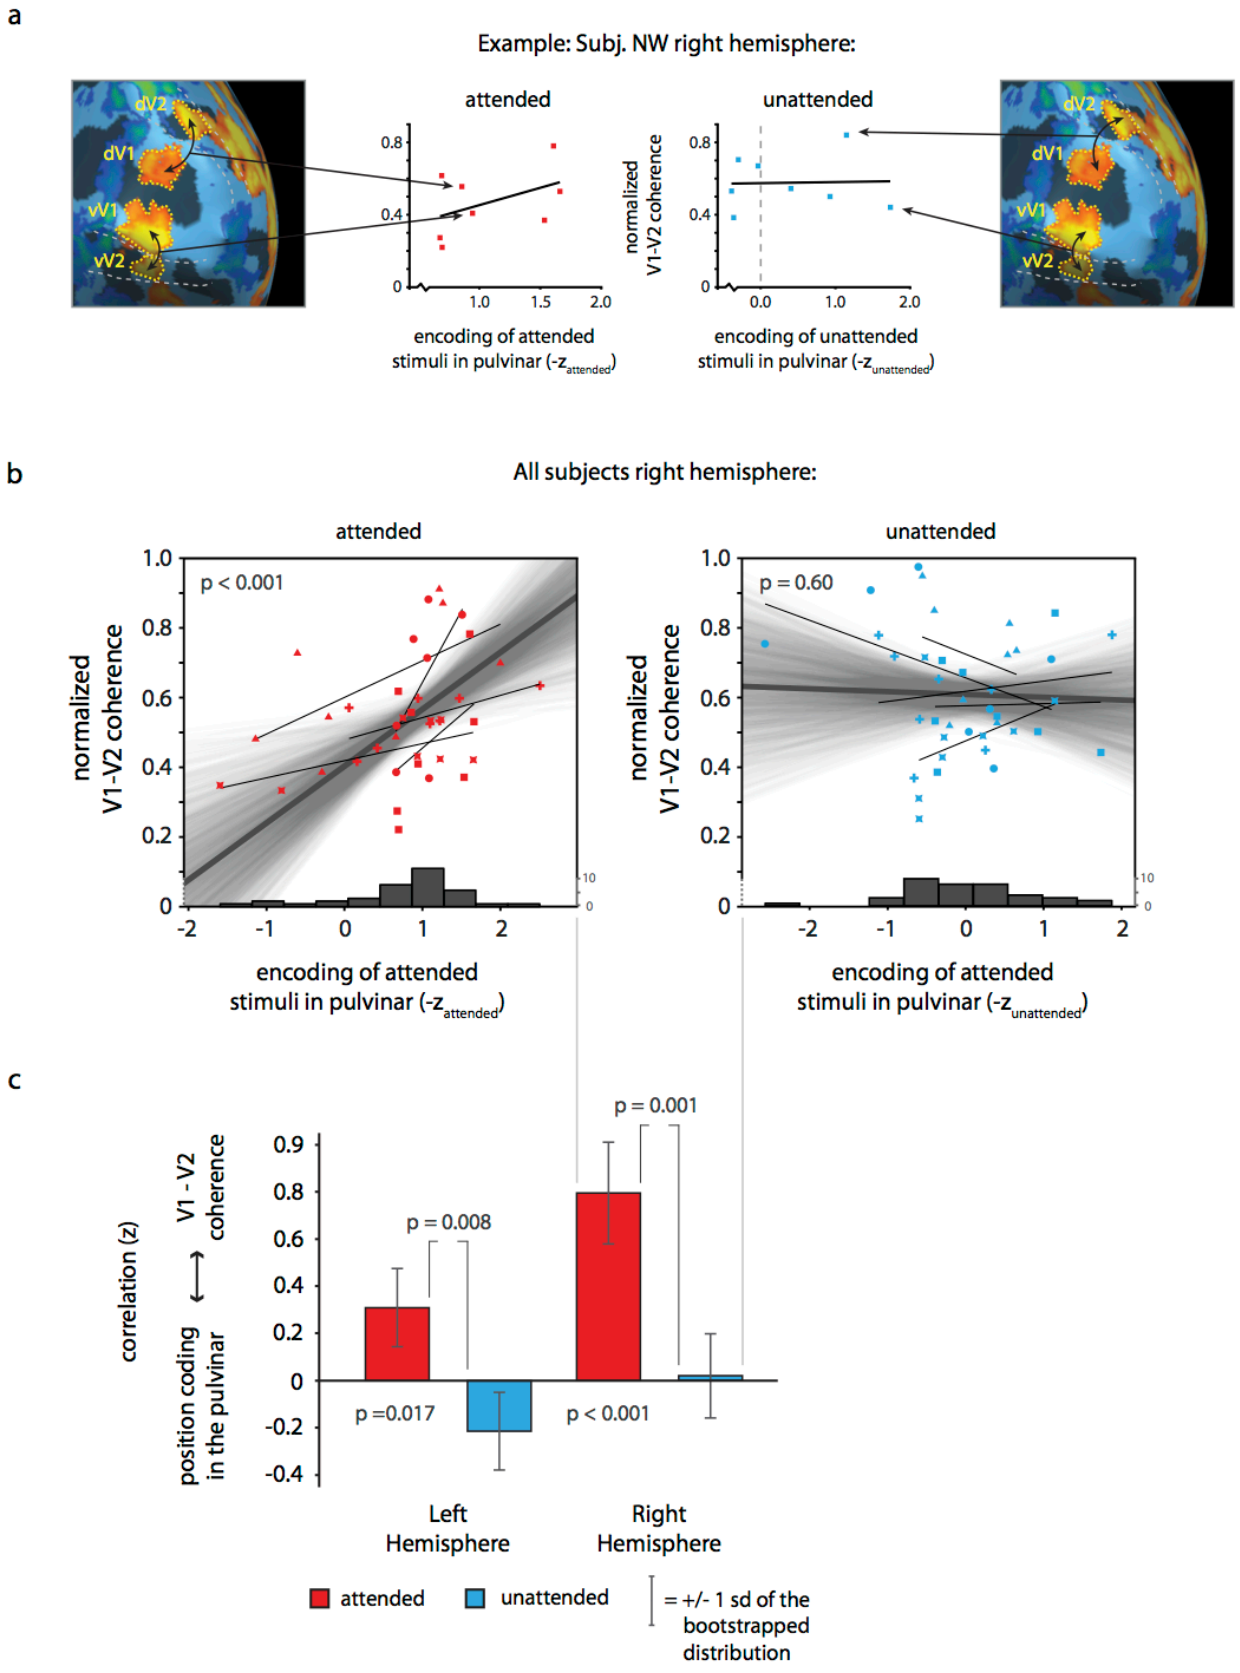
a
Example: Subj. NW right hemisphere:
dV2
dV1
vV1
vV2
attended
unattended
normalized V1-V2 coherence
encoding of attended stimuli in pulvinar (-z attended)
encoding of unattended stimuli in pulvinar (-z unattended)
b
All subjects right hemisphere:
p < 0.001
p = 0.60
c
correlation (z)
V1 - V2 coherence
position coding in the pulvinar
p = 0.001
p = 0.008
p =0.017
p < 0.001
Left Hemisphere
Right Hemisphere
attended
unattended
= +/- 1 sd of the bootstrapped distribution

