

Comparison of Frontal and Parietal Cortices in the Control of Visual Attention

by

Timothy J. Buschman

B.S. Biology
California Institute of Technology, 2001

SUBMITTED TO THE DEPARTMENT OF BRAIN AND COGNITIVE
SCIENCE IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF

DOCTOR OF PHILOSOPHY IN NEUROSCIENCE
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

FEBRUARY 2008

© 2008 Massachusetts Institute of Technology. All rights reserved.

Signature of Author: _____

Brain and Cognitive Science
September 30, 2007

Certified by: _____

Earl K. Miller, Ph.D.
Picower Professor of Neuroscience
Thesis Supervisor

Accepted by: _____

Matthew A. Wilson, Ph.D.
Professor of Neurobiology
Chairman, Committee for Graduate Students

Comparison of Frontal and Parietal Cortices in the Control of Visual Attention

by

Timothy J. Buschman

Submitted to the Department of Brain and Cognitive Science
on January 24, 2008 in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy in Neuroscience

Abstract

The ability to switch between tasks reflects a fundamental part of our intelligence. A foundation of this ability lies in perceiving and processing information pertinent to the situation at hand. It is our capacity to attend to specific objects and, more importantly, our ability to switch our attention from object to object, that supports complex cognitive behavior. Therefore, by understanding the neural mechanisms involved in directing attention we hope to better understand cognition. Previous work investigating the ability to control attention has suggested that attention is influenced from two sources – attention can either be driven from external sources in an bottom-up, exogenous manner or directed internally in an top-down, endogenous manner. This project will utilize two different forms of visual search in order to emphasize these two different types of attentional control. Both the prefrontal and parietal regions are implicated as the source of this control. In order to investigate their relative roles we recorded simultaneously from both parietal cortex (specifically, the lateral intraparietal cortex) and prefrontal cortex (specifically, the frontal eye fields and dorsolateral prefrontal cortex).

We address four main questions. First, we contrast the respective roles of frontal and parietal cortex in the direction of attention when it is under either top-down or bottom-up control. We use the timing of attention signals between frontal and parietal cortex to establish that frontal cortex directs top-down attention back into parietal cortex, while bottom-up attention is reflected first in parietal cortex, flowing forward to frontal cortex. Secondly, we investigated the role of synchrony and the inter-areal relationships underlying top-down and bottom-up control of attention. Our results suggest synchrony between areas shifts as the task shifts, likely aiding in the selection of the network best suited to the current task. Third, we compare the neural mechanisms between internal and external control of attention. Finally, we investigate the neural correlates of the putative parallel and serial mechanisms underlying visual search, finding support for the existence of a serial search and for the role of the frontal eye fields in the direction of spatial attention.

Thesis Supervisor: Earl K. Miller

Title: Picower Professor of Neuroscience

Acknowledgements

I would like to thank my mentor, Earl Miller, for his guidance throughout my graduate career. Earl has taught me to be creative, how to write, how to think, how to design experiments, and how to interpret results. Earl has created an environment of unwavering support and independent thinking while fostering a strong sense of lab camaraderie. I appreciate his devotion to science and to our lab – he truly enjoys being a scientist; an enthusiasm that is contagious. Earl is a great mentor and role model, but just as importantly, he is a good friend.

I would also like to thank my committee members for their help and support. I am grateful to Bob Desimone for giving me the opportunity to jump into systems neuroscience in my post-baccalaureate year and he is responsible for kick-starting my interest in visual attention. Matt Wilson has given me a depth of understanding of neuroscience that I am still probing – it wasn't until the third time I heard his 9.011 lectures that I think I even began to understand them. John Maunsell has fostered my interest in visual attention and has selflessly given his time on such short notice.

I am grateful to my lab mates. There is not a single person in Earl's lab who I have not learned from. Everyone has provided me with illuminating discussions – acting as a sounding-board, supporting me when I was still learning, and (of course) feeling free to correct me when I am wrong. I would like to especially thank Jefferson Roy with his help in learning to manage life as a primate researcher. Without his wisdom and guidance I doubt I would have lasted this long. Kristin MacCully has provided support (and laughter) since the day I joined the lab. Wael Asaad has taught me much about neuroscience and is ultimately responsible for pushing me to do large-scale simultaneous recordings. I would also like to thank Michelle Machon, Dave Freedman, and Jon Wallis; all of whom have taught me more than they probably realize. Finally, I would like to thank the animal care techs as well as the veterinary staff for their support of our research – it would be impossible without them yet they receive so little credit.

I would also like to acknowledge the support of my family. My parents imbued in me a love of science; they taught me to be curious and always encouraged my questions; and, most importantly, they have loved me and supported me throughout my entire life. I am also thankful for my wonderful siblings with whom I spent my childhood laughing, playing, fighting, and learning. They were my first best-friends. I would like also like to thank my grandfather for his love and support; he gave me my interest in computers and it has often proven to be my most valuable skill.

Lastly, I am thankful for the love and support of my family. My son Ryan is a constant reminder of the joy and wonder in life. Finally, and most importantly, I am eternally grateful for the love and support of my wife, Sarah. She has taught me more than anyone, has been a better friend that I have ever deserved, and has loved me more than I thought possible. Without her none of this would have been possible and so it is only fitting that this thesis be devoted to her, as is all of my life.

Table of Contents

Abstract.....	2
Acknowledgements.....	3
Table of Contents.....	4
Chapter 1: General Introduction.....	10
<i>The Balance of Cognition</i>	10
<i>The Advantage of Attention</i>	11
Psychophysical Evidence for Endogenous and Exogenous Control of Attention.....	12
Loss of Attentional Control	14
The Effect of Attention on Neural Representations	15
<i>Biased Competition Model</i>	16
Synchronous Activity as a Biasing Signal.....	17
<i>Brain Regions Involved in the Control of Attention</i>	19
Neural Responses in Lateral Intraparietal (LIP) Region.....	20
Neural Responses in Dorsolateral Prefrontal Cortex (dlPFC).....	21
Neural Responses in Frontal Eye Fields (FEF).....	23
<i>Pre-Motor Theory of Attention</i>	25
<i>Simultaneous Recording</i>	26
<i>Goals of the Current Study</i>	26
Chapter 2: Role of Parietal and Frontal Cortex in the Control of Attention.....	28
<i>Introduction</i>	28
<i>Visual Pop-out and Visual Search</i>	28
<i>Flow of Information about the Locus of Attention</i>	29
Figure 1.....	30
Figure 2.....	32

Figure 3.....	33
Figure 4.....	35
Figure 5.....	36
Table 1 & 2	38
Figure 6.....	39
Figure 7.....	41
Figure 8.....	42
<i>Aligning Data on Array Onset versus Saccade Initiation</i>	<i>43</i>
Figure 9.....	44
<i>Dynamics of Selectivity for Single Units</i>	<i>47</i>
<i>Correlation of Neural Activity with Behavior</i>	<i>54</i>
Figure 10.....	51
Figure 11.....	52
Figure 12.....	53
<i>Conclusion</i>	<i>55</i>
Chapter 3: The Role of Synchrony in the Control of Attention	57
<i>Introduction.....</i>	<i>57</i>
<i>Synchrony between Parietal and Frontal Cortex</i>	<i>57</i>
Figure 13.....	58
Figure 14.....	60
<i>Correcting for a Common Ground.....</i>	<i>61</i>
Figure 15.....	62
<i>Correcting for Reaction Time Differences between Tasks</i>	<i>63</i>
Figure 16.....	64
Figure 17.....	65

Figure 18.....	66
Figure 19.....	67
<i>Inter- and Intra-Areal Coherence</i>	<i>68</i>
Figure 20.....	69
Figure 21.....	70
<i>Synchrony Varies with Location Preference.....</i>	<i>74</i>
Figure 22.....	75
Figure 23.....	77
Figure 24.....	78
Figure 25.....	79
<i>Correlation of Synchrony with Reaction Time.....</i>	<i>80</i>
Figure 26.....	81
Figure 27.....	83
<i>Phase Relationships between Local Field Potentials.....</i>	<i>84</i>
Figure 28.....	86
Figure 29.....	87
Figure 30.....	89
Figure 31.....	90
Figure 32.....	91
Figure 33.....	92
<i>Synchrony between Frequency Bands.....</i>	<i>94</i>
Figure 34.....	96
<i>Synchrony between Spikes and Local Field Potentials</i>	<i>97</i>
Figure 35.....	98
<i>Conclusion</i>	<i>99</i>

Chapter 4: Neural Coding for Task Parameters	100
<i>Introduction</i>	100
<i>Information about Task Parameters over Time</i>	100
Figure 36.....	101
Figure 37.....	103
Figure 38.....	104
Figure 39.....	105
<i>Conclusion</i>	106
Chapter 5: Common & Disparate Mechanisms Underlying Control of Attention	108
<i>Introduction</i>	108
<i>Attention Signals Across Tasks</i>	108
Figure 40.....	111
Figure 41.....	112
Figure 42.....	113
Figure 43.....	114
Figure 44.....	115
Figure 45.....	116
<i>Information about the Stimulus Parameters</i>	117
Figure 46.....	118
Figure 47.....	120
Figure 48.....	121
<i>Intersection of Attention & Stimulus Information</i>	122
<i>Conclusion</i>	122
Figure 49.....	123
Chapter 6: Parallel and Serial Mechanisms for the Direction of Attention	126

<i>Introduction</i>	126
<i>Behavioral Evidence for Parallel and Serial Search Mechanisms</i>	127
Figure 50.....	128
Table 3.....	130
<i>Neural Evidence for Serial and Parallel Mechanisms in Search</i>	131
Figure 51.....	134
Figure 52.....	136
<i>Strength of Serial Search and Reaction Time</i>	137
Figure 53.....	138
<i>Neural Evidence for Clockwise Search</i>	139
Figure 54.....	141
Figure 55.....	142
<i>Serial Search on a Single Neuron Level</i>	143
Figure 56.....	145
<i>Serial Search on a Single Trial Level</i>	146
Figure 57.....	147
<i>FEF leads dIPFC during Search</i>	148
Figure 58.....	149
<i>Psychophysical Estimate of the Time to Shift Attention</i>	150
<i>Controlling for Eye Movement Differences</i>	150
Figure 59.....	151
<i>Conclusions</i>	152
Chapter 7: General Conclusions.....	154
Figure 60.....	157
Chapter 8: Experimental Design and Methods.....	161

<i>Subjects</i>	161
<i>Behavioral Task</i>	161
Figure 61.....	163
<i>Recording Locations and Isolation of Neural Activity</i>	164
<i>Analysis Techniques</i>	166
<i>Mutual Information Statistic</i>	166
Figure 62.....	168
<i>Comparing the Mean vs. Leading Edge of a Distribution</i>	169
<i>Coherence Statistic</i>	170
<i>Estimating Coherence of Single Trials</i>	171
<i>Circular Statistics</i>	173
References	178

The Balance of Cognition

Cognition is the balance between external stimulation and internal motivation. It is important to quickly perceive important (salient) changes in our environment so that we can respond to them when necessary. However, it is also necessary to avoid becoming locked to external stimuli – only being able to represent and respond to the brightest, loudest, flashiest thing would result in simple, boring behaviors. Much of intelligent behavior relies on internal motivations such as goals and rules. If our behavior is only dictated by the most salient object in the environment it is difficult to maintain the consistent representations necessary for long-term, directed actions that are fundamental to intelligent behavior.

Understanding how the brain balances these competing interests will provide a basis for gaining insight into cognition. One manner in which to approach this more general question is through the study of attention. Attention is the ability to select specific neural representations, either externally or internally generated, and focus upon them. Importantly, for the current study, there is psychophysical, imaging, and electrophysiological evidence that attention can be driven both internally (endogenously) or externally (exogenously), making it an ideal platform from which to investigate our ability to balance these competing influences.

An additional advantage of studying attention is that there is an expansive and rich literature, with most of it devoted to visual attention, providing a strong foundation on which to build the current study. Many of the major results will be highlighted below. Furthermore, attention is involved in many cognitive behaviors, including possibly playing a role in the binding of features within an object (Treisman and Gelade, 1980; Friedman-Hill et al., 1995; Treisman, 1996, 1998; Treisman and Kanwisher, 1998), in working memory (Smith and Jonides, 1997; Cabeza and Nyberg, 2000; Mayer et al., 2007; van Swinderen, 2007), and even as a possible gateway into the study of consciousness (Crick and Koch, 1998, 2003). Therefore, in addition to furthering our understanding of cognitive control of behavior, deeper insight into attention will continue to expand our understanding of many other parts of cognition.

The Advantage of Attention

Paralleling the cognitive advantages of attention, visual attention provides a similar advantage to our visual perception. In order to perceive an object with greater resolution we naturally focus on it, drawing our attention to it, usually immediately followed by a saccade. One reason that attention plays a central role in visual perception is the incredibly rich nature of our visual environments. Although a natural visual scene has some structure, there are still a large variety of scenes, each with a large amount of data which needs to be processed. The retina itself does a fairly good job of capturing this information, as it is estimated to send 10 million bits of information into the thalamus every second (Koch et al., 2004; Koch et al., 2006). The computational power to process and analyze such complex scenes is tremendous; this is one of the many problems facing computer vision. The visual system appears to approach the problem through a hierarchical (although heavily interconnected) network of visual areas (Felleman and Van Essen, 1991). Areas lower in the hierarchy respond to simple visual parameters such as orientation (Hubel and Wiesel, 1968), while areas farther along in the hierarchy usually respond to more complex stimuli. For example, visual categorization seems to exist only at the highest levels of the visual system, after several stages of visual processing (Freedman et al., 2001; Freedman and Assad, 2006). Additionally, neurons in higher level regions respond to stimuli presented over a much larger visual area – neurons in V1 have receptive fields 0.2° in diameter while neurons in IT and FST have receptive fields spanning 10° or more (Hubel and Wiesel, 1968; Gross et al., 1969; Gattass et al., 1981; Gattass et al., 1988). Although this is costly in terms of resolution, it allows for the computational load to be reduced by paring down the visual scene to something more manageable. Visual attention compensates for this reduced resolution by allowing specific portions of the visual scene to be processed to a greater degree. This spatial enhancement is often referred to as the ‘spotlight’ of attention (Posner et al., 1980). This can compensate for the loss of acuity of higher visual areas by collapsing the receptive field of a neuron to the attended region (Reynolds et al., 1999; Womelsdorf et al., 2006a). These perceptual advantages of directing visual attention can be measured psychophysically. For example, focusing visual attention to a particular stimulus can increase acuity for simple, low-level parameters such as orientation, color, and spatial frequency as well as complex visual features such as form and categorization (Corbetta et al., 1990; Solomon et al., 1997). Directing attention to a location also speeds up the reaction time to a stimulus appearing at that location, even across modalities (Posner et al., 1980; Lange and Roder, 2006).

Psychophysical Evidence for Endogenous and Exogenous Control of Attention

Visual attention is believed to be under both endogenous (internal) and exogenous (external) control. The majority of evidence for this disassociation comes from visual search experiments. Visual search experiments ask the subject to find a particular stimulus amongst a display that is cluttered with distractors. Experimental evidence has shown that the speed at which subjects can detect and find the target stimulus is determined by the number and type of distractors. For example, under most conditions, if the distractors are visually consistent across the entire visual array - differing from the target in a single, simple dimension (such as color, orientation, or spatial frequency) - reaction times are fast and do not depend on the number of items in the display (Treisman and Gelade, 1980). So, there are no costs, in either reaction time or performance, to adding another distractor to the visual array (Treisman and Gelade, 1980). Under these conditions the target is said to “pop-out” from the distractors. Objects that suddenly appear, are brighter, or faster moving can also automatically draw attention to themselves (Jonides and Yantis, 1988; Treisman and Gormican, 1988).

In contrast, if the distractors are not visually uniform, reaction time to find the target increases linearly with the number of distractors in the display (Treisman and Gelade, 1980; Treisman and Gormican, 1988). In the classic example, if the target is a conjunction of two features (e.g., colored red and orientated vertically) and the distractors consist of stimuli that share one or the other dimension (e.g., a red horizontal bar, or a blue vertical bar), then adding another distractor will increase the reaction time linearly. Under these conditions, the target no longer pops out and therefore the subject must exert effort in searching for it.

The neural mechanisms underlying these two types of search have been the subject of intense debate. Two major theories have been put forth. The work of Ann Treisman and colleagues has suggested that these two forms of visual search are due to two distinct mechanisms with the brain. They propose that the visual system is able to process large portions of the visual field in parallel up to a certain point, after which attention must be used to select a particular portion for further analysis (Treisman and Gelade, 1980; Treisman and Gormican, 1988). Under this model the initial, more basic

analysis, which is carried out in parallel, forms the basis for pop-out search. More difficult search requires top-down selection of specific stimuli, necessitating a serial approach as the spotlight of attention is shifted from object to object. This model predicts the empirically observed reaction time differences and suggests that there exists a specific level of complexity after which subjects can no longer perform searches in a parallel, pop-out manner.

An alternative model championed by John Duncan, Jeremy Wolfe, and colleagues, suggests that true serial search is a rarity, limited to very specific circumstances not common in everyday scenes. Instead, they suggest a gradation of parallel search – as in Treisman’s model, easy search is done quickly in early visual areas, however, more complicated search does not require serial mechanisms but rather still occurs in parallel with internal, top-down influences guiding the parallel search towards items related to the target (Duncan and Humphreys, 1989; Wolfe et al., 1989; Cave and Wolfe, 1990; Duncan et al., 1997; Wolfe and Horowitz, 2004). Evidence for Duncan’s model comes from experimental results showing the reaction time cost of an added distractor does not fall into two limited categories (parallel or serial) but falls along a continuum from fast, efficient search to slower, costly search (McLeod et al., 1988; Duncan and Humphreys, 1989; Wolfe and Horowitz, 2004).

Although these two models differ in many important ways, there are several commonalities that provide the basis for the current set of experiments. First, both models predict a parallel component that allows for the selection of the target stimulus under pop-out conditions. Computational modeling has suggested that this parallel search can be done via a saliency map (Itti and Koch, 2000, 2001). The saliency map charts the degree of salience of each object within the entire visual field. Salience can be thought of as the degree of ‘pop-out-ness’ of an object in the visual field and can be computed for each object from its features such as its relative brightness or movement (Itti and Koch, 2000, 2001). These models have even been extended to spiking networks (Niebur and Koch, 1994; de Breecht and Saiki, 2006) and several brain regions have been proposed as the anatomical instantiation of the saliency map including visual cortex (Li, 2002), parietal cortex (Gottlieb et al., 1998; Colby and Goldberg, 1999; Kusunoki et al., 2000; Bisley and Goldberg, 2003) and FEF (Thompson et al., 1997; Schall and Bichot, 1998). These areas are discussed in greater detail below. Since salience is purely defined by the

features of the objects in the visual field, it is a bottom-up process. In other words, salience is defined externally, with attention automatically drawn to the most salient stimulus.

A second commonality between the two models is the requirement for internal direction of attentional selection under difficult search conditions. The selection process uses goal directed information about the identity of the target to either guide the parallel search under Duncan's model or to direct the attentional spotlight and compare a putative object with the target under Treisman's model. This selection relies on the memory of the target during search. Therefore attention must arise internally, from within the brain, not the external world.

This dichotomy of internal and external control of visual attention forms the basis for the current study – both models suggest that visual search results from the interaction between a bottom-up, externally, driven process and a top-down, internally, driven process. Leveraging this for the current study allows us to use a comparison of pop-out search with conjunctive, difficult, search in order to contrast how the brain implements bottom-up and top-down control of behavior.

Loss of Attentional Control

Deficits in attentional control lead to diverse behavioral issues. The most common disease directly effecting visual attention is often the result of a stroke affecting the parietal cortex. These patients typically exhibit hemi-spatial neglect (Pouget and Driver, 2000; Hillis, 2006; Nachev and Husain, 2006). Initially, they are unable to perceive any objects in the affected hemisphere, but over time their visual perception slowly recovers to near-normal levels. However, even after recovery, while the patient is able to perceive objects presented to affected hemisphere when presented alone, if a competing stimulus is put into the unaffected hemisphere the patients are no longer able to perceive the original object in the affected hemisphere. This has been suggested to reflect the patients' inability to select objects within the affected hemisphere.

A more common disease involving the loss of attentional control is Attention Deficiency-Hyperactivity Disorder (ADHD) or Attention Deficit Disorder (ADD). An estimated 8 to 10% of school-age children are diagnosed as having ADHD/ADD in the United States, making it one of the most prevalent diseases among children (Committee on Quality Improvement, 2000). The deficits underlying ADHD are not well understood, but psychophysical studies of children with ADHD suggests they have trouble disassociating attention from one location to in order to shift to a new location (McDonald et al., 1999). Current theories about the cause of ADHD have focused on differential activation in prefrontal cortex and basal ganglia, possibly leading to deficits in cognitive control (Rubia et al., 1999; Zang et al., 2005).

The Effect of Attention on Neural Representations

Many of the psychophysical advantages conveyed by directing attention to a particular stimulus have parallel benefits on neural selectivity when attention is directed towards a neuron's receptive field. Typically neurons in the visual cortex respond in a varied manner across stimulus space – they have a tuning curve for some stimulus dimension. When attention is directed into a neuron's receptive field, the gain of its tuning curve is increased. In other words, the neuron becomes better able to discriminate between stimuli across its preferred dimension. This is true for selectivity in V2 and V4 (Motter, 1993; Luck et al., 1997), IT (Chelazzi et al., 1993; Miller et al., 1993; Chelazzi et al., 1998), and MT and MST (Treue and Maunsell, 1996; Treue and Trujillo, 1999; Martinez-Trujillo and Treue, 2002). Attentional effects in striate cortex are under debate, with some evidence that they exist (Motter, 1993; Sharma et al., 2003) and some evidence they do not exist (Moran and Desimone, 1985). The exact mechanism of attentional modulation on neural firing rates is not clear. For orientation and direction selectivity there is evidence for a multiplicative effect – the tuning curve of a single neuron is amplified when attention is directed towards its receptive field (McAdams and Maunsell, 1999; Treue and Trujillo, 1999). In contrast, the gain in contrast sensitivity seen when attention is directed towards a stimulus appears to be the result of an additive effect, such that the tuning curve is boosted by a constant (Reynolds et al., 2000). However, some have argued for a multiplicative effect in contrast gain as well (Martinez-Trujillo and Treue, 2002; Carrasco, 2006; Williford and Maunsell, 2006). Attentional effects can also be described as 'shrinking' of the receptive field of a neuron such that the neuron only responds to the selected stimulus (Reynolds et al., 1999; Womelsdorf et al., 2006a). Finally, the effect of attention has also been suggested to be equivalent to the addition of contrast to the attended stimulus (Reynolds et al., 2000; Martinez-Trujillo and Treue, 2002). However, this analogy does not appear to fully explain the

effect of attention on neural activity as contrast is known to be directly correlated with earlier response latency, while direction of attention to a stimulus does not appear to decrease the response latency of a neuron (Lee et al., 2007).

Biased Competition Model

The biased competition model is currently the most accepted theory on how attentional control is implemented in the visual stream. The model, proposed by Desimone and Duncan, suggests that selection occurs at every level of the visual system through lateral competition (Desimone and Duncan, 1995). These competitive interactions are able to select the object most strongly represented, whether due to the physical nature of the stimulus or due to a top-down biasing signal selecting a specific object.

For example, when a single stimulus is presented it drives selective neurons strongly, which pass information on to the next layer. This process continues throughout the visual stream, resulting in representation of that stimulus throughout the entire brain. Similarly, when two or more stimuli are presented alongside one another, they are all represented early on. However, as multiple sets of neurons respond, each selective for one of the visual stimuli, lateral inhibition between these groups of neurons causes competition for representation. Similar to the psychophysical models, this competition can be biased from one of two different sources: either externally driven or internally motivated. A more salient stimulus in the visual field leads to a stronger neural representation (Allman et al., 1985; Desimone et al., 1985; Reynolds and Desimone, 2003), allowing that stimulus to win the competition for representation in downstream cortical regions. Alternatively, if attention is directed to a stimulus then top-down signals can bias competition towards the selected stimulus by giving the representation a boost. The model allows for the top-down influence to be either through an attentional spotlight, through boosting a particular stimulus parameter (such as color), or even a combination of both methods.

Electrophysiological evidence for the biased competition model comes from experiments done by Desimone's group (Reynolds et al., 1999). In this study, single neurons in V2 and V4 were recorded while either one or two stimuli were presented in the neuron's receptive field. When either stimulus

was presented alone, the neuron responded preferentially to one or the other stimulus. When both were presented simultaneously and attention was directed away from the receptive field, the neural response was in between the response to each stimulus alone – as if the two stimuli were competing with one another, resulting in an average response. However, when attention was directed to either stimulus within the receptive field, the response was as if only the attended stimulus was presented. In this case, the competition has been biased through top-down selection towards one of the stimuli, allowing for that stimulus to win representation. Further support for the model has come in the form of computational studies (Szabo et al., 2004; Deco and Rolls, 2005) and anatomical evidence suggesting a canonical cortical circuit that may support similar computation (Moldakarimov et al., 2005).

Although the biased competition model captures much of the electrophysiological and behavioral effects of attention, neither the source nor the manner of the biasing signal is currently fully understood. Although this question is the focus of this thesis, previous work has provided evidence for the role of synchronous activity as a mechanism for biasing signals, and for the fronto-parietal network's involvement in the direction of attention. These results are reviewed below.

Synchronous Activity as a Biasing Signal

Synchronous firing has been proposed to be the manner through which the representation of a selected object can be increased (Aertsen et al., 1989; Engel et al., 2001; Salinas and Sejnowski, 2001; Fries, 2005; Womelsdorf and Fries, 2007). As the postsynaptic potential associated with an incoming action potential has limited duration, coincident inputs will have increased efficacy compared to non-coincident inputs. Therefore, if a population of neurons representing a single stimulus synchronizes their firing activity, that stimulus will be more strongly represented in downstream areas. This has been proposed to be the manner through which specific representations are biased (Aertsen et al., 1989; Engel et al., 2001; Salinas and Sejnowski, 2001; Fries, 2005; Womelsdorf and Fries, 2007).

Electrophysiological evidence for a role of synchrony in attention comes from Desimone's lab (Fries et al., 2001; Womelsdorf et al., 2006b). Monkeys were directed to either attend to an object inside or outside the receptive field of a V4 neuron. They were required to monitor the stimulus for a

slight color change, to which they responded as quickly as possible. Both the spiking activity of multi-units (which is spiking activity not sorted into specific neurons) and local field potentials (oscillatory neural activity at low frequencies) were recorded. The authors found that when attention was directed into the receptive field of a V4 neuron the combined spiking activity of multiple neurons became increasingly synchronized with the underlying local field potential in an upper frequency band (30 – 70 Hz), while becoming less synchronized to low frequency band (5 – 12 Hz). Further analysis suggested that when the synchrony between spikes and LFPs was highest in the upper frequency band, the animals' reaction time was faster, suggesting that synchrony affects the ability of the animal to perceive and respond to changes. A more recent experiment from Bichot, Rossi and Desimone (Bichot et al., 2005) using free-viewing visual search found similar effects in V4 during free search – when searching for objects synchrony between local field potentials increased significantly in the upper frequency band for objects sharing a feature of the target. Taken together, these results suggest that the gain in response with attention could be implemented as an increase in the synchronous firing of neurons. Theoretical work has suggested that this gain modulation is optimal within the upper frequency band (Salinas and Sejnowski, 2000, 2001; Tiesinga et al., 2002; Tiesinga and Sejnowski, 2004; Mishra et al., 2006).

Synchrony may also play a more generalized role in regulating the flow of information (Bressler, 1996; Engel et al., 2001; Salinas and Sejnowski, 2001). Similar to its role in attention, synchronous activity may help to strengthen different pathways, acting as a way to dynamically shift the effective connectivity between regions as needed for a particular task. Support for this model is sparse, but experiment evidence from cats suggests that lower frequencies tend to be more strongly synchronized when cats are purposefully controlling their behavior, while higher frequencies are more strongly synchronized when the animal is passively viewing the same stimulus (von Stein et al., 2000). Further experimental evidence has suggested that different frequency bands can lock to one another (Schanze and Eckhorn, 1997; von Stein et al., 2000; Palva et al., 2005). Finally, synchrony has also been shown to increase between neurons with similar selectivity, suggesting a mechanism for selecting specific sub-networks of neurons (Womelsdorf et al., 2007). Overall, these results not only point to a role for synchrony in providing the biasing signal underlying attention, but also as a more general mechanism for dynamically modulating the effective connectivity between areas.

Brain Regions Involved in the Control of Attention

While electrophysiological recordings are very well suited for understanding how single neurons from a particular region is involved in specific behaviors, until recently, there has been difficulty in recording from multiple neurons simultaneously. However, with the introduction of functional MRI (fMRI), it is now possible to image activity across the entire brain simultaneously, albeit at a much lower resolution.

Several groups have used fMRI to determine what brain regions are involved in the control of attention. Almost all studies have isolated regions in the parietal cortex (specifically within the intraparietal sulcus) as well as regions in prefrontal cortex, including the human analogue of the frontal eye fields in the precentral sulcus (Corbetta and Shulman, 2002). When human subjects were asked to attend to a location in order to wait for the presentation of a stimulus (Corbetta et al., 1993; Coull and Nobre, 1998), or were instructed to attend to either motion or color, regions in the precentral sulcus (human FEF) and intraparietal sulcus were activated (Liu et al., 2003). This pattern of results were true for several other studies attempting to isolate the control of attention, leading to the suggestion of the existence of a fronto-parietal network involved in the control of attention (Corbetta et al., 1993; Corbetta et al., 1998; Coull et al., 1998). Furthermore, these results have been used to suggest a common framework within the brain for different modes of attentional selection (i.e. attending to location, color, motion, etc) (Liu et al., 2003; Yantis and Serences, 2003).

The fMRI literature provides us with a good idea of which brain regions are involved in controlling attention. However, due to the long timecourse of fMRI, it is impossible to establish the relative timing of selectivity during the allocation of attention. Without the time-course, it is difficult to determine the flow of information, and therefore impossible to ascertain which brain regions are in control of behavior under different circumstances. In order to determine the flow of information, one must record neural activity simultaneously from all of the brain regions potentially involved in the control of attention. In primates, the analogous regions seem to be the lateral intraparietal area (LIP), the lateral prefrontal cortex (specifically dorsolateral prefrontal cortex, dlPFC), and the frontal eye fields (FEF). Electrophysiological findings from these regions are reviewed below.

Neural Responses in Lateral Intraparietal (LIP) Region

The lateral intraparietal (LIP) region lies along the lateral bank of the intraparietal sulcus and has been found to be involved in a large variety of visual and visuo-motor tasks. The parietal cortex lies between visual cortex and frontal cortex and is believed to be one of the first areas to begin the transformation of sensory input from purely reflective in nature into more complex, goal, and movement directed representations. LIP lies along the dorsal stream of the visual system in the monkey and receives feed-forward inputs from V3, V3A, V4, MT, MST, and TEO as well as being reciprocally connected with other parietal regions and frontal regions FEF and dIPFC (Schwartz and Goldman-Rakic, 1984; Andersen et al., 1990; Blatt et al., 1990; Stanton et al., 1995; Lewis and Van Essen, 2000; Ferraina et al., 2002). When LIP is stimulated the monkey will make a saccade to the response field of the neuron, suggesting a possible role in saccadic eye movements (Schiller and Tehovnik, 2001). Projections to the superior colliculus are believed to underlie the microstimulation effects found in LIP (Schiller and Tehovnik, 2001; Ferraina et al., 2002). In addition, LIP neurons have also been found to be selective for smooth pursuit eye movements (Bremmer et al., 1997).

Electrophysiological evidence from LIP has shown the area to play several roles. Early work demonstrated that the area shows delay activity during a delayed spatial memory task. During this task a monkey is required to fixate, during which a brief stimulus is flashed in their periphery. The animal must maintain fixation during a delay until instructed to release, at which point it makes a single saccade to the remembered flashed location. LIP neurons show high activity throughout the entire delay when the stimulus is flashed in the neuron's receptive field (Barash et al., 1991). This delay activity is unique to LIP amongst other regions in the parietal cortex, making it an effective way to isolate LIP.

Recent experiments have found LIP to also show activation with the direction of attention. LIP neurons respond transiently to a flashed stimulus in the periphery, an occurrence that is known to temporarily draw attention to it (Bisley and Goldberg, 2006). LIP neurons will also reflect the current location of attention, regardless of whether the animal makes a saccade to that location or not (Bisley and Goldberg, 2006). In visual search tasks, LIP neurons have largely been studied under pop-out conditions, and have been shown to reflect the target location fairly early (at approximately 90 ms)

(Ipata et al., 2006; Thomas and Pare, 2007). Finally, LIP neurons have been shown to reflect attentional priority of the represented location (Bisley and Goldberg, 2003).

LIP also shows flexibility in its representations, morphing to represent task-related information. For example, LIP neurons have been found to carry color information when it is a task-relevant parameter, but not when it is irrelevant to the current task (Toth and Assad, 2002). Furthermore, LIP neurons have been shown to encode shape information (Sereno and Maunsell, 1998) and correlate strongly with the perceived direction of motion of a random-dot stimulus (Williams et al., 2003). They have even shown category related information for random-dot stimuli (Freedman and Assad, 2006). LIP activity has also been related to more cognitive variables such as time (Leon and Shadlen, 2003), decision making (Platt and Glimcher, 1999; Shadlen and Newsome, 2001), reasoning (Yang and Shadlen, 2007) and reward (Bendiksbj and Platt, 2006) when the animal's behavior is reported with a saccade. These results suggest that LIP is a very flexible region and, taken along with its suggested role in the allocation of attention, have led to it being proposed to be the anatomical instantiation of the saliency map (Gottlieb et al., 1998; Colby and Goldberg, 1999; Kusunoki et al., 2000; Bisley and Goldberg, 2003). This role for LIP fits well with previous results, and, we believe, fits well with the results of this thesis.

Neural Responses in Dorsolateral Prefrontal Cortex (dlPFC)

The prefrontal cortex sits at the anterior pole of the brain and reaches its greatest elaboration and relative size in the primate, especially human, brain (Fuster, 1995), suggesting it plays a role in our ability for advanced cognition and goal-directed behaviors. Recent imaging work has also shown the size of prefrontal cortex is directly correlated with intelligence in adult humans (Haier et al., 2004). The PFC seems anatomically well situated to play a role in cognitive control. It is in the position to be a microcosm of cortical processing, able to synthesize a wide range of external and internal information and also exert control over much of the cortex. The PFC receives and sends projections to most of the cerebral cortex (with the exception of primary sensory and motor cortices) as well as all of the major subcortical systems such as the hippocampus, amygdala, cerebellum, and the basal ganglia (Porrino et al., 1981; Amaral and Price, 1984; Amaral, 1986; Selemon and Goldman-Rakic, 1988; Barbas and De Olmos, 1990; Eblen and Graybiel, 1995; Croxson et al., 2005). Different PFC subdivisions have distinct patterns of interconnections with other brain systems (e.g., lateral PFC – sensory and motor cortex,

orbital PFC – limbic), but there are prodigious connections both within and between PFC subdivisions (Pandya and Barnes, 1987; Barbas and Pandya, 1989; Pandya and Yeterian, 1990; Barbas et al., 1991; Petrides and Pandya, 1999). The anatomical architecture suggests an infrastructure ideal for learning, one that can act as a large associative network for detecting and storing patterns between diverse events, experiences, internal states, etc. This anatomical structure makes it ideal to act as a center of control.

Indeed, neurophysiological studies in animals and imaging studies in humans have shown that the PFC is highly multimodal. Neurons from within PFC are responsive to a wide range of information and have other properties useful for cognitive control (Miller, 2000). Furthermore, PFC neurons sustain their activity to maintain information across short, multi-second, memory delays (Pribram et al., 1952; Fuster and Alexander, 1971; Fuster, 1973; Funahashi et al., 1989; Miller et al., 1996). This property is crucial for goal-directed behavior, which unlike “ballistic” reflexes, typically extend over time. After training on a wide range of operant tasks many PFC neurons (from $\frac{1}{3}$ to $\frac{1}{2}$ of the population) reflect the learned task contingencies: the logic or rules of the task (White and Wise, 1999; Asaad et al., 2000; Wallis et al., 2000; Mansouri et al., 2006). For example, a neuron might be selectively activated by a given cue (e.g., a green light) whereas another neuron might be activated when that cue means something different (like “stop”). Some neurons might activate in anticipation of a forthcoming expected reward or a relevant cue (Watanabe, 1996; G et al., 1999; Wallis and Miller, 2003; Padoa-Schioppa and Assad, 2006), and neurons have even been found to reflect whether a monkey is currently following the abstract principles “same” or “different” (White and Wise, 1999; Wallis et al., 2000). In short, the PFC does indeed act like a brain area that absorbs and reflects the rules needed to guide goal-directed, volitional behavior.

There is strong experimental evidence that the prefrontal cortex is involved in the direction of attention. Human patients with lesions to the PFC not only have difficulty controlling their behavior but show deficits in visual search tasks (Eglin et al., 1991b; Knight et al., 1995; Knight, 1997). Inactivating dlPFC with muscimol in primates reduces their ability to do visual search tasks, but not detection tasks (Iba and Sawaguchi, 2003). This deficit was present for both pop-out and conjunction search task conditions, and the reaction time difference observed between the two tasks was lost after inactivation.

This suggests the dlPFC plays a role in both forms of search (an implication that fits well with our results). Electrophysiological recordings within the dlPFC during a visual pop-out task shows selectivity for the target location at approximately 130 ms after the onset of the search array (Hasegawa et al., 2000). These results suggest a role of PFC in the control of visual attention, possibly through the provision of the top-down signals.

Miller and Cohen (2001) argued that all this indicates that the cardinal PFC function is to acquire and actively maintain patterns of activity that represent goals and the means to achieve them (rules) *and* the cortical pathways needed to perform the task (“maps”, hence “rulemaps”). Under this model, activation of a PFC rulemap sets up bias signals that propagate throughout much of the rest of the cortex, affecting sensory systems (visual attention) as well as the systems responsible for response execution, memory retrieval, emotional evaluation, etc. The aggregate effect is to guide the flow of neural activity along pathways that establish the proper mapping between inputs, internal states, and outputs to best perform the task. Establishing the proper mapping is especially important whenever stimuli are ambiguous (i.e. they suggest more than one behavioral response), or when multiple responses are possible and the task-appropriate response must compete with stronger, more habitual, alternatives. In short, task information is acquired by the PFC, which provides support to related information in posterior brain systems, effectively acting as a global attentional controller.

Under the Miller and Cohen model, the prefrontal cortex provides a biasing signal to posterior cortex, not only for attentional selection, but also acts to direct all cognitive functions. Although there is good circumstantial evidence that the PFC is the ‘top’ in top-down biasing signals, in order to demonstrate that PFC directs internal control, it is important to measure the flow of information during internally controlled behavior compared to externally controlled behavior (Miller and D'Esposito, 2005).

Neural Responses in Frontal Eye Fields (FEF)

The frontal eye fields (FEF) lie on the anterior bank of the arcuate sulcus. This area is involved in the production of eye movements and is typically defined as the region of prefrontal cortex that when stimulated with currents < 50 μ A will elicit a saccade (Bruce and Goldberg, 1985). FEF is

intereconnected with the eye plant in the brainstem (which gives it direct control over eye movements) as well as other eye movement related areas, such as the superior colliculus, LIP, and the supplementary eye fields (Stanton et al., 1988; Stanton et al., 1993, 1995). Additionally, FEF sends projections to visual areas V2, V3, V4, MT, MST, and TEO (Stanton et al., 1995). While ablating the FEF does not cause the complete loss of saccades, there is evidence that lesioned animals have difficulty making purposeful saccades (Bruce and Goldberg, 1985; Gaymard et al., 1998; Tehovnik et al., 2000).

Recent research has suggested that FEF plays a role in the control of visual attention. Work from Schall's group has shown that FEF neurons respond to the location of the target stimulus in a pop-out task approximately 50 ms before saccade onset (Schall and Hanes, 1993; Schall et al., 1995; Bichot et al., 2001; Sato et al., 2001). FEF neurons will respond to the target location regardless of whether the animal makes a saccade to the target's location (Thompson et al., 1997; Murthy et al., 2001; Sato et al., 2003; Thompson et al., 2005c). Furthermore, neurons in FEF will show target selectivity under both easy and difficult search conditions (Bichot et al., 2001; Sato et al., 2001). These results have led to Schall's group suggesting that FEF carries a saliency map that integrates both top-down and bottom-up information (Thompson and Bichot, 2005a; Thompson et al., 2005a). This model of FEF functionality is very similar to the saliency map suggested by the guided search model of visual search; it allows for the combination of bottom-up, physical attributes of stimuli with the top-down selection parameters used for search.

The role of FEF in the control of attention has also been supported by the experiments by T. Moore's laboratory. Moore and colleagues stimulated in FEF at subthreshold levels (i.e. levels that would not elicit a saccadic response) while recording from V4 neurons. The selected V4 neurons had receptive fields which overlapped with the motor field of the FEF stimulation site. The response of V4 neurons to visual stimuli were enhanced when FEF was stimulated, as if attention was being driven to that location (Moore and Armstrong, 2003). Further studies have shown that when directing attention to subparts of the receptive field the response of the V4 neuron will collapse around the selected region, in the same manner as found by Reynolds, Chellazi, and Desimone (Reynolds et al., 1999; Armstrong et al., 2006). Microstimulation in FEF will also boost the animal's behavioral discriminability at the target location (Moore and Fallah, 2001, 2004). Finally, recent experiments have shown that when stimulating

in FEF, neurons with overlapping receptive fields in V4 will show an increase in gain of their tuning curve, similar to when the animal attends to that location (Armstrong and Moore, 2007). These results suggest that FEF may play a direct role in directing attention to specific portions of the receptive field and that this direction may come from the same regions that actually induce the eye to move. Naturally, these results have lent strong support to the pre-motor theory of attention, discussed below.

Pre-Motor Theory of Attention

Humans make several saccades a second in order to focus objects onto their fovea. As the fovea is dense with photoreceptors, this allows for fine discrimination of the stimuli. Since the eye movement is explicit, this type of orientation is also known as overt attention. However, as outlined above, we are also able to covertly direct our attention to specific objects or locations allowing for deeper, finer analysis without having to move our eyes. Psychophysical experiments have found strong evidence that saccades cannot be made without first covertly allocating attention to the intended target location – shifts in attention always precede saccades (Rayner et al., 1978; Deubel and Schneider, 1996a; Peterson et al., 2004). This result has led to the pre-motor theory of attention which proposes that covert shifts of attention are due to a sub-threshold activation of the saccadic system (Rizzolatti et al., 1987). As noted above, recent experiments in FEF have supported this model – sub-threshold stimulation in FEF appears to bias posterior cortex in a manner similar to attention. Furthermore, all three regions found to be involved in the control of saccadic eye movements are also found to be activated by the covert direction of attention.

A recent electrophysiological study attempted to disassociate motor related activity from visual (or visuo-motor) activity in the frontal eye fields of monkeys by training them respond manually to a pop-out search task, as compared to the more common eye movement response (Thompson et al., 2005c). Thompson et al found that when no saccadic response was required, the neurons that responded in a manner that was purely motor (i.e. they did not respond to visual stimuli in their motor fields) were not activated by the visual search task. In contrast, visuo-motor neurons, which were activated both by the onset of visual stimuli and by saccadic movements, showed a strong selectivity for the pop-out target stimulus. In fact, the authors found that the purely movement related neurons were

suppressed during search. Based on these results, and previous results, the FEF is a possible region underlying the pre-motor signals involved in directing covert attention.

Simultaneous Recording

As we are interested in the relative roles of the three cortical regions putatively involved in the control of attention (LIP, dIPFC, and FEF), it is important that we expand upon typical electrophysiological techniques by recording simultaneously from all three regions of interest. This is necessary for several reasons. Absolute timing of neural activity can vary with further training/experience, between tasks, and with the statistical criterion used. Therefore, in order to accurately measure the flow of information between regions, our main interest was in the *relative* timing differences between areas. Simultaneous recording from multiple electrodes aids in detecting them because it reduces the influence of extraneous variables such as differences in performance across sessions. This is especially true in our case as the effects we are interested in contrasting are possibly subtle (such as timing differences between populations), and therefore simultaneously recording from all three regions will allow for improved discrimination. It is crucial to resolve the flow of information between different brain regions as it is the only way to truly determine which brain regions seem to be in control of different behaviors (Miller and D'Esposito, 2005). Simultaneous recordings also allow for relational timing effects, such as synchrony, to be investigated.

Goals of the Current Study

The first goal of this study, and the focus of the first part of the thesis, will be to determine the relative role of prefrontal and parietal cortices in the endogenous and exogenous control of attention. Animals were trained to perform both a covert visual search and covert visual pop-out task, utilizing internal and external control of attention, respectively. Simultaneously recorded neurons in LIP, FEF, and dIPFC will then be compared in order to determine when each region find the location of the target stimulus. The ordering of this latency will give insight into the direction of information flow between the two tasks and illuminate which region(s) is involved in the control of attention.

The second goal of the study will be to investigate the differences in the synchronous relationship between/with-in each region(s) during the tasks. As noted in the Introduction, there is extensive evidence that local synchrony plays a role in attention, and theoretical work suggesting synchrony could provide a mechanism for the dynamic change in effective connectivity between brain regions. Therefore, we aim to investigate the relationship between the local field potential of different regions and within regions, as well as the relationship between spiking activity and LFPs.

Our third aim is to resolve the similarities and differences in the neural networks underlying exogenous and endogenous control of attention. While there is considerable evidence from fMRI experimentation that both parietal and frontal regions are involved in both types of attentional control, due to the low spatial resolution of fMRI it is not possible to determine at the network level whether these two functions are separate or combined. We hope to investigate the commonalities and differences in the networks supporting behavior in each task.

Our fourth goal is to investigate the relative role of serial and parallel search mechanisms supporting endogenous and exogenous control of attention. Furthermore, we hope to determine which brain regions underlie these differing neural computations.

Introduction

Visual attention is thought to be controlled from both endogenous (internal) and exogenous (external) sources. Experiments with visual search paradigms provide the majority of evidence for the existence of these two forms of attentional control. For example, under most conditions, when a visual scene is presented attention will be automatically drawn to the most salient object in the scene, where saliency is defined by a variety of parameters, including relative brightness and movement (Jonides and Yantis, 1988; Treisman and Gormican, 1988). However, when a visual scene is presented that does not have any stimulus of particularly high saliency, in order to select a visual object for greater scrutiny, one must internally direct attention to the stimulus (this is classically shown in a conjunction search, Treisman and Gelade, 1980; Treisman and Gormican, 1988).

The best working model of both the psychophysical and neural effects of attention is the biased competition model (Desimone and Duncan, 1995). The model proposes that lateral inhibition between the neural representation of visual stimuli leads to competition between stimuli for representation in the cortex. Under this model the stimulus with the highest saliency will have the strongest representation (or even possibly the earliest representation), allowing it to win this competition. However, internal direction of attention can select specific stimuli and bias those representations to win the competition. While physiological support exists for the model's representation of competing stimuli (Reynolds et al., 1999), two important points remain with little to no experimental evidence: whether the competition proceeds in a bottom-up, feed-forward manner and the origin of the biasing signal. We hope to answer these questions by utilizing the exogenous and endogenous nature of visual pop-out and visual search in order to investigate the relative roles of parietal and prefrontal cortex in the control of visual attention. Most of this chapter has previously appeared in print (Buschman and Miller, 2007).

Visual Pop-out and Visual Search

In order to determine the respective roles of prefrontal and parietal cortices in the endogenous and exogenous control of attention, we recorded from multiple electrodes simultaneously implanted in frontal and parietal cortices as monkeys found a visual target under two conditions (Figure 1). Under

both conditions the target was randomly located in an array of four stimuli, and only differed in how the distractors related to the target. As with the classic human studies overviewed above, the ‘pop-out’ condition was designed to rely on exogenous, external, bottom-up control of attention. Under this condition the distractors were all identical and differed from the target along two dimensions (color and orientation), so the target’s salience automatically draws attention (Treisman and Gelade, 1980). In contrast, during the ‘search’ condition, each distractor independently differed from the target. Because the target matched some of the distractors in each dimension, it was not salient and had to be sought just using the memory of the sample stimulus. As noted above, this form of visual search requires endogenous, internal, top-down control of attention.

The monkeys showed the behavioral hallmarks of bottom-up versus top-down attention. As with humans, under these two conditions, we can measure the cost of an added distractor by varying the number of distractors in the search array. There was a shallower increase in reaction time (RT) with more distractors during pop-out than search (6 ms/item for pop-out; 22 ms/item for search; $p < 0.001$, t-test of least-squares linear regression, Figure 2). This difference suggests a mechanistic difference underlying the two types of search. Furthermore, under recording conditions, when there were always four items in the search array (three distractors, one target), the monkeys’ reaction time was significantly longer and more variable for search than pop-out (Figure 3). Average RTs for search (272 ms) and pop-out (233 ms) differed significantly ($p < 10^{-5}$, t-test). The variance in RT also differed significantly (standard deviations of 43 ms for search and 33 ms for pop-out, $p < 10^{-5}$, χ^2 test). This difference is also typical for visual search and pop-out in humans (Treisman and Gelade, 1980; Duncan and Humphreys, 1989; Wolfe et al., 1989) and further reflects different mechanisms underlying the behavior.

Flow of Information about the Locus of Attention

We were interested in the respective roles of prefrontal and parietal cortex in the direction of attention under these two tasks. In order to determine these roles, we leveraged the behavioral design: the animal always made its choice with a saccade to the target location and previous psychophysical work has shown saccades are preceded by shifts in attention. Therefore, in order to establish when each anatomical area directs attention to the target location, we can determine the time-point at which

Figure 1




Figure 1. Schematic of Task Design. Red circle indicates eye position throughout the task. Both Search and Pop-out required the animal to begin the task fixating, followed by the sample stimulus, which was the target to be found in the visual array. After a short delay, the visual array was presented and the animal was required to make a single, direct, saccade to the target location in order to receive a reward. Visual Search and Pop-out tasks only differed in how the distractors related to the target in the search array.

neurons in each area signaled the target location. This activity reflects both the pre-saccadic shift in attention to the target as well as the eye movement itself. Since the saccade acts as a behavioral marker which informs us of when the animal knew the target location, we can look backwards in time, relative to the saccade, in order to find the neural correlates of directing attention.

We determined when each neuron first “found” (reflected) the target location by computing when the amount of information in its firing rate about target location first reached significance. Significance was determined via the mutual information statistic (see Methods). We were interested in the amount of information about the target location carried by the firing rate of an individual cell. Trials were broken down into groups by condition (pop-out or search) and by target location. This grouped across differences in target identity (color and orientation) as well as distractor identity, controlling for any selectivity for these parameters. As attention is known to precede a saccade (Rayner et al., 1978; Deubel and Schneider, 1996b; Peterson et al., 2004), measuring selectivity in this manner ensures any selectivity observed is only related to the allocation of attention or to the eye movement itself.

The degree of selectivity was calculated across the entire timecourse of the trial in non-overlapping windows of 25 ms, where the amount of information about the target location contained in the firing activity of each cell was determined within each bin. Significance was determined within each bin independently by comparing the observed amount of information against a null distribution created through randomization tests (see Method). The amount of observed information was denoted as significant only if it exceeded 95% of the null distribution ($p < 0.05$ in each bin). The point at which an individual neuron began to significantly reflect the target location was defined as the time point at which significant information was found for two consecutive bins (chance level = 0.05^2 or $2.5 \cdot 10^{-3}$). This is known as the *time to first significance* for each neuron. Figure 4, top and bottom, show the firing rates (for pop-out only) as well as the mutual information over time (for both pop-out and search) for an example dIPFC and LIP neuron, respectively. The starred position shows the time to first significance for each cell. As will be seen on the population level, the LIP neuron shows selectivity for the target location well before dIPFC during pop-out, while dIPFC shows earlier selectivity than LIP during search.

Figure 2

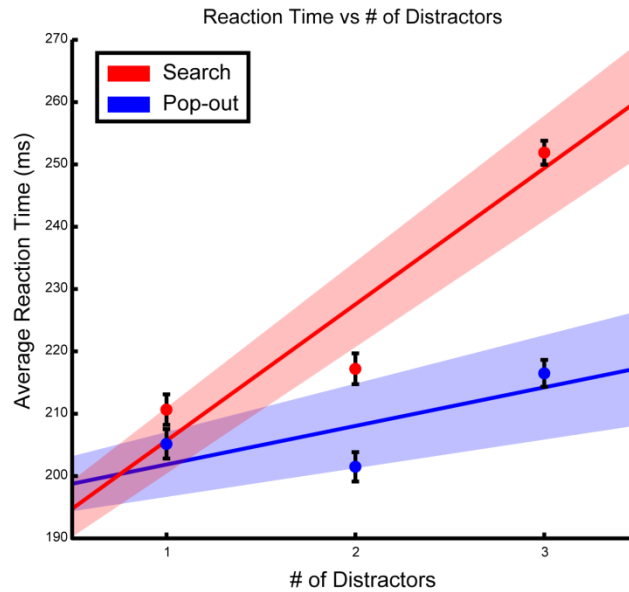


Figure 2. Reaction Time Increase with Increased Distractors. Psychophysical testing with variable number of distractors in the search array showed differences between search (in red) and pop-out (in blue). The average reaction time for the animal to find the target is plotted as a circle for visual arrays with 1, 2, and 3 distractors. Black lines show standard deviation. The solid colored lines show the linear fit for both search (red) and pop-out (blue). Shaded regions show 95% confidence interval about these fits. There was a shallower increase in reaction time (RT) with more distractors during pop-out than search (6 ms/item for pop-out; 22 ms/item for search; $p < 0.001$, t-test of least-squares linear regression).

Figure 3

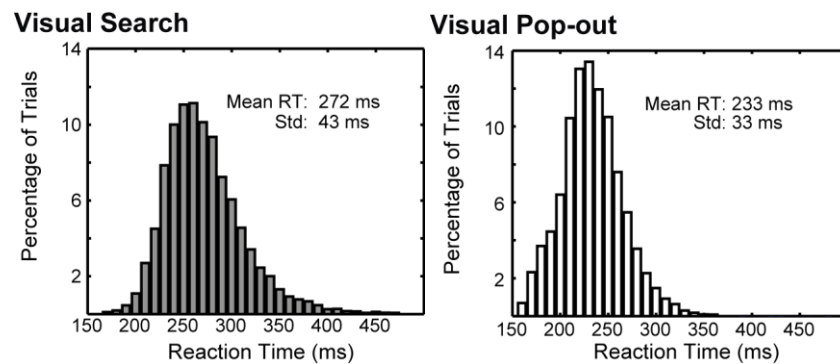


Figure 3. Overall reaction time distributions for Search and Pop-out. Reaction time distributions for Search (left) and Pop-out (right) from both animals during recording sessions (with 4 items in visual array). Visual search showed both an average increase in reaction time as well as a more variable reaction time. Average RTs for search (272 ms) and pop-out (233 ms) differed significantly ($p < 10^{-5}$, t-test). The variance in RT also differed significantly (standard deviations of 43 ms for search and 33 ms for pop-out, $p < 10^{-5}$, χ^2 test).

The time to first significance was determined for each neuron recorded, generating a distribution of when neurons from each anatomical region first found the target location relative to the start of the saccade. These distributions are shown in Figure 5. During pop-out, there was a bimodal distribution for all three regions (Figure 5, left). For each area there was a population of neurons that first found the target well before the saccade (i.e., shortly after visual array onset) and a separate population that found the target after the saccade. The early population consisted of 35% of all target location selective neurons in LIP (24/68), 51% in dIPFC (40/78), and 31% in FEF (17/54). During pop-out there were clear differences in timing across these early populations: LIP neurons found the target first, followed by dIPFC neurons, and then FEF neurons.

In contrast, during the search task, neurons began finding the target later than in pop-out, just before the saccade, and in the reverse order: the frontal areas (dIPFC and FEF) showed selectivity first, followed by LIP (Figure 4B). About 1/3 of all selective neurons in FEF and dIPFC began to reflect the target location before the saccade (19/60 and 21/70, respectively), while only 14% (8/58) of selective LIP cells did so. This is greater than expected by chance for FEF and dIPFC but not for LIP ($p = 8 \times 10^{-5}$, $p = 4 \times 10^{-5}$, and $p = 0.41$, respectively, tested against binomial distribution).

The distribution of when individual neurons in each area first showed significant information for the target location appeared to be multimodal. In order to quantify the number and location of modes we fit models with a mixture of one, two, or three Gaussians to the data. We used the Bayesian Information Criterion (BIC, Schwarz, 1978) to determine how well each model fit the observed data. The BIC is a combination of the residuals between the model and data and the number of free parameters in the model, correcting for the advantage of more complex models. The model with the lowest BIC is the model that has the best fit to the data without over-parameterization. As is shown in Table 1 we found that a bimodal distribution fit the best for the observed data in all three anatomical regions during both search and pop-out. The resultant R^2 of the fits (Table 2) all show a large proportion of the variance in the observed data was captured by the mixture of Gaussians. The bimodal fit estimated both the mean and variance of the distributions. The deviation about each parameter was estimated using the Fisher Information and was used to both calculate a confidence interval for each parameter and to test for significant differences between areas (similar to a t-test).

Figure 4

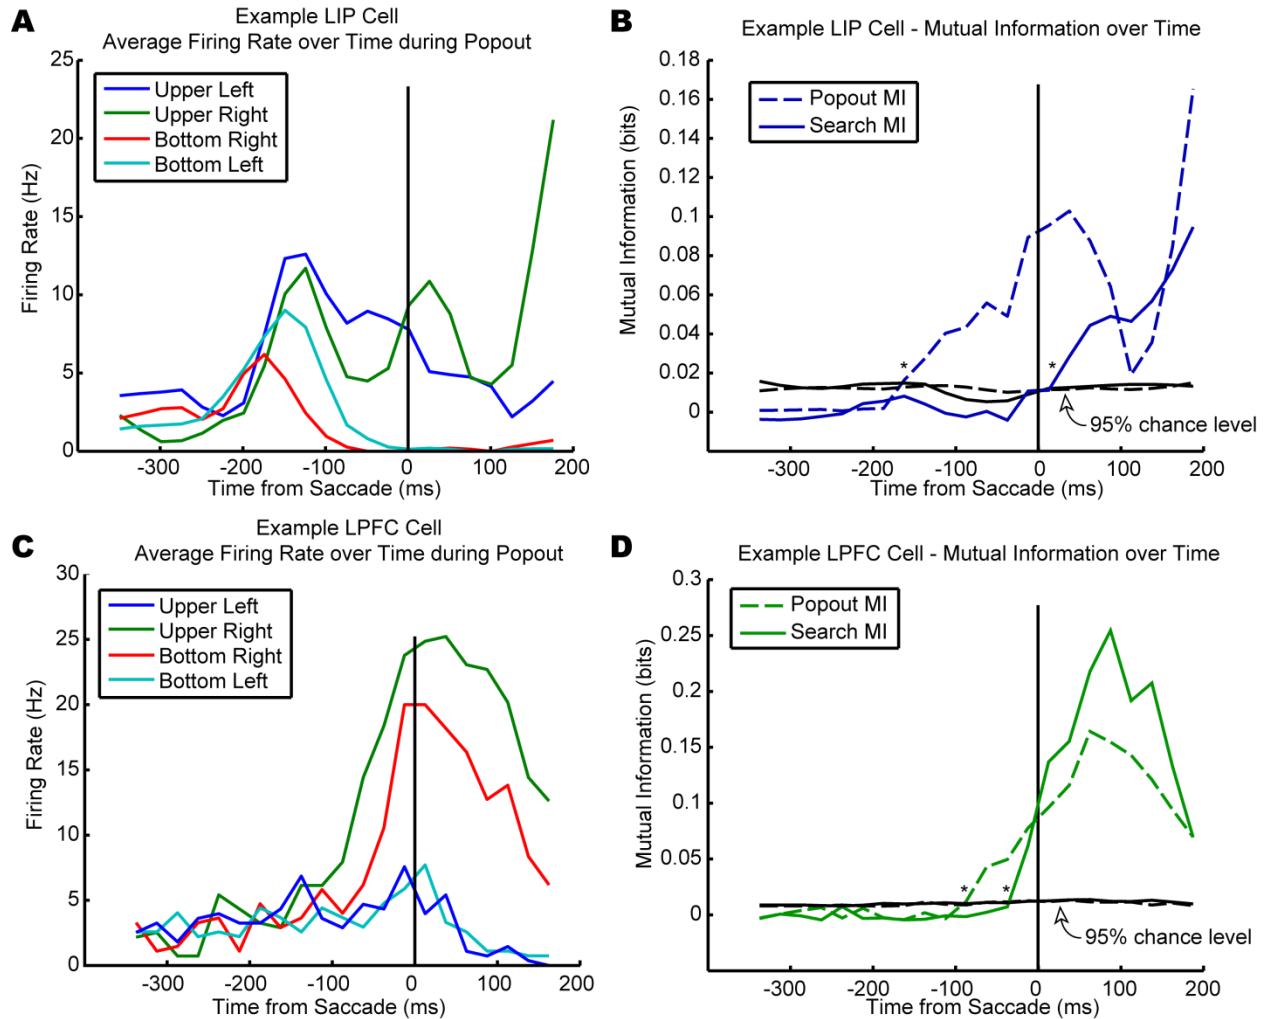


Figure 4. The firing rate of example LIP and LPFC neurons during the visual pop-out task (left column, top and bottom, respectively). Trials were aligned on saccade in order to capture the shift of attention before the saccade. Selectivity for target location, regardless of the identity of the target, can be clearly seen in the firing rate histograms and was captured in the mutual information statistic. Panels in the right column show the amount of information about the target location carried in the firing rate for the example LIP and LPFC neurons, respectively. The amount of information over time is shown for both the visual search and pop-out tasks. The asterisk indicates the time-point at which the observed amount of information was significant for two consecutive bins at $p < 0.05$. This marks the time to first significance. The population effects shown in the main text are reflected in these example neurons: during pop-out, selectivity in LIP precedes LPFC, while during search, LPFC precedes LIP.

Figure 5

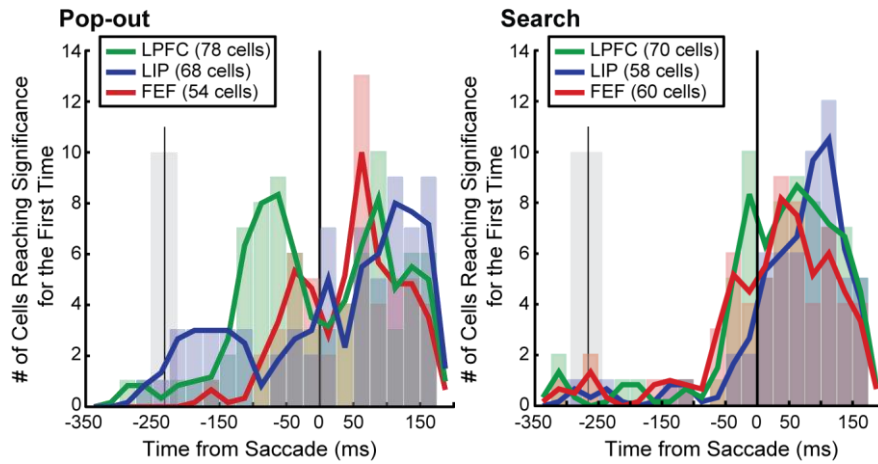


Figure 5. Distribution of time to first show significant target information during pop-out (left) and search (right) relative to the saccade. Vertical black line indicates saccade, grey shaded regions indicate mean and +/- one standard deviation of distribution of visual array onset.

For pop-out, fits of bimodal Gaussians (Figure 6, left column) indicated that the early population of LIP neurons was centered at 162 ms before the saccade (bottom row; 95% confidence interval or CI: 200 ms – 124 ms), followed by the early populations in dIPFC and FEF, 77 ms (middle row; 95% CI: 84 ms – 70 ms) and 40 ms (top row; 95% CI: 56 ms – 23 ms) before the saccade, respectively (LIP < PFC, $p < 10^{-25}$; LIP < FEF, $p = 6 \cdot 10^{-8}$; PFC < FEF, $p = 6 \cdot 10^{-5}$; t-test). The distribution of the neurons that found the target after the saccade was overlapping in all three areas and centered about 100 ms after saccade (Figure 6, left column).

Curves were also fit to the distributions from the search condition (Figure 6, right column) and showed the reverse ordering. Estimates of the early populations showed that FEF and dIPFC preceded LIP: FEF and dIPFC had early modes at 46 ms (top row; 95% confidence interval: 75 – 17 ms) and 19 ms (middle row; 25 - 13 ms) before the saccade, respectively, followed by LIP at 19 ms (bottom row; 8 – 30 ms) *after* the saccade (FEF < LIP, $p = 8 \cdot 10^{-13}$; dIPFC < LIP, $p = 6 \cdot 10^{-8}$; all comparisons by t-test). The distributions of neurons that found the target after the saccade were overlapping and centered at approximately 100 ms after saccade (Figure 6, right column).

In order to determine when a neural population as a whole began to represent the target location, we compared the observed cumulative distribution to the distribution that is expected by chance (see Figure 7 for raw cumulative histograms and chance distribution). Given our significance criterion of two consecutive bins of $p < 0.05$, we calculated the average number of neurons expected to reflect the target location by chance in each 25 ms time bin, along with the standard deviation about that mean. As each comparison was conducted independently for each 25 ms step, the number of comparisons increases with time. But as a neuron can only ‘first’ show selectivity once, the chance level does not increase linearly. Therefore, a Monte-Carlo analysis was used to estimate the number of significant neurons in each time bin by chance. Use of a binomial distribution yielded similar results to the Monte-Carlo analysis (although the binomial distribution allows for an analytical estimate of significance levels, it does not correct for the absence of multiple ‘firsts’, and therefore provides a theoretical upper-bound to the significance criterion). By subtracting the average number of neurons expected by chance from the observed distribution, and normalizing by the standard deviation, we constructed a z-score for the entire population (Figure 8). This corrected for multiple comparisons

Table 1 & 2

		SINGLE MODE	DOUBLE MODE	TRIPLE MODE
LIP	Pop-out	-563.9	-596.2	-583.5
	Search	-518.0	-569.2	-493.6
LPFC	Pop-out	-651.5	-794.7	-781.7
	Search	-641.7	-713.2	-616.2
FEF	Pop-out	-424.1	-439.3	-427.3
	Search	-562.2	-572.5	-537.7

Table 1. Table of Bayesian Information Criterion (BIC) values of model fit to distribution of time to first significance for each area and task (Figures 5 and 6). The BIC describes the goodness of fit for varying models while correcting for model complexity. The best fitting model is the one with the lowest BIC value, which we found to be a mix of two Gaussians for all three regions on both tasks.

	POP-OUT	SEARCH
LIP	0.92	0.98
LPFC	0.94	0.97
FEF	0.88	0.94

Table 2. Table of r^2 values for fit of mixture of two Gaussians on distributions of time to first significance for each area and task (Figures 5 and 6). The r^2 provides a measure of how much of the variance in the data is captured by the model fit.

Figure 6

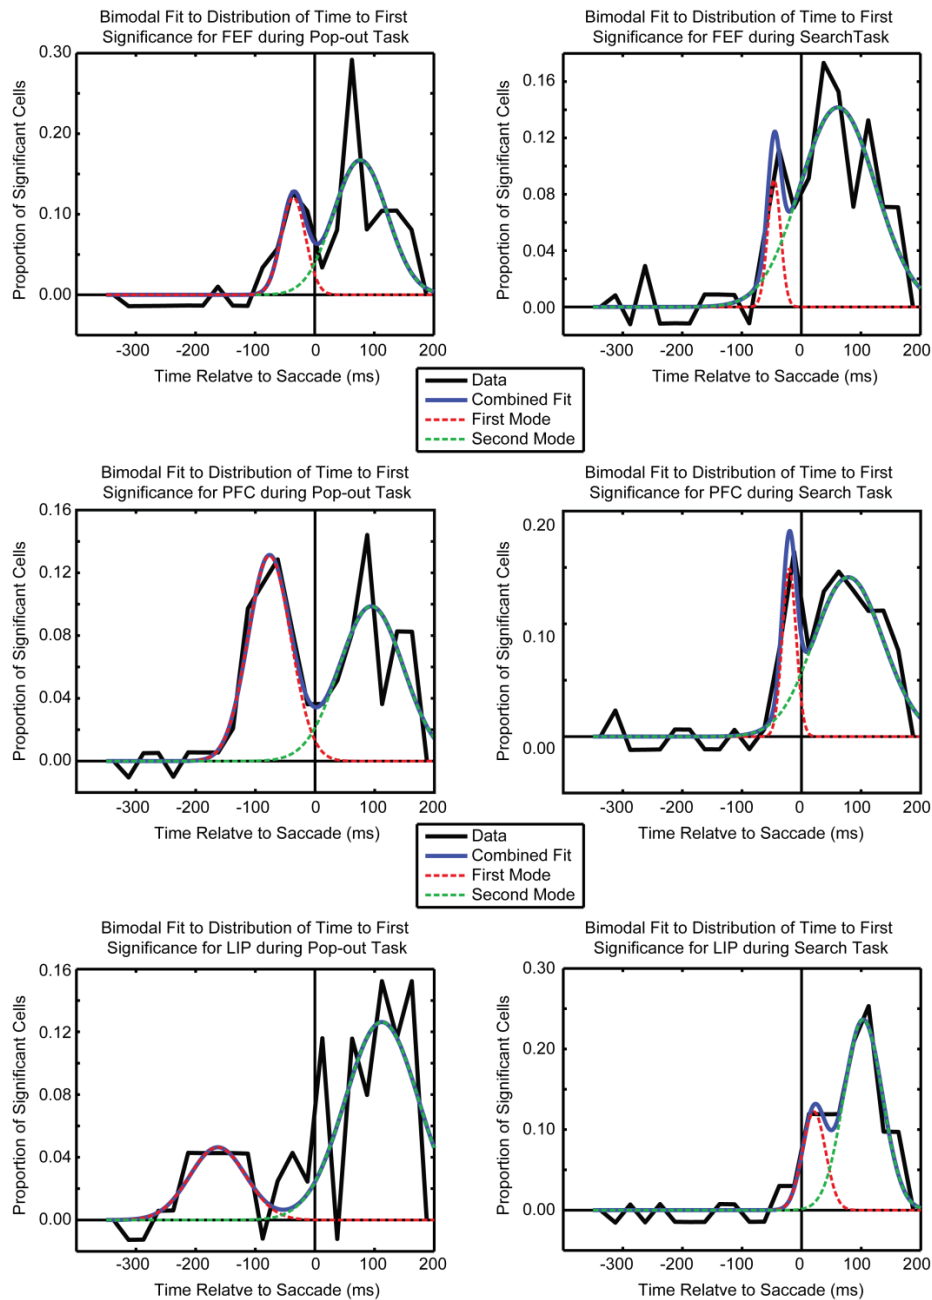


Figure 6. Fits of bimodal gaussians to distribution of time to first find target location, relative to saccade. Vertical black line indicates saccade. Black lines show the distribution of time to first find the target in each of the three recorded areas (corrected for chance). Data is shown for pop-out (left column) and search (right column) and for FEF, LPFC, and LIP (top row, middle row, and bottom row, respectively). As shown in Tables 1 and 2, the bimodal gaussian distributions fit the observed data well.

across time bins and gave us a specific latency at which the entire observed population carried significant information about the target location. Differences between anatomical areas were tested for significance through a randomization method: we randomly assigned neurons to different anatomical areas and re-calculated the difference in the time for each population to reach significance. This allowed us to construct a null distribution and determine the p-value of the observed difference in latency between anatomical areas.

The same ordering as seen in the fits of bimodal distributions was seen in the cumulative distributions of the time to first significance data (Figure 8). During pop-out, LIP, dIPFC and FEF neurons began finding the target 170 ms, 120 ms, and 35 ms before saccade, respectively (Figure 6, left column; LIP < PFC, $p < 0.05$; LIP < FEF, $p = 0.009$; PFC < FEF, $p = 0.002$, randomization test). A quarter of the selective LIP neurons (17/68) began encoding the target location before the dIPFC population first carried significant information. This is more than expected by chance ($p = 2 \cdot 10^{-5}$, tested against binomial distribution). When aligning trials on visual array onset instead of saccade, we found similar results (see below).

For search, the cumulative distributions in Figure 8, right column, show that target location information reached significance in the FEF and dIPFC at 50 and 40 ms before the saccade, respectively, followed by LIP (dIPFC and FEF were earlier than LIP, $p = 0.027$ and $p = 0.006$, respectively; randomization test). In fact, although during pop-out, LIP neurons found the target first, and well before the saccade, during search target location information in LIP did not reach significance until 32 ms *after* the saccade. As with pop-out, similar results were observed when trials were aligned on visual array onset (see below).

Figure 7

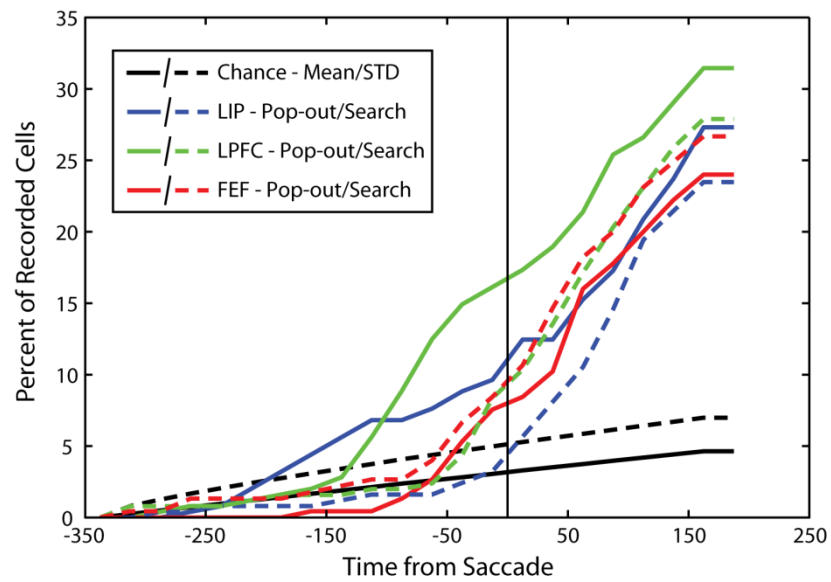


Figure 7. Raw cumulative sum of time to first significance histogram for LIP, LPFC, and FEF in both the search and pop-out conditions are shown in colors. The black solid (dashed) lines show the mean (standard deviation) of the number of significant cells expected by chance.

Figure 8

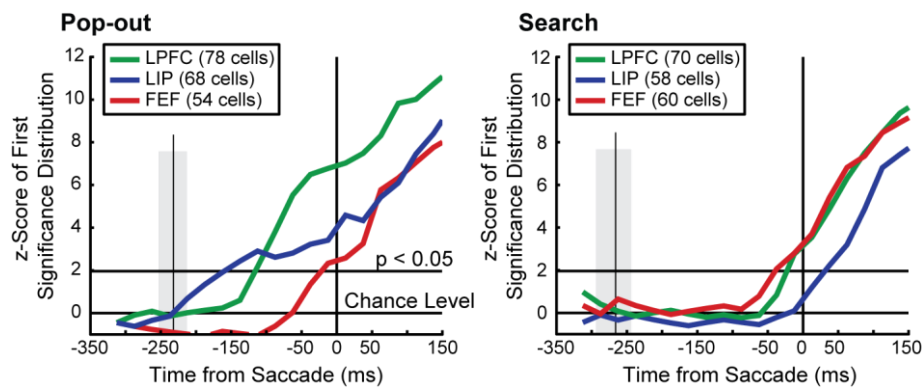


Figure 8. Timing of selectivity for each area aligned on saccade. Normalized cumulative sum of the histogram shown in Figure 5. A z-score for the observed distribution was calculated through randomization tests and was corrected for multiple comparisons. The same ordering as seen in the fits of bimodal distributions was seen in the cumulative distributions of the time to first significance data (Figure 5 and 6).

Aligning Data on Array Onset versus Saccade Initiation

Although aligning trials on saccade initiation allows us to look at the process of attending to and selecting a target location relative to the key behavioral response, it is also possible to investigate the ordering of selectivity for each area when trials are aligned on visual array onset. Array onset marks the onset of the search and therefore any deterministic process involved in processing the visual array and finding the target location must proceed from this point. It is not entirely clear which parts of visual search are locked to visual array onset or precede the target selection (i.e. are locked to the saccade). For example, the pop-out task, since it seems to be parallel in nature and rely mostly on bottom-up signals, might show a greater locking to the onset of the visual array (as reflected by the tighter reaction time distribution). In contrast, a completely serial search task will only show attention selection relative to the saccade since it takes a variable, unknown, amount of time to find the target. Since it is not clear whether search tasks are deterministically determined by visual array onset or by a generative process leading to a saccade, we investigated the timecourse of selectivity for each area when aligning trials on both markers.

Selectivity relative to visual array onset was determined in exactly the same manner as relative to saccade, with the obvious exception that trials were aligned on array onset. Mutual information was calculated in independent 25 ms bins and significance was determined in each bin alone through randomization tests. The time to first significance for individual neurons was again denoted as the time at which a neuron carried significant information about the target location in two consecutive bins. The distribution of times to first significance for all three areas with trials aligned on onset are relatively similar to trials aligned on saccade. Most importantly, when trials were aligned on visual array onset for both the pop-out and search conditions, the ordering of selectivity across areas was the same as when trials were aligned on saccade. Furthermore, bimodal distributions are again suggested, although both modes for each area's distribution appear to have greater variance (possibly indicated aligning on saccade is the more appropriate choice).

Cumulative histograms were again used to determine when an entire area began to carry significant information (or alternatively, when an area had a significant number of cells indicating the target location). Results were similar to those found when aligned on saccade (Figure 9). For the pop-

Figure 9

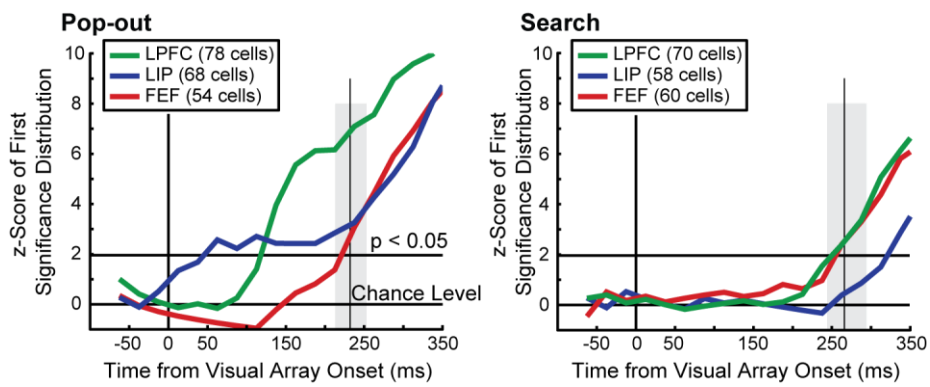


Figure 9. Timing of selectivity for each area aligned on array onset. Cumulative histogram of times to first showing significant selectivity when trials are aligned on visual array onset for pop-out (left) and search (right). Vertical black line indicates visual array onset, grey shaded regions indicate mean and +/- one standard deviation of distribution of saccade.

out condition (Figure 9, left), while the distribution was more variable due to the variability in reaction time (as indicated by a shallower slope), LIP showed selectivity for the target location approximately 50 ms after array onset, followed by dIPFC and then FEF (after 120 and 220 ms, respectively). All these differences were also significant (LIP < PFC, $p = 0.038$; LIP < FEF, $p = 0.013$; PFC < FEF, $p = 0.001$; randomization tests, see Methods). When aligning visual search trials on visual array onset, dIPFC and FEF carried significant information at 250 ms after array onset, significantly preceding selectivity in LIP, which began at 320 ms after onset (Figure 9, right; $p = 0.044$ and $p = 0.047$, respectively, by randomization test). The overall longer neural latency to find the target during search (approximately 250 ms after array onset) compared to pop-out (between 50 and 100 ms after array onset) parallels the differences in RT for the two conditions (Figure 3).

Controlling for Visual Array Response and Target Selectivity

One concern might be that the observed differences between visual search and pop-out are due to different response rates to the visual array during the two tasks. In order to ensure that this is not the case, Figure 10, left column, shows the normalized response of location-selective neurons during both tasks. There are no significant differences in response between the two tasks for any of the three anatomical areas of interest. This suggests that the difference in location information is not due to a difference in the overall responsiveness of neurons to the visual array, but rather in how the responses differentiate between targets and distractors. An alternative way to compare the response to the visual array during search and pop-out is to compare the difference in firing rate of individual neurons in each task to zero. The difference in response to the visual array during visual search and pop-out did not differ from zero in any of the three areas (until late in the trial, again due to differences in reaction time).

The overall average response of location-selective LIP neurons shows a modest increase overtime, without a large visual array onset response. This is due to the variable nature of LIP responses, as shown in Figure 10, right column. All three of these example LIP neurons have very similar responses to the visual array onset (although they do tend to differentiate to a greater degree later in the trial as the reaction time to the tasks differs). The dynamics of single units are also further discussed below.

Although we do not observe a difference in response to the visual array between the two tasks, the existence of such a difference does not necessarily lead to a difference in location selectivity. As the location selectivity is measured between conditions within a task, a generalized increase in response would not necessarily lead to information about the target location. One advantage of using the mutual information technique is that it does not suffer from the boundary condition issues of other tests that assume an underlying distribution not well suited for spiking activity (such as t-tests, ANOVAs, etc). Furthermore, since significance is determined using a randomization procedure, it is not the specific values of neural activity that are being tested, but rather whether their distribution is more informative about the target location than random chance.

Another important question is how much the location selectivity of individual neurons in each of our three areas were also selective for target identity. In order to determine the degree of target identity selectivity we examined the activity of all location selective neurons from each area and for both tasks when the target lay in its preferred direction (as defined as the target location to elicit the greatest response from the neuron). As can be seen in Figure 11, very few neurons carried information about the target identity; rather they only indicated the location of the target. While Figure 11 only tests for target identity information in the same time bin as the location information, we can also test for target identity information proceeding or following the location selective information and over the entire trial, relative to the saccade. For no time bin relative to when an individual neuron carried significant information about the target location, nor for any time bin relative to the saccade, were there a significant number of neurons carrying target identity information. These results support the model that these neurons are involved in the spatial direction of attention, not in the comparison of currently selected objects with the memory of the target stimulus. However, as will be shown later in this thesis, there are many neurons in all three regions that are involved in maintaining information about the target stimulus, likely for this purpose. The number of neurons showing an interaction between the target location and its identity is very small – perhaps unexpectedly so. An interaction effect would not necessarily change the interpretation of our results (as the neuron’s response would still carry significant information about the target location, all that is needed to follow the location of attention), it is nonetheless interesting that the brain appears to completely differentiate the neural populations support these two functions.

Figure 10

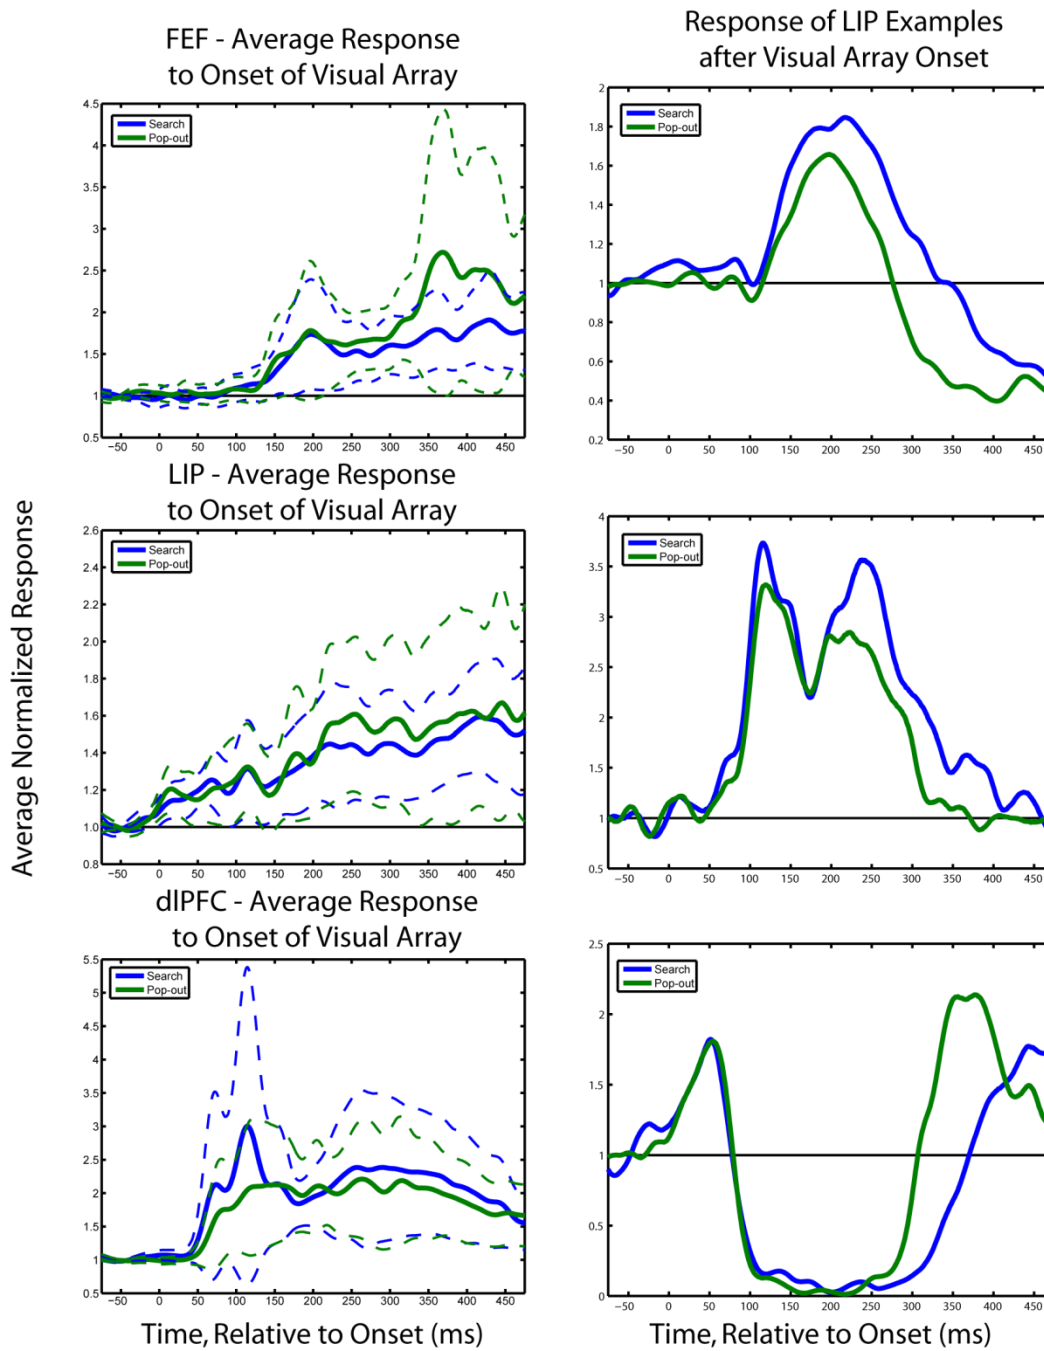


Figure 10. Left column: Normalized firing rate across time (relative to visual array onset) for all location-selective neurons in FEF (top), LIP (middle), and dIPFC (bottom). Average normalized firing rate is shown for both visual search (blue line) and pop-out (green line). 95% CI for mean is shown with dashed line. Right column: Normalized firing rate of three example LIP neurons. Response to visual search is shown in blue, pop-out in green. These examples show the variety of responses in LIP (more-so than either dIPFC or FEF).