Figure 5.1: V1-V2 functional connectivity is predicted by the strength of visual encoding in the pulvinar. a) Example cortical regions of interest and functional connectivity data from the right hemisphere of one subject. Each point is one scanning run. Subjects attended to the positions of cued objects while ignoring non-cued objects. Position discrimination measurements reflect the degree to which the positions of the attended and unattended stimuli could be decoded from the pulvinar, and coherence was computed as the magnitude of $Coh_{xy}(\lambda) = \frac{\left|f_{xy}(\lambda)\right|^2}{f_{xx}(\lambda) f_{yy}(\lambda)}$ – the cross-spectrum of the two fMRI timecourses, normalized by the individual spectrum of each. Panel **b** shows data from the same (right) hemisphere for all subjects, with a separate regression line shown for each subject and a thicker group average regression line. The shaded region on each plot shows 2000 bootstrapped estimates of the group average regression line. Also shown along the abscissa is a histogram of the position discrimination values for all subjects. **c**) We used a bootstrapping procedure to test the significance of the groupwise regression of V1-V2 coherence on visual encoding in the pulvinar. The fit was significant and positive for attended locations (left hemisphere: $p = 0.017$, right hemisphere: $p < 0.001$; two-tailed), but not for unattended locations (left hemisphere: $p = 0.14$, right hemisphere: $p = 0.60$). The correlation for attended locations was significantly greater than the correlation for unattended locations in both hemispheres (left hemisphere: $p = 0.008$, right hemisphere: $p = 0.001$).

5.5 Conclusions

The results presented here collectively show that position coding in the brain can be highly flexible, particularly under conditions of selective attention and raise a number of additional questions. If position coding in high-level, object-selective visual areas is closely linked to perception, does the brain possess position invariant object representations? Our results, together with those of Afraz et al. (2010), suggest that it may not. Position coding may be a special case, acting as a medium in which invariance in other dimensions is computed, but not actually achieving invariance itself. More research is needed before such a strong conclusion could be drawn. Do the effects reported here operate differently for foveally vs. peripherally presented stimuli? In everyday life, we quickly foveate objects of interest and perform most of our fine-scale spatial judgments using our high-resolution central vision. Performing our experiments with peripherally-presented stimuli allowed us to maximize the magnitude of the perceptual effects of interest and likewise maximize our chances of finding corresponding effects in the BOLD response. It will be important to perform the same sorts of experiments using stimuli that subjects are allowed to foveate, perhaps even in a free viewing scenario. These are just a few outstanding questions related to perceptual localization and spatial mapping, but they point in the direction that I hope to take this line of research in the future.

References

Afraz, A. and P. Cavanagh (2009). "The gender-specific face aftereffect is based in retinotopic not spatiotopic coordinates across several natural image transformations." Journal of Vision 9(10).

Afraz, A., M. V. Pashkam, et al. (2010). "Spatial Heterogeneity in the Perception of Face and Form Attributes." Current Biology 20(23): 2112-2116.

Altmann, C. F., A. Deubelius, et al. (2004). "Shape saliency modulates contextual processing in the human lateral occipital complex." J Cogn Neurosci 16(5): 794-804.

Alvarez, G. A. and P. Cavanagh (2005). "Independent Resources for Attentional Tracking in the Left and Right Visual Hemifields." Psychological Science 16(8): 637-643.

Andersen, R. A., G. K. Essick, et al. (1985). "Encoding of spatial location by posterior parietal neurons." Science 230(4724): 456-458.

Andersen, R. A. and V. B. Mountcastle (1983). "The influence of the angle of gaze upon the excitability of the light- sensitive neurons of the posterior parietal cortex." The Journal of Neuroscience 3(3): 532-548.

Andersen, R. A., L. H. Snyder, et al. (1997). "Multimodal representation of space in the posterior parietal cortex and its use in planning movements." Annu Rev Neurosci 20: 303-330.

Ashby, F. G., W. Prinzmetal, et al. (1996). "A formal theory of feature binding in object perception." Psychological Review 103(1): 165-192.

Avidan, G., U. Hasson, et al. (2002). "Analysis of the neuronal selectivity underlying low fMRI signals." Curr Biol 12(12): 964-972.

Balz, G. W. and H. S. Hock (1997). "The effect of attentional spread on spatial resolution." Vision Res 37(11): 1499-1510.