Figure 11

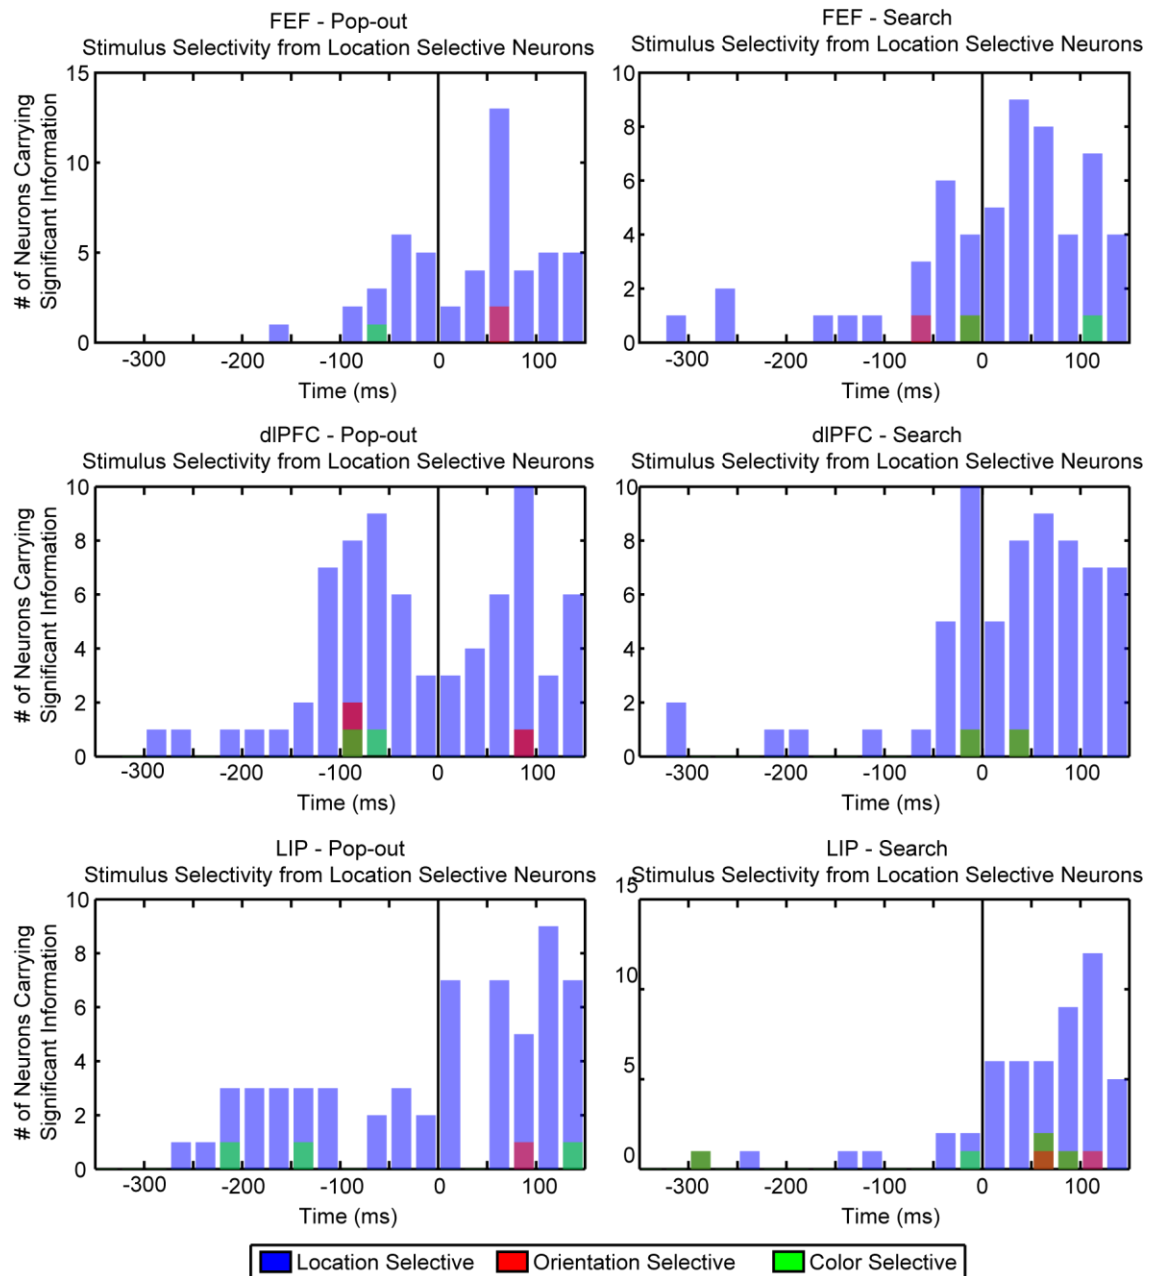


Figure 11. Information about target identity in location-selective neurons from FEF (top row), dIPFC (middle row), and LIP (bottom row) for both pop-out (left column) and visual search (right column). Selectivity for the target orientation (red) and color (green) are shown overlapping location selectivity. Very few location selective neurons show target identity information, suggesting the process of directing attention to a spatial location is separate from the identity of the target at that location.

Dynamics of Selectivity for Single Units

The previous analysis focused on when each individual neuron, and then when the population of neurons from a given area, first carried significant information about the target location (and therefore the direction of attention to it). Figures 12, 13, and 14 show the selectivity for target location for every neuron recorded in LIP, dIPFC, and FEF, respectively. The top plot in each figure shows the selectivity over time for individual cells during visual pop-out, while the bottom plot shows the same for visual search. As previously described, the neural activity from each cell was binned into 25 ms independent bins across time, relative to saccade. The color of each point indicates the observed information about the target location relative to the 95% level of the null distribution (which, as described above, was derived from a randomization procedure). The neurons have been sorted by their time to first significance (determined, as above, as the point in time where the observed information about the target location exceeded 95% of the null distribution for two consecutive time bins). Non-significant cells are randomly ordered above the selective ones. The heavy blue line follows the time to first significance for all selective cells.

As was seen previously, selectivity in prefrontal and parietal cortex was doubly dissociated between the pop-out and search tasks. Parietal cortex found the target early in bottom-up, exogenous attention and prefrontal cortex driving target selection in top-down, endogenous attention. Again, although it is more difficult to see in this format, all three regions also show a bimodal distribution of times to first significance.

The dynamics of target related information not only varied with time across the entire anatomical area but also for each cell alone. Although a variety of temporal responses exists for neurons in all three regions across both tasks, it is interesting to note the relative temporal-sparseness of LIP responses in both tasks when compared to the other two regions. Although there are a few neurons that show sustained information, LIP neurons appear to transiently represent the target location. Once one neuron's response becomes reduced another neurons response starts. This effect is reflected in the relatively sparse color underneath the blue line in the LIP plots. In contrast, for both dIPFC and FEF, neurons appear to represent maintain their representation across the behavioral response.

These differences may be reflective of the nature of the information each area is maintaining – LIP appears to be representing the most salient object in the visual field, likely enhancing its representation and passing it forward. However, while it is important to represent a salient object in the visual field, it is also important to remain flexible to responding to newer, more salient objects. In fact, over-representing a particular stimulus would lead to it becoming more difficult to later disengage when a newer, more salient, stimulus is presented. An area which is representing the saliency of the visual field should maintain a balance, possibly reflected in the staccato representation observed in LIP. In contrast, dlPFC and FEF are regions involved in making and executing the decision, and therefore the response to the target location should be strongly maintained through the behavioral response.

Figure 12

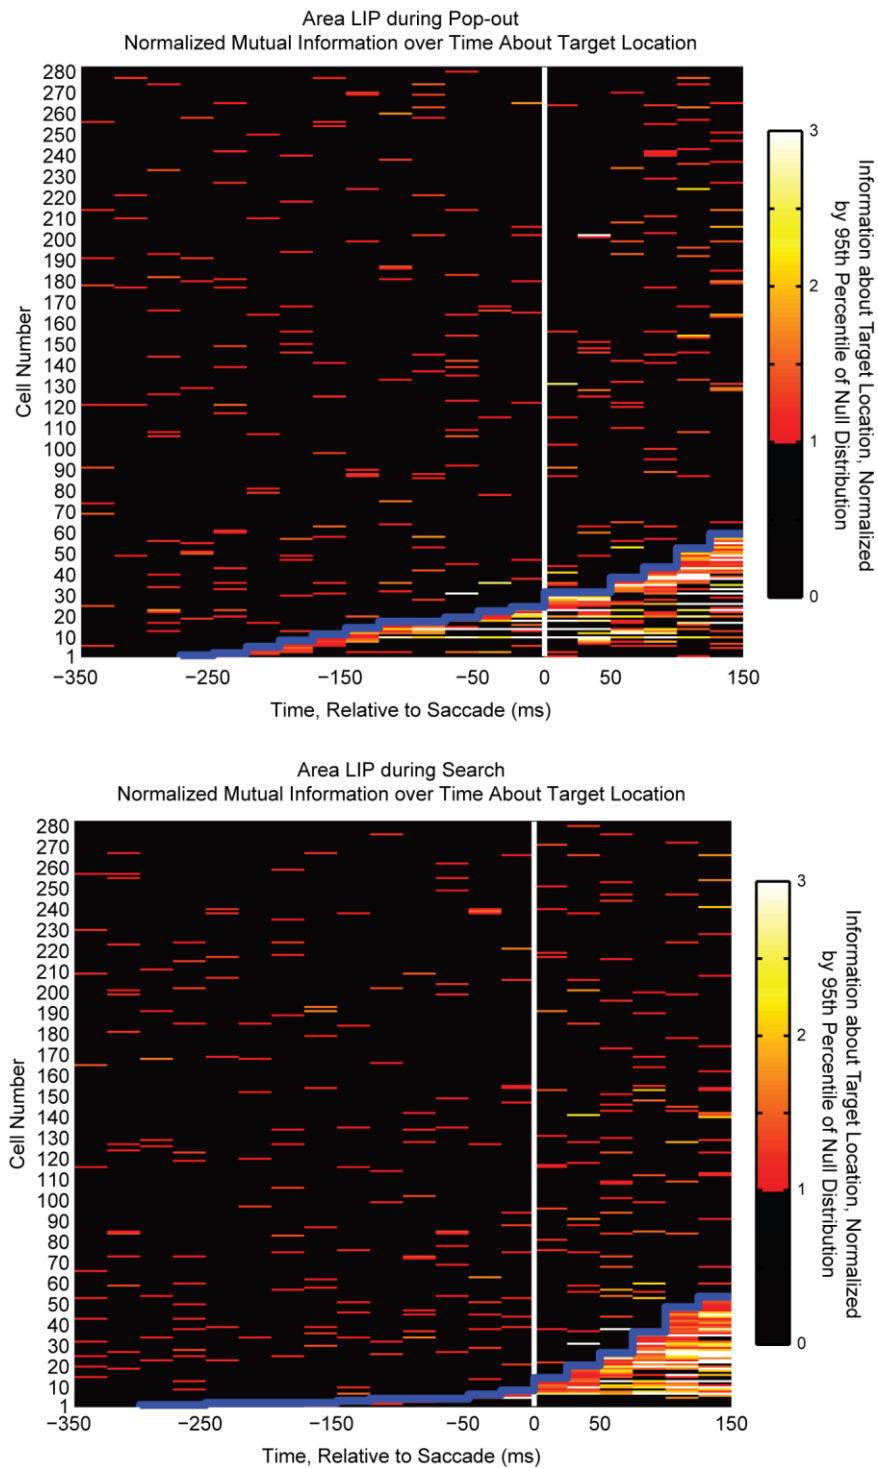


Figure 12. Normalized mutual information (MI) across time for all recorded neurons in **LIP** for both pop-out (top) and search (bottom). Neurons are sorted by their time to first significance, with the blue line following trend over time. MI was calculated in independent 25 ms time bins and neurons were only taken to carry significant information when above chance for two consecutive bins.

Figure 13

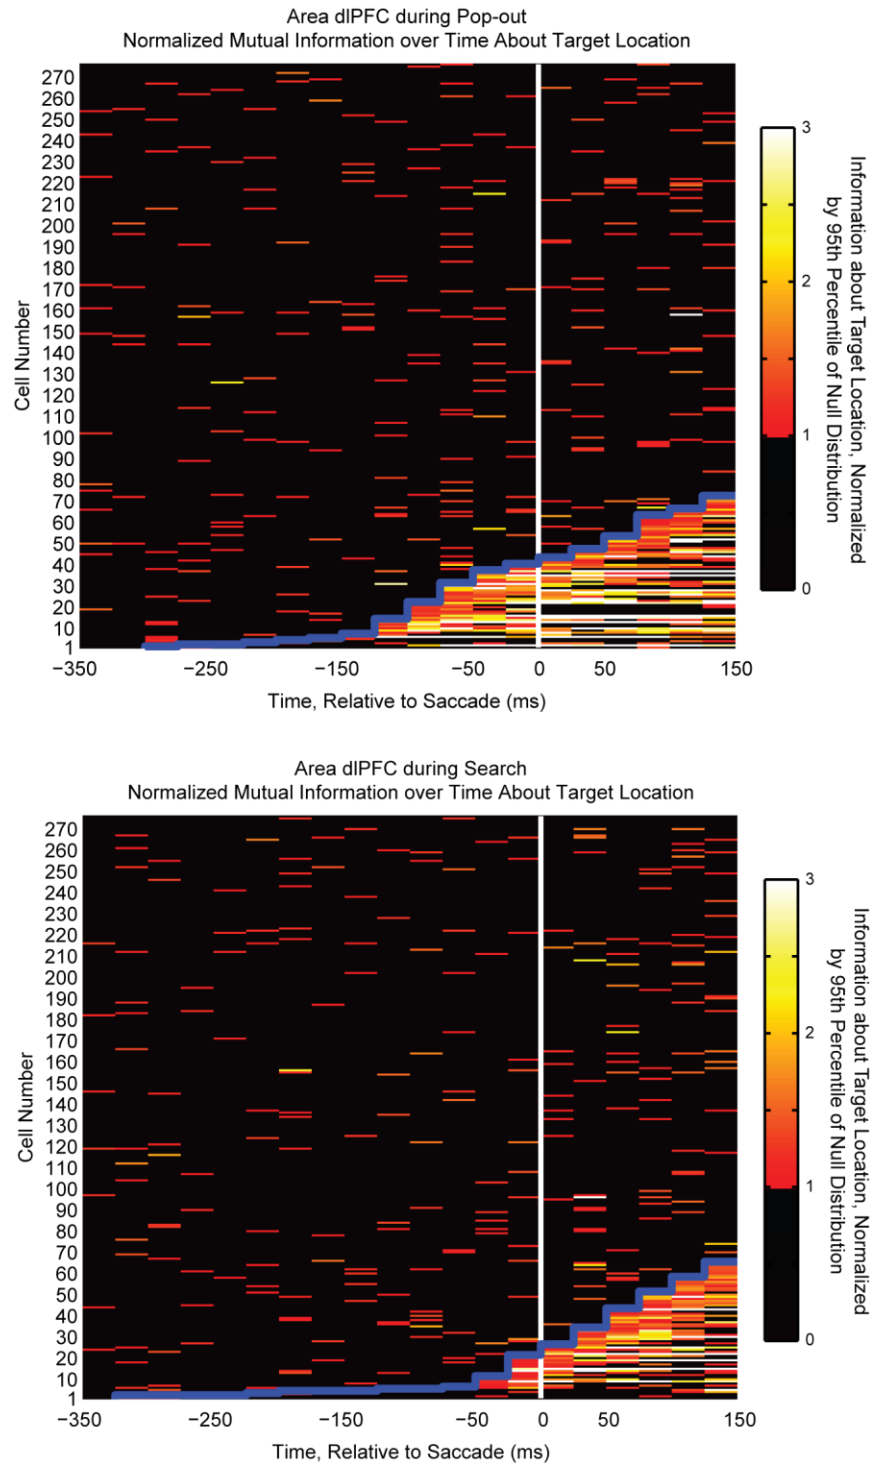


Figure 13. Normalized mutual information (MI) across time for all recorded neurons in **dIPFC** for both pop-out (top) and search (bottom). Neurons are sorted by their time to first significance, with the blue line following trend over time. MI was calculated in independent 25 ms time bins and neurons were only taken to carry significant information when above chance for two consecutive bins.

Figure 14

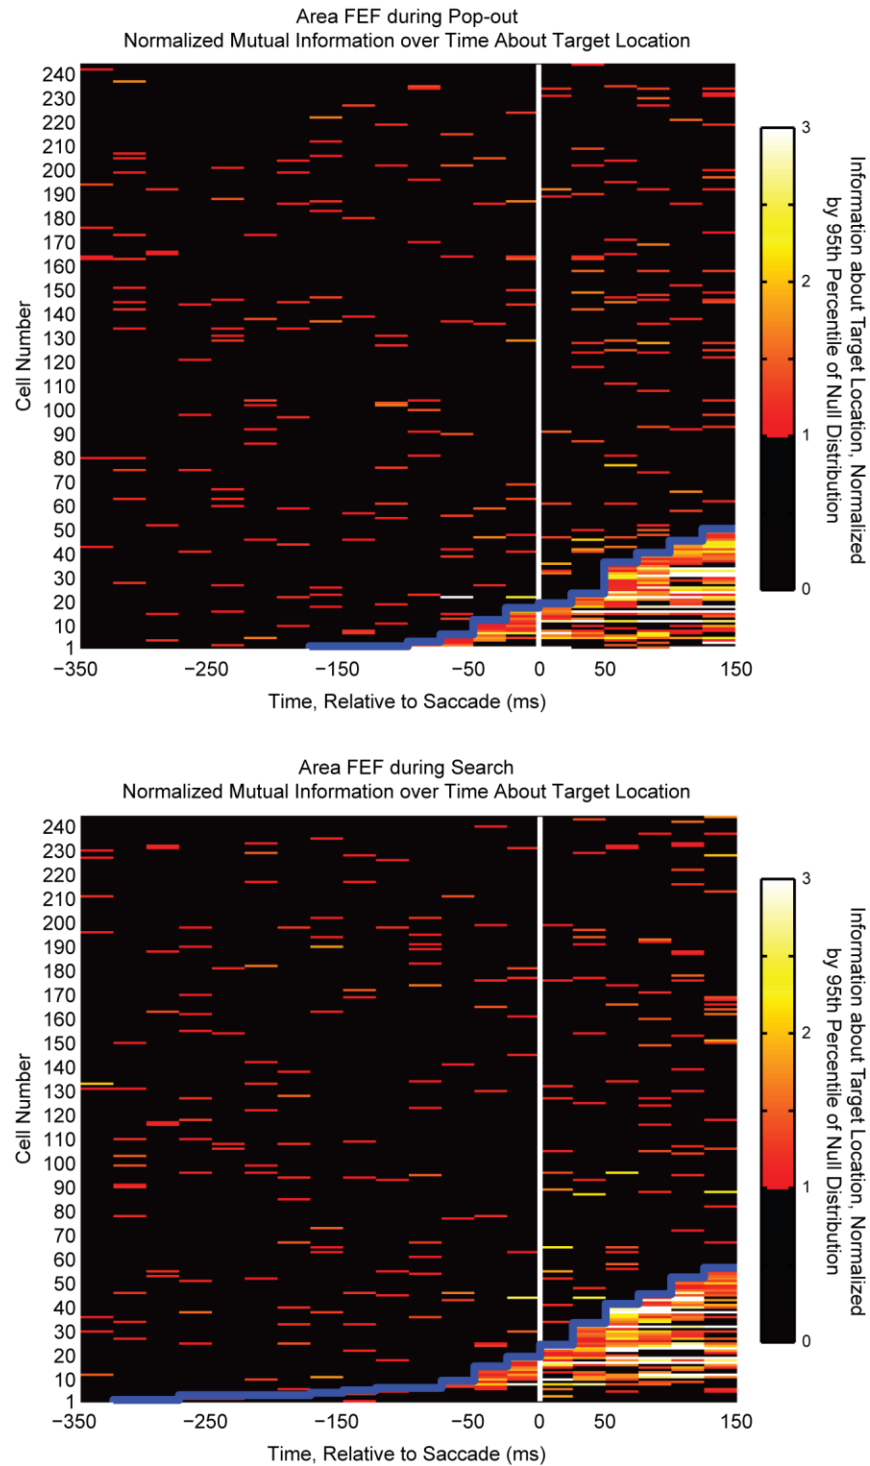


Figure 14. Normalized mutual information (MI) across time for all recorded neurons in **FEF** for both pop-out (top) and search (bottom). Neurons are sorted by their time to first significance, with the blue line following trend over time. MI was calculated in independent 25 ms time bins and neurons were only taken to carry significant information when above chance for two consecutive bins.

Correlation of Neural Activity with Behavior

All three regions show neural selectivity for the target location at some point during both the pop-out and search tasks. However, selectivity measures can be passive in nature – while they show neural activity is related to a variable of interest (target/attention location in this case), it is not necessary that the observed selectivity is actively used in the behavior. Although previous research has shown activity in these regions across a large spectrum of visual attention tasks (reviewed above), we can bolster these arguments by showing that neural activity is directly correlated with behavior. A number of prior studies have shown a relationship between neural activity in the frontal and parietal cortices and behavior during focal attention tasks (Moore and Fallah, 2001; Sato et al., 2001; Schall, 2002b; Schall, 2002a; Bisley and Goldberg, 2003; Moore and Fallah, 2004). Here, we also found that single neuron activity within each area correlated with how quickly the animal found the target.

In order to determine whether there was a direct relationship between firing activity of single neurons recorded in all three regions and the behavioral response of the animal we computed the correlation between firing rate and reaction time. This was done in a sliding window manner similar to neural selectivity for target location. Each recorded neuron that met the requirements to be included in the selectivity analysis was tested for significant correlation ($p < 0.05$) in 25 ms windows, stepped 25 ms over the trial. A total of 6 steps were used, limiting the analysis to the first 150 ms after visual array onset. This was done to avoid any potential contamination from signals directly moving the eyes – this is well before the reaction time distribution for both tasks (in fact, it is at least 50 ms faster than at least 90% of both distributions). Additionally, correlation between the level of neural activity and the reaction time was only done on trials where the target location was in the neuron's preferred direction. For this analysis preferred direction was taken to be the target location which showed the greatest firing rate over a 200 ms peri-saccadic window (starting 150 ms prior to saccade, ending 50 ms post-saccade). This was done in order to control for saccade metrics and any location-reaction time correlations (which were strong, as discussed below, and will become important for later analyses).

During the pop-out task all three areas had a significant number of neurons with significant correlation between neural activity and reaction time in at least one of the 6 windows. LIP had 23 significantly correlated neurons (9%, $p = 0.01$), dIPFC had 35 (14%, $p = 7 \cdot 10^{-8}$), and FEF had 31 (14%, $p =$

7×10^{-7}). However, during the search condition, while both dlPFC and FEF still had a significant number of selective neurons (25, $p = 2.4 \times 10^{-3}$, and 37, $p = 3.3 \times 10^{-10}$, respectively) there was not a significant number of selective neurons in LIP (only 16, $p > 0.5$). This parallels our findings on when these areas found the target location: while all 3 regions showed selectivity for the target before the saccade during pop-out, only dlPFC and FEF carried target information before the saccade in the search task. Furthermore, by demonstrating a correlation between single unit activity and the speed of the behavioral response this confirms that our regions of interest were directly involved in the task and supports the results on their roles in attention.

Conclusion

The source of top-down signals has largely been inferred from indirect evidence such as patterns of anatomical connections (Miller and D'Esposito, 2005). In the case of visual attention, previous research has shown its neural correlates throughout the cortex, with control attributed to parieto-frontal networks (Desimone and Duncan, 1995; Kastner et al., 1999; Reynolds and Chelazzi, 2004). Our results suggest that within this network, fast, bottom-up target selection occurs first in LIP, whereas longer latency top-down selection occurs first in the frontal cortex. This supports the hypothesis that parietal neurons form “saliency maps” for bottom-up selection (Itti and Koch, 2001; Bisley and Goldberg, 2003; Constantinidis and Steinmetz, 2005) as well as studies showing that stimulation of the frontal cortex causes attention-like effects in the extrastriate cortex (Moore and Fallah, 2001). It also fits with attenuation of top-down effects in the posterior cortex after PFC damage (Eglin et al., 1991a; Chao and Knight, 1997; Tomita et al., 1999). Although both the frontal and parietal cortex are involved in attention, our results illustrate that bottom-up signals appear first in LIP and top-down signals appear first in the frontal cortex.

The issue of the control of directing attention is a subset of the greater problem of cognitive control. Our results suggest that the top-down influences crucial for many complex behaviors likely originate from the prefrontal cortex and are imposed on posterior cortex. There are a plethora of results demonstrating that neurons in lateral prefrontal cortex carry information important to maintaining the rules and structure of the current task, as well as the most abstract information necessary for decision making associated with the behavior. As suggested by Miller and Cohen (2001)

the lateral PFC is in a good position anatomically to provide top-down direction necessary for behavior. However, our results are the first direct neural evidence that these top-down biasing signals do appear to originate within the prefrontal cortex. We believe that these results provide an important piece of information supporting current models of cognitive control.

Introduction

As noted in the Introduction, synchrony of neural activity may be useful for dynamically shifting neural networks by increasing the effectiveness of connections between brain areas (Aertsen et al., 1989; Engel et al., 2001; Salinas and Sejnowski, 2001; Fries, 2005; Womelsdorf and Fries, 2007). This might aid in direction of top-down signals as well to enhance the representation of attended stimuli (Aertsen et al., 1989; Engel et al., 2001; Salinas and Sejnowski, 2001; Fries, 2005; Womelsdorf and Fries, 2007). In order to investigate the role of synchrony in the internal and external control of attention, we compared the inter- and intra-areal synchrony between local field potentials (LFPs) as well as the synchrony between frequency bands within the LFP and the degree of synchronization between spiking activity and the LFP at different frequency bands. Some of this chapter has previously appeared in print (Buschman and Miller, 2007).

Synchrony between Parietal and Frontal Cortex

Based on the relative timing of selectivity across areas for visual pop-out and visual search, we have proposed that the prefrontal cortex (dlPFC and FEF) provide the top-down signal during endogenous attention, while activity flows ‘up’ from posterior parietal cortex into prefrontal cortex during exogenous attention. Therefore, the relationship between the prefrontal cortex and parietal cortex changes between endogenous visual search and exogenous visual pop-out.

To investigate the relationship between synchrony of the parietal and prefrontal cortices and this shifting relationship, we quantified the degree of synchrony between local field potentials (LFPs) in the parietal and frontal cortices during both visual search and visual pop-out. The local field potential is a more spatially generalized neural signal, thought to be related to the average input or activity of a region around the tip of the electrode (Legatt et al., 1980). Therefore, by measuring synchrony between these signals we are able to get a more general view of coupling amongst these regions and how this coherence might vary with the task at hand (and therefore with attentional control). Synchrony was measured between all pairs of simultaneously recorded LIP and frontal electrodes. Only electrodes that had at least one neuron selective for target location were used (resulting in 282 LIP-frontal pairs). The

Figure 15

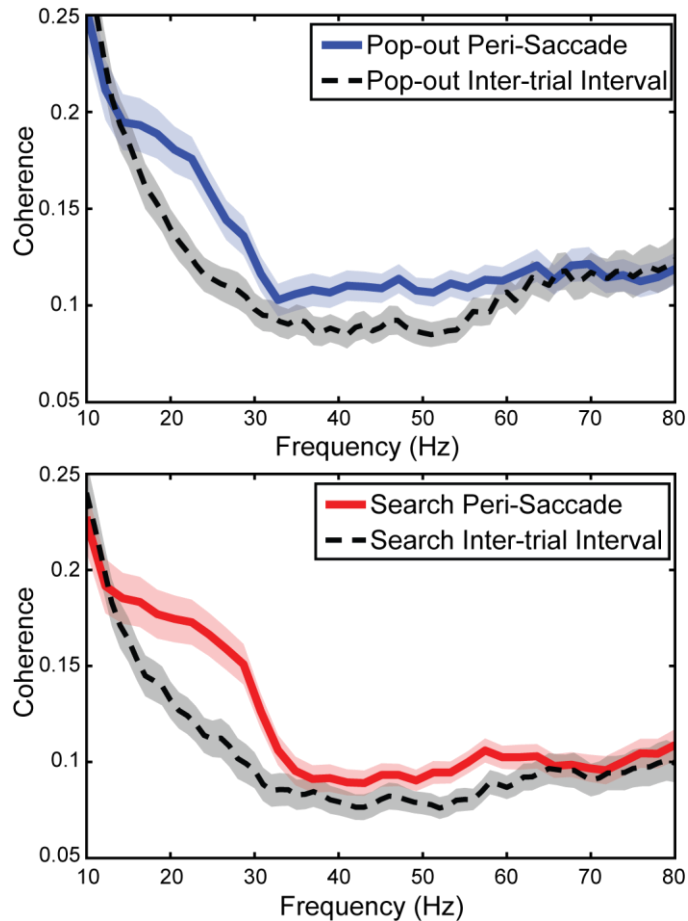


Figure 15. Local field potential (LFP) coherence between LIP and frontal cortex (LPFC and FEF) across frequencies for both Pop-out (top) and Search (bottom) tasks. Coherence was calculated around the time of the attention shift (in a peri-saccadic period, beginning 150 ms before saccade to 50 ms afterwards) and compared to a baseline, an inter-trial interval (ITI) epoch (a 200 ms window starting 500 ms before trial onset). Shaded regions are 95% confidence intervals around average coherence. Frequencies below 10 Hz are not meaningful (thus, not shown) because of the relatively short time epochs used.

degree of synchrony was captured in the coherence statistic, a measure of the co-spectrum between two signals, normalized for the power (Jarvis and Mitra, 2001). Significance was determined by randomization tests (see Methods for more information). Because similar results were found between LIP and dIPFC and between LIP and FEF, we combined data from the dIPFC and FEF.

In order to capture the process of attending to the target, coherence was estimated over a 200 ms peri-saccadic window starting 150 ms before the saccade and continuing until 50 ms after the saccade. The baseline period was taken from the inter-trial interval (a 200 ms window starting 500 ms before fixation spot display). The inter-trial interval was chosen to be the baseline specifically because it is a completely uncontrolled period of time, minimizing the amount of structure in the signal. Coherence during the fixation period was similar to that of the ITI and could also have been used as a baseline. In order to determine the evolution of coherence over the trial, we also estimated the coherence between LIP and frontal cortex during the “sample” period (taken to be a 400 ms window beginning 200 ms after sample onset) and the memory delay (a 450 ms window beginning 50 ms after sample offset).

During both search and pop-out, there was an increase in coherence between LIP and frontal cortex in a middle (22 – 34 Hz) and upper (35 -55 Hz) frequency band (Figure 15) that peaked during the peri-saccadic period, i.e. around the time of the attention shift (Figure 16A). However, the relative increase in coherence for each frequency band differed between bottom-up and top-down control of attention. This difference in coherence between task conditions can be highlighted by subtracting the coherence during pop-out from search. Figure 16B shows the z-score of this difference. We found a greater increase in middle frequency coherence (22 -34 Hz) between LIP and frontal cortex during top-down search than during bottom-up pop-out. By contrast, the increase in upper frequency (35 – 55 Hz) coherence was greater during pop-out than search. Thus, bottom-up and top-down attention may rely on different frequency bands of coherence between the frontal and parietal cortex.

As noted in the Introduction, localized synchrony of activity within a brain area may help resolve competition for attentional selection and inter-areal synchrony may aid in long-range communication

Figure 16

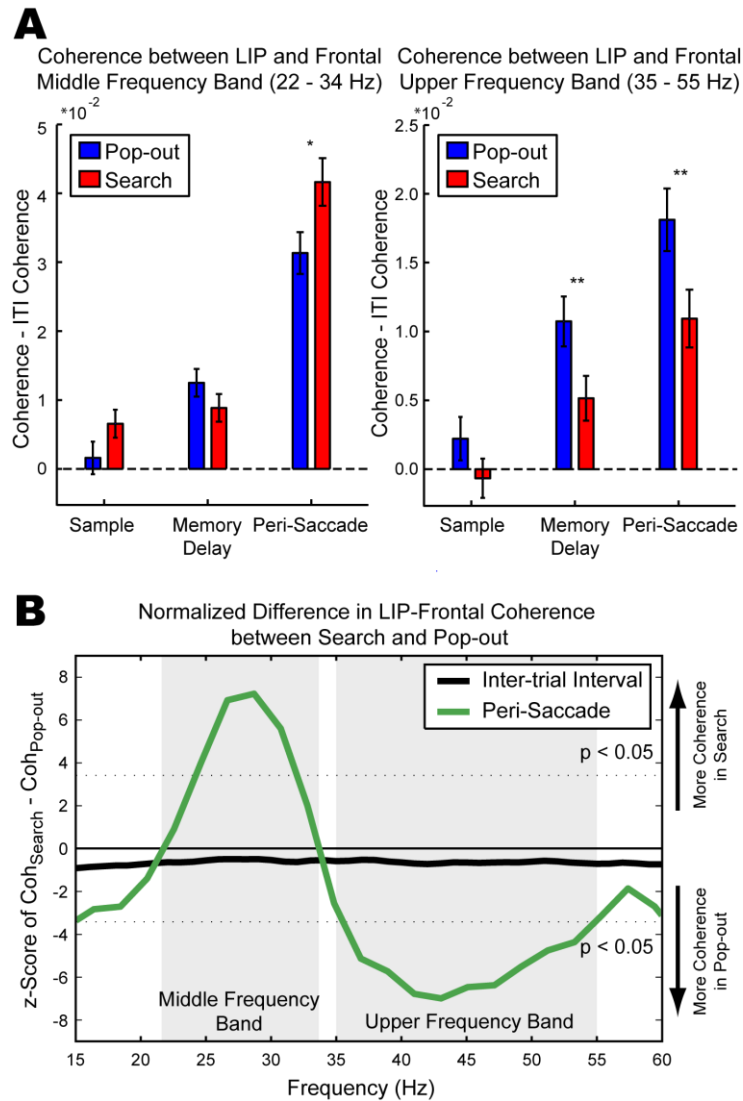


Figure 16. (A) Level of coherence for pop-out and search during the middle (left, 22 – 34 Hz) and upper (right, 35 – 55 Hz) frequency bands in different trial epochs. Significant differences of $p < 0.05$ are marked with *; $p < 0.01$ as **, as determined by t-test. **(B)** Differences in LFP coherence between LIP and frontal cortex during pop-out and search for the peri-saccadic period (green) and ITI (black). Pop-out coherence was subtracted from search coherence (Fig. 3). Dashed lines indicate significance levels ($p < 0.05$, corrected for multiple comparisons). Differences above the upper dashed line indicate significantly more coherence during search than pop-out, and below the lower dashed line significantly more coherence during pop-out than search.

between areas. Our results suggest that the flow of top-down and bottom-up information is aided by coherence emphasizing different frequency bands. Lower frequency bands are more robust to spike timing delays and thus may be better suited for longer-range coupling between multiple, distant areas (Kopell et al., 2000; von Stein and Sarnthein, 2000; Engel et al., 2001). The increase in low frequency synchrony during search could reflect a ‘broadcast’ of top-down signals on a larger anatomical scale. Synchrony at higher frequency bands might support the local interactions needed to enhance stimulus representations (Kopell et al., 2000; von Stein and Sarnthein, 2000; Engel et al., 2001). The emphasis of higher-frequency synchrony during pop-out could reflect local enhancement of stimulus representations that are passed forward from parietal to frontal cortex. This suggests that the brain may emphasize coherence at different frequency bands for the dynamic modulation of inter-areal connections that engages the network suited for the current task.

Correcting for a Common Ground

In order to ensure that the coherence results observed were not due to task dependent fluctuations in the ground potential, but rather due to true task-related differences in synchrony amongst these regions, we corrected for a common ground amongst the electrodes. Since during recording all electrodes were differentiated to a common ground, this induces a shared, synchronous signal in all electrodes. Therefore, if the ground signal fluctuated significantly from task to task then this might induce a spurious shift in the coherence statistic. To remove this potential confound, we re-referenced all of the electrodes within an area to the average potential for that area (minus the electrode under consideration). This subtracts any common signal across all electrodes, including the common-referenced ground potential. After this subtraction was made, we re-calculated the coherence signal between LIP and frontal cortex. Figure 17 shows the results. The pattern of coherence between LIP and frontal cortex between the search and pop-out tasks was virtually the same pattern as the original results. Coherence was greater for the middle frequency band during search over pop-out ($p = 3.5 \cdot 10^{-8}$, by paired t-test) and greater for pop-out over search in the upper frequency band ($p = 3.5 \cdot 10^{-8}$, by paired t-test).

Furthermore, as Figure 18 shows, the raw power of the local field potentials within the LIP and within the frontal cortex did not significantly vary between the search and pop-out task conditions. Any

Figure 17

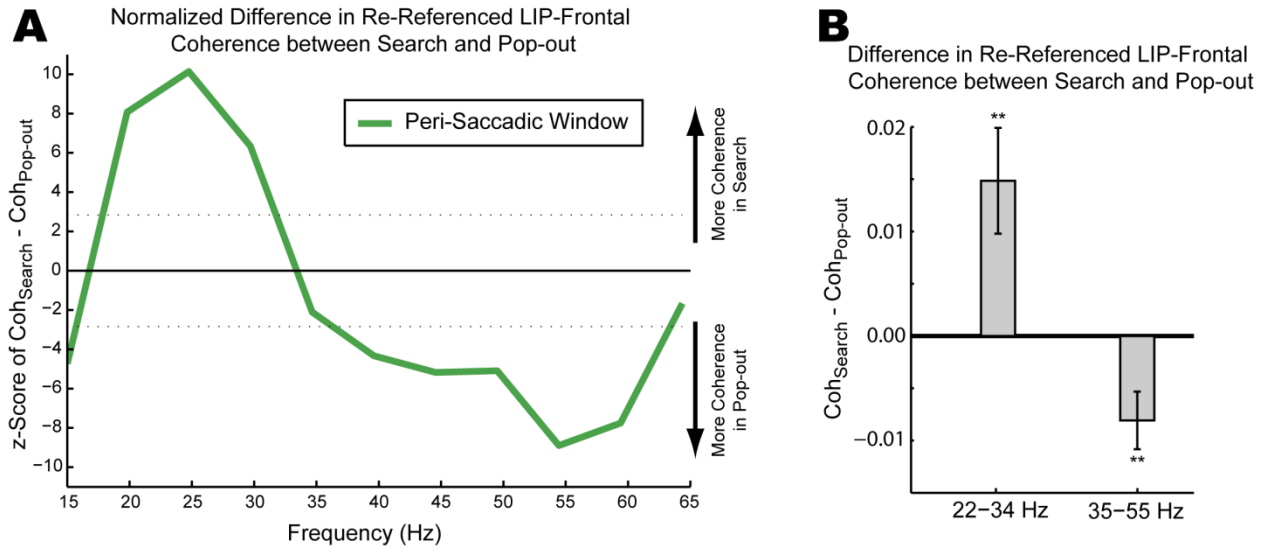


Figure 17. Coherence between LIP and both Frontal regions is different based on task condition for re-referenced field potentials. **(A)** Shows the same normalized difference in LIP-frontal coherence during search and pop-out as Fig 11B, but with the field potentials from each area re-referenced to the area's average potential. This was done in order to remove any possible role of a common ground in driving the coherence effects. **(B)** Raw difference in LIP-frontal coherence between search and pop-out. The coherence between each pair of LIP and Frontal electrodes was determined for both search and pop-out and the difference within each pair of electrodes is shown. Error bars indicate 95% confidence interval around the mean difference. During search LIP-frontal coherence was significantly greater than during pop-out ($p = 3.5 \cdot 10^{-8}$ by paired t-test), while the coherence in the upper frequency band was significantly greater during pop-out than search ($p = 3.5 \cdot 10^{-8}$, by paired t-test).

fluctuations in the ground potential possibly underlying the coherence differences would also be reflected as differences in the average power between task conditions. As no differences in power were observed, this indicates that increases in power per se did not underlie the changes in inter-area coherence. Taken together, these corrections demonstrate that the observed differences in coherence are not due to a common ground or due to fluctuations in the raw power of the field potentials but reflect true changes in coherence between areas.

Correcting for Reaction Time Differences between Tasks

The monkeys showed differing reaction times (RTs) during search versus pop-out and, because we calculated coherence over a fixed time interval, in principle, this might cause the differences in coherence that we observed. To confirm that this was not the case we performed two control analyses. First, we recalculated coherence between LIP and frontal areas over a variable window starting 75 ms after visual array onset and ending 50 ms after saccade (the window thus varied with the animal's reaction time). This time period was chosen to avoid any initial visual response and to capture as much of the process of attending and selecting the target as possible. As shown in Figure 19, we obtained the same results as with a fixed time interval. In the middle frequency band search was significantly greater than pop-out ($p = 5.9 \cdot 10^{-3}$, by paired t-test), while in the upper frequency band pop-out was significantly greater than search ($p = 9.0 \cdot 10^{-6}$, by paired t-test).