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

Benjamini, Y. and Y. Hochberg (1995). "Controlling the False Discovery Rate - a Practical and Powerful Approach to Multiple Testing." Journal of the Royal Statistical Society Series B-Methodological 57(1): 289-300.

Berry, M. J., I. H. Brivanlou, et al. (1999). "Anticipation of moving stimuli by the retina." Nature 398(6725): 334-338.

Blohm, G. and J. D. Crawford (2009). "Fields of gain in the brain." Neuron 64(5): 598-600.

Boynton, G. M. (2004). "Adaptation and attentional selection." Nat Neurosci 7(1): 8-10.

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

Bremmer, F., A. Schlack, et al. (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.

Bressler, D., N. Spotswood, et al. (2007). "Negative BOLD fMRI response in the visual cortex carries precise stimulus-specific information." PLoS ONE 2(5): e410.

Bressler, D. W. and M. A. Silver (2010). "Spatial attention improves reliability of fMRI retinotopic mapping signals in occipital and parietal cortex." NeuroImage 53(2): 526-533.

Bridgeman, B., S. Peery, et al. (1997). "Interaction of cognitive and sensorimotor maps of visual space." Percept Psychophys 59(3): 456-469.

Buchsbaum, M. S., B. R. Buchsbaum, et al. (2006). "Thalamocortical circuits: fMRI assessment of the pulvinar and medial dorsal nucleus in normal volunteers." Neurosci Lett 404(3): 282-287.

Buracas, G. T. and G. M. Boynton (2007). "The Effect of Spatial Attention on Contrast Response Functions in Human Visual Cortex." The Journal of Neuroscience 27(1): 93-97.

Cai, R. H., A. Pouget, et al. (1997). "Perceived geometrical relationships affected by eye-movement signals." Nature 386(6625): 601-604.

Cameron, E. L., J. C. Tai, et al. (2002). "Covert attention affects the psychometric function of contrast sensitivity." Vision Res 42(8): 949-967.

Carlson, T. A., P. Schrater, et al. (2003). "Patterns of activity in the categorical representations of objects." J Cogn Neurosci 15(5): 704-717.

Carrasco, M. (2011). "Visual attention: The past 25 years." Vision Res 51(13): 1484-1525.

Carrasco, M., S. Ling, et al. (2004). "Attention alters appearance." Nat Neurosci 7(3): 308-313.

Carrasco, M. and Y. Yeshurun (2009). Covert attention effects on spatial resolution. Attention. S. Narayanan, Elsevier. Volume 176: 65-86.

Casanova, C. (2004). The Visual Functions of the Pulvinar. The visual neurosciences. J. S. Werner and L. M. Chalupa. Cambridge, Mass., MIT Press: 592-608.

Cavanagh, P., A. R. Hunt, et al. (2010). "Visual stability based on remapping of attention pointers." Trends in Cognitive Sciences 14(4): 147-153.

Cohen, M. R. and J. H. R. Maunsell (2009). "Attention improves performance primarily by reducing interneuronal correlations." Nature Neuroscience 12(12): 1594-1600.

Cotton, P. L. and A. T. Smith (2007). "Contralateral visual hemifield representations in the human pulvinar nucleus." J Neurophysiol.

Crespi, S., L. Biagi, et al. (2011). "Spatiotopic Coding of BOLD Signal in Human Visual Cortex Depends on Spatial Attention." PLoS ONE 6(7): e21661-e21661.

d'Avossa, G., M. Tosetti, et al. (2007). "Spatiotopic selectivity of BOLD responses to visual motion in human area MT." Nat Neurosci 10(2): 249-255.

Dakin, S. C., J. Cass, et al. (2010). "Probabilistic, positional averaging predicts object-level crowding effects with letter-like stimuli." Journal of Vision 10(10).

Danziger, S., R. Ward, et al. (2001). "The effects of unilateral pulvinar damage in humans on reflexive orienting and filtering of irrelevant information." Behav Neurol 13(3-4): 95-104.
Danziger, S., R. Ward, et al. (2004). "Contributions of the human pulvinar to linking vision and action." Cogn Affect Behav Neurosci 4(1): 89-99.
Dassonville, P., B. Bridgeman, et al. (2004). "The induced Roelofs effect: two visual systems or the shift of a single reference frame?" Vision Res 44(6): 603-611.
De Valois, R. L. and K. K. De Valois (1991). "Vernier acuity with stationary moving Gabors." Vision Res 31(9): 1619-1626.
Desimone, R. and J. Duncan (1995). "Neural mechanisms of selective visual attention." Annu Rev Neurosci 18: 193-222.
Desimone, R., M. Wessinger, et al. (1990). "Attentional control of visual perception: cortical and subcortical mechanisms." Cold Spring Harbor symposia on quantitative biology 55: 963-971.
DiCarlo, J. J. and D. D. Cox (2007). "Untangling invariant object recognition." Trends in cognitive sciences 11(8): 333-341.
DiCarlo, J. J. and J. H. R. Maunsell (2003). "Anterior Inferotemporal Neurons of Monkeys Engaged in Object Recognition Can be Highly Sensitive to Object Retinal Position." Journal of Neurophysiology 89(6): 3264-3278.
DuBois, R. M. and M. S. Cohen (2000). "Spatiotopic Organization in Human Superior Colliculus Observed with fMRI." NeuroImage 12(1): 63-70.
Duhamel, J.-R., F. Bremmer, et al. (1997). "Spatial invariance of visual receptive fields in parietal cortex neurons." Nature 389(6653): 845-848.
Dukelow, S. P., J. F. DeSouza, et al. (2001). "Distinguishing subregions of the human MT+ complex using visual fields and pursuit eye movements." J Neurophysiol 86(4): 1991-2000.
Dumoulin, S. O., R. G. Bittar, et al. (2000). "A new anatomical landmark for reliable identification of human area V5/MT: a quantitative analysis of sulcal patterning." Cereb Cortex 10(5): 454-463.
Eldridge, L. L., B. J. Knowlton, et al. (2000). "Remembering episodes: a selective role for the hippocampus during retrieval." Nat Neurosci 3(11): 1149-1152.
Epstein, R., K. S. Graham, et al. (2003). "Viewpoint-specific scene representations in human parahippocampal cortex." Neuron 37(5): 865-876.
Epstein, R., A. Harris, et al. (1999). "The parahippocampal place area: recognition, navigation, or encoding?" Neuron 23(1): 115-125.
Eriksen, C. W. and J. F. Collins (1969). "Temporal course of selective attention." Journal of Experimental Psychology 80(2, Pt.1): 254-261.
Eriksen, C. W. and J. W. Rohrbaugh (1970). "Some Factors Determining Efficiency of Selective Attention." The American Journal of Psychology 83(3): 330-342.
Eurich, C. W. and H. Schwegler (1997). "Coarse coding: calculation of the resolution achieved by a population of large receptive field neurons." Biol Cybern 76(5): 357-363.
Ezzati, A., A. Golzar, et al. (2008). "Topography of the motion aftereffect with and without eye movements." Journal of Vision 8(14).