As a further test to ensure that RT differences between task conditions did not induce our coherence differences, we used a stratification procedure to match the RTs trials between task conditions. This approach has been used in similar situations to help normalize reaction time effects between tasks (Tallon-Baudry et al., 1998). Stratification was done by pairing trials from each task condition together if they had reaction times close to one another (within 7 ms). If no trial from the other task was within this buffer then that trial was discarded. The result of this procedure was RT distributions between task conditions that were 93% overlapping (compared to 60% overlapping for the original data) and no longer were significantly different ($p > 0.05$ for corrected compared to $p = 5 \cdot 10^{-4}$ uncorrected; by mutual information, see Figure 20). While this procedure reduced the total number of trials, the pattern of coherence results was the same (Figure 21). Similar to the original data, coherence

Figure 18

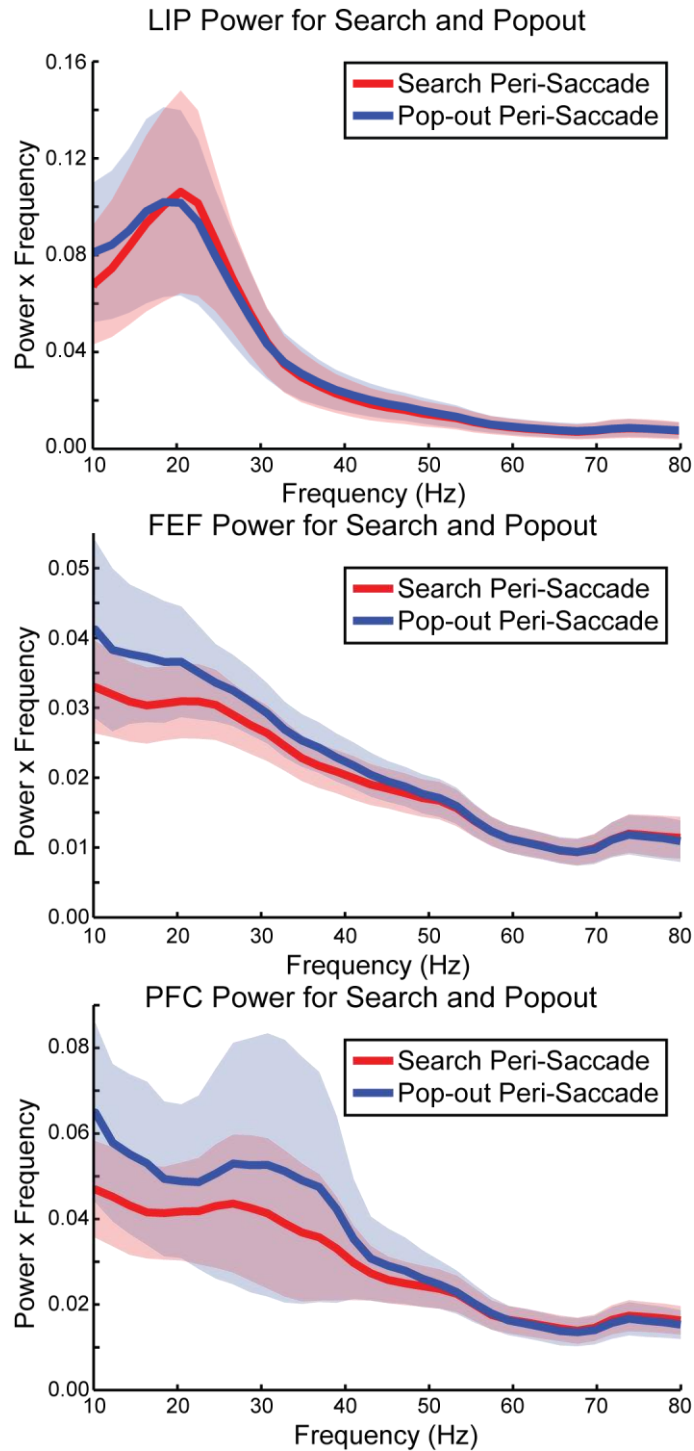


Figure 18. Spectral power over frequency during search and pop-out for LIP (top), FEF (middle) and dIPFC (bottom). The power spectrum was multiplied by frequency in order to make the data easier to visualize. Shaded regions indicate 95% confidence interval about the mean power for each area. As all three regions show overlapping power spectrum between the two task conditions, there is no significant difference in power between the two conditions.

Figure 19

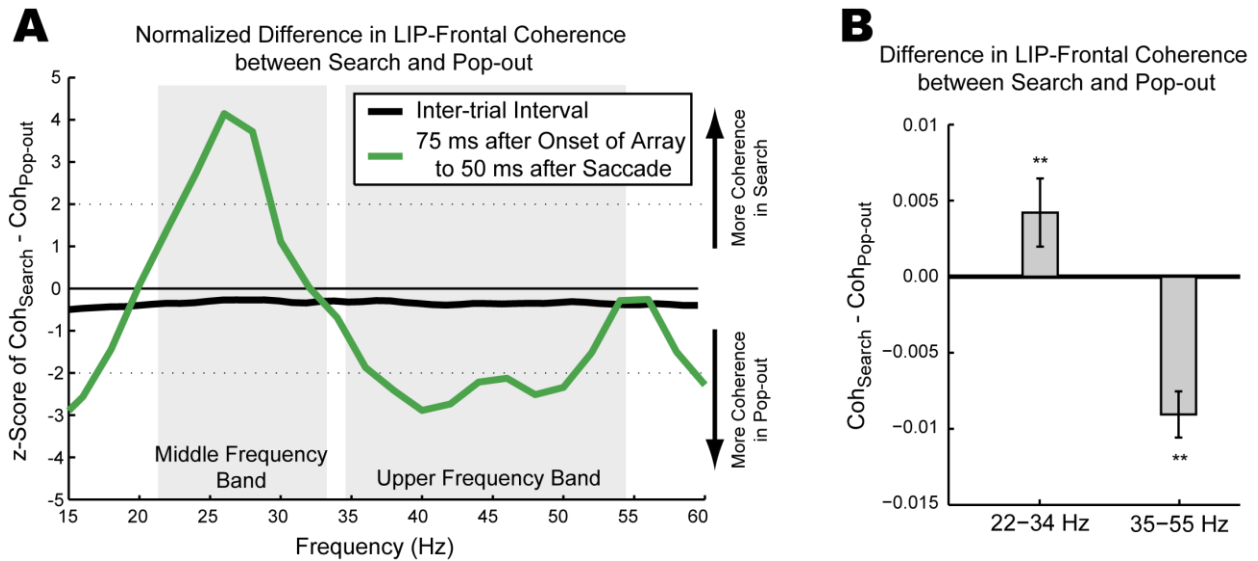


Figure 19. Coherence between LIP and both Frontal regions is different based on task condition in variable time window. **(A)** Shows the same normalized difference in LIP-frontal coherence during search and pop-out as Fig 14B, but over a variable time window beginning 75 ms after visual array onset and ending 50 ms after saccade. **(B)** Raw difference in LIP-frontal coherence between search and pop-out for same variable window. The coherence between each pair of LIP and Frontal electrodes was determined for both search and pop-out and the difference within each pair of electrodes is shown. Error bars indicate 95% confidence interval around the mean difference. During search LIP-frontal coherence was significantly greater than during pop-out ($p = 5.9 \cdot 10^{-3}$ by paired t-test), while the coherence in the upper frequency band was significantly greater during pop-out than search ($p = 9.0 \cdot 10^{-6}$, by paired t-test).

Figure 20

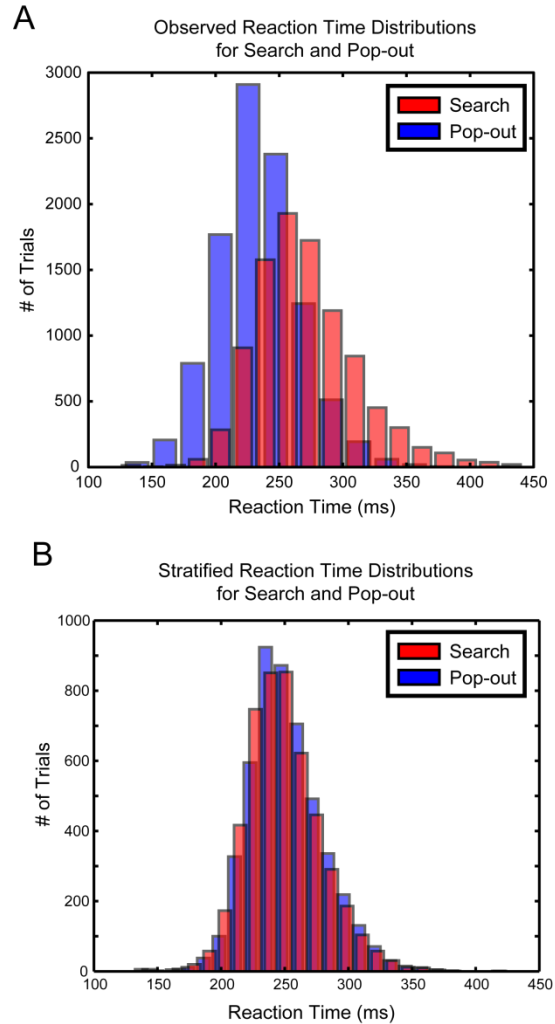


Figure 20. Reaction time distributions before and after stratification procedure. In order to remove any reaction time effects on the observed coherence, the reaction time distributions were stratified. Trials were paired across tasks (to within 7 ms of one another) in order to equalize the reaction time distributions. The result of this procedure was RT distributions between task conditions that were 93% overlapping (compared to 60% overlapping for the original data) and no longer were significantly different ($p > 0.05$ for corrected compared to $p = 5 \cdot 10^{-4}$ uncorrected; by mutual information).

Figure 21

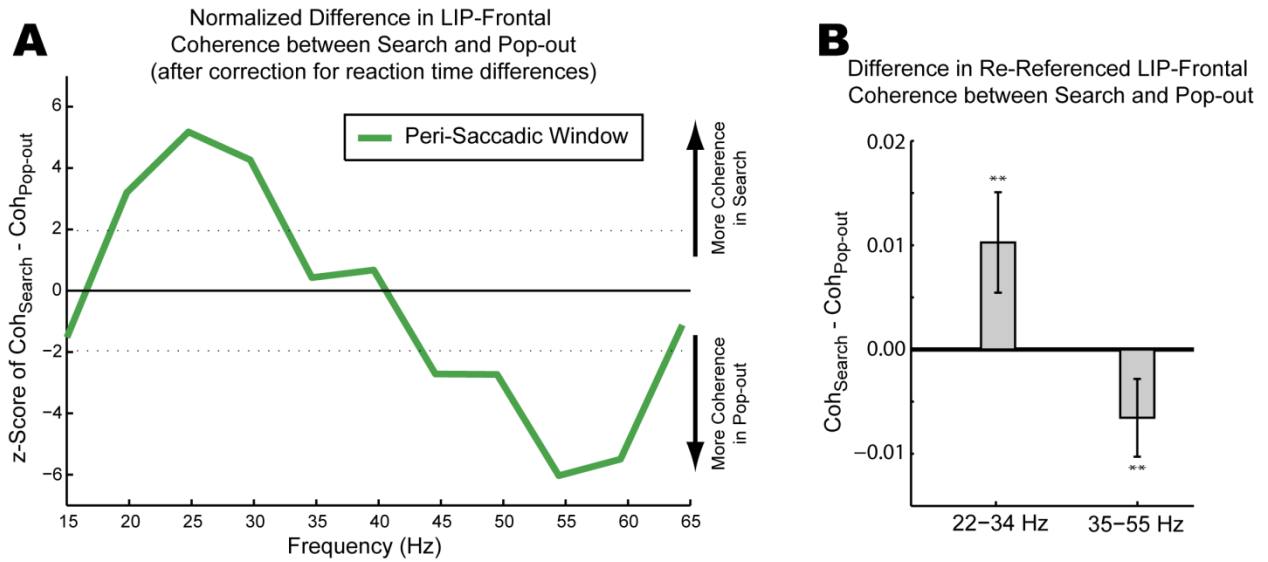


Figure 21. Coherence between LIP and both Frontal regions is different based on task condition after task stratification. **(A)** Shows the same normalized difference in LIP-frontal coherence during search and pop-out as Fig 14B, but with trials stratified by reaction time. The stratification process corrects for reaction time differences by creating overlapping distributions for search and pop-out **(B)** Raw difference in LIP-frontal coherence between search and pop-out. The coherence between each pair of LIP and Frontal electrodes was determined for both search and pop-out and the difference within each pair of electrodes is shown. Error bars indicate 95% confidence interval around the mean difference. During search LIP-frontal coherence was significantly greater than during pop-out ($p = 4.5 \cdot 10^{-5}$ by paired t-test), while the coherence in the upper frequency band was significantly greater during pop-out than search ($p = 7.1 \cdot 10^{-4}$, by paired t-test).

between LIP and frontal cortex was significantly greater in the middle frequency band during search ($p = 4.5 \cdot 10^{-5}$, by paired t-test) and significantly greater in the upper frequency band during pop-out ($p = 7.1 \cdot 10^{-4}$, by paired t-test). These results confirm the differences in coherence between search and popout while controlling for differences in reaction time.

Inter- and Intra-Areal Coherence

The previous sections have analyzed the synchrony between parietal and frontal cortex during both visual search and pop-out. This focus has been due to the timing differences observed between tasks, motivating an analysis of how this change in relative relationship might be reflected in changes in the degree of synchrony. Here we present a more complete view of the degree of coherence between and within all of the studied anatomical regions.

Coherence was calculated in the same manner as above (with a detailed description available in Methods). A randomization procedure was used in order to determine a null distribution. Confidence intervals about the mean coherence were calculated from the standard error of the mean. Figure 22 shows the raw coherence between all possible pairings of anatomical regions (including both intra- and inter-areal coherence). The coherence between LIP and both dIPFC and FEF alone show very similar patterns to one another, which prompted us to combine the results, as previously described. Figure 23 shows the difference in coherence during search and pop-out for each pair of areas. As was done previously, the null distribution was used to determine significance levels for both an increased coherence during search (in which the difference is positive) or increased coherence during pop-out (for which the difference is negative). Again, the task-difference in coherence between LIP and both dIPFC and FEF show very similar results.

While synchrony between the frontal and parietal regions showed differences in two distinct bands, coherence between and within other regions was not always as clear-cut. Synchrony within dIPFC electrodes showed a very similar effect as LIP-dIPFC (Figure 23) – coherence was found to be increased during visual search in a middle band of approximately 20 to 35 Hz and increased during visual pop-out for an upper frequency band of 35-60 Hz. Coherence within FEF is much noisier due to a lower

Figure 22

Inter- and Intra-Areal LFP-LFP Coherence during Peri-Saccade Time Epoch

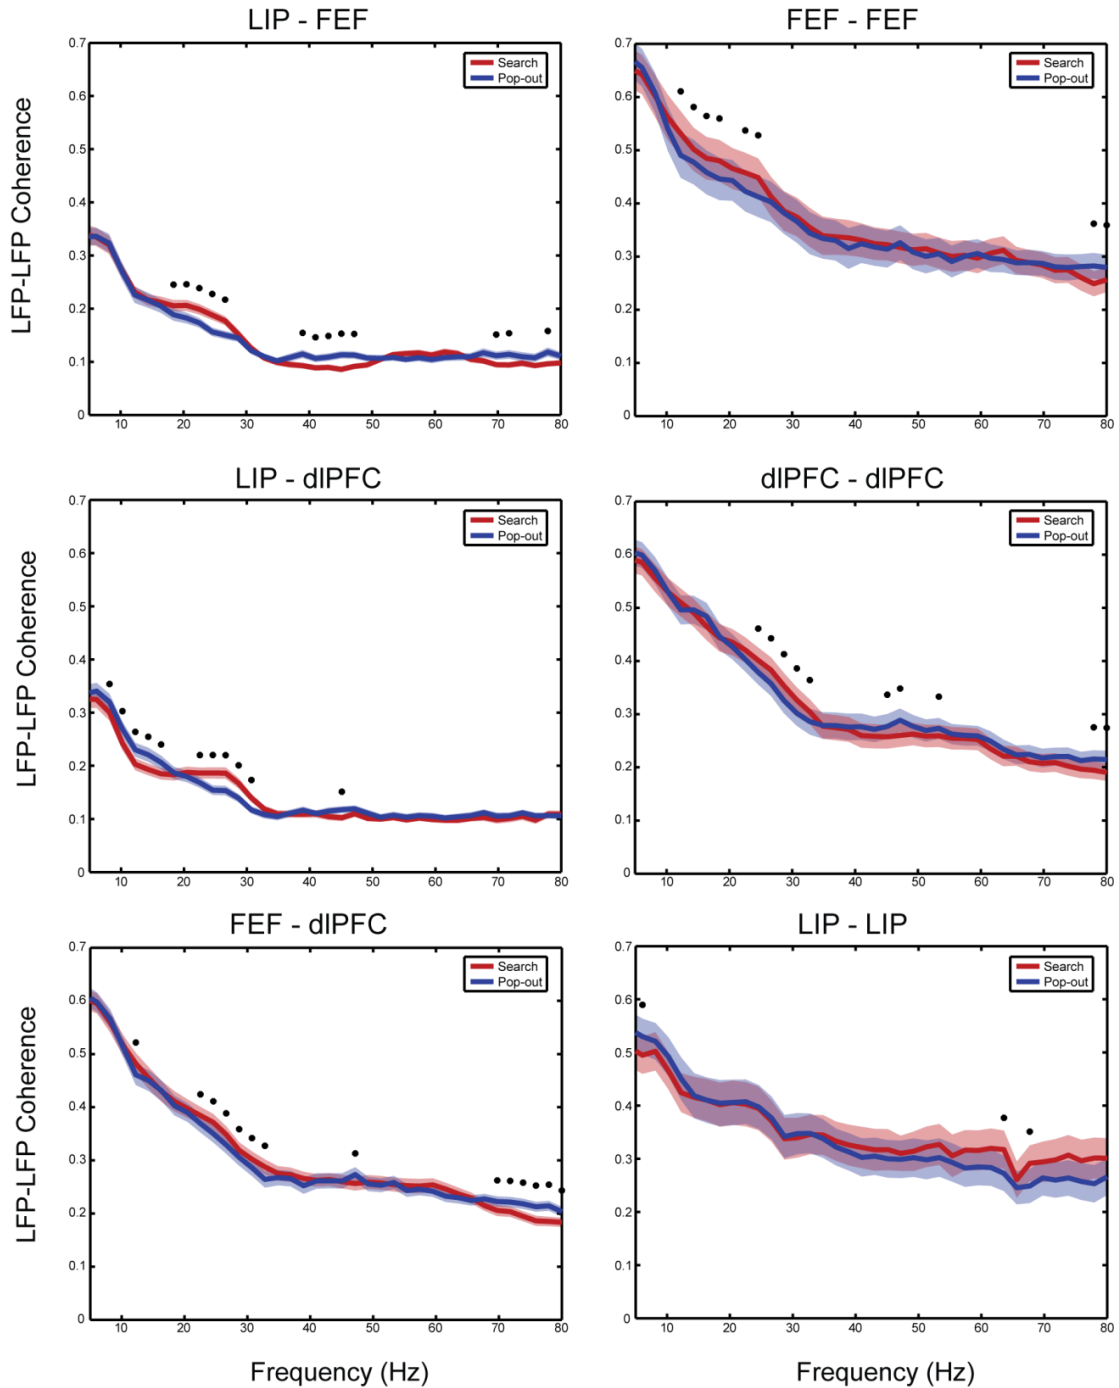


Figure 22. The raw coherence measured between areas (left column) and within areas (right column). The coherence is shown for both search (red) and pop-out (blue) with the shaded region indicating the standard deviation. Asterisks indicate a significant difference between search and pop-out ($p < 0.05$, t-test). The randomization procedure used in Figure 23 is a more valid test, but the t-test does provide similar results.

Figure 23

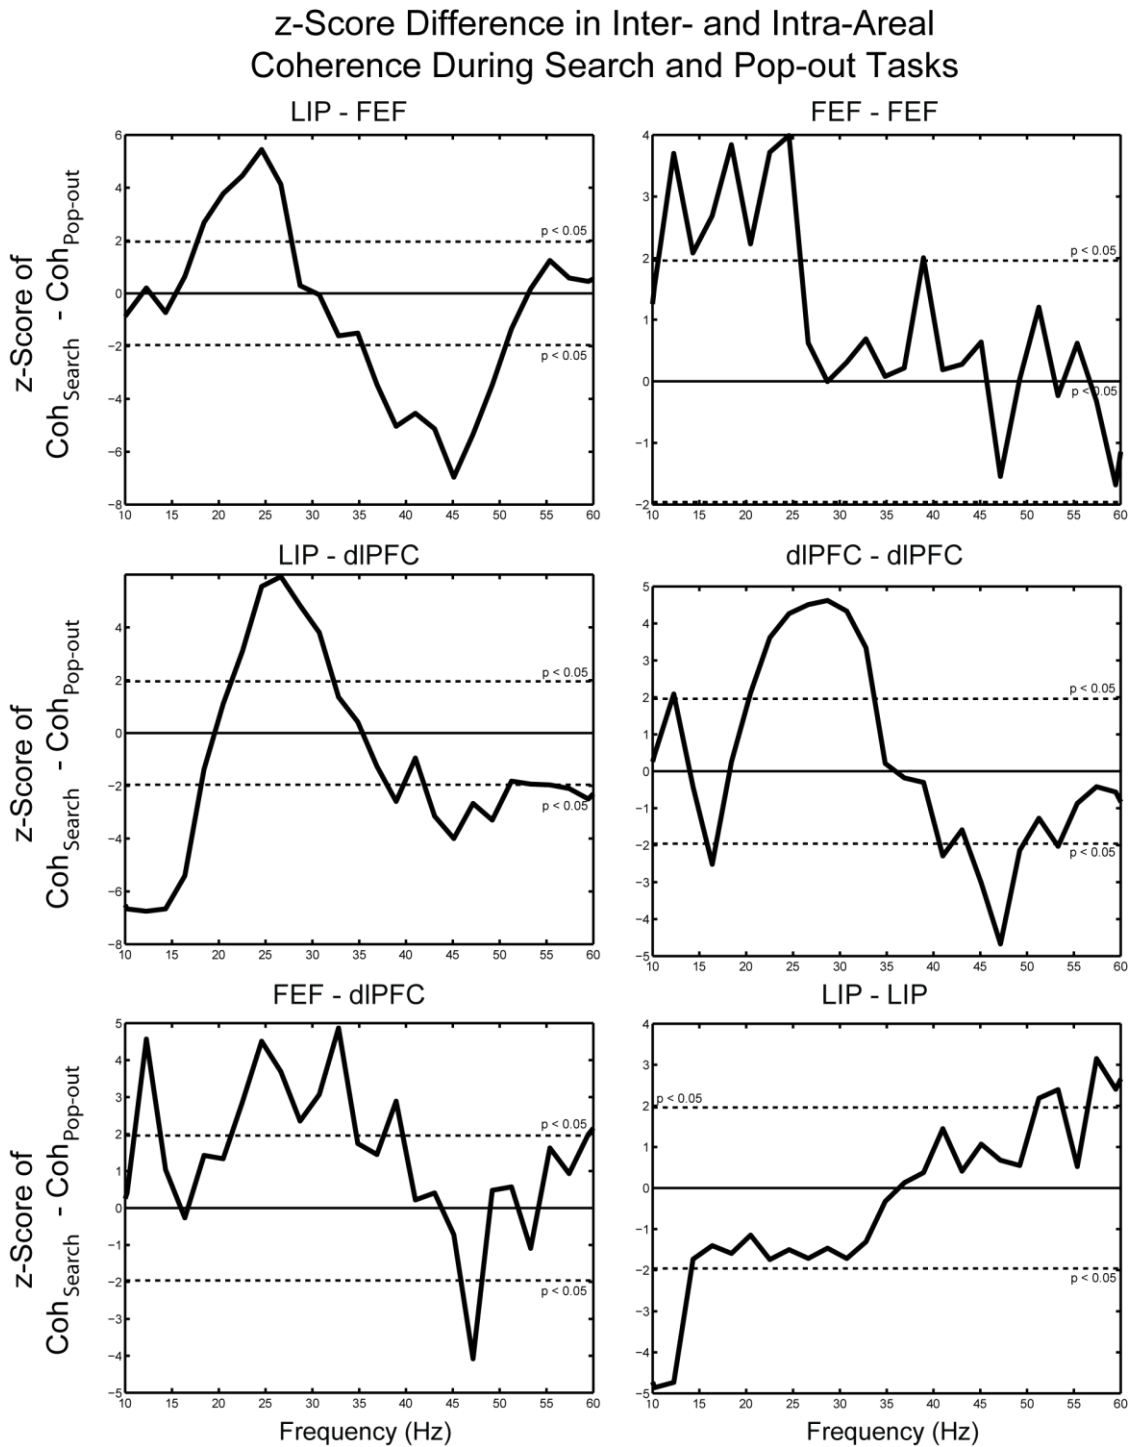


Figure 23. The difference in coherence during visual search and pop-out. Coherence is measured between areas (left column) and within areas (right column). Significance was determined by a randomization test (see main text). Differences above the upper dotted line indicate significantly more coherence during visual search; below the lower dotted line indicates more coherence during pop-out.

number of simultaneously recorded pairs, but shows a generally similar pattern. Coherence within FEF was relatively higher during visual search for a lower-middle band from approximately 10 to 25 Hz and trended towards an increase in coherence during visual pop-out for higher frequencies. Coherence between the two frontal regions (FEF-dIPFC) again shows an increase in middle frequency range coherence during visual search relative to visual pop-out, but shows a muddle difference in the upper frequency range. Later analysis will show a possible upward trend during visual search of FEF-dIPFC in a frequency band of 35-65 Hz (see LFP to LFP phase-locking section).

One possible concern is that these effects within and between frontal regions might reflect intrinsic changes within frontal cortex or they may also reflect the second-order synchronization of frontal electrodes in response to a first-order synchronization between frontal and parietal regions (although dIPFC electrodes show a pattern of coherence differences much more similar to LIP-Frontal than FEF electrodes do). However, LIP-LIP coherence did not show the same pattern of results, rather not showing a significant increase in either band between tasks. Although this argues against the strength of second-order effects, it is difficult to tease apart the sources of these changes in synchrony. It is important to note that the synchrony observed between LIP and the frontal regions cannot be result of a second-order for effect.

In order to provide a more complete picture of the distribution of coherence, Figure 24 shows the raw coherence between pairs of LIP and Frontal electrodes, as well as the raw distribution with standard deviation instead of 95% C.I. as shown in Figure 15. Although the effects across the population were highly significant, there is a good deal of variation in the degree of coherence between pairs of electrodes. Part of this variance is due to differences in synchrony with respect to the receptive fields of electrodes, as discussed in the following section. Figure 24 shows the standard deviation in coherence across individual trials for both visual search and pop-out (along with the 95% C.I. of that standard deviation). The technique for estimating single trial coherence is discussed in Chapter 8. With respect to the average coherence effects (> 0.1), the standard deviation across trials was relatively small. This suggests that the observed synchrony between pairs of electrodes in parietal and frontal cortex is fairly consistent from trial to trial. Later sections will attempt to explain some of this variance across trials with behavioral parameters.

Figure 24

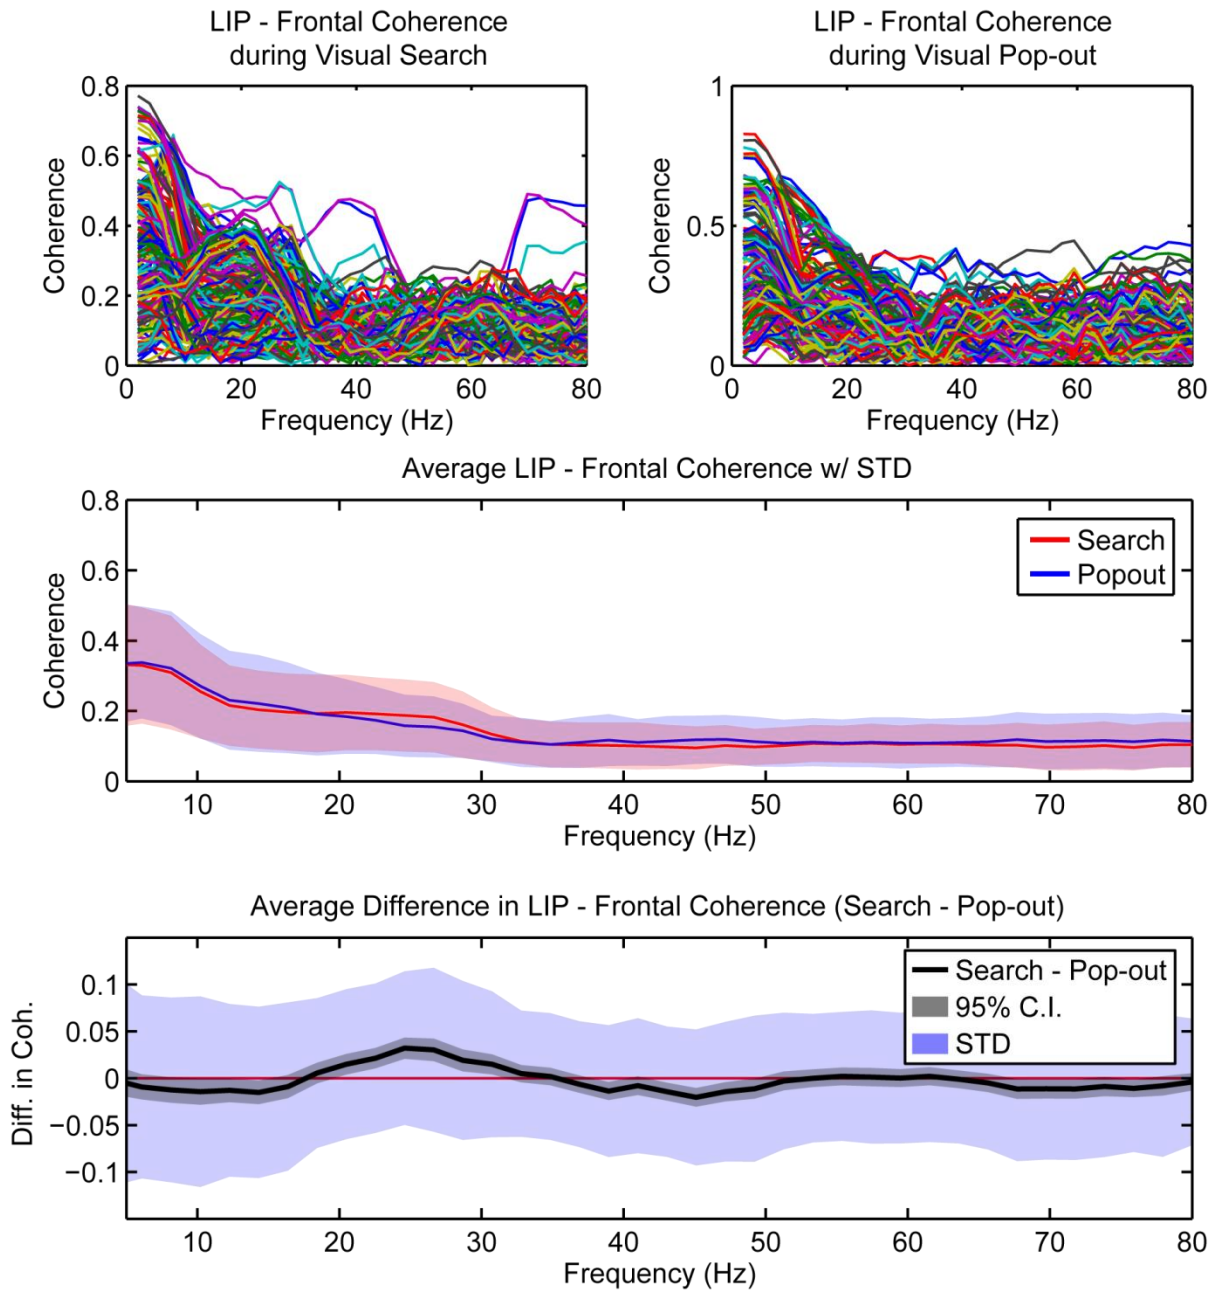


Figure 24. Raw coherence traces for pairs of LIP and Frontal electrodes during visual search (top row, left) and pop-out (top row, right). The average coherence is shown in the middle row, along with the standard deviation of the distribution. The difference in coherence between visual search and pop-out is shown in the bottom row, now showing both the corrected 95% CI and the standard deviation. All three plots show a large spread of values, but with highly significant differences due to the large number of pairs. Part of the variance in coherence across pairs is explained through receptive field effects as explained in the next section.

Figure 25

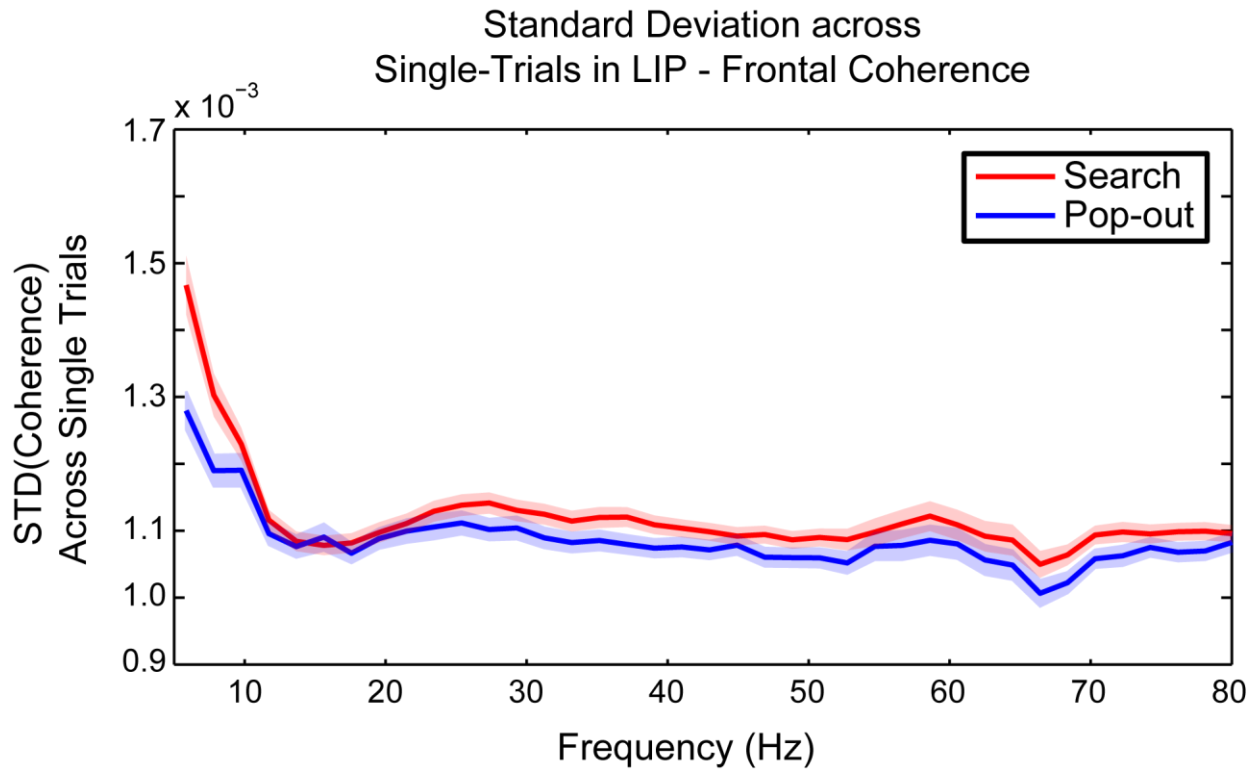


Figure 25. Standard deviation in coherence across single trials over frequency. The variance in coherence across trials was roughly the same for both visual search and pop-out and was fairly small relative to the overall coherence. This suggests that the observed effects were not strongly driven by a few outlying trials, but were fairly consistent.

Synchrony Varies with Location Preference

So far we have demonstrated large scale interactions between entire anatomical areas, but an important question is whether synchrony between regions happens on a finer scale. If synchrony between two brain regions does facilitate communication then it might be advantageous for the interactions to happen between subregions sharing similar properties. In our task the main variable of interest is the target location. Therefore, in order to test this hypothesis we investigated whether the observed coherence between two brain regions is modulated by the difference in location preferences for each electrode pair.

The location preference was determined for each individual electrode by defining a vector by the average response of all single neurons isolated from that electrode (the average response of isolated single neurons was taken over the same window as the coherence measurement – a 200 ms window starting 150 ms before the saccade). The difference in preferred direction for a pair of electrodes was then taken to be the difference in the angle of these two vectors. The first area listed in each pair acts as the reference vector, with the second area either rotated clockwise or counter-clockwise from that direction. Clockwise rotations were defined as negative angle values, while counter-clockwise rotations were positive. Angle differences were binned into 4 equally sized bins of 90° as only four target locations were used during recording. Coherence was calculated in the same manner as reported above. We examined both the middle frequency band (22-34 Hz) and the upper frequency band (35-55 Hz) found to be task related in the previous analysis. The average coherence across all pairs is reported for both search trials (shown in red) and pop-out trials (shown in blue). An unpaired t-test was used to determine whether the observed differences in coherence between tasks was significant.

Figure 26 shows the influence of location preference on coherence between all possible pairings of anatomical areas for the middle frequency band (22-34 Hz). Inter-areal coherence is shown in the left column, while intra-areal coherence is shown in the right column. Coherence between FEF and LIP was strongest during visual search for electrodes that were offset by -90° (this response was significantly greater than that observed in visual pop-out ($p = 0.015$, by t-test)). As the angle is negative, electrodes in FEF with a location preference 90° clockwise of those in LIP showed the greatest coherence. If the role of synchrony is to enhance communication, then this would suggest that sub-regions within FEF are

Figure 26

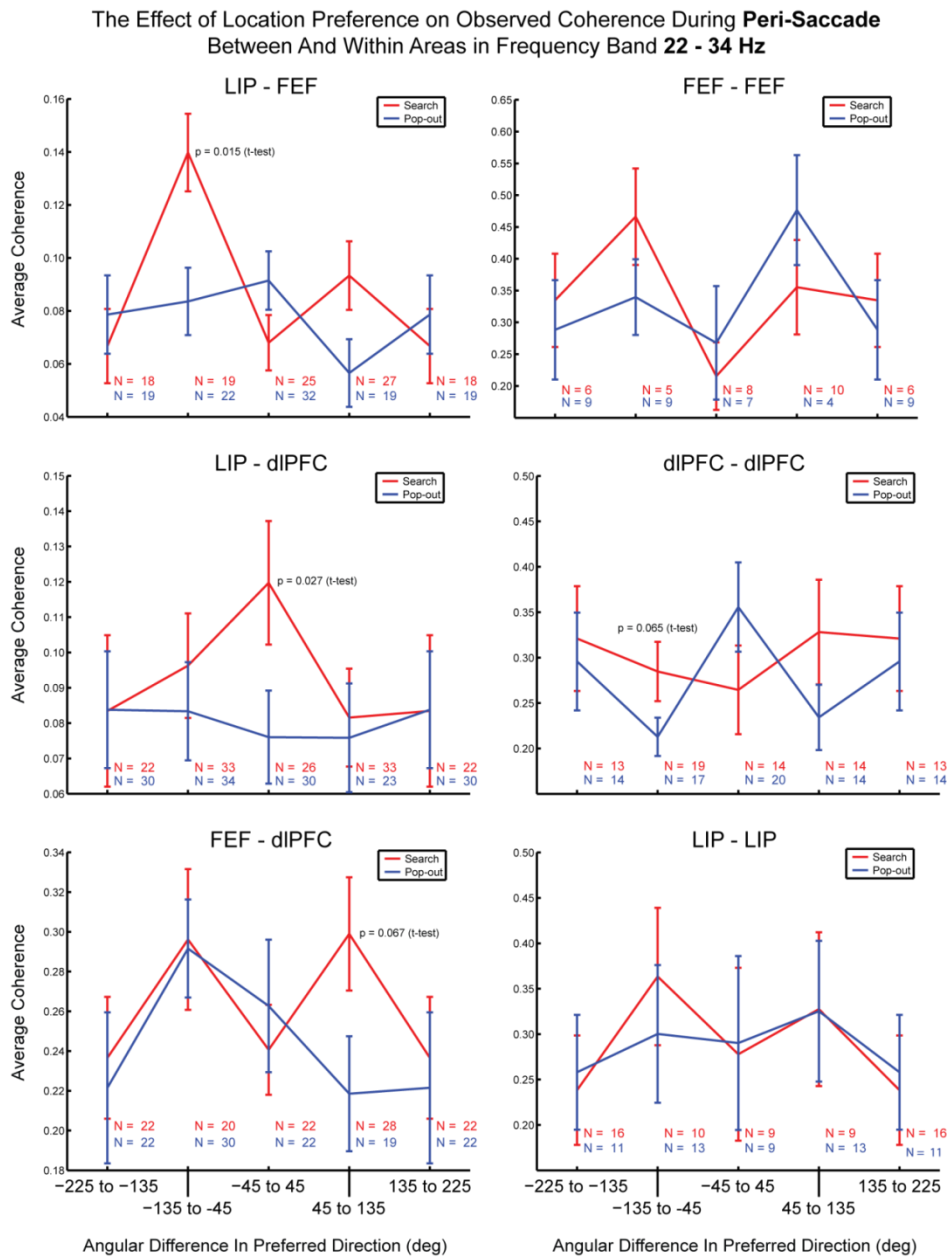


Figure 26. 22 – 34 Hz coherence between pairs of electrodes based on the difference in their preferred target locations during the **peri-saccade** time period. The coherence is compared between areas (left column) and within areas (right column) for both search (red) and pop-out (blue). The number of pairs of electrodes for each location preference offset is given for both search (red) and pop-out (blue). Errorbars indicate the 95% confidence interval about the mean. Significant differences between search and pop-out were determined using an unpaired t-test, with the p-value indicated.

receiving the most information from regions within LIP that are 90° counter-clockwise. Interestingly, this effect is not true for coherence between LIP and dlPFC in the 22-34 Hz band. Instead, coherence between these two regions was highest for electrodes within each area that shared the same location preference – as if the greatest communication occurred between sub-regions carrying common information.

While coherence between LIP and FEF was greatest in the 22-34 Hz frequency range when electrodes were offset by on target stimulus, coherence in the upper frequency range (35-55 Hz) was greatest for visual pop-out when the electrodes shared the same preferred direction (Figure 27). This dichotomy in the dependency of synchrony on location preferences between the middle and upper frequency bands further supports the differential role they play in visual search and pop-out. Coherence between LIP and dlPFC in the 35-55 Hz range was also maximal when the electrodes shared a preferred direction, although there was no significant difference between search and pop-out. However, a significant difference in 35-55 Hz coherence was observed for pairs of electrodes within dlPFC that share a preferred location. No such intra-areal increase in coherence was observed for the middle frequency band, bolstering the suggestion that higher frequency synchronization may play a larger role in the localized facilitation of communication. Although similar trends are observed for coherence during the inter-trial interval, there was no significant difference between tasks at any frequency band or location preference offset (Figure 28 for the middle, 22-34 Hz, frequency band; Figure 29 for the upper, 35-55 Hz, band). This suggests that the brain may begin to establish these synchronous relationships early in the trial in anticipation of performing the search task, however the effective connectivity between brain regions is maximally different during the execution of the task.

Figure 27

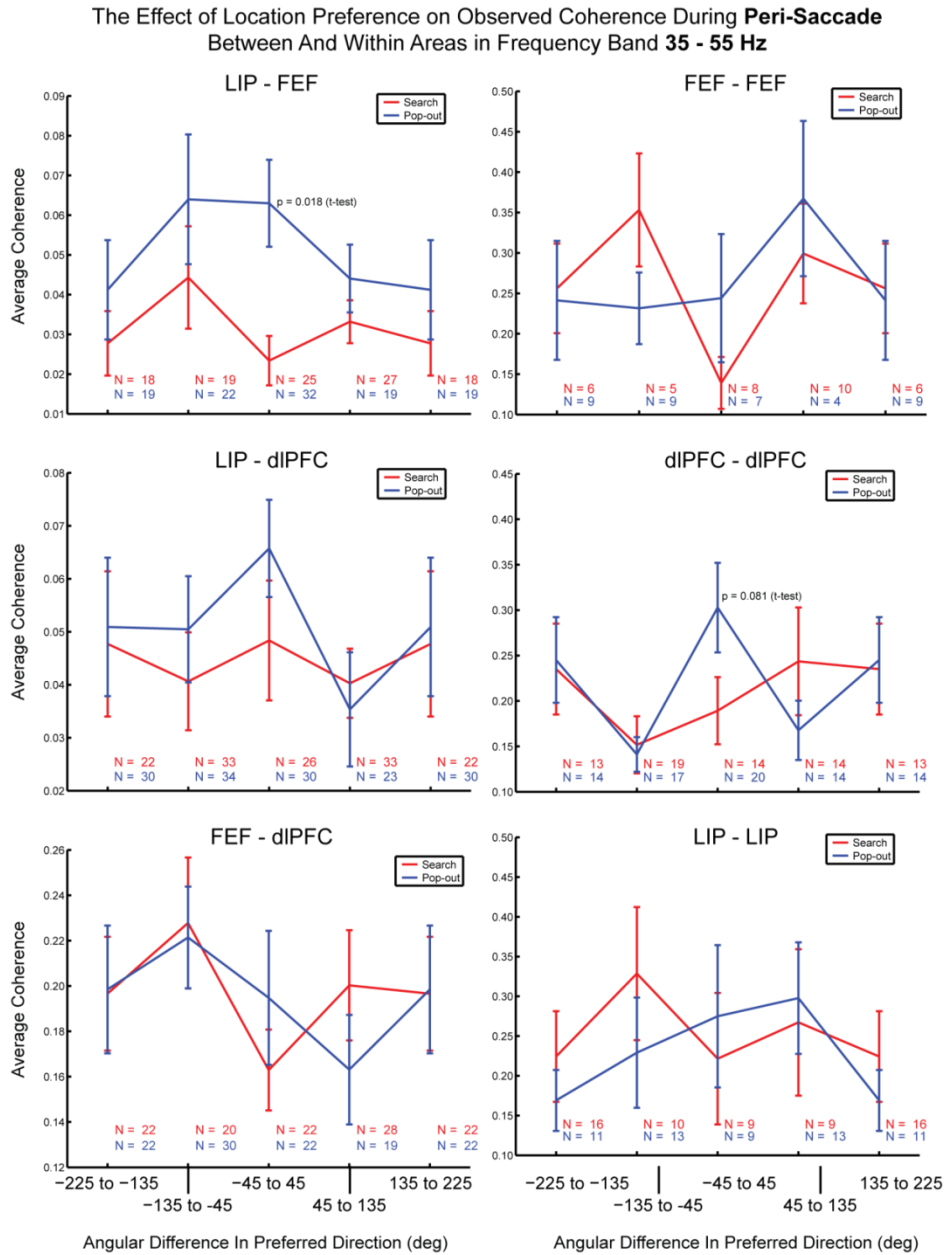


Figure 27. 35 – 55 Hz coherence between pairs of electrodes based on the difference in their preferred target locations during the **peri-saccade** time period. The coherence is compared between areas (left column) and within areas (right column) for both search (red) and pop-out (blue). The number of pairs of electrodes for each location preference offset is given for both search (red) and pop-out (blue). Errorbars indicate the 95% confidence interval about the mean. Significant differences between search and pop-out were determined using an unpaired t-test, with the p-value indicated.

Figure 28

The Effect of Location Preference on Observed Coherence During **Inter-trial Interval** Between And Within Areas in Frequency Band **22 - 34 Hz**

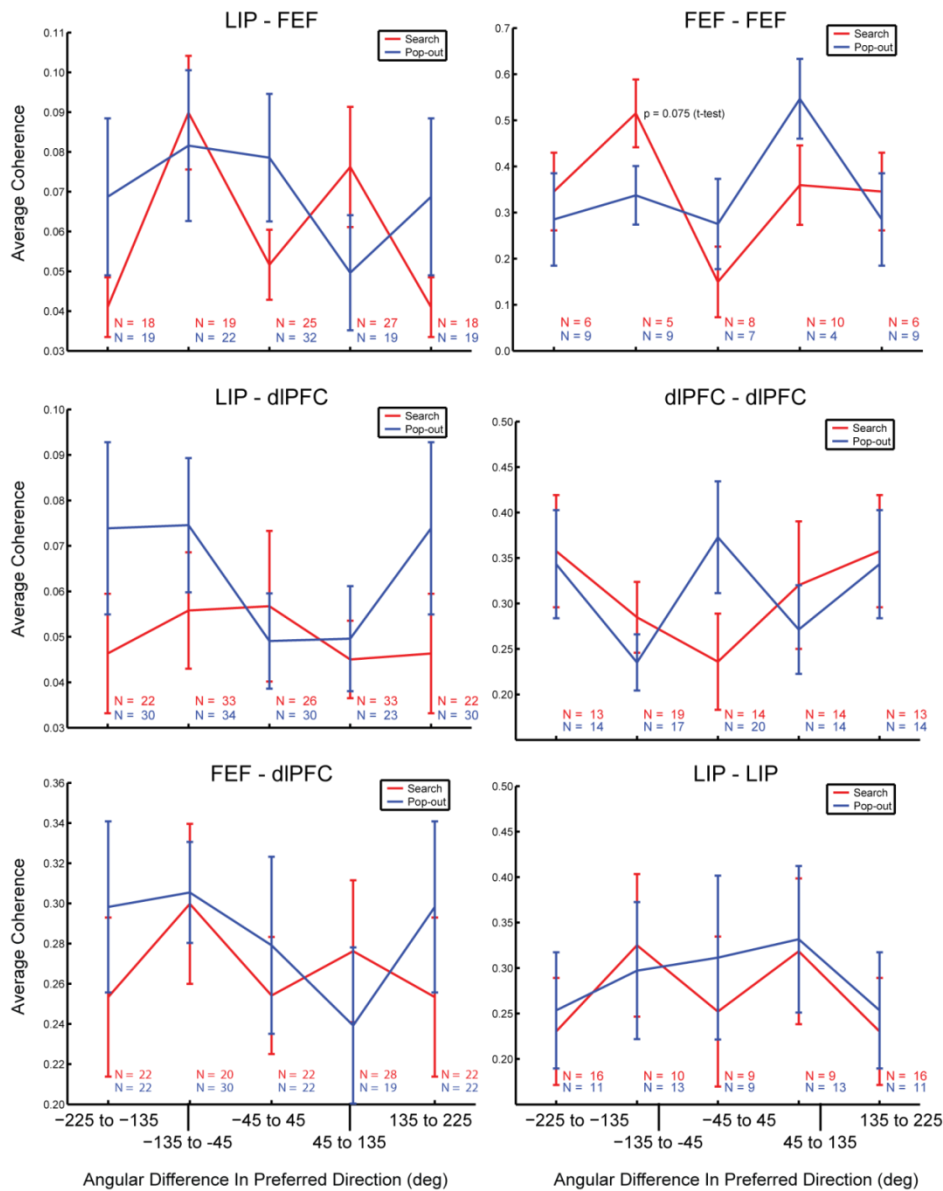


Figure 28. 22 – 34 Hz coherence between pairs of electrodes based on the difference in their preferred target locations during the **inter-trial interval**. The coherence is compared between areas (left column) and within areas (right column) for both search (red) and pop-out (blue). The number of pairs of electrodes for each location preference offset is given for both search (red) and pop-out (blue). Errorbars indicate the 95% confidence interval about the mean. Significant differences between search and pop-out were determined using an unpaired t-test, with the p-value indicated.

Figure 29

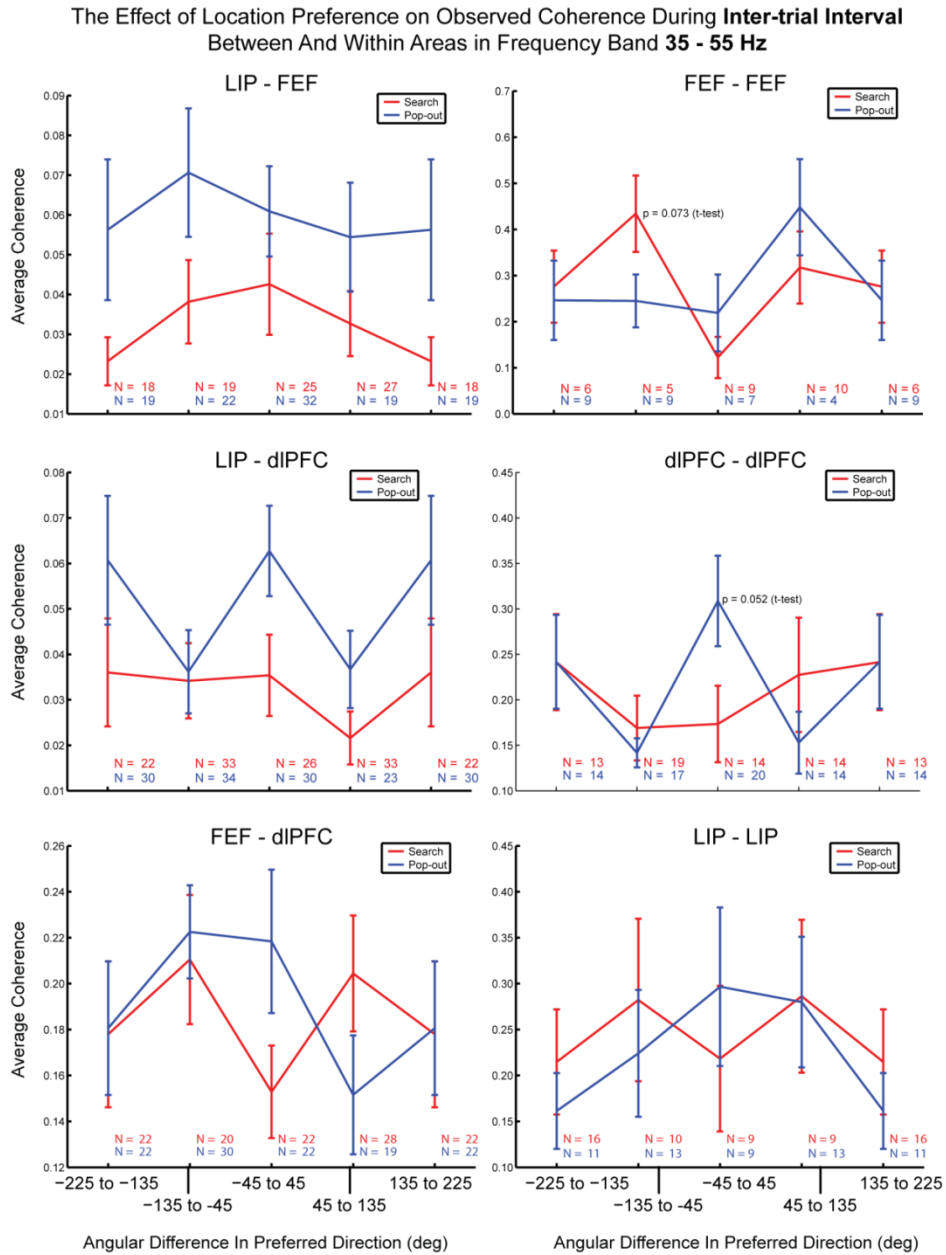


Figure 29. 35 – 55 Hz coherence between pairs of electrodes based on the difference in their preferred target locations during the **inter-trial interval**. The coherence is compared between areas (left column) and within areas (right column) for both search (red) and pop-out (blue). The number of pairs of electrodes for each location preference offset is given for both search (red) and pop-out (blue). Errorbars indicate the 95% confidence interval about the mean. Significant differences between search and pop-out were determined using an unpaired t-test, with the p-value indicated.

Correlation of Synchrony with Reaction Time


Although we have described the relative strength of coherence across visual search and pop-out, it is also interesting to ask whether the strength of the coherence on a trial-by-trial basis is correlated with how well the animal is able to perform the task. In order to determine the trial-by-trial coherence we used a jackknife method as described in the methods. We then used the estimated coherence from a single trial to attempt to predict whether the reaction time from that trial would be high or low. Trials were binned into 10 separate groups based on their reaction times. A linear correlation was then computed between the coherence of a given pair of electrodes and the reaction time. This was done for all possible pairings of areas in our two frequency bands of interest (the middle frequency band, 18-34 Hz, and the upper frequency band 35-55 Hz). Trials from visual search and visual pop-out were analyzed separately in order to avoid task-related correlations in reaction time and coherence.

There was a negative correlation between reaction time and coherence in the 18-34 Hz frequency band for FEF-FEF, FEF-dIPFC, dIPFC-dIPFC, and LIP-LIP (Figure 30). Since the correlation is negative, an increase in coherence on a single trial leads to a decrease in the reaction time. Therefore, when synchrony within frontal cortex and within parietal cortex is increased in the middle frequency band the animal is able to respond more quickly. This finding supports the importance of the middle frequency band to the search task. Interestingly, the largest difference in coherence observed between search and pop-out was in the middle frequency between the frontal regions and LIP, however an increased level of coherence between these areas does not cause a significant improvement in reaction time. This may be due to a saturation of the coherence between these regions during search.


During visual pop-out there were negative correlations between the FEF-FEF coherence in the 18-34 Hz range as well as LIP-LIP coherence in the 35-55 Hz range (Figure 31). As the FEF-FEF, 18-34 Hz range also showed a significant negative correlation with reaction time during search it may play a more general role in speeding up eye movements regardless of the task. Coherence within LIP was the only pairing that showed a significant correlation in the upper frequency band, demonstrating a very strong effect – over 20% of all electrode pairs within LIP were negatively correlated with reaction time. As the correlation is negative, increased coherence within LIP in the upper frequency band leads to a faster

Figure 30


Distribution of Correlation Coefficient between RT and FEF - LIP Coherence during Search




Distribution of Correlation Coefficient between RT and FEF - FEF Coherence during Search




Distribution of Correlation Coefficient between RT and LIP - PFC Coherence during Search




Distribution of Correlation Coefficient between RT and PFC - PFC Coherence during Search



Distribution of Correlation Coefficient between RT and FEF - PFC Coherence during Search



Distribution of Correlation Coefficient between RT and LIP - LIP Coherence during Search



□ Sig., $p < 0.05$ ○ Not Sig.

* = $p < 0.05$ ** = $p < 0.01$


Figure 30. The distribution of correlation coefficient's between coherence and reaction time during **visual search**. The distribution is shown for all possible area pairs in both the middle (18-34 Hz) and upper (35-55 Hz) frequency bands. The left sub-figure for each area pair shows the distribution of correlation coefficients for each pair of electrodes. Significant correlations are indicated with a red square ($p < 0.05$). The right sub-figure shows the percent of all pairs demonstrating a significant correlation in either the Up (blue) or Down (red) direction. The proportion of pairs of electrodes with significant reaction time correlations was tested against the binomial distribution. A single asterisk indicates $p < 0.05$, a double asterisk indicates $p < 0.01$.

reaction. Similar effects were previously shown in earlier visual areas (Womelsdorf et al., 2006b). These results suggest that when local synchrony is enhanced the animal is able to respond more quickly.


Neither of the tested frequency ranges was positively correlated with reaction time for either task, suggesting that while increases in synchronization at a particular band may aid in the task performance, increasing synchrony does not hinder behavior. Whether this is a general truth is an open question and warrants further investigation.

Figure 31


Distribution of Correlation Coefficient between RT and FEF - LIP Coherence during Pop-out




Distribution of Correlation Coefficient between RT and FEF - FEF Coherence during Pop-out




Distribution of Correlation Coefficient between RT and LIP - PFC Coherence during Pop-out




Distribution of Correlation Coefficient between RT and PFC - PFC Coherence during Pop-out



Distribution of Correlation Coefficient between RT and FEF - PFC Coherence during Pop-out



Distribution of Correlation Coefficient between RT and LIP - LIP Coherence during Pop-out



□ Sig., $p < 0.05$ ○ Not Sig.

* = $p < 0.05$ ** = $p < 0.01$

Figure 31. The distribution of correlation coefficient's between coherence and reaction time during **visual pop-out**. The distribution is shown for all possible area pairs in both the middle (18-34 Hz) and upper (35-55 Hz) frequency bands. The left sub-figure for each area pair shows the distribution of correlation coefficients for each pair of electrodes. Significant correlations are indicated with a red square ($p < 0.05$). The right sub-figure shows the percent of all pairs demonstrating a significant correlation in either the Up (blue) or Down (red) direction. The proportion of pairs of electrodes with significant reaction time correlations was tested against the binomial distribution. A single asterisk indicates $p < 0.05$, a double asterisk indicates $p < 0.01$.

Phase Relationships between Local Field Potentials

Synchrony between the local field potentials can also be computed using relative phase relationships. Once frequency bands of interest are known, we can filter the LFP signal from each electrode into those bands. After locating the peaks and valleys of the filtered oscillation it is possible to determine the relative phase of another signal. For all of the results shown here the transitive property was maintained, so only one pair is shown (e.g. LIP-dIPFC is the same as dIPFC-LIP, but with reversed phase relationships). Although this method does require one to previously identify the frequency bands of interest, it does confer two major advantages. First, the signals can be filtered across the entire recording time, allowing for the estimation of synchrony in limited time windows for frequency bands much lower than is possible using the spectral method. Second, since the signal is no longer assumed to be stationary within our window (as is required of the spectral method) the relative phase between our filtered reference signal and a probe signal can be more accurately determined. This is vitally important when the reference signal contains phase-jumps – a common occurrence in neural signals, especially around stimulus presentations and behavioral responses.

For this analysis we will use phase-locking between two LFP signals as a measure of their degree of synchrony. Circular statistics were used to determine whether the observed phase distribution was significantly different from the null distribution (see Methods for detailed description of these analyses). We used 4 different frequency bands of interest – a low frequency band of 4-8 Hz, a low-middle band of 8-16 Hz, a middle band of 18-34 Hz, and an upper frequency band of 35-65 Hz. All filtering was done in a phase-conserving manner with high order Butterworth filter applied across the entire signal from each trial (all filters had a minimum -10 db attenuate within 1 Hz outside the filter range). These frequency bands were chosen to match well with existing literature, as well as our previous results from the coherence analysis.

Figures 32 and 33 show the difference in observed phase-locking between visual search and visual pop-out for inter- and intra-areal pairings, respectively. Each figure shows the difference in phase-locking between LFPs recorded from all possible pairs of areas. Differences in synchrony are plotted across time for 4-8Hz (left-most column), 8-16 Hz (middle-left column), 18-34 Hz (middle-right

column), and 35-65 Hz (rightmost column). All possible pairings are shown for completeness and because all 4 frequency ranges will be important for later analysis.

Generally, the phase-locking statistic follows the pattern of results found using the coherence statistic when comparisons are possible. Inter-areal effects (Figure 32) appear to be very consistent with the coherence statistic, while intra-areal effects (Figure 33) appear to show a stronger second-order effect than the coherence measure. For example, synchrony in the 18-34 Hz band was increased during search for all pairing groups, an effect that was also seen in the coherence statistic with the exception of LIP-LIP which previously showed no significant effect. The phase-locking in the upper frequency band (35-65 Hz) showed an increase during pop-out for FEF-LIP and LIP-dIPFC as did the coherence statistic, as well as additional increases in FEF-dIPFC synchrony during visual search. Again, intra-areal synchrony was slightly different than the coherence statistic as LIP-LIP showed a significant increase in phase-locking during pop-out, FEF-FEF showed an increase in synchrony for search, and dIPFC-dIPFC showed no significant change between tasks (coherence was increased for pop-out).

Both of the lower frequency bands (4-8 Hz and 8-16 Hz) showed a temporal pattern of increasing their overall synchrony from the fixation onward for both visual search and pop-out (data not shown). The 4-8 Hz band tended to peak around the memory delay, while 8-16 Hz synchrony tended to show a flat or slight upwards trend as the trial proceeded. These effects existed across all 6 pairs of areas. The difference in synchrony between search and pop-out showed that for the lower frequency band (4-8 Hz) visual pop-out was generally more synchronous across most of the trial. However, during

Figure 32

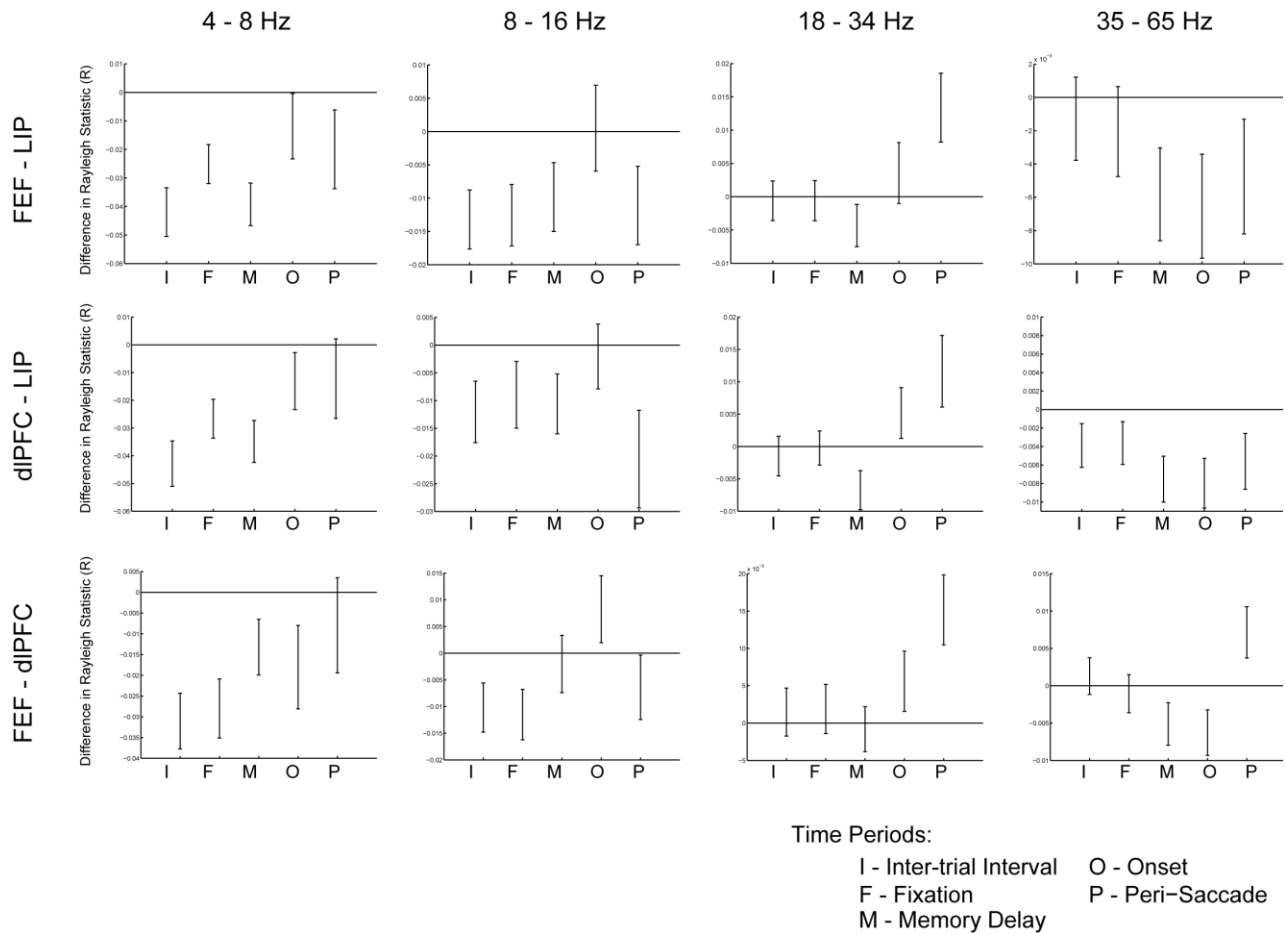


Figure 32. Difference in **inter-areal** synchronization between visual search and pop-out. The difference is plotted across time (x-axis), for each inter-areal pairing (across rows), and for our 4 frequency bands of interest: 4-8 Hz (leftmost column), 8-16 Hz (middle-left column), 18-34 Hz (middle-right column), and for 35-65 Hz (rightmost column). The degree of synchronization is measured with the Rayleigh statistic, R. If a circular distribution is farther away from being uniform, then R will increase. Since a non-uniform distribution suggests synchronization through phase-locking, a high R means more synchronization. The mean difference in R between search and pop-out is plotted with the errorbars indicating a 95% confidence interval. Therefore, a mean below, and not overlapping with, zero indicates increased phase-locking during pop-out. Conversely, a R significantly above zero indicates increased phase-locking during visual search.

Figure 33

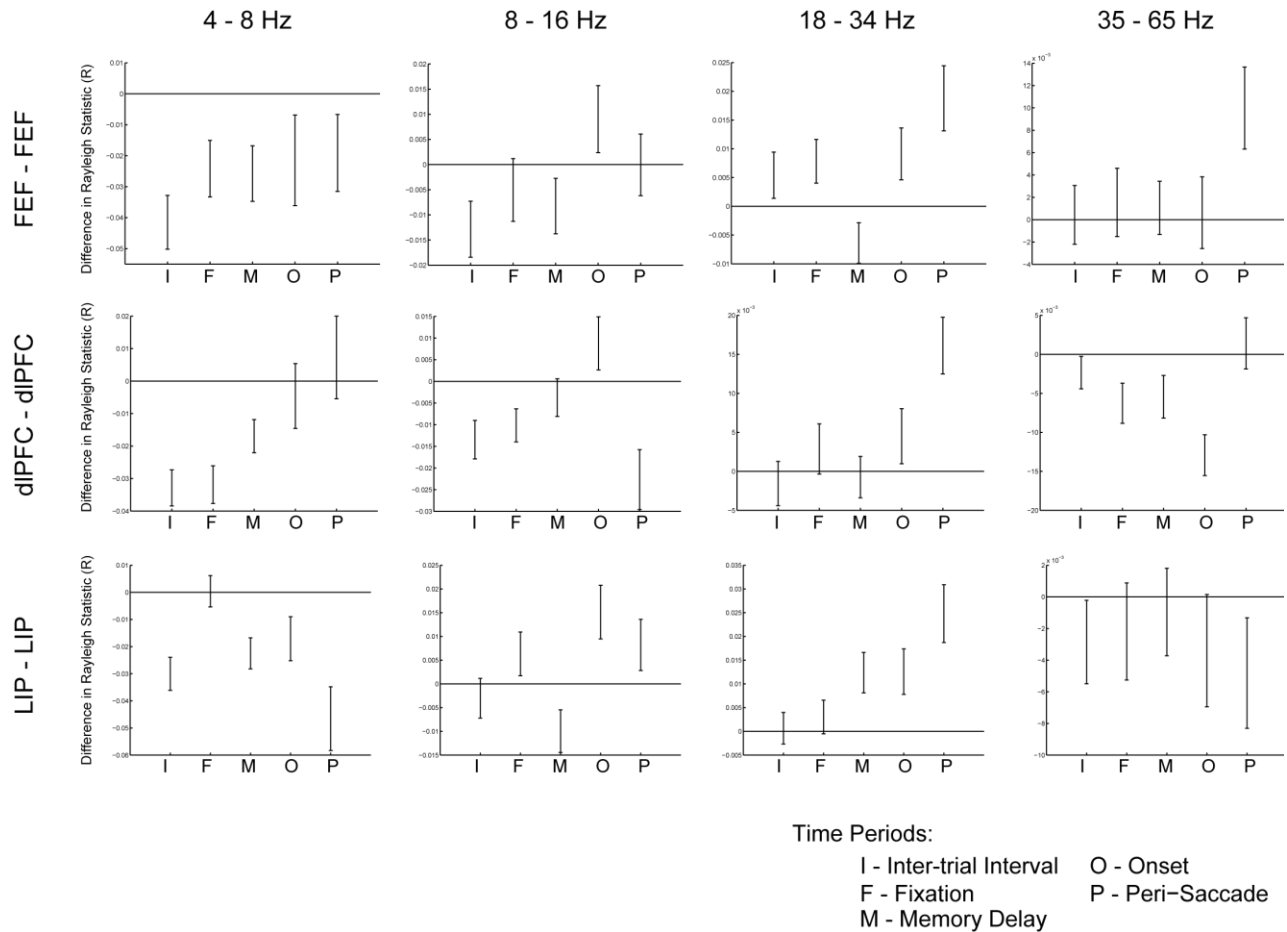


Figure 33. Difference in **intra-areal** synchronization between visual search and pop-out. The difference is plotted across time (x-axis), for each intra-areal pairing (across rows), and for our 4 frequency bands of interest: 4-8 Hz (leftmost column), 8-16 Hz (middle-left column), 18-34 Hz (middle-right column), and for 35-65 Hz (rightmost column). The degree of synchronization is measured with the Rayleigh statistic, R. If a circular distribution is farther away from being uniform, then R will increase. Since a non-uniform distribution suggests synchronization through phase-locking, a high R means more synchronization. The mean difference in R between search and pop-out is plotted with the errorbars indicating a 95% confidence interval. Therefore, a mean below, and not overlapping with, zero indicates increased phase-locking during pop-out. Conversely, a R significantly above zero indicates increased phase-locking during visual search.

the peri-saccadic period, for which the differences in task are maximal, interactions between dIPFC and any other region show no difference in synchrony, while interactions within and between LIP and FEF show a continued increase in synchrony for visual pop-out.

Measuring the degree of synchrony through phase-locking also allows us to determine the phase offset of the LFP signals. For this analysis all of the phases are relative to the peak of the reference electrode (which is always listed first in a grouping). The phase difference between the signals is determined by the phase of the reference signal at the peak of the probe signal (which is listed second). The phase relationships for X-Y are symmetric to Y-X, so only X-Y are plotted for succinctness. Due to this fact, all relationships between signals within an area (X-X) showed no phase differences. The phase relationships between areas are plotted on concentric circles. The two semicircles plotted for each area pair represent the 95% confidence interval of the mean phase difference during search (red) and pop-out (blue). The phase difference is plotted over time for each frequency band (for this analysis the response in fixation is omitted as it showed similar effects to the ITI, which serves as a good baseline). Average phases that lie in the upper half of the circle suggest that the reference electrode is leading the probe electrode (for X-Y, X would be leading Y). Phases in the bottom half suggest the probe leads the reference (Y leads X).

Figure 34 shows the inter-areal LFP phase differences for the lowest frequency band, 4-8 Hz, across the trial. Initially, all three regions show a phase lag around zero – the two areas are oscillating in phase. However, once the visual array is presented the relationship between these regions changes dramatically – both LIP-dIPFC and LIP-FEF show a significant negative phase shift, suggesting the frontal regions are leading the parietal region. Later in the trial, during the peri-saccade time period, this large shift disappears and instead we have a double dissociation between the current task and which area is leading. During pop-out, parietal regions lead frontal regions but during search frontal regions are leading parietal regions in this band. This difference mirrors the difference in our flow of information about the target location and suggests a role for the 4-8 Hz band in controlling the overall flow of information between brain regions.

Figure 34

Phase Relationship between Inter-Areal LFP Signals in the 4 - 8 Hz Frequency Band

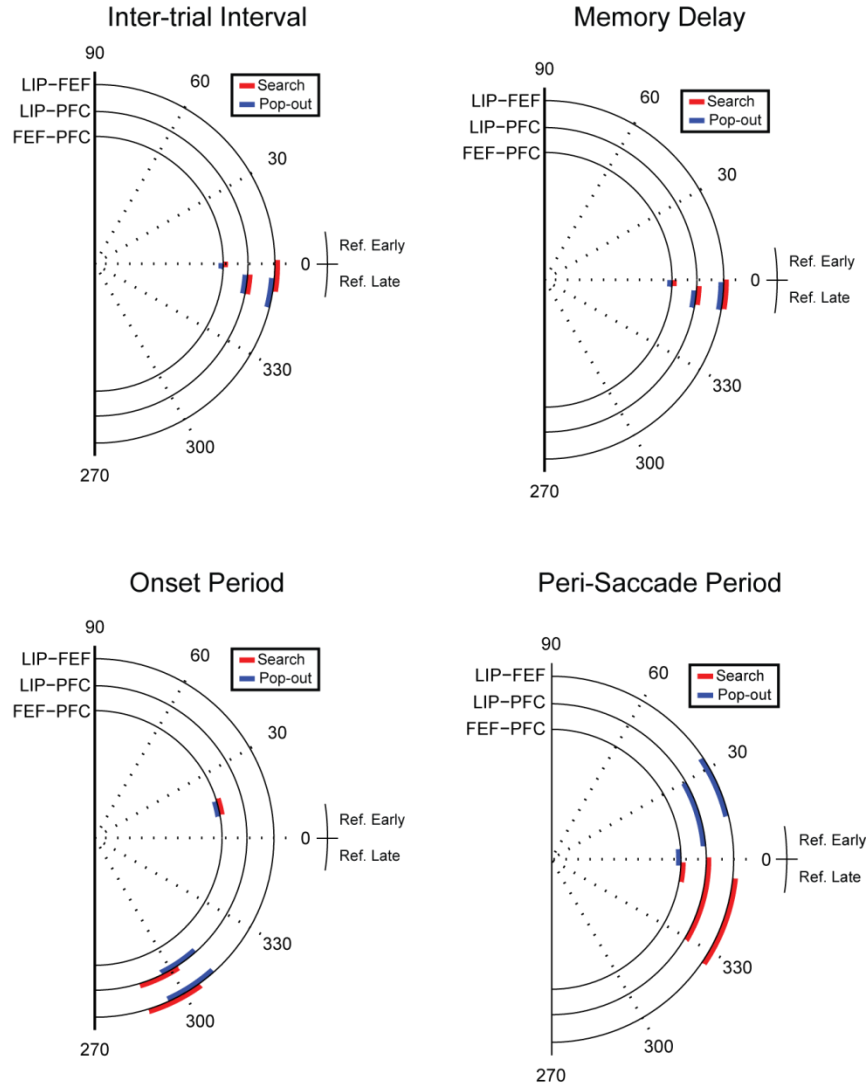


Figure 34. Phase difference in the LFP signals of different brain regions for the 4 – 8 Hz frequency band. The phase difference is plotted across time for the inter-trial interval (upper left), the memory delay (upper right), after the onset of the visual array (lower left), and in the peri-saccade period around the animal’s decision (lower right). The relative phase is measured with respect to the first, reference, signal (e.g. for the outermost right, LIP). Therefore a positive angle suggests the reference signal preceded that of the probe signal, while a negative angle suggests the probe signal preceded the reference signal. The average angle was determined for both search (red) and pop-out (blue) trials. The spread of the semi-circle indicates the 95% confidence interval around the average angle.

Figure 35

Phase Relationship between Inter-Areal LFP Signals in the 8 - 16 Hz Frequency Band

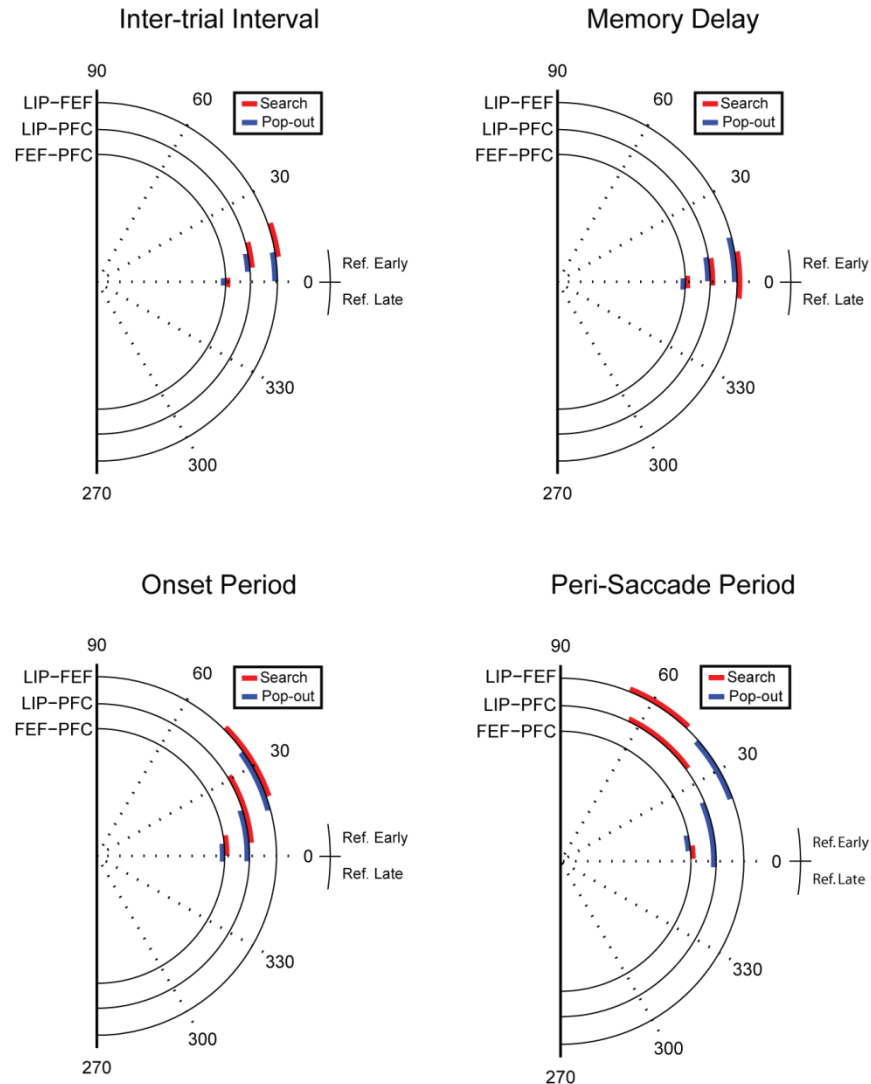


Figure 35. Phase difference in the LFP signals of different brain regions for the 8 – 16 Hz frequency band. The phase difference is plotted across time for the inter-trial interval (upper left), the memory delay (upper right), after the onset of the visual array (lower left), and in the peri-saccade period around the animal’s decision (lower right). The relative phase is measured with respect to the first, reference, signal (e.g. for the outermost right, LIP). Therefore a positive angle suggests the reference signal preceded that of the probe signal, while a negative angle suggests the probe signal preceded the reference signal. The average angle was determined for both search (red) and pop-out (blue) trials. The spread of the semi-circle indicates the 95% confidence interval around the average angle.

Figure 36

Phase Relationship between Inter-Areal LFP Signals in the 18 - 34 Hz Frequency Band

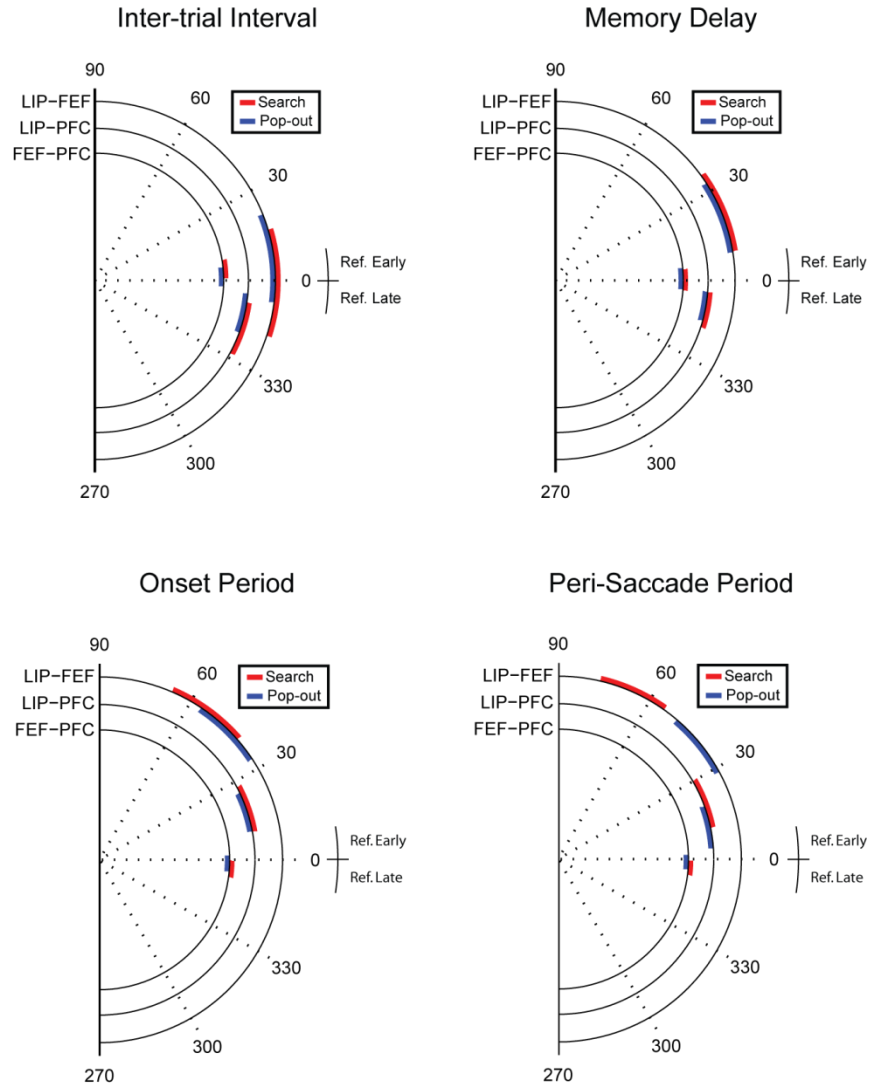


Figure 36. Phase difference in the LFP signals of different brain regions for the 18 - 34 Hz frequency band. The phase difference is plotted across time for the inter-trial interval (upper left), the memory delay (upper right), after the onset of the visual array (lower left), and in the peri-saccade period around the animal's decision (lower right). The relative phase is measured with respect to the first, reference, signal (e.g. for the outermost right, LIP). Therefore a positive angle suggests the reference signal preceded that of the probe signal, while a negative angle suggests the probe signal preceded the reference signal. The average angle was determined for both search (red) and pop-out (blue) trials. The spread of the semi-circle indicates the 95% confidence interval around the average angle.

Figure 37

Phase Relationship between Inter-Areal LFP Signals in the 35 - 65 Hz Frequency Band

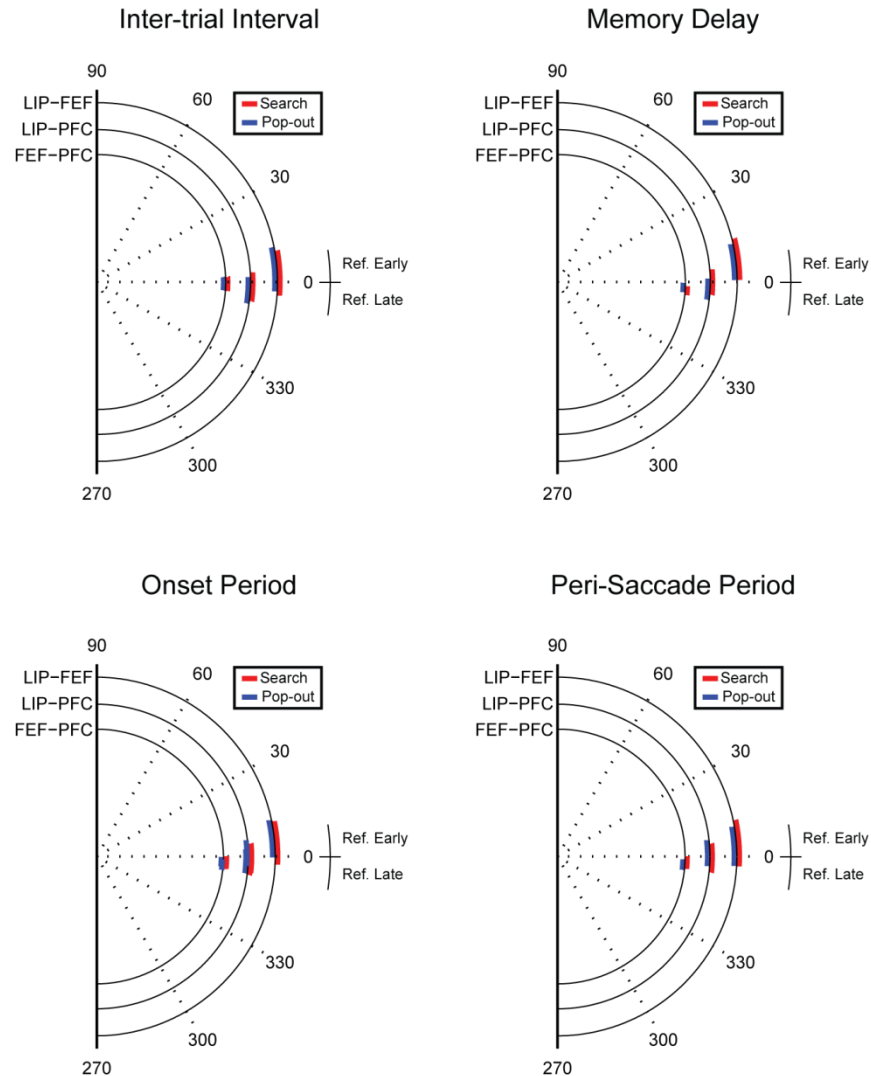


Figure 37. Phase difference in the LFP signals of different brain regions for the 35 - 65 Hz frequency band. The phase difference is plotted across time for the inter-trial interval (upper left), the memory delay (upper right), after the onset of the visual array (lower left), and in the peri-saccade period around the animal's decision (lower right). The relative phase is measured with respect to the first, reference, signal (e.g. for the outermost right, LIP). Therefore a positive angle suggests the reference signal preceded that of the probe signal, while a negative angle suggests the probe signal preceded the reference signal. The average angle was determined for both search (red) and pop-out (blue) trials. The spread of the semi-circle indicates the 95% confidence interval around the average angle.