Fischer, J., N. Spotswood, et al. (2011). "The Emergence of Perceived Position in the Visual System." J Cogn Neurosci 23(1): 119-136.
Fischer, J. and D. Whitney (2009). "Attention narrows position tuning of population responses in V1." Curr Biol 19(16): 1356-1361.
Fischer, J. and D. Whitney (2009). "Precise discrimination of object position in the human pulvinar." Human Brain Mapping 30(1): 101-111.
Fischer, J. and D. Whitney (2012). "Distractor suppression: Attention gates visual coding in the human pulvinar." Nature Communications *in press*.
Fortenbaugh, F. and L. Robertson (2011). "When here becomes there: attentional distribution modulates foveal bias in peripheral localization." Attention, Perception, & Psychophysics 73(3): 809-828.
Fu, Y. X., Y. Shen, et al. (2004). "Asymmetry in visual cortical circuits underlying motion-induced perceptual mislocalization." Journal of Neuroscience 24(9): 2165-2171.
Galletti, C., P. P. Battaglini, et al. (1993). "Parietal neurons encoding spatial locations in craniotopic coordinates." Experimental Brain Research 96(2): 221-229.
Galletti, C., P. P. Battaglini, et al. (1995). "Eye Position Influence on the Parieto-occipital Area PO (V6) of the Macaque Monkey." European Journal of Neuroscience 7(12): 2486-2501.
Gandhi, S. P., D. J. Heeger, et al. (1999). "Spatial attention affects brain activity in human primary visual cortex." Proc Natl Acad Sci U S A 96(6): 3314-3319.
Gardner, J. L., E. P. Merriam, et al. (2008). "Maps of visual space in human occipital cortex are retinotopic, not spatiotopic." Journal of Neuroscience 28(15): 3988-3999.
Gauthier, I., M. J. Tarr, et al. (1999). "Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects." Nat Neurosci 2(6): 568-573.
Golomb, J. and N. Kanwisher (2010). "Location information in category-selective areas: retinotopic or spatiotopic?" Journal of Vision 10(7): 950-950.
Golomb, J. D., M. M. Chun, et al. (2008). "The Native Coordinate System of Spatial Attention Is Retinotopic." The Journal of Neuroscience 28(42): 10654-10662.
Golomb, J. D. and N. Kanwisher (2011). "Higher Level Visual Cortex Represents Retinotopic, Not Spatiotopic, Object Location." Cerebral cortex (New York, N.Y.: 1991).
Golomb, J. D. and N. Kanwisher (2012). "Retinotopic memory is more precise than spatiotopic memory." Proceedings of the National Academy of Sciences.
Golomb, J. D., A. Y. Nguyen-Phuc, et al. (2010). "Attentional Facilitation throughout Human Visual Cortex Lingers in Retinotopic Coordinates after Eye Movements." The Journal of Neuroscience 30(31): 10493-10506.
Golomb, J. D., V. Z. Pulido, et al. (2010). "Robustness of the retinotopic attentional trace after eye movements." Journal of Vision 10(3): 19.11-12-19.11-12.
Goodale, M. A. and A. D. Milner (1992). "Separate visual pathways for perception and action." Trends in Neurosciences 15(1): 20-25.
Graziano, M. S. (2006). "Progress in understanding spatial coordinate systems in the primate brain." Neuron 51(1): 7-9.

Graziano, M. S. and C. G. Gross (1998). "Spatial maps for the control of movement." Curr Opin Neurobiol 8(2): 195-201.

Greenwood, J. A., P. J. Bex, et al. (2009). "Positional averaging explains crowding with letter-like stimuli." Proceedings of the National Academy of Sciences 106(31): 13130-13135.

Grieve, K. L., C. Acuña, et al. (2000). "The primate pulvinar nuclei: vision and action." Trends in Neuroscience 23: 35-39.

Grill-Spector, K., N. Knouf, et al. (2004). "The fusiform face area subserves face perception, not generic within-category identification." Nat Neurosci 7(5): 555-562.

Grill-Spector, K., T. Kushnir, et al. (1999). "Differential processing of objects under various viewing conditions in the human lateral occipital complex." Neuron 24(1): 187-203.

Grill-Spector, K., T. Kushnir, et al. (1998). "A sequence of object-processing stages revealed by fMRI in the human occipital lobe." Hum Brain Mapp 6(4): 316-328.

Grill-Spector, K. and R. Malach (2001). "fMR-adaptation: a tool for studying the functional properties of human cortical neurons." Acta Psychol (Amst) 107(1-3): 293-321.