The inter-areal phase relationship of the 8-16 Hz filtered LFP signal is shown Figure 35. Similar to the lowest frequency band, the signals are initially in phase until the presentation of the visual array. After array onset the parietal LFP begins to precede the frontal signals. This trend continues into the peri-saccade time period. Interestingly, in contrast to the lowest frequency band, the 8-16 Hz band shows a larger phase offset between LIP and the frontal regions during visual search than visual pop-out. As this frequency band shows a general synchrony increase during pop-out, and the observed phase relationship suggests parietal leading frontal cortex, one possible role for the 8-16 Hz band may be the forward flow of information.

Figure 36 presents the phase relationships over time for inter-areal phase-locking in the 18-34 Hz, middle, band. Unlike the two lower bands, there is a consistent phase relationship between parietal and dIPFC during the ITI and memory delay – the LFP from dIPFC proceeds that of LIP, suggesting a downward flow of information. However, after the visual array onset and during the peri-saccade period, this ordering reverses with dIPFC and FEF LFPs following those of LIP. We have demonstrated that the degree of synchrony between the frontal and parietal regions increases in a middle frequency band of approximately 18-34 Hz during visual search over pop-out. As the visual search paradigm requires a more integrated approach, we suggest this lower frequency band allows for integration over much large temporal differences, allowing the brain to gather all of the information necessary to find the target and complete the task. The phase differences between parietal and frontal regions suggest that although this frequency band may be emphasized when top-down influences are maximized, the signal originates in posterior cortex flowing into frontal regions.

The local field potentials of all three regions are in phase throughout the entire task for the highest frequency band, 35-65 Hz (Figure 37). This is different than the other frequency band in that the relationship appears to be fairly stationary across time. Synchrony between parietal and frontal cortex was shown to be enhanced in this frequency band during visual pop-out when compared to search. This, along with previous work, suggests a role for high-frequency oscillations in the local enhancement of attended stimuli, with this information being ‘passed’ along to downstream areas. The lack of phase differences between LFP signals in parietal and frontal cortex suggests that while information is passed from one region to the next, the high frequency oscillations are very local and are themselves not

propagated forward, but rather phase-lock to one another, possibly through cross-frequency phase locking.

Synchrony between Frequency Bands

Up to this point we have discussed the relative phase-locking of different anatomical areas across different frequency bands. We have suggested that communication between regions is facilitated through this synchronization and that task-specific enhancement of particular frequencies may aid in the flow of pertinent information. However, if one were to hypothesize that different types of information are carried on different frequency bands it becomes important to integrate across these frequencies. Here we investigate the possibility that the different frequency bands within a signal may phase-lock between themselves in order to facilitate cross-band communication.

As with the comparison of phase relationships between LFP signals, each LFP signal was filtered into four bands of interest. However, instead of comparing the relative phase between electrode pairs, we are now comparing the relative phase of different filtered signals within a single electrode. Due to the inherent multi-modal nature of casting higher frequencies onto lower frequencies, all circular analysis were done with a non-parametric test (see Methods).

The only observed phase-locking to occur between any frequency ranges from any of the three areas were between the lowest frequency band (4-8 Hz) and the other three bands. Furthermore, the cross-band synchronization only occurred after the visual array onset in all three regions (Figure 38, the 8-16 Hz band is not plotted in order to save space). Although all three regions showed similar increases in locking across the trial, they did vary in their relative increase for visual search and pop-out. Within LIP both the middle and upper frequency band showed significantly more synchronization with the 4-8 Hz band during pop-out, whereas for both dIPFC and FEF the opposite was true – cross-band synchronization was increased during search (Figure 38, top two rows). These results suggest a possible mechanism for brain regions to integrate across frequency bands – maybe allowing for the collection of different information carried on each stream. This hypothesis is supported by the differences in cross-band synchronization by task – LIP is more directly involved in finding the target during visual pop-out

and therefore is more likely to need increased cross-modal interactions to support the integration of information. In contrast, during visual search, a task requiring prefrontal control, both prefrontal regions show an increase in cross-band synchronization, as if to facilitate the increased need for integration.

Figure 38

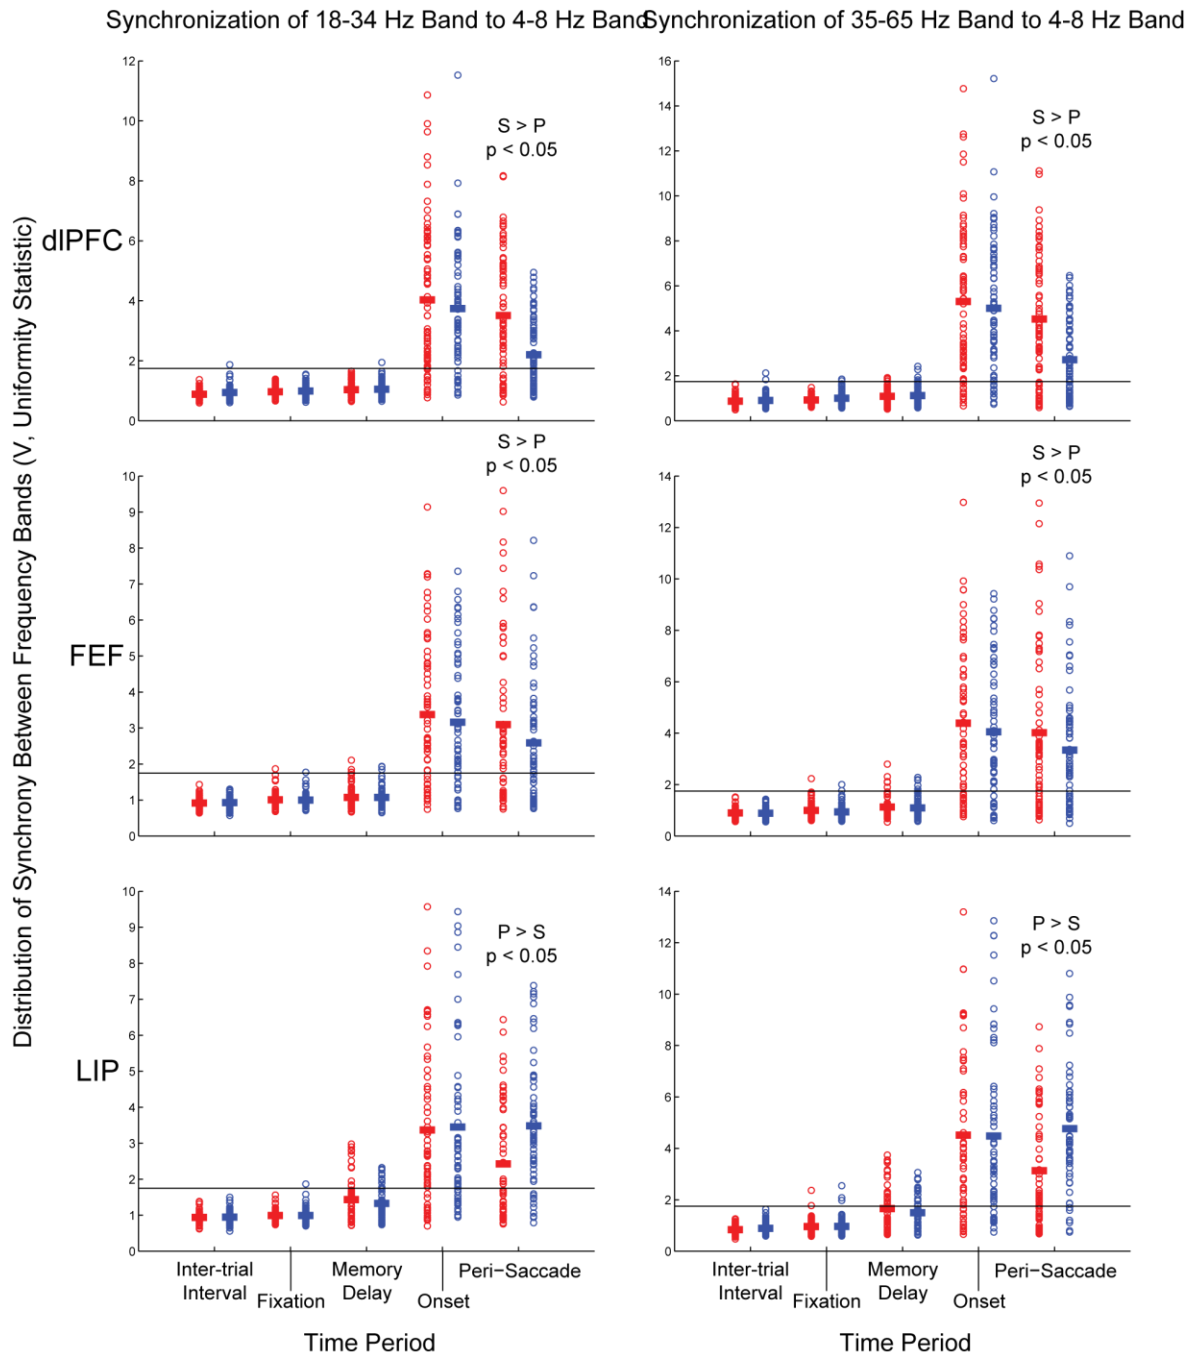


Figure 38. Degree of cross-band synchronization over time. A non-parametric test for uniformity (V) was used to determine significant phase-locking between the 4-8 Hz band and both the 18-34 Hz band (left column) and the 35-65 Hz band (right column). The degree of synchronization is shown for all three areas. The distribution of our synchrony statistic is shown over time for both visual search (red) and pop-out (blue). The solid colored line indicates the average V value and values above 1.747 (black line) are significant at $p < 0.05$. Significant differences between tasks are indicated as search greater than pop-out ($S > P$) or vice versa ($P > S$).

Synchrony between Spikes and Local Field Potentials

Although we have extensively investigated the synchronization of local field potentials, the brain ultimately communicates via spiking activity. Therefore, in order to ensure that the synchronization of local field potentials is relevant to the information currency of the brain we investigated the synchronization of spiking activity to the local field potential.

The phase of each spike was determined by interpolating the phase along the filtered LFP signal. Synchronization between the neuron and LFP was assessed by comparing the resultant distribution of phases from all spikes from a single neuron to the null distribution (see Methods for more information). Spiking activity and the LFP signal were both taken from a peri-saccade period of 200 ms, starting 150 ms before the saccade. As with all of the coherence measurements only electrodes that had at least one neuron with target location information were used. Spiking activity was drawn from isolated neurons that showed selectivity for the target location in either task. All possible neuron-electrode pairings were used with the exception of the neuron to its own electrode. This was avoided in order to prevent any contamination of the LFP signal with low-frequency components of the spike waveform.

The proportion of neuron-electrode pairs that showed a significant phase-locking is plotted in Figure 39 for all four frequency bands of interest. There are a significant proportion of neurons that synchronize with all four frequency bands and do so in a fairly diverse manner (within areas and across areas). This is not surprising as all three areas are heavily interconnected. Furthermore, as has been thoroughly demonstrated, there are high degrees of synchronization between LFP signals (and even between different frequency bands). Therefore, many of the effects we see across anatomical areas may be second-order effects due to LFP-LFP synchronization.

Figure 39




Figure 39. The proportion of single neurons within each area [denoted by (s)] phase-locked to the LFP signal from either the same, or another, area [denoted by (l)]. The asterisk (*) indicates a significant proportion of neurons were phase-locked ($p < 0.05$, tested against binomial). A small 'b' marks a significant intersection of neurons selective in both search and pop-out ($p < 0.05$, binomial test).

The strongest spike-field synchronization occurs within the 4-8 Hz range (although this is true for the peri-saccade time period, it is not always the case earlier in the trial – for example, during the memory delay, the 35-65 Hz frequency band is emphasized). Interestingly, for both the middle and upper frequency bands, the strongest synchronization between spiking activity and the local field potential occurs within the same region. This is especially true for the upper frequency band (Figure 39, bottom row), supporting the hypothesis that this high frequency oscillation is mostly a local one.

All spike-field pairings that have a population of neurons that are synchronized in both visual search and pop-out show a large proportion of overlap between the two populations (all $p < 0.05$, with populations showing 30-60% overlap). This suggests that when there is significant synchronization between a neuron and a particular neuron then they tend to maintain that synchrony across tasks. This may allow for the same neuron to consistently represent its information on particular frequency bands.

Conclusion

As noted in the Introduction, localized synchrony of activity within a brain area may help resolve competition for attentional selection while inter-areal synchrony may aid in long-range communication between areas. Our results suggest that the flow of top-down and bottom-up information is aided by coherence emphasizing different frequency bands. Lower frequency bands are more robust to spike timing delays and thus may be better suited for longer-range coupling between multiple, distant areas (Kopell et al., 2000; von Stein and Sarnthein, 2000; Engel et al., 2001). The increase in middle frequency synchrony during search could reflect a ‘broadcast’ of signals on a larger anatomical scale. Synchrony at higher frequency bands might support the local interactions needed to enhance stimulus representations (Kopell et al., 2000; von Stein and Sarnthein, 2000; Engel et al., 2001). The emphasis of higher-frequency synchrony during pop-out could reflect local enhancement of stimulus representations that are passed forward from parietal to frontal cortex. Our results suggest that the brain may emphasize coherence at different frequency bands in order to dynamically modulate inter-areal connections such that the network best suited for the current task is always engaged.

Introduction

This thesis is focused on the relative relationship between parietal and frontal cortex in the control of attention and therefore our behavioral paradigm was designed in order to compare and contrast exogenous attention in visual pop-out and endogenous attention in visual search. However, despite the main variable of manipulation being whether or not attention is being directed in an external or internal manner, there are several other pieces of information necessary to complete the task. It is clear from the first chapter that while a good proportion of neurons in each area reflect the allocation of attention to, and subsequent selection of, the target, approximately 2/3 of recorded neurons from each area are not involved in carrying information about the target location. This chapter is focused on the other parameters involved in completing the task, and the neural representations of these parameters over time.

Information about Task Parameters over Time

Although we used information about the target location as a window into the allocation of attention, the main manipulation was to investigate differences in how this information is represented across parietal and frontal cortex under two different tasks – visual pop-out and visual search. We needed to block the two tasks (into blocks of ~35) in order to get the reaction time differences necessary to make a psychophysical argument about the type of attentional control employed by each task. Because the tasks were presented in blocks of contiguous trials, the animal was able to fully expect the type of task that would be used on the next trial – information that is likely useful to the animal in order to prepare for handling the two different types of task. As can be seen in Figure 40, information about the type of task used during the trial is strongly represented in all three areas throughout the entire task. The number of selective neurons from each area are shown for several time periods across the trial: the Inter-trial Interval (ITI, 1000 ms window starting 500 ms before the presentation of the fixation point); the fixation period (300 ms window, starting 300 ms before the presentation of the sample stimulus in order to avoid saccade artifacts); the sample presentation (1000 ms window across the entire time the sample was presented); the memory delay (500 ms window of the entire memory delay between sample stimulus and visual array onset); the onset epoch (the first 200 ms

Figure 40




Figure 40. Percent of recorded neurons that carried significant information about the task type (either search or pop-out) over the length of the trial. Data is shown from LIP (red), dIPFC (green) and FEF (blue). All three regions carried significant information about the task type during all time periods of the task ($p < 0.05$, tested against binomial).

of the visual array presentation); and a peri-saccade period (a 300 ms window starting 200 ms before the saccade – this was chosen to encompass as much of the process of finding the target as possible).

As the animal is searching for a particular target stimulus matching that of a sample stimulus presented earlier in the trial, information about sample stimulus is also vitally important to correctly completing the task. Figures 41, 42, and 43 show the percent of neurons in dlPFC, FEF, and LIP, respectively, selective for stimulus parameters (and target location for comparison) across time during both visual pop-out (top) and visual search (bottom). The same time windows are used as for the analysis of task selectivity. Information was tested for the color of the target, the orientation of the target, and the identity of the target (the conjunction of both color and orientation that makes up each unique stimulus).

All three regions carry significant information about the target stimulus in at least one time period, emphasizing the integrative role of all three brain regions. No region is devoted to a single task alone, whether it is directing attention, providing working memory, or moving the eyes. However, the temporal characteristics of the selectivity does vary between areas. Neurons in dlPFC carry complete information to uniquely identify the target in both tasks and do so from the sample stimulus presentation. In contrast, FEF neurons do not encode sample information until later in the trial, as the animal is finding the target and LIP neurons do not carry sufficient information to uniquely identify the target stimulus, but rather are only selective for the target orientation in both tasks.

It is interesting to note that, in general, neurons in all three areas tended to carry information about the target orientation and the target identity, but were relatively unresponsive to information about the target color. In fact, only the dlPFC carried significant information about the target color and only during visual search, not during pop-out. It is possible that as the target stimulus was unique in both color and orientation during visual pop-out, both animals tended to focus on a single parameter in order to make their judgment (in this case orientation). As the two parameters were not explicitly controlled for difficulty, it might be that the animals preferred to use orientation because it was the easier of the two parameters to perceive. This dichotomy may also be represented in the responses of

Figure 41




Figure 41. Percentage of recorded neurons in **dIPFC** that show selectivity for several task parameters over the length of the trial. Selectivity for target location, color, orientation and identity were tested. Highlighted bars indicate a significant percentage of cells were selective for that parameter during that time window ($p < 0.05$, tested against binomial distribution).

Figure 42

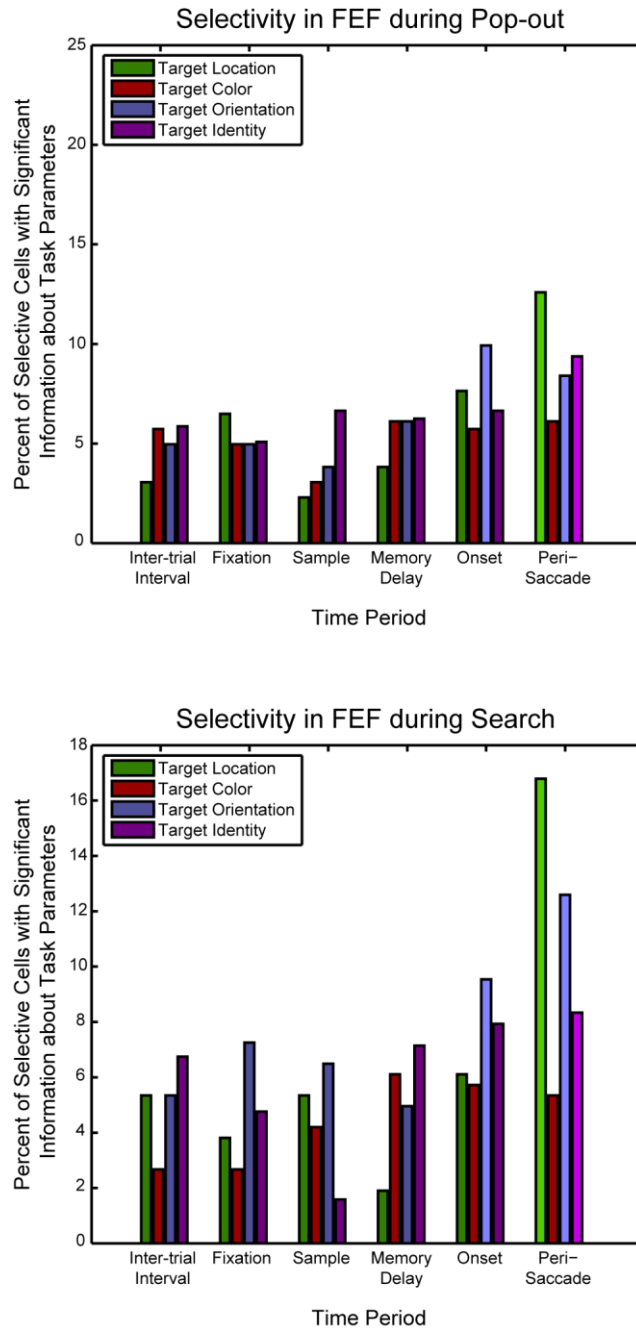


Figure 42. Percentage of recorded neurons in FEF that show selectivity for several task parameters over the length of the trial. Selectivity for target location, color, orientation and identity were tested. Highlighted bars indicate a significant percentage of cells were selective for that parameter during that time window ($p < 0.05$, tested against binomial distribution).

Figure 43

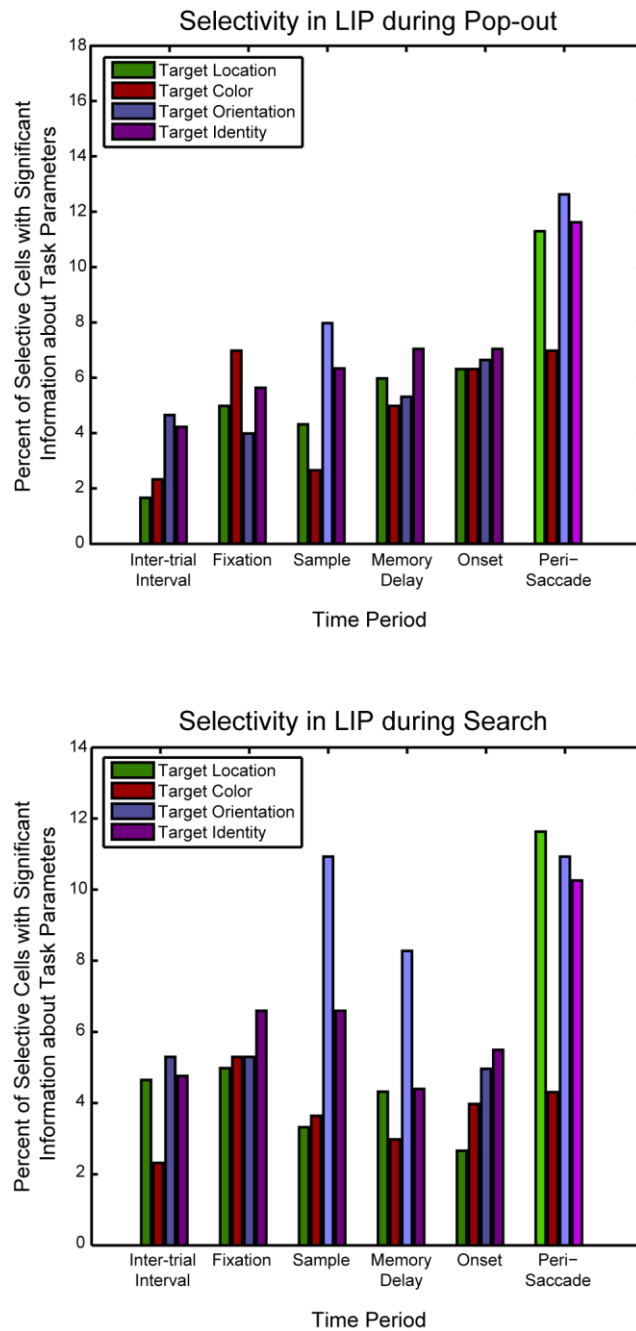


Figure 43. Percentage of recorded neurons in LIP that show selectivity for several task parameters over the length of the trial. Selectivity for target location, color, orientation and identity were tested. Highlighted bars indicate a significant percentage of cells were selective for that parameter during that time window ($p < 0.05$, tested against binomial distribution).

dIPFC neurons between the two tasks – while dIPFC neurons carry information about the target color during visual search, they do not carry the same information during visual pop-out, possibly because that information does not convey any further advantage to the animal (above orientation alone).

dIPFC was the only area of the three sampled to show significant information about the sample stimulus in both search and pop-out during the memory delay. This is not unexpected as the working memory capabilities of lateral PFC are well documented (Miller and Cohen, 2001). Interestingly, in the visual search task LIP neurons do continue to carry information about the target orientation throughout the memory delay. This may play a role in establishing pre-attentive ‘filters’ along the visual stream in order to emphasize stimuli in the visual search array with orientation matching the target stimulus. However, further experiments are necessary to directly test this issue.

Conclusion

Based on the flow of information about the allocation of attention across LIP, FEF, and dIPFC we have proposed that during conditions requiring endogenous control of attention, the two frontal regions provide top-down direction, while during conditions emphasizing exogenous control of attention, parietal cortex appears to feed information into frontal cortex. Our analysis of other task parameters supports this basic model while filling out the more general role each of these three regions play in behavior.

All three regions showed selectivity for the target stimulus itself, not just its location. This emphasizes the fact that these regions play complex roles, both directing attention to specific stimuli/locations as well as representing the information needed to identify the target. The strongest representation of stimulus identity was found in the dIPFC. This results fits well with the role of dIPFC acting to gather and integrate information across the rest of the brain in order to direct behavior. Importantly, dIPFC was the only area to contain neurons with enough information to uniquely identify the target stimulus throughout the entire task. Conversely, neurons in FEF carried the least information about the target stimulus, suggesting a more active role in the direction of spatial attention and eye movements than stimulus processing per-se. LIP did carry information about the target stimulus, but

not enough to completely identify the target. Instead, it appears to focus on a single dimension of the two-dimensional stimuli. This result fits well with our own work as well as previous work suggesting that LIP acts as a saliency map (Gottlieb et al., 1998; Colby and Goldberg, 1999; Kusunoki et al., 2000; Bisley and Goldberg, 2003). This made aid in the visual search process by providing a one-dimensional filtering of the visual field.

Introduction

Previous chapters of the thesis have focused on the role of neurons in visual pop-out and visual search separately with comparisons between the two tasks made at the population level. For example, we determined the relative roles of parietal and frontal cortex in directing attention during bottom-up, externally-guided pop-out and compared that to top-down, internally driven search and we have shown that these three regions (LIP, dlPFC, and FEF) are all involved in both forms of attentional control, but with shifts in relative timing. However, the neural networks within each region involved in both tasks has not been described -- a fundamental question is the degree to which these two forms of control have overlapping mechanisms. To try to answer this question we will investigate whether neurons within each area are active in one or both tasks, and if active in both tasks, whether or not the information represented by a single neuron is consistent across tasks. This is an important question as it may offer evidence towards the two task types being independent, as might be argued by Triesman and colleagues, or whether the two tasks lie along a continuum of possibilities, with each individual instance of visual search utilizing underlying mechanisms to different degrees.

For our analysis we will denote the observed proportion of neurons selective in both tasks as significant if it meets two requirements: first, the proportion of selective neurons in each task alone must be significant; second, the observed percentage of neurons responsive in both tasks must be above chance level when chance probability is defined as the product of the observed proportion of selective neurons in search and pop-out alone. This ensures the existence of a population of selective neurons in each task and then determines whether these two populations are significantly overlapping.

Attention Signals Across Tasks

Neurons in all three regions show selectivity for the target location in both tasks, although with different latencies. Based on the differences in timing we have suggested that the flow of information between frontal and parietal cortex shifts between the two tasks. However, it is not clear from these results how the neural mechanisms underlying behavior varies between tasks. One possibility is that the two tasks utilize similar neural mechanisms but with different regions taking the lead. Alternatively, the

neural substrate underlying each behavior may be different. In order to determine the degree of overlap in the neural substrate for each task, we performed several analyses comparing the groups of neurons active in each task.

Figures 44, 45, and 46 show the time to first significance of individual neurons in LIP, dIPFC, and FEF, respectively, during both visual pop-out (y-axis) and visual search (x-axis). Neurons that were selective for the location of the target in both tasks are plotted in black, while individual neurons responsive in only a single task are plotted in blue (pop-out) and red (search) along the respective axis. Cumulative sums are shown outside each axis. These figures highlight the degree of overlap in early (pre-saccadic) and late (post-saccadic) responses in each area. Both LIP and FEF showed disparate pre-saccadic (early) populations between the two tasks (see Figure 47 for summary), suggesting that the neural mechanisms underlying these two tasks differed. In contrast, there are a significant number of neurons in dIPFC that were selective for both tasks before the saccade (4% of overall selective neurons, $p = 2.1 \times 10^{-4}$, corrected for multiple comparisons across regions and tested against binomial distribution as described above). While the populations do show an overlap, there is no significant linear correlation between the time neurons selective in both task first reached significance ($\rho = -0.52$, $p = 0.1$, Figure 45). Although there is no correlation in the time neurons find the target between the two tasks, the location preference of early-responding neurons is conserved across visual search and pop-out (Figure 48).

Neurons that responded after the saccade (late) showed significant overlap for all three areas (see Figure 47 for summary; LIP, $p < 10^{-5}$; FEF, $p < 10^{-5}$; dIPFC, $p < 10^{-5}$; all tested by binomial, corrected for multiple comparisons across area). Additionally, the time to first significance of late responding neurons is significantly correlated for both LIP and FEF (LIP, $\rho = 0.437$, $p = 0.01$; FEF, $\rho = 0.489$, $p = 0.02$) but not for dIPFC ($\rho = 0.260$, $p = 0.31$). We can also determine the preferred direction of each neuron selective late in the trial using the neuron's activity to when the target lay in each of the four possible directions. The difference between the preferred direction for significantly selective neurons is shown in Figure 49 for dIPFC (top), LIP (middle), and FEF (bottom). It is clear from this analysis that the population of neurons responding late in the task conserved their selectivity between visual search and

pop-out, suggesting that the common framework supporting the behavior late in the trial is consistent across tasks.

These results fit well with our understanding of the role of LIP and FEF in the production of saccades. It is likely that the late neurons are reflective of the motor movement itself and not the direction of attention. However, the neurons responsive early in the trial likely play a large role in directing attention. Interestingly, this population of neurons appears to be disparate in FEF and LIP for the two tasks while overlapping for dIPFC.

Figure 44

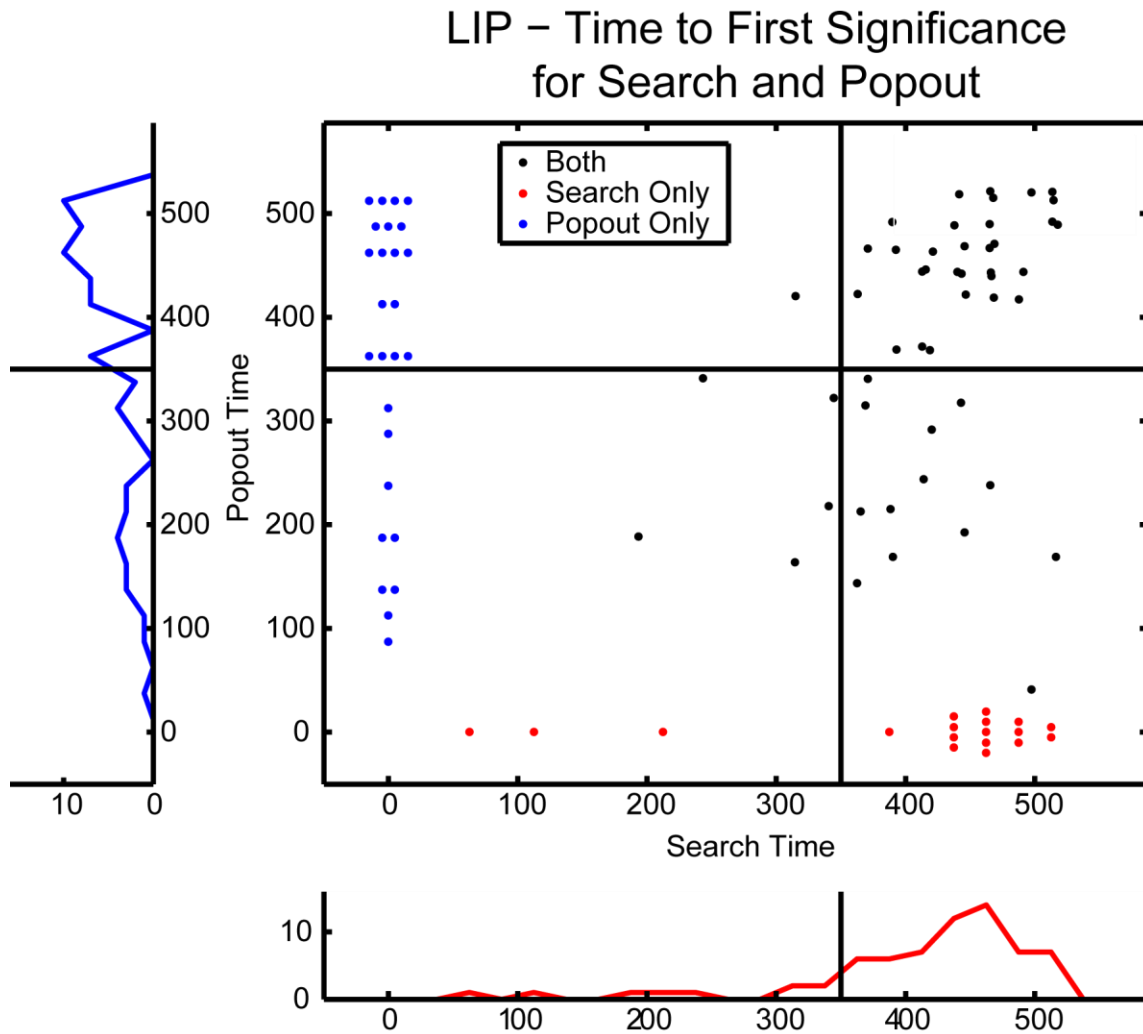


Figure 44. Comparison of the time for LIP neurons to first carry significant information about the target location in visual search (x-axis) and visual pop-out (y-axis). Each dot represents a neuron from LIP selective for target location. Black dots indicate neurons selective during both tasks, while neurons only selective during search and pop-out are red and blue dots respectively. As these neurons do not have timing information available for both tasks they are plotted along the axis. Selectivity was determined by testing the observed mutual information against a null distribution created through a randomization procedure (see Methods). The histogram of when neurons found the target during search and pop-out are plotted outside the x- and y-axis, respectively.

Figure 45

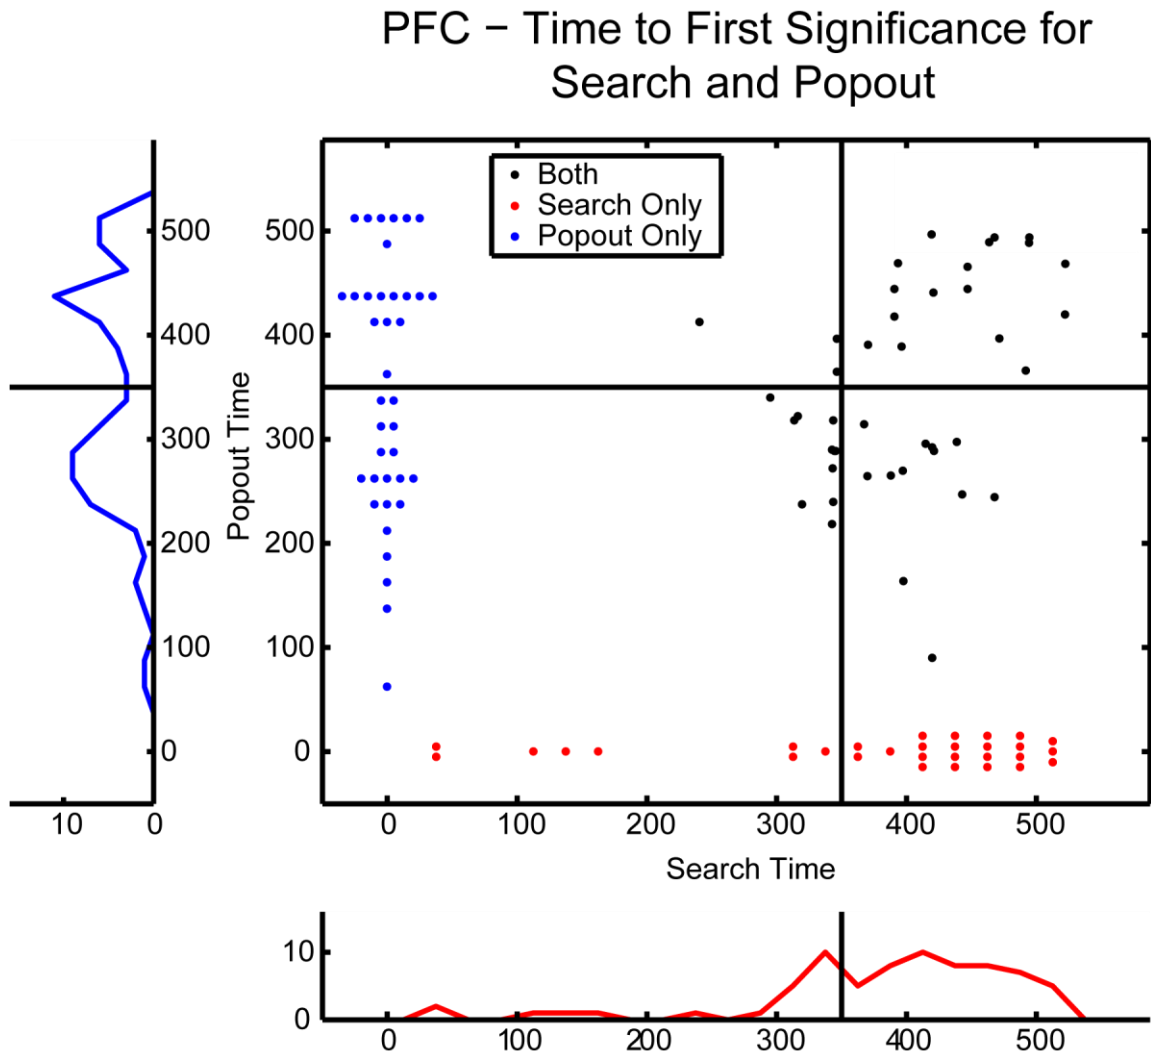


Figure 45. Comparison of the time for **dIPFC** neurons to first carry significant information about the target location in visual search (x-axis) and visual pop-out (y-axis). Each dot represents a neuron from dIPFC selective for target location. Black dots indicate neurons selective during both tasks, while neurons only selective during search and pop-out are red and blue dots respectively. As these neurons do not have timing information available for both tasks they are plotted along the axis. Selectivity was determined by testing the observed mutual information against a null distribution created through a randomization procedure (see Methods). The histogram of when neurons found the target during search and pop-out are plotted outside the x- and y-axis, respectively.

Figure 46

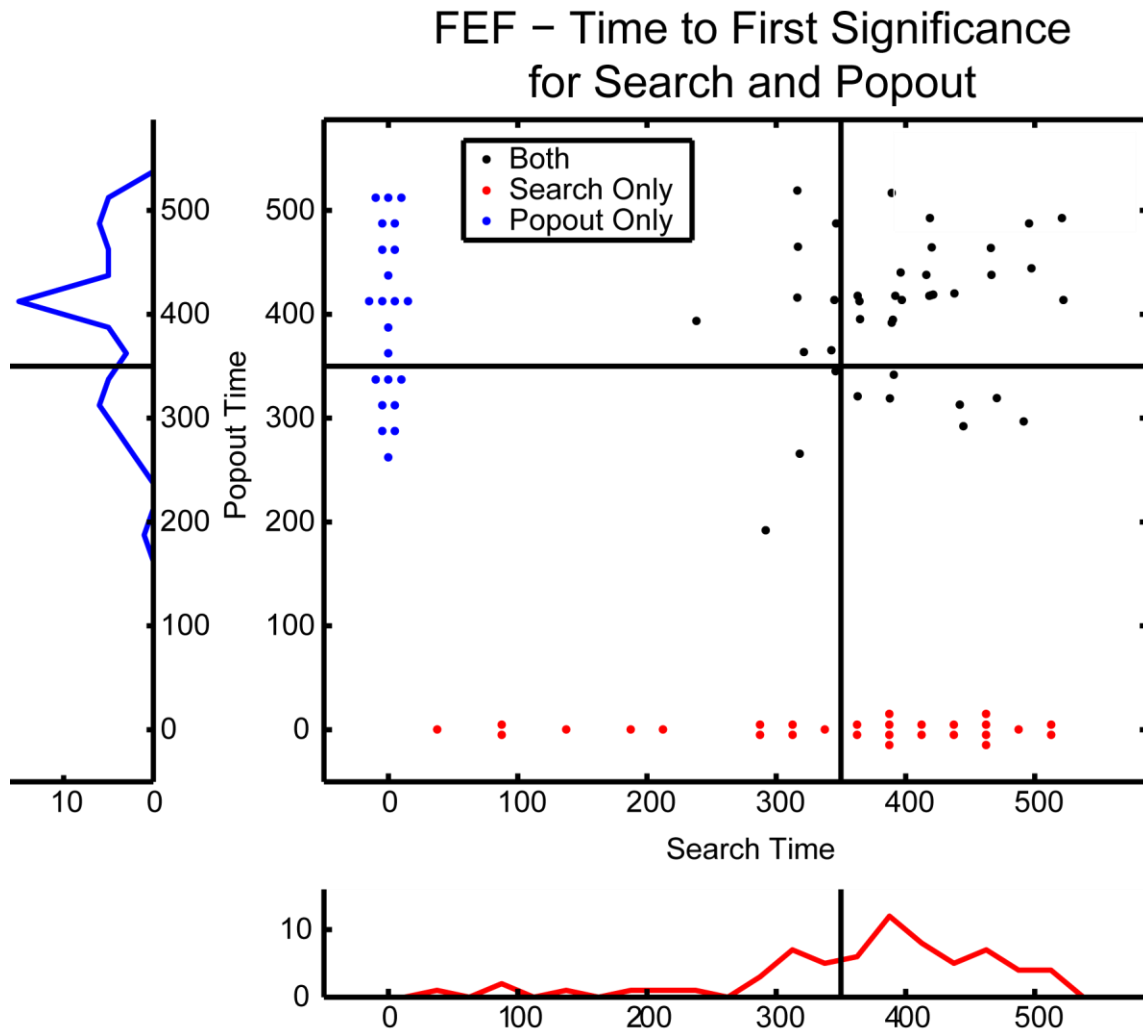


Figure 46. Comparison of the time for FEF neurons to first carry significant information about the target location in visual search (x-axis) and visual pop-out (y-axis). Each dot represents a neuron from FEF selective for target location. Black dots indicate neurons selective during both tasks, while neurons only selective during search and pop-out are red and blue dots respectively. As these neurons do not have timing information available for both tasks they are plotted along the axis. Selectivity was determined by testing the observed mutual information against a null distribution created through a randomization procedure (see Methods). The histogram of when neurons found the target during search and pop-out are plotted outside the x- and y-axis, respectively.

Figure 47




Figure 47. Magnitude and intersection of neural populations carrying information about the **target location**. Diamond size indicates the percentage of neurons selective for target identity during visual search (red) and pop-out (blue). The size of the overlap measures the intersection between the search and pop-out populations of selective neurons. The asterisk denotes a significant proportion of cells carrying sample information in each task or if the intersection of these populations was significant ($p < 0.05$, tested against binomial).

Figure 48

Histogram of Angular Difference between Location Preference of Pre-Saccadic Neurons Responsive in Both Search and Pop-out

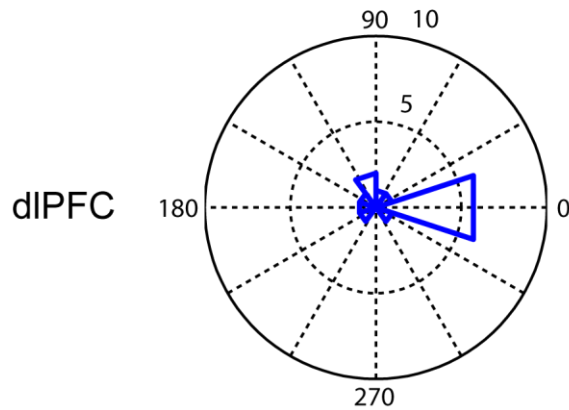


Figure 48. Circular histogram of the difference in preferred target location between visual search and pop-out for pre-saccadic neurons in dIPFC. Only neurons that were selective during both visual search and pop-out before the initiation of the saccade are included. The pre-saccadic response was used to determine a vector for each neuron during both tasks and the angle between these vectors is plotted here. Only dIPFC is shown since it is the only area to have a significant number of neurons active before the saccade in both tasks.

Figure 49

Histogram of Angular Difference between Location Preference of Post-Saccadic Neurons Responsive in Both Search and Pop-out




Figure 49. Circular histogram of the difference in preferred target location between visual search and pop-out for post-saccadic neurons for dIPFC (top), LIP (middle) and FEF (bottom). Only neurons that were selective during both visual search and pop-out after the initiation of the saccade are included. Their post-saccadic response was used to determine a vector for each neuron during both tasks and the angle between these vectors is plotted here.

Information about the Stimulus Parameters

As was shown in the previous chapter, all three regions carried information relating to the identity of the target stimulus at some point during the trial. As the stimuli and stimulus presentation were conserved across both tasks, one might expect the underlying neural framework to be conserved. Alternatively, as the relative importance of the exact target identity differs between the two tasks it may be that different neural substrates support the representations in each task. In order to determine which was the case we employed a very similar analysis to that used above for the signals relating to target location.

Orientation was the most strongly represented stimulus parameter across all three regions and as seen in Figure 50 when it was significantly encoded during each task alone, it tended to do so in an overlapping manner between tasks. During the sample epoch, both LIP and dIPFC show an overlapping representation of target orientation, dIPFC neurons during the memory delay are also overlapping, and when FEF neurons begin to represent the target orientation during the onset and peri-saccade periods, it does so with a common framework between tasks. Only LIP neurons representing the target orientation during the peri-saccade epoch do not show a significant overlap.

The overlap of populations of neurons representing the identity of the stimulus (the conjunction of both color and orientation) is shown in Figure 51. dIPFC is the only region to show a significant population of neurons encoding the target identity in both visual search and pop-out early in the task. As can be seen in the middle column of Figure 51 the neural substrate encoding target identity during the sample presentation and memory epoch are both significantly overlapping, suggesting a common framework between the two tasks. Similarly, neurons in FEF encoding the target identity in the peri-saccade period also show a significant overlap. However, during this same time period LIP neurons do not show an overlap in the population of neurons encoding the stimulus (and although dIPFC neurons encode the identity during the search task, they do not do so during pop-out).

Color information was only sparsely encoded by the three regions of interest, with only dIPFC showing a significant number of neurons encoding the color of the stimulus and only during visual

Figure 50

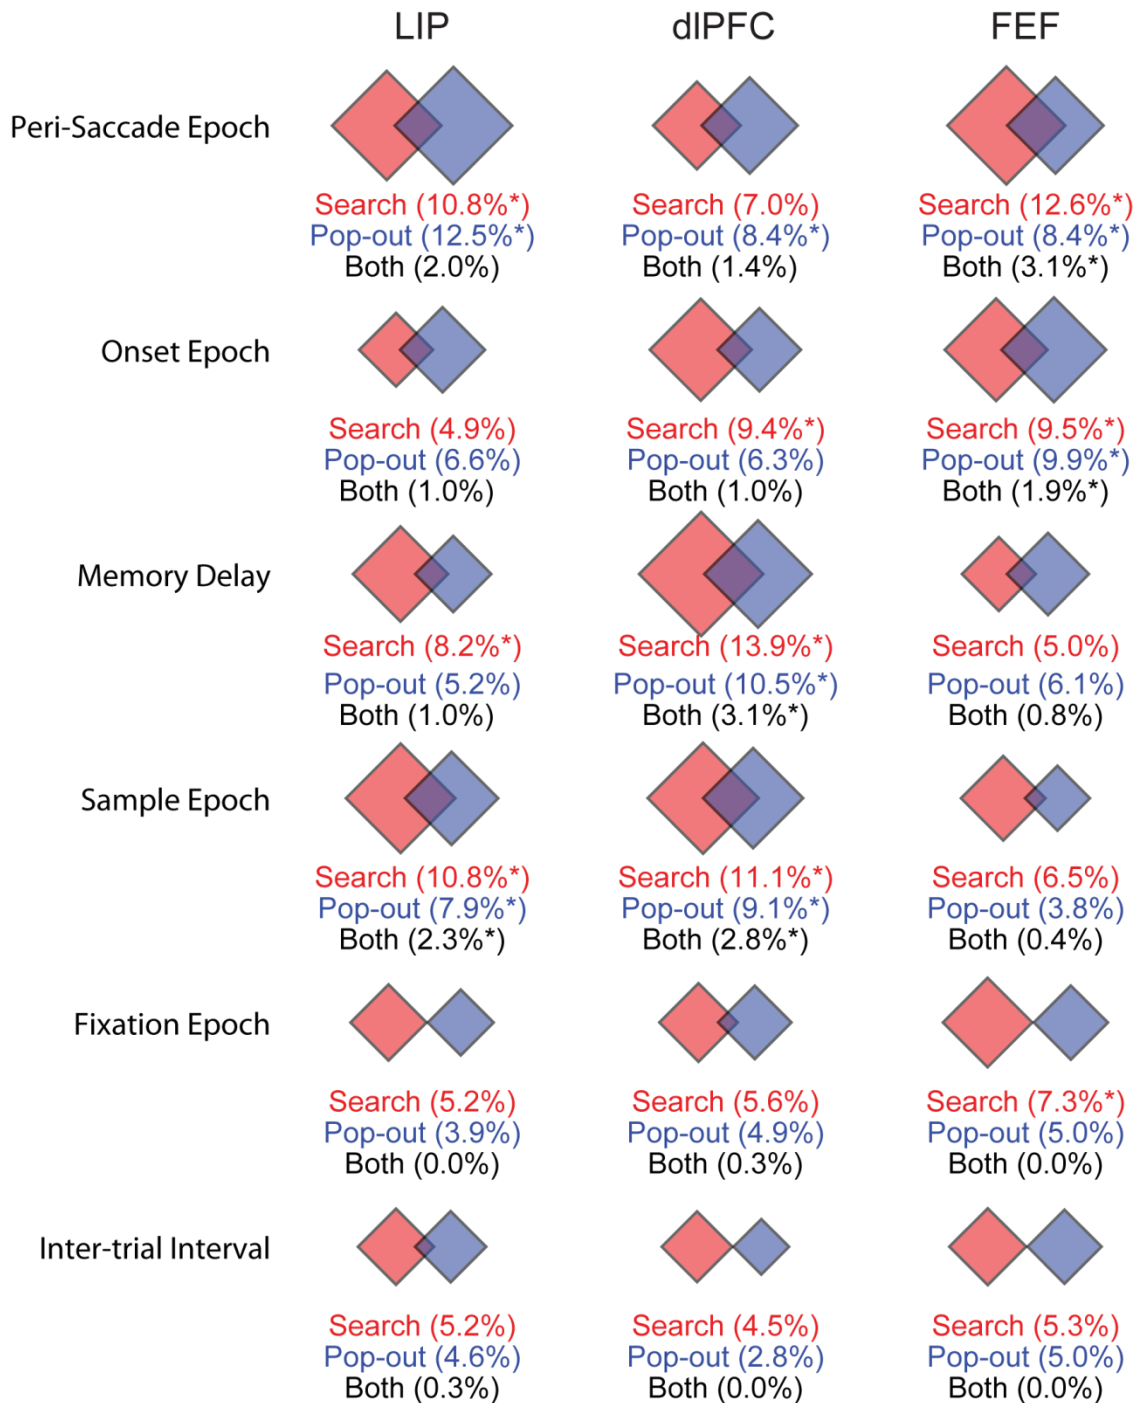


Figure 50. Magnitude and intersection of neural populations carrying information about the **target orientation**. Diamond size indicates the percentage of neurons selective for target orientation during visual search (red) and pop-out (blue). The size of the overlap measures the intersection between the search and pop-out populations of selective neurons. The asterisk denotes a significant proportion of cells carrying sample information in each task or if the intersection of these populations was significant ($p < 0.05$, tested against binomial).

search. Although the total number of neurons encoding color in dlPFC during pop-out is not significant, a significant proportion of the neurons active in visual search are also responsive during pop-out (see Figure 52; this effect is not plotted as significant as it fails our earlier definition).

In general it appears that stimulus information is represented in a fairly consistent manner across tasks. This is in spite of the fact that the manner in which the information about the target stimulus must be used differs between visual search (where the conjunction of both orientation and color information is necessary) and visual pop-out (where the stimulus identity is not needed at all since the array itself selects the target). However, it is important to note that despite the fact that many of the populations of neurons are overlapping between the tasks, the majority of selective neurons are not. This leaves plenty of room for different neural representations to be best utilized for each task.

Figure 51

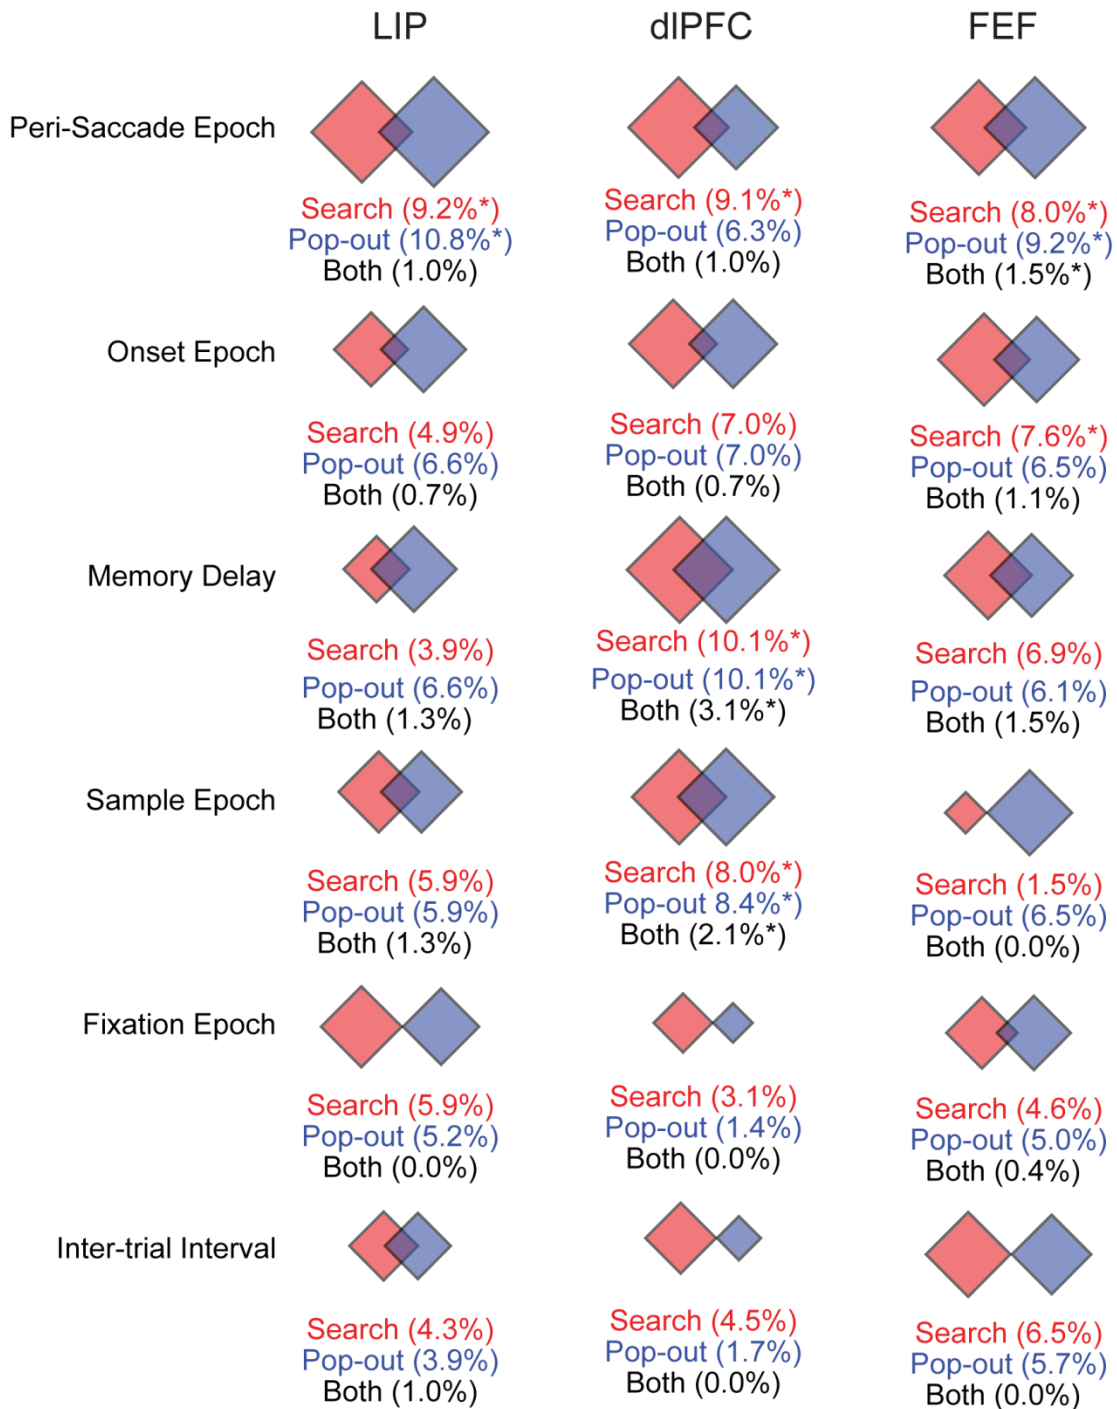


Figure 51. Magnitude and intersection of neural populations carrying information about the **target identity** – the conjunction of color and orientation. Diamond size indicates the percentage of neurons selective for target identity during visual search (red) and pop-out (blue). The size of the overlap measures the intersection between the search and pop-out populations of selective neurons. The asterisk denotes a significant proportion of cells carrying sample information in each task or if the intersection of these populations was significant ($p < 0.05$, tested against binomial).

Figure 52

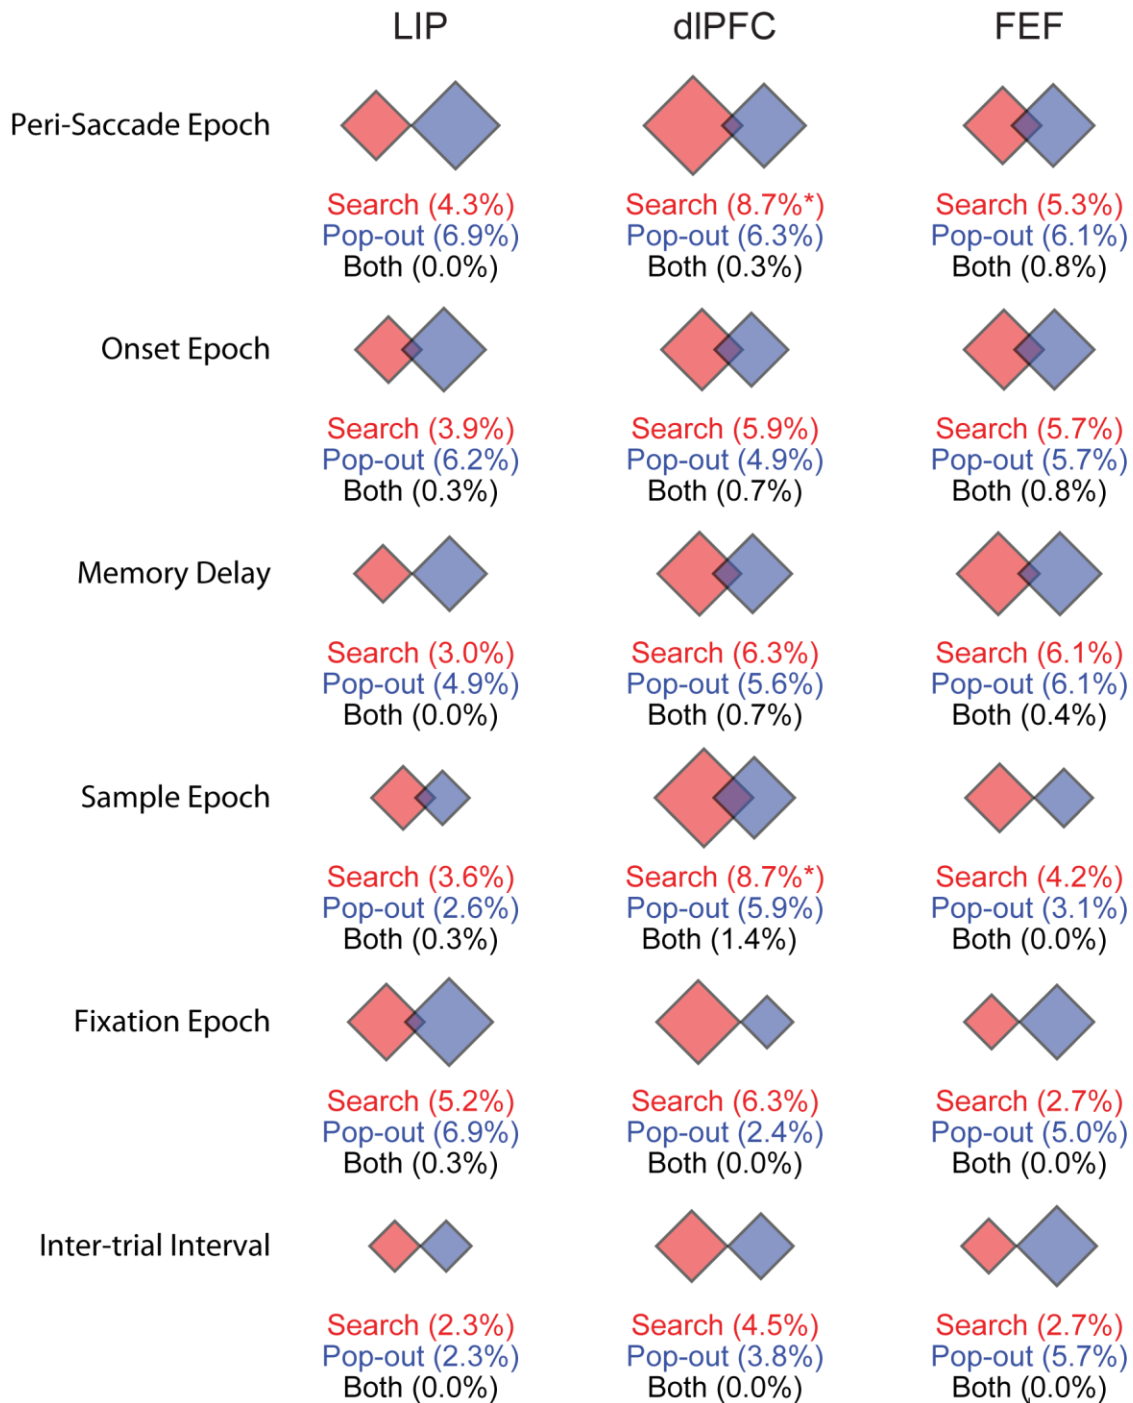


Figure 52. Magnitude and intersection of neural populations carrying information about the **target color**. Diamond size indicates the percentage of neurons selective for target color during visual search (red) and pop-out (blue). The size of the overlap measures the intersection between the search and pop-out populations of selective neurons. The asterisk denotes a significant proportion of cells carrying sample information in each task or if the intersection of these populations was significant ($p < 0.05$, tested against binomial).

Intersection of Attention & Stimulus Information

Up to this point we have demonstrated the existence of neurons that represent the stimulus itself, as well as neurons that carry information about the allocation of attention to the target location. We have also demonstrated that stimulus parameters appear to have a common neural framework supporting their representation, while the allocation of attention between the two tasks is more complex. One final comparison that can be made is to compare the neural substrates representing the allocation of attention with those carrying information about the sample stimulus. As can be seen in Figure 53, no significant overlap was found between the networks carrying information about these two components. This finding suggests that the mechanisms underlying target location are separate from those carrying stimulus information. This result helps to confirm that the target location information found is not a reflection of neurons responding to the specific stimulus parameters of the target. Instead, the target location is likely to be encoded in a more general fashion.

Figure 53




Figure 53. Magnitude and intersection of neural populations carrying information about the **target location (in green) and the sample stimulus (in blue; this included information about color, orientation, or identity)**. Diamond size indicates the percentage of neuron falling into each group across time (forward through the trial from bottom row towards top) and across areas and task. Asterisks denote a significant proportion of cells carrying sample information, target location, or if the intersection of these populations was significant ($p < 0.05$, tested against binomial).

Conclusion

This chapter has investigated the degree of overlap between neural networks representing various parameters related to both tasks. We have shown evidence that neurons selective early in search task are disparate in both FEF and LIP, while dIPFC neurons show a significant overlap. In contrast, late responding, post-saccadic, neurons show a significant overlap between tasks with common direction preferences and response times. The networks representing the stimulus parameters are, for the most part, overlapping, suggesting that the representation of the target stimulus has a common component between tasks. Finally, there was no significant overlap in the networks representing each of these task parameters.

These results provide several important insights into the roles each of these areas plays in performing the behavioral task and how these roles may be represented by the neural network. Neurons in dIPFC are the only ones to show a significant overlap in their representation of the allocation of attention prior to the saccade. This suggests the dIPFC plays a conserved role across both tasks. One possible purpose might be to determine when the correct target has been ‘found’. As it is also the area to show the highest degree of sample stimulus selectivity, it is also well suited to make comparisons between the currently selected target and the remember sample stimulus. In contrast, both FEF and LIP appear to be serving different roles between the two tasks. This is not to say that the underlying computations provided by each region differs between tasks (although it might), rather these computations serve a different purpose within each task. For example, LIP may play a consistent role in the maintenance of a saliency map of the visual field, but show differing effects – during visual pop-out this essentially ‘solves’ the task, while during visual search it might only provide information about how array stimuli relate to the target. All three anatomical areas showed well conserved responses during the post-saccadic period, likely due to the fact that the post-saccadic responses are largely influenced by the eye movement itself. As we attempted to behaviorally restrict eye movements to be equal across tasks to avoid confounds, it is heartening to see this overlap.

The overlap in the representation of the target stimulus suggests a commonality across tasks in the maintenance and use of this information. As the stimuli used as the target were the same across both tasks, it is not surprising to see some representation of that stimulus alone. However, as noted

above, the majority of stimulus selective neurons were only active during a single task. This dichotomy may allow the brain to both maintain a common representation and a disparate one, allowing the brain to manipulate and utilize part of the representation while conserving a basic, common representation. For example, changes in network connectivity during learning can improve visual search behavior without disrupting visual pop-out.

Introduction

Current models of visual attention agree that attention is under the influence of both internal and external forces, however it is currently unclear what form these influences take. As noted in the Introduction, there are currently two models of attention favored in the field – one that suggests bottom-up and top-down interactions occur within a single map, while the other suggests the existence of exogenous parallel mechanisms and endogenous serial mechanisms. Although both models predict that the pop-out task is controlled by external, bottom-up factors, they make very different predictions for the pattern of visual search. We hope to investigate the neural correlates of shifting attention during the visual search task in order to determine whether we can find support for either model. It is important to note that these two models are not mutually exclusive, and in all likelihood the brain uses a combination of the two models to direct attention.

By definition, one prediction of a truly serial search would be that the animal must visit each possible target location in turn during the search. Theoretically, it might be possible to follow this process of serial search by recording from single neurons in an area that directs attention. However, in order to determine whether the deduced location of attention was correct would require some knowledge of where the animal's attention was throughout the task. Fortunately for this analysis, both animals showed a consistent pattern when searching the visual array – both animals searched the array in a clockwise manner. When searching a visual array repeatedly for a target that requires effort to find, it is likely the most efficient approach to adopt a specific search pattern. This not only creates a consistent behavior but also ensures that the animal does not accidentally visit the same location twice, slowing the search process. This consistency not only aids the animal in performing difficult visual searches, but allows us to reliably predict the location of attention. In this chapter we hope to utilize this fact in order to investigate the serial nature of top-down attention.

The premotor theory of attention states that a covert shift of attention to a specific location utilizes the same neural mechanisms as an overt saccadic movement to that location and has gained support from both psychophysical and electrophysiology studies (see Introduction). Here we hope to

test this model by defining the direction of selectivity for each individual neuron in a post-saccadic window and then attempting to decode the locus of attention before the target was found (i.e. pre-saccadic). Since we know the location of attention preceding the saccade (it is at the selected location) we can test our decoding procedure by predicting the location of attention just prior to the saccade. Additionally, we can predict the locus of attention earlier in the trial during both the pop-out and search tasks for each of the three regions. Using this type of analysis, in combination with the clockwise tendency of both animals, we will demonstrate a serial component to visual search evident in the firing activity of FEF neurons that is absent in the other two regions and is absent from pop-out.

Behavioral Evidence for Parallel and Serial Search Mechanisms

Based on the observed reaction times both animals demonstrate a proclivity to search from a preferred target location. Monkey S showed the quickest reaction time to the bottom-right (location 2) for 8 out of the 10 days of recording ($p = 3.0 * 10^{-5}$), while Monkey W appeared to prefer upper right (location 1, 8 out of 15, $p = 0.0042$). However, even without making the assumption that the animal always initiated search in the same location, there may still be a tendency to start at one position or in one hemifield, which would be reflected in an ordering of reaction times to different target locations. Figure 54 shows an example day from Monkey S demonstrating this effect. Monkey S was quickest to respond when the target was at location 2 (lower right), next quickest at the lower left, and so-on. Overall 8 out of 25 days showed a rank order of location preference that suggested a clockwise search order. This is significantly more than expected by chance ($p = 0.016$, by binomial test where chance is 4/24 possible orderings). In contrast, there were not a significant number of days showing a counter-clockwise rank order for search ($N = 2/25$, $p = 0.81$, by binomial) nor were there a significant number of days showing clockwise or counter-clockwise rank order during pop-out ($N = 4/25$, $p = 0.40$ for both, by binomial).

In order to determine whether the observed ordering was more clockwise in nature or more counter-clockwise in nature we performed a cost analysis. Each observed rank-order pattern of target locations was 'mutated' into the closest clockwise and counter-clockwise pattern through a series of swaps between pairs of locations in the order. The cost of such a swap was defined as the reaction time difference between the locations. Therefore, if two target locations showed very similar reaction times

Figure 54

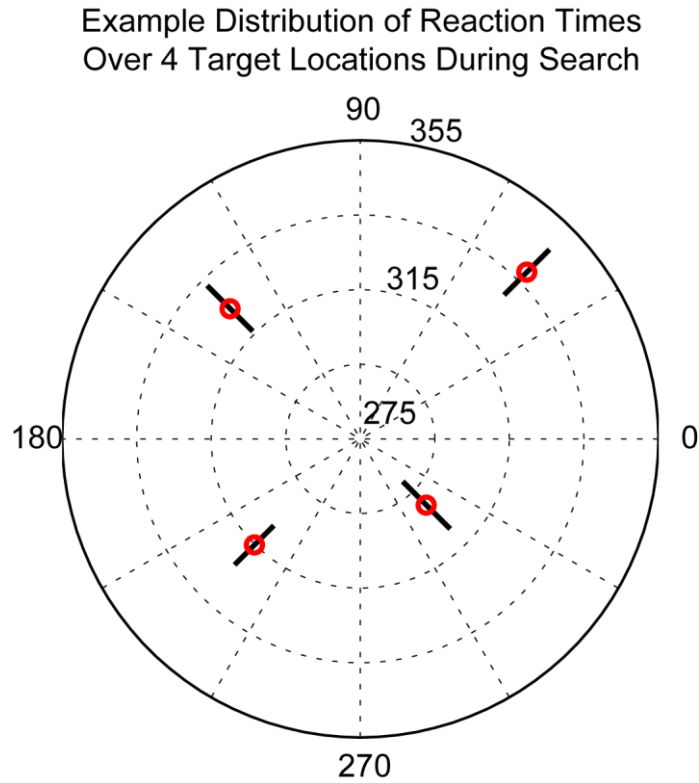


Figure 54. Example reaction times to target at each of the four locations. The mean time to find the target at each of the four possible locations is plotted as a red circle (black lines indicate 95% confidence interval) for a single day of recording from Monkey S (day 4 in Table 3). Monkey S was quickest to respond when the target was at location 2 (lower right), next quickest at the lower left, and so-on in a clockwise manner. This effect is also true for the entire population, as shown in Table 3.

it would cost very little to switch their ordering. In contrast, two locations with disparate reaction times would be very costly to switch. The closest pattern for both clockwise and counter-clockwise search was determined. Table 3 shows the original reaction times for each target location and the rank-ordered location preference, as well as the cost to convert that order into a clockwise pattern or counter-clockwise pattern. As can be seen from the table, a large majority of days show a pattern that more closely resembles clockwise search than counter-clockwise search (16 out of 22 days, with 3 days having equal distances to clockwise and counter-clockwise). This is significantly greater than chance ($p = 0.0085$, by binomial test where chance is 50/50 for clockwise/counter-clockwise patterns). Alternatively, by shuffling the reaction time/location pairings of trials and re-calculating the distance to CW/CCW ordering as above, we can generate a null distribution of how many days were closest to CW versus CCW. The observed difference in days with CW over CCW ordering (10 more CW days than CCW ones) lies significantly outside the null distribution ($p = 0.042$, two-tailed test). Based on these results, we believe both animals were searching the visual array in a fairly consistent clockwise manner.

We also fit several models to the reaction time data for both search and pop-out in order to test what was the most parsimonious description of the distribution of responses. All models were fit using a normal generalized linear model (GLM) and the quality of the fit was described using the Akaike Information Criterion (corrected for small data size, AIC_c). Six different models were tested. The 'Constant' model assumes that there is no difference in reaction time to when the target was at each of the four locations and was used as a baseline (only 1 free parameter). The 'Hemi-field' model allows for a hemi-field effect in the reaction time where each hemi-field (either vertical or horizontal) is allowed a different reaction time (possibly due to damage to the cortex; model has 3 free parameters). The 'Location' model makes no assumptions about the ordering of location reaction times (fitting a distribution to each location's reaction times independently; model has 5 free parameters). The 'CW' model assumes a clockwise ordering to the animal's search pattern – the lowest reaction time location is taken to be the starting point of the search with other locations being n clockwise steps away from the initial stimulus (model has 2 free parameters). The 'CCW' model is the same as the CW model but assumes a counter-clockwise ordering to the animal's search pattern. The "Pop-out" model assumes a CW search pattern, but relies on the fact that each visual search array had a "pop-out" stimulus that differed from the target and distractors in at least one dimension (this model is really only valid under

Table 3

		Reaction Times (Median, ms)				Rank Order Of Reaction Times				Distance to CW (in ms)	Distance to CCW (in ms)	CW/CCW Closest?
		Upper Right	Lower Right	Lower Left	Upper Left							
Monkey S		268.5	261	274.5	276	2	1	3	4	7.5	1.5	CCW
		267.5	258	265	268	2	3	1	4	0.5	5.5	CW
		326	317	313	311.5	4	3	2	1	10.5	0	CCW
		327.5	298.5	309	315	2	3	4	1	0	23	CW
		282	262	270.5	275	2	3	4	1	0	15.5	CW
		283	263.5	277	283	2	3	1	4	0	12	CW
		295.5	289	270.5	278	3	4	2	1	6.5	7.5	CW
		288	266	280	281	2	3	4	1	0	16	CW
		289	259	273	272.5	2	4	3	1	0.5	27.5	CW
		329	290	302	301	2	4	3	1	1	23	CW
Monkey W		249	246	240	242	3	4	2	1	3	2	CCW
		252	247	253	258	2	1	3	4	5	5	--
		261	269	270.5	273.5	1	2	3	4	0	9	CW
		250	252	254.5	265	1	2	3	4	0	9	CW
		255	254	247	240.5	4	3	2	1	7.5	0	CCW
		251	253.5	245.5	254.5	3	1	2	4	4.5	2.5	CCW
		249	249	245.5	249	3	1	2	4	0	0	--
		250	261.5	257	259	1	3	4	2	7	2	CCW
		243	253	248	260	1	3	2	4	5	15	CW
		256	258	258	269	1	2	3	4	0	4	CW
		258	259	239.5	256	3	4	1	2	0	6	CW
		244	249	247	240	4	1	3	2	2	4	CW
		231	240.5	231	236	1	3	4	2	5	5	--
		250.5	267	265	279	1	3	2	4	2	26	CW
		262.5	270.5	269	273.5	1	3	2	4	1.5	7.5	CW

Table 3. Reaction times table of target appearing at each location for each day of recording. The left 4 columns show the median reaction time for the animal to find the target in each of the 4 possible locations. The next 4 columns show the rank order of these reaction times. 8 out of 25 days showed a clockwise rank order. In order to determine the distance to clockwise (CW) or counter-clockwise (CCW) ordering, we estimated the cost (in reaction time) to re-order the observed rank order to a clockwise or counter-clockwise one (described in detail in methods). Whether the rank order is closer to CW or CCW ordering is displayed in the far right column.

the visual search condition since it is the same as the Constant model for the pop-out task, but with an extra free parameter). One possible process underlying the animal's behavior might be that it always began its search from the pop-out position and then moved in a clock-wise pattern (model has 2 free parameters).

As can be seen in Table 4.1, the best fitting model across all of the behavioral data was the CW model for visual search and the Location model for pop-out. The CW model for visual search is consistent with earlier behavioral analyses that suggest the animal's tended to start their search in a fairly consistent location and then shifted their attention around the visual array in a clock-wise manner. We can highlight the CW ordering of the search by comparing the CW and CCW models for each individual day's reaction times during both visual search and pop-out (Table 4.2). There is a significant tendency for the CW model to better fit the visual search reaction time data over the CCW model. However, during pop-out, the two models fit equally well. In fact, for visual search, the CW model is the best fitting model of all 6 tested on a significant number of recording days, with no other model garnering a significant proportion (Table 4.3). The "Pop-out" model is not a good fit to the behavioral data, nor were we able to find any neural effect suggesting such a strategy. As the 'pop-out' stimulus was never the target, it might be that the animals learned to ignore it in order to avoid confusion. The much better fitting model is that the animals started at a fairly consistent location and then search in a clock-wise manner.

Interestingly, there also appears to be a significant location-based effect during visual pop-out, although the majority of that effect is also captured in the hemi-field model. This suggests that the damage with recording (either acute or chronic) tended to increase the animal's reaction time to pop-out stimuli in certain hemi-fields. As the number of electrodes lowered into LIP were slightly higher this might have yielded slightly more damage. Another possibility is that somehow parietal damage is more disruptive to the underlying neural processes and therefore is more disruptive to pop-out.

Table 4

Table 4.1 Comparison of Different Models Averaged Across All Days

	Constant	Hemi-field	Location	CW	CCW	"Pop-out"
Search	2755.1	2751.8	2749.3	2747.7	2753.9	2755.5
Pop-out	2620.5	2539.7	2531.7	2584.8	2587.8	2622.5

Table 4.2 Comparison of CW & CCW Models Across All Recording Days

	# CW Best	# CCW Best	p-value
Search	18	7	0.0073
Pop-out	14	11	0.2122

Table 4.3 Comparison of All Models Across All Recording Days

	Days w/ Best Fit	Percent of Days w/ Best Fit	p-Value
Search:			
Constant	4	16%	0.4063
Hemi-field	4	16%	0.4063
Location	3	12%	0.6184
CW	10	40%	0.0012
CCW	2	8%	0.8113
"Pop-out"	2	8%	0.8113
Pop-out:			
Constant	0	0%	0.9895
Hemi-field	8	32%	0.0157
Location	12	48%	0.0001
CW	4	16%	0.4063
CCW	1	4%	0.9371
"Pop-out"	0	0%	0.9895

Table 4. Table 4.1 shows the corrected AIC (AIC_c) for several different models of behavior (models are described in text). Table 4.2 shows a direct comparison of the clockwise (CW) and counter-clockwise (CCW) models for each day's reaction times during both visual search and pop-out. Table 4.3 shows a comparison of the best fitting models for both visual search and pop-out using all of the possible models.

Neural Evidence for Serial and Parallel Mechanisms in Search

As noted in the Introduction, the premotor theory of attention states that the neural framework that supports the allocation of attention to a particular spatial location is the same framework that encodes the motor response during a saccade. We can leverage this theory by using the post-saccadic time period to define a motor field for each neuron. Based on the premotor theory of attention, a neuron should respond when attention is being driven into its motor field. We can then utilize this fact to attempt to decode the location of the ‘spotlight’ of attention during the task. For all of the analysis used in this chapter, the motor field was determined from a post-saccadic time period was taken to be the first 75 ms after the saccade. As the animal was required to maintain fixation for 150 ms after the saccade in order to receive a reward, this time period was chosen to capture the majority of the saccadic response while avoiding any responses due to reward or post-reward eye movements.

Initially we will take an agnostic view on the exact nature of the search order (although both animals showed a tendency to search the visual array in a clockwise manner based on their reaction times). Instead, we will only assume that an ordering to the search did exist and therefore that the animal was more likely to ‘visit’ a location near the target before it found the target itself. So, we will investigate the activity of FEF neurons when the target was in its preferred direction, when the target fell in either location adjacent to the preferred direction (i.e. either clockwise or counter-clockwise), and, finally, when the target fell in the anti-preferred direction. Figure 55 shows the activity across the entire population of location selective neurons in FEF during both search (top) and pop-out (bottom). For this analysis the activity from each neuron was binned into independent 40 ms bins relative to the saccade (i.e. the first pre-saccadic bin was from 40 ms before the saccade until the time of the saccade). In order to facilitate equal representation of each neuron in the population average, the activity of each neuron was normalized by converting its response in each bin to a z-score based on the response of that neuron, in that time bin, across all target locations. This also facilitates the determination of significance, as we can test our observed population average against that of a standard normal distribution with zero mean. The z-score of this difference is what is plotted in Figure 55. Significance across target locations within each time bin was also determined using an ANOVA. A significant effect across target location was marked with a black circle in each bin ($p < 0.05$, corrected for multiple comparisons across time). Post-saccadic responses are significant due to the selection procedure.

Figure 55

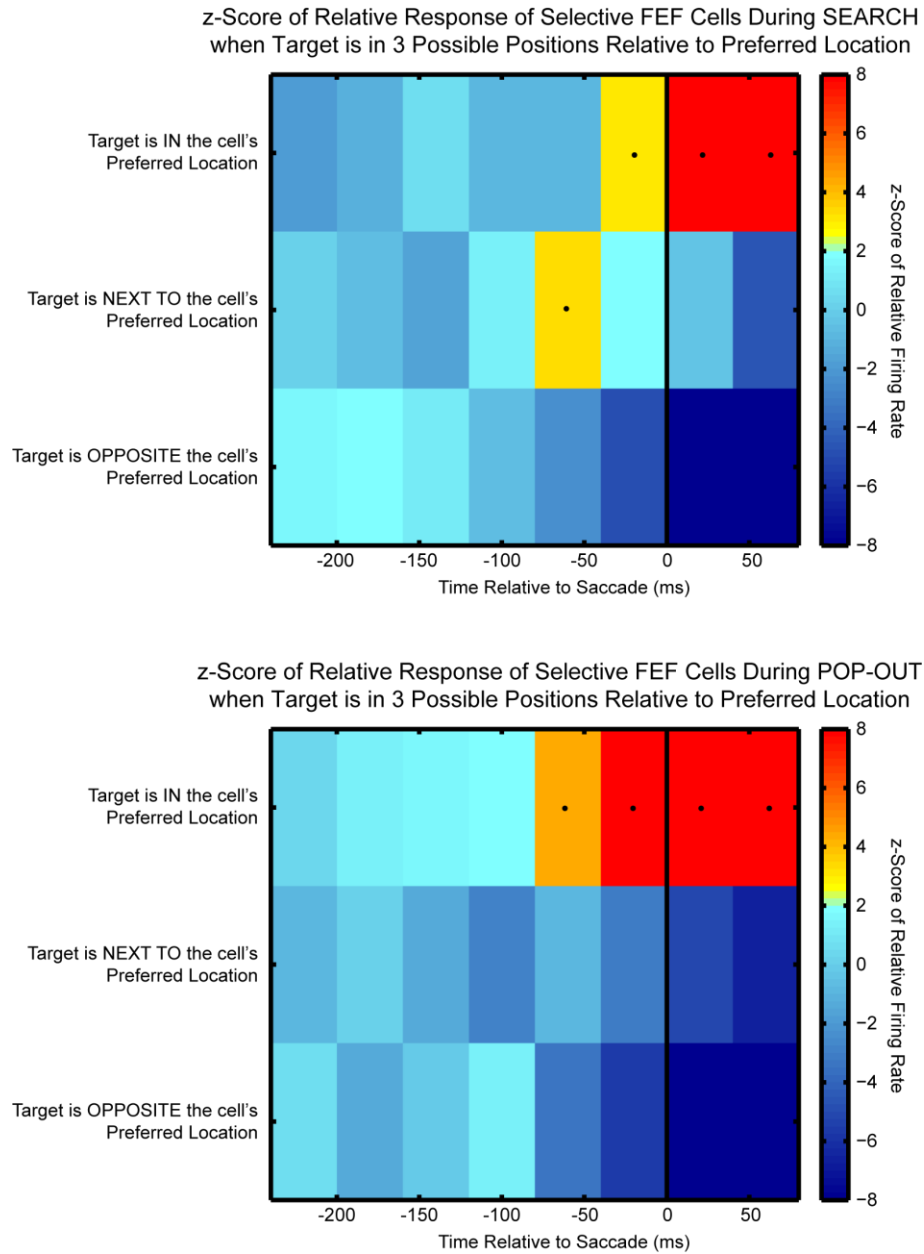


Figure 55. Average, normalized firing rate over time for location selective neurons in FEF during search (top) and pop-out (bottom). Trials within each task are sorted by the *displacement* of the target position from that of the neuron's preferred target location (which is defined in a 75 ms post-saccadic period). Activity was analyzed in independent bins of 40 ms. Color indicates the z-score of the average response above chance (which is 0 since the activity is relative to the overall average). Dots indicate when the activity across bins was significant by ANOVA at $p < 0.05$ after Bonferroni correction for multiple comparisons.