Guillery, R. W. and S. M. Sherman (2002). "Thalamic Relay Functions and Their Role in Corticocortical Communication - Generalizations from the Visual System." Neuron 33(2): 163-175.

Haenny, P. E. and P. H. Schiller (1988). "State dependent activity in monkey visual cortex. I. Single cell activity in V1 and V4 on visual tasks." Exp Brain Res 69(2): 225-244.

Hagler, D. J., Jr., L. Riecke, et al. (2007). "Parietal and superior frontal visuospatial maps activated by pointing and saccades." NeuroImage 35(4): 1562-1577.

Hagler Jr, D. J. and M. I. Sereno (2006). "Spatial maps in frontal and prefrontal cortex." NeuroImage 29(2): 567-577.

Hardy, S. G. and J. C. Lynch (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.

Haushofer, J., M. S. Livingstone, et al. (2008). "Multivariate patterns in object-selective cortex dissociate perceptual and physical shape similarity." PLoS Biol 6(7): e187.

Haxby, J. V., M. I. Gobbini, et al. (2001). "Distributed and overlapping representations of faces and objects in ventral temporal cortex." Science 293(5539): 2425-2430.

Hayhoe, M., J. Lachter, et al. (1991). "Integration of form across saccadic eye movements." Perception 20(3): 393¬†,Äì¬†402-393¬†,Äì¬†402.

Haynes, J. D. and G. Rees (2005). "Predicting the orientation of invisible stimuli from activity in human primary visual cortex." Nat Neurosci 8(5): 686-691.

Haynes, J. D. and G. Rees (2006). "Decoding mental states from brain activity in humans." Nature Reviews Neuroscience 7(7): 523-534.

Hemond, C. C., N. G. Kanwisher, et al. (2007). "A preference for contralateral stimuli in human object- and face-selective cortex." PLoS ONE 2(6): e574.

Henson, R. N., M. Hornberger, et al. (2005). "Further dissociating the processes involved in recognition memory: an FMRI study." J Cogn Neurosci 17(7): 1058-1073.
Holmes, A. P. and K. J. Friston (1998). "Generalisability, random effects and population inference." Neuroimage 7(S754).
Huk, A. C., R. F. Dougherty, et al. (2002). "Retinotopy and functional subdivision of human areas MT and MST." J Neurosci 22(16): 7195-7205.
Jehee, J. F. M., D. K. Brady, et al. (2011). "Attention Improves Encoding of Task-Relevant Features in the Human Visual Cortex." The Journal of Neuroscience 31(22): 8210-8219.
Jonides, J., D. E. Irwin, et al. (1982). "Integrating visual information from successive fixations." Science 215(4529): 192-194.
Kaas, J. H. and D. C. Lyon (2007). "Pulvinar contributions to the dorsal and ventral streams of visual processing in primates." Brain Res Rev.
Kamitani, Y. and F. Tong (2005). "Decoding the visual and subjective contents of the human brain." Nat Neurosci 8(5): 679-685.
Kanwisher, N., J. McDermott, et al. (1997). "The fusiform face area: a module in human extrastriate cortex specialized for face perception." J Neurosci 17(11): 4302-4311.
Kastner, S., D. H. O'Connor, et al. (2004). "Functional imaging of the human lateral geniculate nucleus and pulvinar." J Neurophysiol 91(1): 438-448.
Kastner, S., M. A. Pinsk, et al. (1999). "Increased activity in human visual cortex during directed attention in the absence of visual stimulation." Neuron 22(4): 751-761.
Kay, K. N., T. Naselaris, et al. (2008). "Identifying natural images from human brain activity." Nature 452(7185): 352-355.
Kerzel, D. (2003). "Attention maintains mental extrapolation of target position: irrelevant distractors eliminate forward displacement after implied motion." Cognition 88(1): 109-131.
Kourtzi, Z. and N. Kanwisher (2000). "Activation in human MT/MST by static images with implied motion." J Cogn Neurosci 12(1): 48-55.
Kravitz, D. J., N. Kriegeskorte, et al. (2010). "High-Level Visual Object Representations Are Constrained by Position." Cerebral Cortex 20(12): 2916-2925.
Kriegeskorte, N., M. Mur, et al. (2008). "Representational similarity analysis - connecting the branches of systems neuroscience." Front Syst Neurosci 2: 4.
LaBerge, D. and M. S. Buchsbaum (1990). "Positron emission tomographic measurements of pulvinar activity during an attention task." J Neurosci 10(2): 613-619.
Laberge, D., M. Carter, et al. (1992). "A Network Simulation of Thalamic Circuit Operations in Selective Attention." Neural Computation 4(3): 318-331.
Larsson, J. and D. J. Heeger (2006). "Two retinotopic visual areas in human lateral occipital cortex." J Neurosci 26(51): 13128-13142.

Lauritzen, T. Z., M. D'Esposito, et al. (2009). "Top-down flow of visual spatial attention signals from parietal to occipital cortex." Journal of Vision 9(13): 18.11-14-18.11-14.
Levi, D. M. (2008). "Crowding - an essential bottleneck for object recognition: a mini-review." Vision Research 48(5): 635-654.
Levy, I., U. Hasson, et al. (2001). "Center-periphery organization of human object areas." Nat Neurosci 4(5): 533-539.
Lin, C.-J. and C.-C. Chang (2011). "LIBSVM : a library for support vector machines." ACM Transactions on Intelligent Systems and Technology 2:27:1--27:27.
Liverence, B. M. and B. J. Scholl (2011). "Selective Attention Warps Spatial Representation Parallel but Opposing Effects on Attended Versus Inhibited Objects." Psychological Science 22(12): 1600-1608.
Luck, S. J., L. Chelazzi, et al. (1997). "Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex." J Neurophysiol 77(1): 24-42.
Luck, S. J., S. A. Hillyard, et al. (1989). "Independent hemispheric attentional systems mediate visual search in split-brain patients." Nature 342: 543-545.
MacEvoy, S. P. and R. A. Epstein (2007). "Position selectivity in scene- and object-responsive occipitotemporal regions." J Neurophysiol.
MacKay, D. M. (1961). Interactive processes in visual perception. Sensory Communication. W. A. Rosenblith. Cambridge, MA, MIT Press: 339–355.
MacMillan, N. A. and C. D. Creelman (2004). Detection Theory: A User's Guide, Lawrence Erlbaum Associates.
Mangun, G. R. (1995). "Neural mechanisms of visual selective attention." Psychophysiology 32(1): 4-18.
Martinez-Trujillo, J. C. and S. Treue (2004). "Feature-based attention increases the selectivity of population responses in primate visual cortex." Curr Biol 14(9): 744-751.
Matin, L., E. Picoult, et al. (1982). "Oculoparalytic illusion: visual-field dependent spatial mislocalizations by humans partially paralyzed with curare." Science 216(4542): 198-201.
Maunsell, J. H. and C. J. McAdams (2000). Effects of attention on neuronal response properties in visual cerebral cortex. The New Cognitive Neurosciences. M. S. Gazzaniga. Cambridge, Mass., MIT Press: 315-324.
Maunsell, J. H. R. (2004). The Role of Attention in Visual Cerebral Cortex. The visual neurosciences. J. S. Werner and L. M. Chalupa. Cambridge, Mass., MIT Press: 1538-1545.
Maus, G. W., J. Fischer, et al. (2009). Dragging it out: Motion-sensitive areas code perceived positions in the motion-induced flash mislocalization illusion. Society for Neuroscience Annual Meeting.
McAdams, C. J. and J. H. Maunsell (1999). "Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4." J Neurosci 19(1): 431-441.
McGraw, P. V., N. W. Roach, et al. (2012). "Size-induced distortions in perceptual maps of visual space." Journal of Vision 12(4).

McKee, S. P., G. H. Silverman, et al. (1986). "Precise velocity discrimination despite random variations in temporal frequency and contrast." Vision Res 26(4): 609-619.
McKyton, A. and E. Zohary (2007). "Beyond retinotopic mapping: the spatial representation of objects in the human lateral occipital complex." Cereb Cortex 17(5): 1164-1172.
Melcher, D. (2005). "Spatiotopic Transfer of Visual-Form Adaptation across Saccadic Eye Movements." Current Biology 15(19): 1745-1748.
Melcher, D. and M. C. Morrone (2003). "Spatiotopic temporal integration of visual motion across saccadic eye movements." Nature neuroscience 6(8): 877-881.
Mishkin, M., L. G. Ungerleider, et al. (1983). "Object vision and spatial vision: two cortical pathways." Trends in Neurosciences 6(0): 414-417.
Moran, J. and R. Desimone (1985). "Selective attention gates visual processing in the extrastriate cortex." Science 229(4715): 782-784.
Motter, B. C. (1993). "Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli." J Neurophysiol 70(3): 909-919.
Mourão-Miranda, J., A. L. W. Bokde, et al. (2005). "Classifying brain states and determining the discriminating activation patterns: Support Vector Machine on functional MRI data." Neuroimage 28(4): 980–995.
Muller, N. G. and A. Kleinschmidt (2004). "The attentional 'spotlight's' penumbra: center-surround modulation in striate cortex." Neuroreport 15(6): 977-980.
Murray, S. O., H. Boyaci, et al. (2006). "The representation of perceived angular size in human primary visual cortex." Nat Neurosci 9(3): 429-434.
Murray, S. O. and E. Wojciulik (2004). "Attention increases neural selectivity in the human lateral occipital complex." Nat Neurosci 7(1): 70-74.
Nakayama, K. and M. Mackeben (1989). "Sustained and transient components of focal visual attention." Vision Research 29(11): 1631-1647.
Nijhawan, R. (1994). "Motion extrapolation in catching." Nature 370(6487): 256-257.
Nijhawan, R. (2002). "Neural delays, visual motion and the flash-lag effect." Trends Cogn Sci 6(9): 387.
Norman, K. A., S. M. Polyn, et al. (2006). "Beyond mind-reading: multi-voxel pattern analysis of fMRI data." Trends Cogn Sci 10(9): 424-430.
O'Regan, J. K. and A. Levy-Schoen (1983). "Integrating visual information from successive fixations: does trans-saccadic fusion exist?" Vision research 23(8): 765-768.
Olshausen, B. A., C. H. Anderson, et al. (1993). "A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information." The Journal of Neuroscience 13(11): 4700-4719.
Op De Beeck, H. and R. Vogels (2000). "Spatial sensitivity of macaque inferior temporal neurons." The Journal of Comparative Neurology 426(4): 505-518.
Padmala, S., S. Lim, et al. (2010). "Pulvinar and affective significance: responses track moment-to-moment stimulus visibility." Frontiers in Human Neuroscience 4(64): 1-9.

Pashkam, M. V., W. Shim, et al. (2007). Human motion area MT is not spatiotopic. Society for Neuroscience Annual Meeting. San Diego, CA.
Popple, A. V. and D. M. Levi (2005). "The perception of spatial order at a glance." Vision Research 45(9): 1085-1090.
Pouget, A., P. Dayan, et al. (2000). "Information processing with population codes." Nat Rev Neurosci 1(2): 125-132.
Purushothaman, G. and D. C. Bradley (2005). "Neural population code for fine perceptual decisions in area MT." Nat Neurosci 8(1): 99-106.
Rajimehr, R. and R. B. H. Tootell (2009). "Does Retinotopy Influence Cortical Folding in Primate Visual Cortex?" The Journal of Neuroscience 29(36): 11149-11152.
Ramachandran, V. S. and S. M. Anstis (1990). "Illusory displacement of equiluminous kinetic edges." Perception 19(5): 611-616.
Reynolds, J. H. and R. Desimone (1999). "The role of neural mechanisms of attention in solving the binding problem." Neuron 24(1): 19-29, 111-125.
Reynolds, J. H., T. Pasternak, et al. (2000). "Attention increases sensitivity of V4 neurons." Neuron 26(3): 703-714.
Roberts, M., L. S. Delicato, et al. (2007). "Attention alters spatial integration in macaque V1 in an eccentricity-dependent manner." Nat Neurosci 10(11): 1483-1491.
Robinson, D. L. and S. E. Petersen (1992). "The pulvinar and visual salience." Trends Neurosci 15(4): 127-132.
Roelofs, C. (1935). "Optische Localisation." Arch Augenheilkunde 109: 395-415.
Ross, J., M. C. Morrone, et al. (1997). "Compression of visual space before saccades." Nature 386(6625): 598-601.
Ross, J., M. C. Morrone, et al. (2001). "Changes in visual perception at the time of saccades." Trends Neurosci 24(2): 113-121.
Saalmann, Y. B. and S. Kastner (2009). "Gain control in the visual thalamus during perception and cognition." Current Opinion in Neurobiology 19(4): 408-414.
Saalmann, Y. B., M. A. Pinsk, et al. (2012). "The Pulvinar Regulates Information Transmission Between Cortical Areas Based on Attention Demands." Science 337(6095): 753-756.
Salinas, E. and L. F. Abbott (1997). "Invariant visual responses from attentional gain fields." J Neurophysiol 77(6): 3267-3272.
Sceniak, M. P., D. L. Ringach, et al. (1999). "Contrast's effect on spatial summation by macaque V1 neurons." Nat Neurosci 2(8): 733-739.
Schlag, J. and M. Schlag-Rey (2002). "Through the eye, slowly: delays and localization errors in the visual system." Nat Rev Neurosci 3(3): 191-215.
Schneider, K. A. (2011). "Subcortical Mechanisms of Feature-Based Attention." The Journal of Neuroscience 31(23): 8643-8653.
Schneider, K. A., M. C. Richter, et al. (2004). "Retinotopic Organization and Functional Subdivisions of the Human Lateral Geniculate Nucleus: A High-Resolution Functional Magnetic Resonance Imaging Study." The Journal of Neuroscience 24(41): 8975-8985.
Schwartz, G., S. Taylor, et al. (2007). "Synchronized Firing among Retinal Ganglion Cells Signals Motion Reversal." Neuron 55(6): 958.

Schwarzlose, R. F., J. D. Swisher, et al. (2008). "The distribution of category and location information across object-selective regions in human visual cortex." Proc Natl Acad Sci U S A 105(11): 4447-4452.
Serences, J. T., S. Saproo, et al. (2009). "Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions." Neuroimage 44(1): 223-231.
Sereno, M. I., A. M. Dale, et al. (1995). "Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging." Science 268(5212): 889-893.
Sereno, M. I. and R.-S. Huang (2006). "A human parietal face area contains aligned head-centered visual and tactile maps." Nature Neuroscience 9(10): 1337-1343.
Serre, T., L. Wolf, et al. (2007). "Robust object recognition with cortex-like mechanisms." IEEE Transactions on Pattern Analysis and Machine Intelligence 29: 411-426.
Shipp, S. (2003). "The Functional Logic of Cortico-Pulvinar Connections." Philosophical Transactions: Biological Sciences 358(1438): 1605-1624.
Shipp, S. (2004). "The brain circuitry of attention." Trends Cogn Sci 8(5): 223-230.
Silver, M. A. and S. Kastner (2009). "Topographic maps in human frontal and parietal cortex." Trends in Cognitive Sciences 13(11): 488-495.
Silver, M. A., D. Ress, et al. (2005). "Topographic maps of visual spatial attention in human parietal cortex." Journal of Neurophysiology 94(2): 1358-1371.
Smith, A. T., P. L. Cotton, et al. (2009). "Dissociating vision and visual attention in the human pulvinar." J Neurophysiol 101(2): 917-925.
Smith, A. T., K. D. Singh, et al. (2000). "Attentional suppression of activity in the human visual cortex." Neuroreport 11(2): 271-277.
Snedecor, G. W. and W. G. Cochran (1980). Statistical Methods (7th ed). Ames, Iowa State University Press.
Snow, J. C., H. A. Allen, et al. (2009). "Impaired attentional selection following lesions to human pulvinar: evidence for homology between human and monkey." Proc Natl Acad Sci U S A 106(10): 4054-4059.
Snowden, R. J. (1998). "Shifts in perceived position following adaptation to visual motion." Curr Biol 8(24): 1343-1345.
Somers, D. C., A. M. Dale, et al. (1999). "Functional MRI reveals spatially specific attentional modulation in human primary visual cortex." Proc Natl Acad Sci U S A 96(4): 1663-1668.
Spitzer, H., R. Desimone, et al. (1988). "Increased attention enhances both behavioral and neuronal performance." Science 240(4850): 338-340.
Sundberg, K. A., M. Fallah, et al. (2006). "A motion-dependent distortion of retinotopy in area V4." Neuron 49(3): 447-457.
Suzuki, S. and P. Cavanagh (1997). "Focused attention distorts visual space: an attentional repulsion effect." J Exp Psychol Hum Percept Perform 23(2): 443-463.
Swisher, J. D., M. A. Halko, et al. (2007). "Visual topography of human intraparietal sulcus." The Journal of neuroscience: the official journal of the Society for Neuroscience 27(20): 5326-5337.

Talairach, J. and P. Tournoux (1988). Co-planar stereotaxic atlas of the human brain : 3-dimensional proportional system: an approach to cerebral imaging. Stuttgart ; New York, New York, G. Thieme ; Thieme Medical Publishers.

Tong, F. (2003). "Primary visual cortex and visual awareness." Nat Rev Neurosci 4(3): 219-229.

Tong, F. and M. S. Pratte (2012). "Decoding Patterns of Human Brain Activity." Annual Review of Psychology 63(1): 483-509.

Tootell, R. B., N. Hadjikhani, et al. (1998). "The retinotopy of visual spatial attention." Neuron 21(6): 1409-1422.

Tootell, R. B., J. B. Reppas, et al. (1995). "Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging." J Neurosci 15(4): 3215-3230.

Tootell, R. B., D. Tsao, et al. (2003). "Neuroimaging weighs in: humans meet macaques in "primate" visual cortex." J Neurosci 23(10): 3981-3989.

Treisman, A. and G. Gelade (1980). "A feature-integration theory of attention." Cognitive Psychology 12(1): 97-136.

Treisman, A. M. (1969). "Strategies and models of selective attention." Psychol Rev 76(3): 282-299.

Treue, S. (2001). "Neural correlates of attention in primate visual cortex." Trends Neurosci 24(5): 295-300.

Treue, S. and J. C. Martinez Trujillo (1999). "Feature-based attention influences motion processing gain in macaque visual cortex." Nature 399(6736): 575-579.

Treue, S. and J. H. Maunsell (1996). "Attentional modulation of visual motion processing in cortical areas MT and MST." Nature 382(6591): 539-541.

Ullman, S. and S. Soloviev (1999). "Computation of pattern invariance in brain-like structures." Neural networks: the official journal of the International Neural Network Society 12(7-8): 1021-1036.

Ungerleider, L. G. and M. Mishkin (1982). Two cortical visual systems. Analysis of Visual Behavior. D. J. Ingle, M. A. Goodale and R. J. W. Mansfield. Cambridge, MA, The MIT Press: 549-586.

Volkmann, J. (1932). "The Method of Single Stimuli." The American Journal of Psychology, 44(4): 808-809.

Wallis, G. and E. T. Rolls (1997). "Invariant face and object recognition in the visual system." Progress in neurobiology 51(2): 167-194.

Wandell, B. A., S. O. Dumoulin, et al. (2007). "Visual field maps in human cortex." Neuron 56(2): 366-383.

Ward, R., S. Danziger, et al. (2002). "Deficits in spatial coding and feature binding following damage to spatiotopic maps in the human pulvinar." Nature Neuroscience 5(2): 99-100.

Watson, J. D., R. Myers, et al. (1993). "Area V5 of the human brain: evidence from a combined study using positron emission tomography and magnetic resonance imaging." Cereb Cortex 3(2): 79-94.

Whitaker, D. and P. V. McGraw (1998). "The effect of suprathreshold contrast on stimulus centroid and its implications for the perceived location of objects." Vision Res 38(22): 3591-3599.
Whitaker, D., P. V. McGraw, et al. (1997). "The influence of adaptation on perceived visual location." Vision Res 37(16): 2207-2216.
Whitaker, D., P. V. McGraw, et al. (1995). "Centroid analysis predicts visual localization of first- and second-order stimuli." Vision Res 36(18): 2957-2970.
Whitaker, D., P. V. McGraw, et al. (1996). "Centroid analysis predicts visual localization of first- and second-order stimuli." Vision Res 36(18): 2957-2970.
Whitney, D. (2002). "The influence of visual motion on perceived position." Trends Cogn Sci 6(5): 211-216.
Whitney, D. and D. W. Bressler (2007). "Spatially asymmetric response to moving patterns in the visual cortex: re-examining the local sign hypothesis." Vision Res 47(1): 50-59.
Whitney, D. and P. Cavanagh (2000). "Motion distorts visual space: shifting the perceived position of remote stationary objects." Nat Neurosci 3(9): 954-959.
Whitney, D., H. C. Goltz, et al. (2003). "Flexible Retinotopy: Motion-Dependent Position Coding in the Visual Cortex." Science 302(5646): 878-881.
Whitney, D., I. Murakami, et al. (2010). The utility of visual motion for goal-directed reaching. Space and time in the perception of action. R. Nijhawan and B. Khurana. Cambridge, UK, Cambridge University Press: 121-145.
Wilke, M., K. M. Mueller, et al. (2009). "Neural activity in the visual thalamus reflects perceptual suppression." Proc. Natl. Acad. Sci. 106: 9465-9470.
Williford, T. and J. H. Maunsell (2006). "Effects of spatial attention on contrast response functions in macaque area V4." J Neurophysiol 96(1): 40-54.
Wolford, G. (1975). "Perturbation model for letter identification." Psychological Review 82(3): 184-199.
Yeshurun, Y. and M. Carrasco (1998). "Attention improves or impairs visual performance by enhancing spatial resolution." Nature 396(6706): 72-75.
Yeshurun, Y. and M. Carrasco (1999). "Spatial attention improves performance in spatial resolution tasks." Vision Res 39(2): 293-306.
Yeshurun, Y., B. Montagna, et al. (2008). "On the flexibility of sustained attention and its effects on a texture segmentation task." Vision Research 48(1): 80-95.
Yi, D. J., T. A. Kelley, et al. (2006). "Attentional modulation of repetition attenuation is anatomically dissociable for scenes and faces." Brain Res 1080(1): 53-62.
Yuille, A. and D. Kersten (2006). "Vision as Bayesian inference: analysis by synthesis?" Trends in Cognitive Sciences 10(7): 301-308.
Zimmermann, E., D. Burr, et al. (2011). "Spatiotopic visual maps revealed by saccadic adaptation in humans." Current Biology 21(16): 1380-1384.

Zipser, D. and R. A. Andersen (1988). "A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons." Nature 331: 679-684.