As can be seen in Figure 55 (top), an interesting pattern of FEF activity emerges before the saccade as the animal performs a visual search. Just prior to the saccade (in the time bin from 40 ms to 0 ms before the saccade), FEF activity is strongest when the target is located in the neuron's preferred direction. Presumably this reflects the process of directing attention to, and then making a saccade to, the target. This response carries into the post-saccadic region as this is the time period we use to define each neuron's preferred direction. This is similar to the effect shown in an earlier chapter using mutual information. Interestingly, the FEF activity in the time period 80 to 40 ms before the saccade is strongest for when the target location is adjacent to that neuron's preferred target location. As the target actually lies adjacent to this location, and the eye movement will be made to the target at a later time point, this activity does not reflect a motor response but rather is the allocation of attention into that neuron's receptive field. As the animal performs an ordered search through the visual array (as supported by our behavioral results) attention is directed into the receptive field of different populations of neurons. As the animal makes its way towards the target location, due to the ordered search, it is more likely to visit one of the adjacent locations before finding the target. Therefore, a neuron that has a receptive field at one of the adjacent locations must activate in order to direct attention into that spatial location. The response in the 80-40 ms pre-saccade time period reflects this direction of attention. This pattern of an ordered search continues even earlier in the trial, as there is a (non-significant) tendency for the population to represent the location opposite to the target location approximately 180 ms before the saccade. Although this effect is weaker and more temporally diffuse than the one observed for adjacent target locations, this trend further extends the pattern of the search backwards in time.

A different pattern of results is observed for visual pop-out. Where-as we see an activation for adjacent locations earlier in the trial during search, FEF neurons only respond to the target location during visual pop-out. This makes sense as FEF is receiving input from both LIP and dIPFC with information about where the target lies. Therefore attention has already been drawn to that location, the target has been identified (and so no more shifts of attention are needed), and the animal is ready to make its behavior response.

Figure 56




Figure 56. Plot of average activity in FEF neurons in two pre-saccadic bins across reaction times. The bins were chosen based on the selectivity from Figure 55. The orange line shows the effect of reaction time on the average FEF activity just prior to the saccade when the target is in the neuron’s preferred direction. In contrast, the blue line plots the average FEF activity in a 40 ms window start 80 ms before the saccade when the target is in one of the two locations adjacent to the neuron’s preferred location. A significant correlation was observed between the strength of activity in the blue bin and reaction time, but not for the orange bin and reaction time.

Strength of Serial Search and Reaction Time

If the pattern of activity observed during visual search reflects a true serial component to visual search then we can make a strong prediction – a larger relative response in the 80-40 ms time bin should correlate with a longer reaction. This prediction relies on the fact that the activity in the 80-40 ms time bin is actually drawing attention away from the target location, and therefore ‘impeding’ the animal in finding the target somewhere else. This is most likely an effect that exists across trials – longer reaction time trials will, in general, require a longer serial search, and therefore will, in general, be more likely to visit adjacent locations prior to the target location. However, as reaction times speed up, it becomes increasingly likely that the animal will directly find the target location without requiring a serial search. Therefore, shorter reaction time trials should correlate with reduced activity in the 80-40 ms bin. In contrast, as the animal must always attend to, and initiate an eye movement to, the target location, the pre-saccadic response of neurons to the target in its preferred location should not vary with reaction time.

As shown in Figure 56, this is exactly the pattern of results observed. The blue line shows the (un-normalized) response of FEF neurons in the time period 80 to 40 ms before the saccade when the target is adjacent to its preferred location. The orange line shows the average response of the same FEF neurons during the 40ms just before the saccade on trials during which the target was in the neuron’s preferred direction. Both averages are plotted over four reaction time groupings. Reaction times were binned into four percentile groups, 10-30%, 30-50%, 50-70%, and 70-90%. The fastest and slowest 10% of trials were left out in order to avoid outliers.

There was a strong (significant) correlation between reaction time and the strength of the relative response in the 80-40 ms time bin ($\rho = 0.97$, $p = 0.02$, Figure 56). Therefore, as the relative response increases, drawing attention away from the target location, the reaction time of the animal also increases. In contrast, there is not a significant correlation between the response of neurons 40 ms before the saccade ($\rho = 0.12$, $p = 0.87$). This reflects the necessity of attending to, and selecting, the target location before deciding and making a behavioral response regardless of the animal’s reaction time.

Figure 57

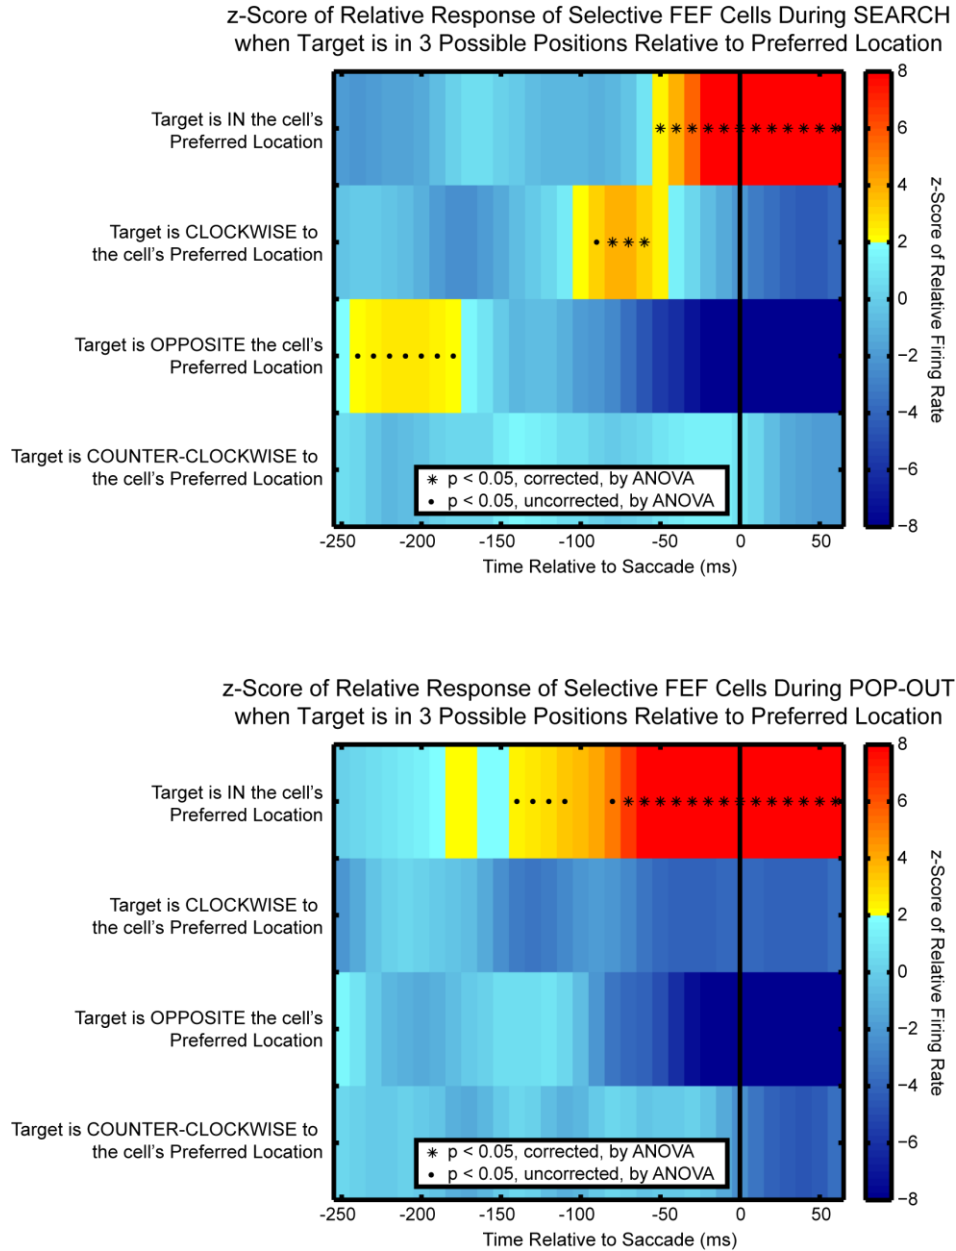


Figure 57. Average, normalized firing rate over time for location selective neurons in **FEF** during search (top) and pop-out (bottom). Trials within each task are sorted by the location of the target relative to the neuron’s preferred target location (which is defined in a 75 ms post-saccadic period). Color indicates the z-score of the average response above chance (which is 0 since the activity is relative to the overall average). Asterisks indicate when the activity across bins was significant by ANOVA at $p < 0.05$ after Bonferroni correction for multiple comparisons, while dots indicate an uncorrected $p < 0.05$.

Neural Evidence for Clockwise Search

Initially, we examined the neural response over time of FEF neurons during both visual search and visual pop-out for three conditions: when the target was in the neuron's preferred direction, when the target was adjacent to the neuron's preferred direction, and when the target was opposite to the neuron's preferred direction. This made no assumption about the directionality of the order search. However, based on the behavioral evidence, during the visual search task both animals not only searched the visual array in a organized manner, but they seemed to search in a clockwise manner as well. In order to determine whether the neural response follows with the clockwise nature suggested by the behavior we can perform a very similar analysis but now 'unwrap' the adjacent condition into two conditions – one where the target was clockwise from the preferred location and one where the target was counter-clockwise to the target location. Figures 57, 58, and 59 show the relative response of location-selective neurons from FEF, dlPFC, and LIP, respectively, during both the visual search (top) and visual pop-out (bottom) tasks. Similar to previous analyses, these figures show the relative difference across target locations (relative to each neuron's preferred target location). Bins of 40 ms were used, but instead of being taken independently as before, the window was slid in steps of 10 ms across the entire time period. In this way we can more accurately capture the temporal dynamics of the response. Significant differences from a zero response is shown in the color of each bin (although the number of multiple comparisons is larger in this analysis, so higher z-scores are much more relevant). Significance within each time bin was determined using an ANOVA across the four possible target locations. Black stars indicate significant responses ($p < 0.05$) after a simple Bonferroni correction for multiple comparisons, while black dots now indicate significant responses at an uncorrected level ($p < 0.05$, uncorrected).

As was seen in Figure 55, FEF activity is ordered during visual search even when the adjacent locations are unwrapped (Figure 57). The ordering now follows in a clockwise manner, suggesting that the animals were performing a serial, clockwise search and that activity in FEF is involved in the direction of attention into that neuron's receptive field. Because of the higher temporal resolution, we can see that the shift between FEF activity representing the target location or FEF activity maximally representing the clockwise position occurs at approximately 50 ms. The neural activity representing this location (presumably directing attention to the clockwise stimulus) lasts for approximately 50 ms. Furthermore, now the response of FEF neurons to directing attention opposite to the target location is

significant (although at an uncorrected level -- it is still weaker and more temporally diffuse than the earlier locations). It is likely that the weakness of this response is due to the lower probability of the animal visiting three possible target locations (as it searches more locations it becomes increasingly likely to find the target thereby ending the trial). During visual pop-out FEF neurons did not show a sequential ordering pattern, but rather only showed significant increases in activity when the target lay in its preferred direction (Figure 57, bottom). At the population level significant differences are first observed at approximately 140 ms before the saccade and reach significant levels 70 ms before the saccade.

Figure 58, top, shows the average activity of dIPFC neurons during the visual search task. Unlike the activity of neurons in FEF, neurons in dIPFC are only responsive when the target stimulus is in its preferred location. This suggests these neurons do not show activity following the 'spotlight' of attention as do the FEF neurons. Instead, the dIPFC neurons begin to show activity significantly deviating from background at approximately 70 ms before the saccade. Although this is earlier than the relative response of FEF neurons we cannot make a direct comparison in timing using this measure as it is only a measure relative to the neuron's response across all target locations. As FEF activity is obviously stronger in the clockwise position for some time periods, any neural response to the target in its receptive field will be reduced when doing relative comparisons. Therefore, in order to determine the order of response, one must do cross-correlations on simultaneously recorded dIPFC and FEF neurons with similar preferred locations when the target was in their preferred location (see below). Similar to FEF, dIPFC activity during visual pop-out does not show an ordered effect, only significantly responding when the target was in the neuron's preferred location (Figure 58, bottom). Average activity across dIPFC neurons began to significantly diverge at approximately 140 ms before the saccade.

Performing a similar analysis is more difficult on LIP neurons as they were not responsive during visual search until very late in the trial (see earlier chapter), making it difficult to reliably determine a preferred target location and also making any pre-saccadic inferences difficult. Figure 59 shows the activity of LIP neuron during visual search (top) and visual pop-out (bottom) utilizing the same method as was employed for FEF and dIPFC. Results for both tasks are similar to the results of dIPFC – LIP neurons only significantly diverge from baseline levels in the preferred location. During search LIP

Figure 58

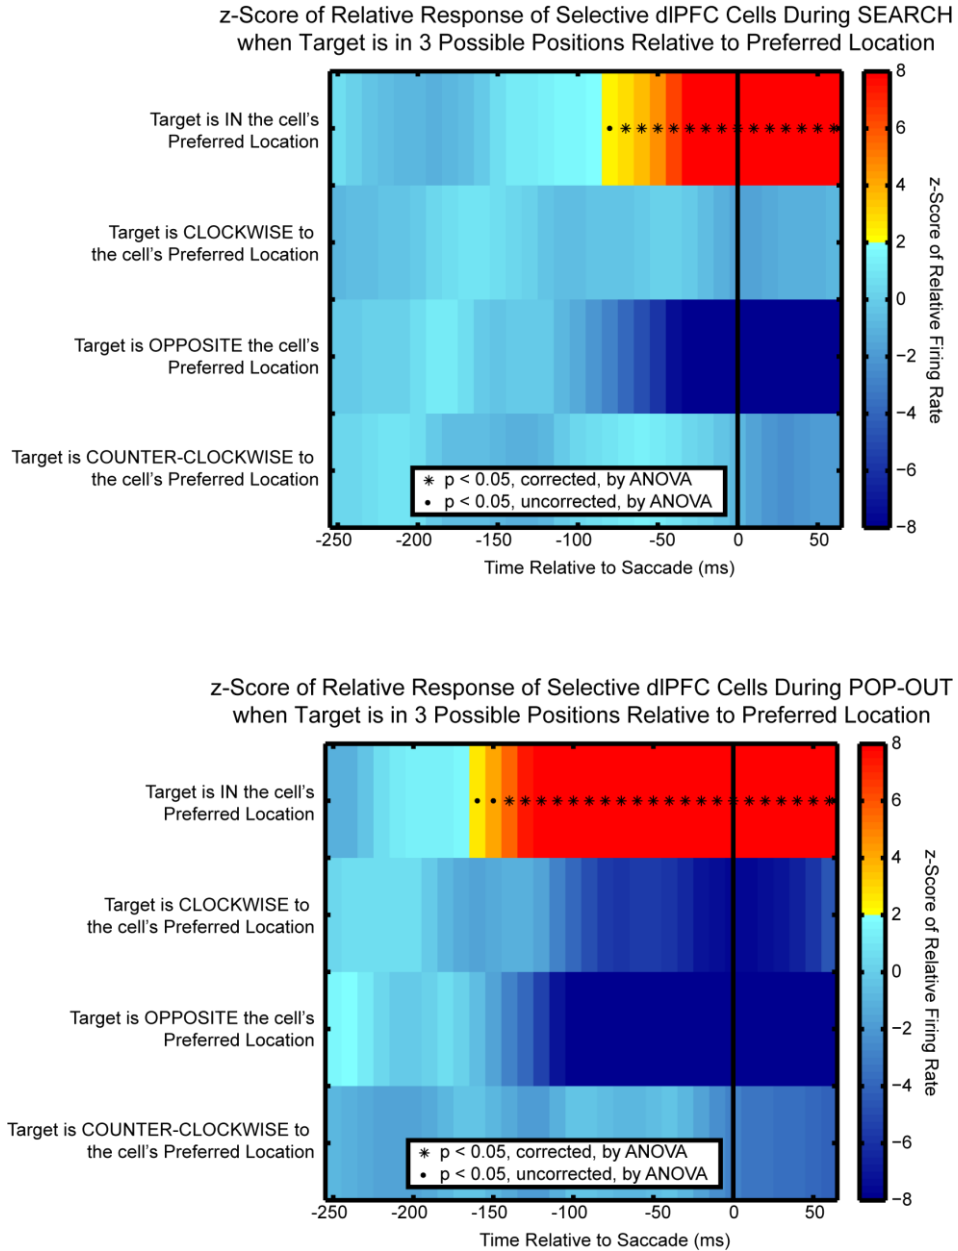


Figure 58. Average, normalized firing rate over time for location selective neurons in **dIPFC** during search (top) and pop-out (bottom). Trials within each task are sorted by the location of the target relative to the neuron's preferred target location (which is defined in a 75 ms post-saccadic period). Color indicates the z-score of the average response above chance (which is 0 since the activity is relative to the overall average). Asterisks indicate when the activity across bins was significant by ANOVA at $p < 0.05$ after Bonferroni correction for multiple comparisons, while dots indicate an uncorrected $p < 0.05$.

Figure 59

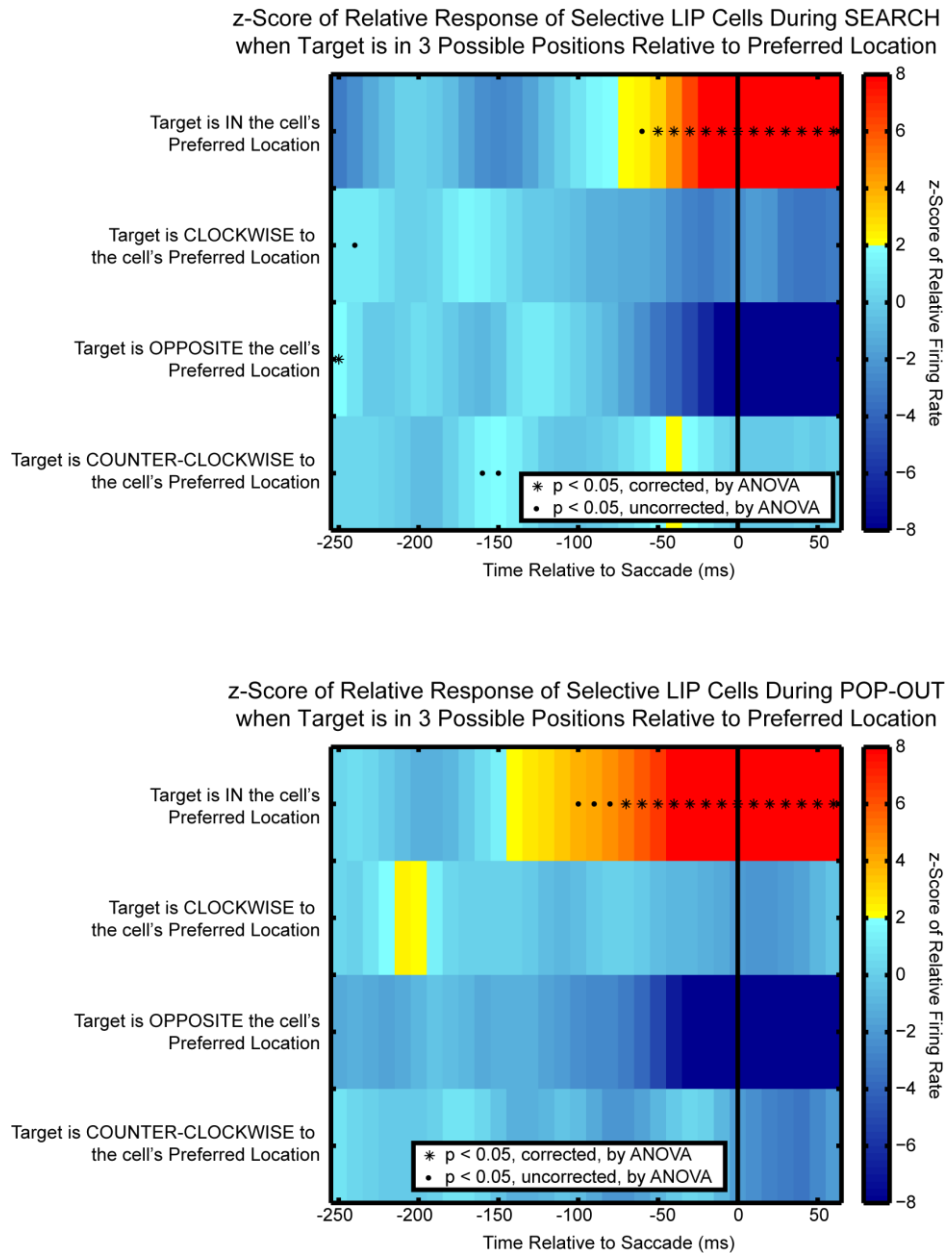


Figure 59. Average, normalized firing rate over time for location selective neurons in LIP during search (top) and pop-out (bottom). Trials within each task are sorted by the location of the target relative to the neuron’s preferred target location (which is defined in a 75 ms post-saccadic period). Color indicates the z-score of the average response above chance (which is 0 since the activity is relative to the overall average). Asterisks indicate when the activity across bins was significant by ANOVA at $p < 0.05$ after Bonferroni correction for multiple comparisons, while dots indicate an uncorrected $p < 0.05$.

neurons diverged at approximately 50 ms before the saccade while during pop-out neural activity diverged at around 70 ms before saccade. It is important to reiterate the difficulty in using this analysis to determine the relative latency of location selectivity across regions. The previously used mutual information analysis is more appropriate for asking that question. The relative neural activity highlights differences between neuronal responses at different target locations and since it is an average response across the entire population it is not good at capturing the exact temporal dynamics of individual neurons, especially when they are non-homogeneous. Another important difference is that the mutual information analysis employed earlier makes no assumptions about how a neuron encodes the target location, whereas this analysis makes a strong assumption that neurons respond preferentially to a single location and do so by increasing their firing rate. Finally, as this analysis uses larger time windows as well as averaging across the entire population, it will tend to under-represent transient selectivity, such as was seen in LIP.

Serial Search on a Single Neuron Level

So far our analysis has been restricted to the population level. While the effects described are strongly suggestive of a role of FEF in serially directing attention to specific locations, one could hypothesize that the effects are merely a manifestation of some strange interactions at the population level. To address these concerns and to bolster our results, we can also examine the responses at the single neuron level.

In order to accomplish this, we will use the response of individual neurons across the four target locations (again relative to the preferred target location) in each time bin to define a ‘vector’ of response. For example, if for a particular time bin the neuron responds only when the target is in its preferred direction then the vector of that response will point in the direction of the preferred location. Essentially we are using the activity of each individual neuron to define the response vector within each time window. From this we are able to determine whether the population of response vectors is non-uniform, and therefore has a significant direction, and what that direction is. Figure 60, shows the average angle, relative to a neuron’s preferred direction, across the population of location-selective units from FEF (top), dIPFC (middle), and LIP (bottom). The mean angle direction is plotted in red for search and blue for pop-out with the shaded regions indicating the 95% confidence interval about the

mean angle when the population shows a significant direction preference (as measured by a Rayleigh test, see Methods). The raw response of each neuron was used to determine the vector of response in the same time bins as were used for the 'unwrapped' analysis above.

Similar to the results from previously presented, Figure 60 shows that average vector of individual FEF neurons shows a clockwise pattern. There is a significant clustering of vectors during search in the direction of the preferred target location until approximately 40 ms before the saccade. At this point the population vector drifts towards the clockwise position, where there is a significant clustering of angles in the clockwise direction from about 70 to 90 ms. This is followed by one last significant clustering in the anti-preferred direction. Again, similar to previous results, individual FEF neurons do not show an ordered effect during visual pop-out. Rather they cluster in the preferred direction until about 80 ms before the saccade.

Results from dIPFC and LIP neurons also reflected the population level analysis. Neither region clustered significantly about any direction other than the preferred direction (Figure 60, middle and bottom figures). During visual search, dIPFC neurons showed significant clustering approximately 70 ms before the saccade, while LIP neurons were slightly later at 50 ms before the saccade. During visual pop-out, dIPFC neurons showed significant clustering about their preferred location at 130 ms prior to the saccade, while LIP neurons were later at 80 ms before the saccade. Similar to above, latency measurements are difficult to interrupt with this method, with previous analyses tailored for that question.

One last noteworthy result from this analysis is the success of post-saccadic clustering towards the preferred direction despite the fact that neurons were not pre-selected for specific receptive fields overlapping the target location. This bolsters our argument that we recorded from a wide variety of neurons each with different response properties and that this diverse sampling provided good coverage of the available neuron types.

Figure 60

Average Angle of Neural Response, Relative to Preferred Direction




Figure 60. The mean angle of the distribution of individual neuron’s represented location, over time. The average angle is plotted for all three regions and for both visual search (red) and pop-out (blue) trials. Shaded regions indicate when the 95% C.I. of the average angle and are only present when the distribution of individual angles was significantly clustered.

Serial Search on a Single Trial Level

In addition to looking at the single neuron level for the effect of a serial, clockwise search we can leverage our simultaneous recordings to examine the relative spiking activity of simultaneously recorded neurons on single trials. If neurons have preferred target locations offset by one clockwise step from one another, and our model of a shifting spotlight of attention is correct, then one could hypothesize that we should see an ordering in our neuron pair's activity – the clockwise neuron should fire *after* the counter-clockwise neuron. The cross-correlation is used to test for this effect.

Pairs of simultaneously recorded FEF neurons with preferred directions offset by one target location step (90°) were compared. The neuron with the clockwise preferred direction was taken to be the primary neuron and only trials in which the target was at the primary neuron's preferred direction were analyzed. This would ensure the animal 'visits' the secondary neuron's preferred location before the primary's preferred direction (see inset of Figure 61 for diagram). The spiking activity from a peri-saccadic period starting 100 ms prior to the saccade and ending 50 ms after the saccade was used. This was chosen to maximally capture the time period during which we have previously found the attention shift occurring. All spiking activity was binned into 5 ms bins, and a shuffle-corrector was used in order to remove any saccade-locked covariations in response probability.

The red line in Figure 61 shows the cross-correlation between spiking activity in pair of neurons in FEF that have preferred target locations offset by one clockwise step. There is a significant positive correlation at +40 ms. ($p = 0.012$, two-tailed t-test against no correlation), demonstrating that FEF neurons with a preferred direction that is clockwise to the preferred direction of another FEF neuron will show neural activity that *follows* the other neuron. This supports our previous findings at the population and single neuron level. There is also a significantly negative correlation between these pairs of neurons at +65 ms ($p = 0.0425$, two-tailed t-test against no correlation). The blue line shows the cross-correlation between selective neurons in FEF and dIPFC when, again, those neurons had the same preferred target location and for trials in which the target was in that direction. During visual pop-out, no significant correlation between neurons with offset preferred directions were observed. This supports the previous results demonstrating the parallel nature of visual pop-out. Similar results are also observed for time periods relative to visual array onset.

Figure 61




Figure 61. Shuffle-corrected cross-correlelogram between FEF neurons with preferred direction's offset by 90° clockwise. Correlation was calculated in a 150 ms window starting 100 ms before the saccade for both search (red) and pop-out (blue). Only trials in which the animal correctly found the target in the preferred direction of the primary neuron were used. Temporal offsets > 0 indicate the activity of the secondary neuron proceeded the primary neuron, while offsets < 0 suggest the primary neuron leads the secondary one. Significance is indicated with an asterisk and was tested against zero with a two-sided t-test. The only significantly higher correlation occurred at +40 ms. This suggests the activity of a neuron follows the activity of counter-clockwise neurons by 40 ms.

FEF leads dIPFC during Search

An interesting question is whether FEF or dIPFC is directing attention during visual search. In our initial analysis of information about the target location, the latency for both regions was at approximately 40-50 ms before the saccade (see Figure 8). Although information about the locus of attention was slightly earlier than dIPFC during visual search, the difference was not significant, making it difficult to assign a direction of information flow. In contrast, the work presented in this chapter would suggest that dIPFC responds to a target in its preferred location earlier than FEF does. However, as noted above, since the measurement is relative, this is a difficult conclusion to draw from that analysis. In order to accurately determine the latency difference between the two areas, one must perform a cross-correlation analysis on simultaneously recorded neurons with shared location preferences. Since dIPFC neurons do not seem to respond reliably unless the target is in its preferred direction, only trials for which the target lies in that direction are used. Figure 62 shows the cross-correlation between all pairs of simultaneously recorded neurons in FEF and dIPFC on trials in which the target lay in their shared preferred direction. The cross-correlation was done in a 300 ms time window from 50 ms after the onset of the visual array in order to capture the entire process of finding the target location. As before, spiking activity was binned into 5 ms bins. The shuffle-corrector was used in order to remove any stimulus-locked covariations in response probability.

The red line in Figure 62 shows the cross-correlation between spiking activity in FEF and spiking activity in dIPFC when the target was in their shared preferred direction. There is a significant correlation at approximately -25 ms. ($p = 0.031$, two-tailed t-test against no correlation), suggesting that FEF is leading dIPFC during search. No other temporal offset show a significant correlation in either a positive or negative direction. The blue line shows the cross-correlation between selective neurons in FEF and dIPFC when, again, those neurons had the same preferred target location and for trials in which the target was in that direction. During visual pop-out, the only significant correlation occurred at +80 ms ($p = 0.041$, two-tailed t-test against no correlation), suggesting that dIPFC leads FEF by 80 ms. This results fits well with the offset determined by the mutual information analysis (85 ms, see first chapter) and with the difference in a significant effect for the normalized population response (70 ms, see above).

Figure 62

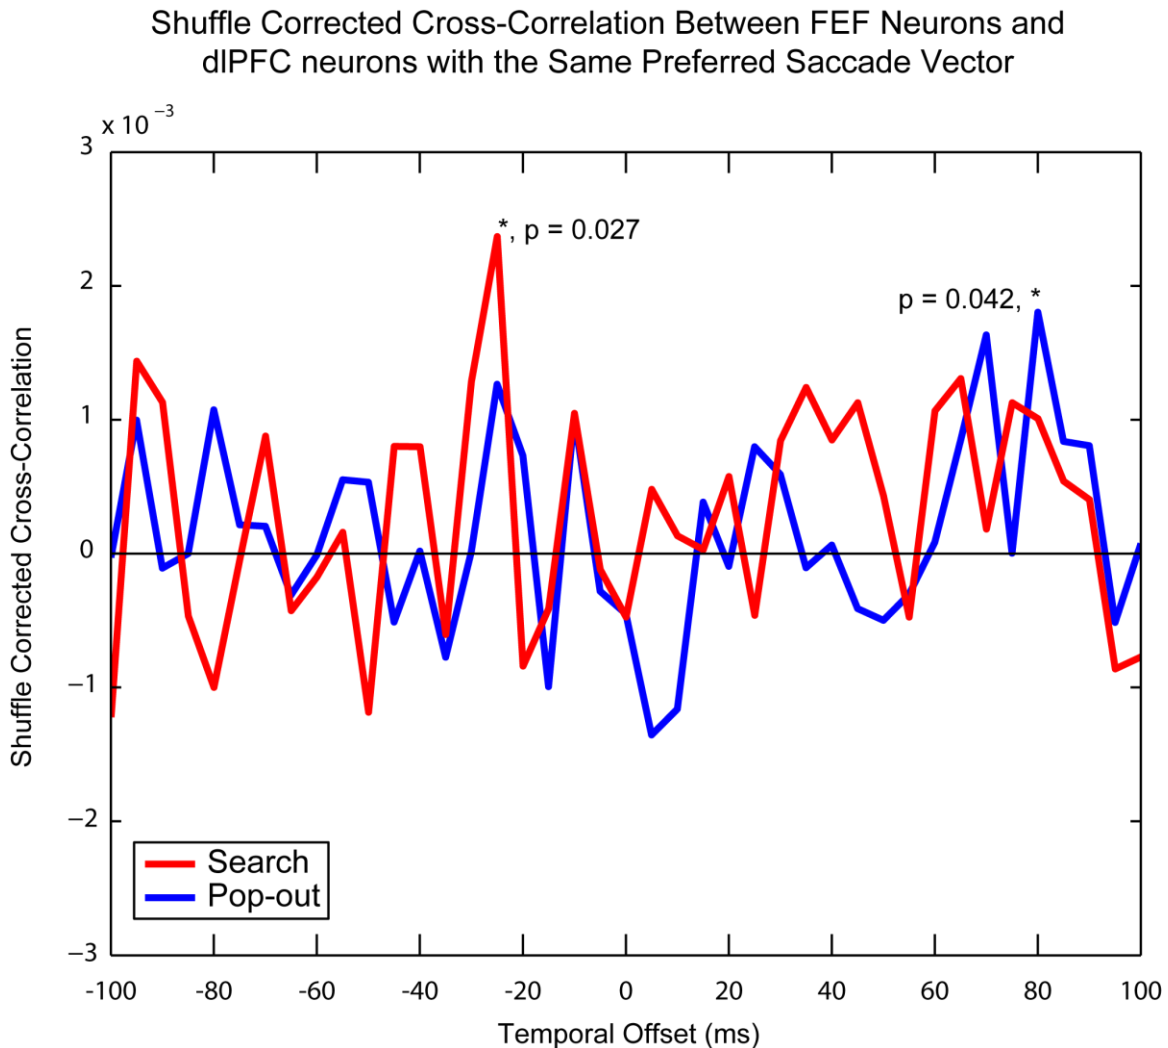


Figure 62. Shuffle-corrected cross-correlogram between FEF and dIPFC neurons with overlapping preferred directions. Correlation was calculated in a 300 ms window starting 50 ms after visual array onset for both search (red) and pop-out (blue). Only trials in which the animal correctly found the target in the preferred direction of both neurons were used. Temporal offsets > 0 indicate the activity of FEF neurons follow that of dIPFC neurons, while offsets < 0 suggest FEF activity precedes dIPFC. Significance is indicated with an asterisk and was tested against zero with a two-sided t-test. The only significantly higher correlation during visual search occurred at -25 ms. This suggests that during search FEF activity representing the target location proceeds that of dIPFC by 25 ms. However, during visual pop-out a significant increase in correlation was observed for +80 ms. Therefore, during pop-out, the effect reverses and dIPFC neurons proceed FEF neurons by 80 ms.

Psychophysical Estimate of the Time to Shift Attention

Using the psychophysically estimated cost of adding a distractor to the search array during visual search it is possible to estimate the time to shift the attentional spotlight from one putative target to another. We can model the time to search the visual array as coming from two sources, one fixed and one variable. The fixed component of the reaction time is due to things that stay consistent across the task – visual perception, decision making, initiation of eye movement, etc. The variable reaction time component is associated with the cost of actively finding the target – as the target is located randomly in the array the animal has a fixed chance of finding the target on each subsequently attended location. If we assume a strong inhibition of return, then we can directly model the observed reaction time as a combination of fixed reaction time and the scaled cost of shifting attention when there are n stimuli to search through:

$$RT_n = RT_{fixed} + \frac{1}{n} \sum_{i=1}^n i RT_{att}$$

The cost of adding a distractor to the search array is therefore:

$$\Delta RT = RT_n - RT_{n-1} = R_{att} \left[\frac{1}{n} \sum_{i=1}^n i - \frac{1}{n-1} \sum_{i=1}^{n-1} i \right] = R_{att} \left[1 - \frac{1}{n(n-1)} \sum_{i=1}^{n-1} i \right] = \frac{R_{att}}{2}$$

Using the reaction time cost derived from our psychophysical experiments (22 ms/item), we can estimate the time needed for a shift in attention to be roughly 44 ms/item. This matches well with the observed switch time between the representation of the counter-clockwise target to the clockwise target across the population (Figures 55 and 57), as well as the observed within-trial cross-correlation offset of 40 ms (Figure 61).

Controlling for Eye Movement Differences

We have presented data suggesting that both of our animals searched the visual search array in an ordered, clockwise manner. Furthermore, we have demonstrated that the neural activity in FEF directs the spotlight of attention from locus to locus before finding the target. However, one potential confound with our observance would be if the eye position varied across the task in an ordered manner, perhaps following where the animal is currently attending. For example, the animal's eye might have a tendency to drift towards a particular stimulus as the animal is attending to it. FEF activity might then

Figure 63

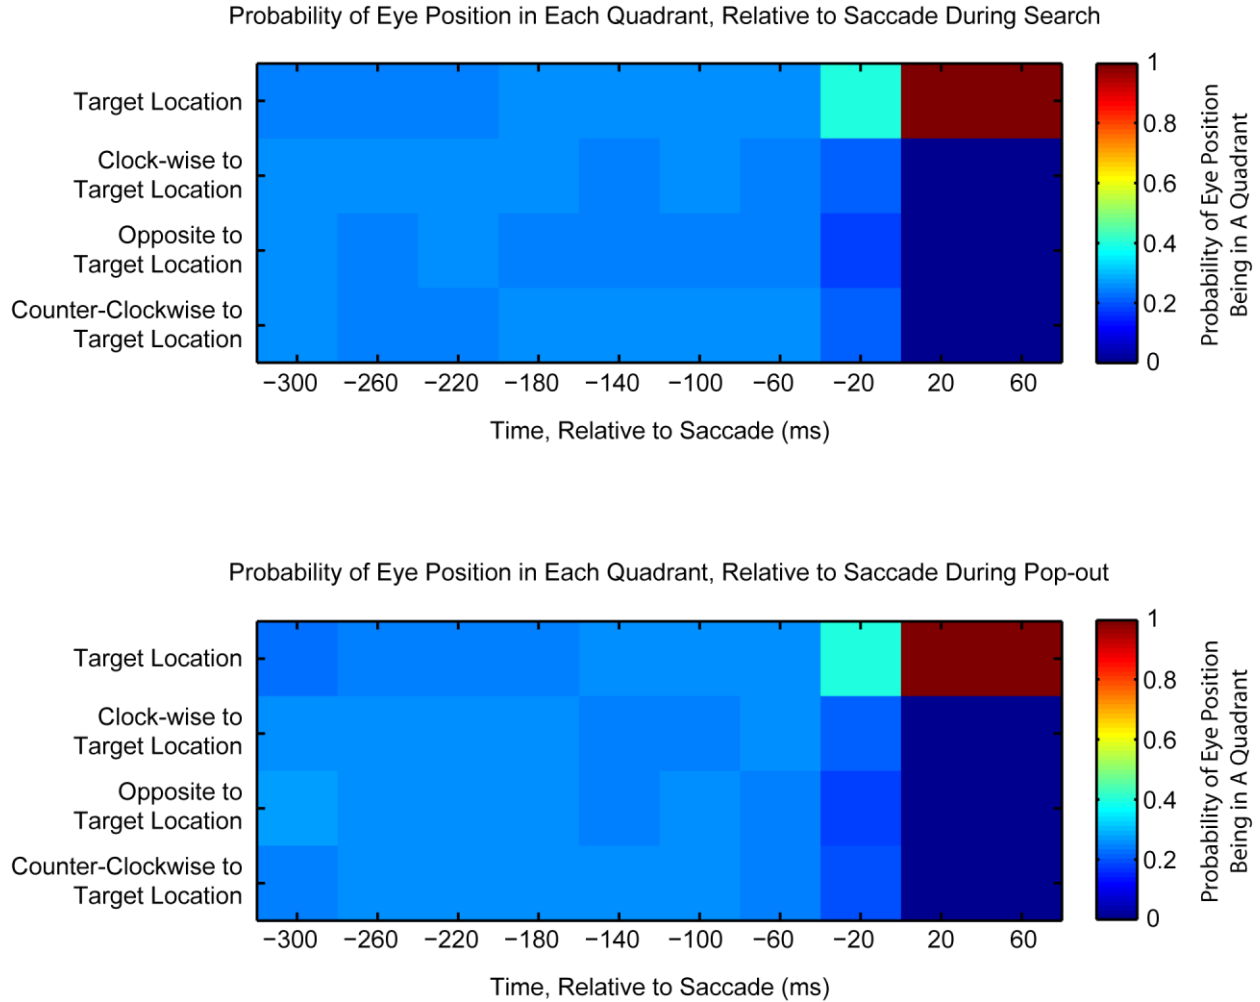


Figure 63. The probability of the eye falling in each quadrant, relative to the target location, over the entire trial. The probability is shown for both visual search (top) and pop-out (bottom). Unlike with the neural activity (Figure 57), the eye position did not show any change in its position based on the target location. The only significant clustering of eye position occurred in the last 40 ms before the saccade (for both visual search and pop-out). This effect captured the initial eye movement towards the target location that occurred before the saccade was triggered.

reflect this shift in eye position and not the direction of attention itself. Figure 63 shows the probability of the eye position falling in each of the four quadrants (relative to target location) over time. Unlike with the neural activity in FEF, there appears to be no pattern in the eye position over time. As with the neural activity, independent bins of 40 ms were used. The post-saccadic eye position reflects our sorting procedure. In the time bin just prior to the saccade there is a preference for the eye position to be in the target quadrant for both search and pop-out (40% of the time, $p < 0.01$ by ANOVA, for both tasks). This is a result of the eye beginning to move towards the target as the saccade is triggered. Crucially, unlike the neural activity, there is no preference for eye location in the 40 ms from 80 ms to 40 ms pre-saccade ($p = 0.21$ for search and $p = 0.95$ for pop-out, both by ANOVA). The (unsignificant) maximally likely position of the eye in this bin was in target quadrant for search (very slightly above the opposite quadrant) and in the quadrant clockwise from the target for pop-out. The probability of the eye being located in any quadrant for any time period prior to 40 ms before the saccade was never significantly different from chance. These results suggest that there is no confounding pattern between eye movement and target location over time, supporting the conclusion that the FEF activity reflects the shift in the locus of attention.

Conclusions

In this chapter we hope to have presented data demonstrating the existence of a serial component to visual search as well as a parallel component to visual pop-out. Both animals show behavioral evidence that they searched the visual array in an organized manner during visual search (with both animals showing a preference to search the visual array in a clockwise manner). Neural activity in FEF shows a distinct clockwise pattern. This effect was demonstrated on the population level (Figures 55 and 57), the individual neuron level (Figure 60), between pairs of FEF neurons simultaneously recorded on single trials (Figure 61). Furthermore, the size of the effect was demonstrated to correlate with the reaction time of the animal during search (Figure 56), suggesting that the effect is weakest when the animal found the target quickly (as it didn't have to search for as long) and strongest when the animal found the target slowly (presumably due to a more involved search process). Finally, this pattern of results was not observed for either dIPFC or LIP (Figures 58 and 59); nor was it observed for any area during visual pop-out (Figures 57-59, bottom).

A long-standing discussion within the field of visual search has been over the existence of a serial component and, if a serial search does exist, whether or not it is a common occurrence. There is good psychophysical evidence for both a parallel and serial component to visual search tasks (Treisman and Gelade, 1980; Treisman and Gormican, 1988). However, visual search clearly does not fall into two mutually exclusive categories, but rather exists across a continuum between purely parallel and exclusively serial search. Although this is well studied in the human psychophysical literature, there is limited electrophysiological evidence for a serial component to search. Our results provide the first direct electrophysiological evidence of a truly serial component to visual search and a parallel mechanism underlying visual pop-out. However, previous results suggest FEF carries a mixture of top-down and bottom-up information (Thompson and Bichot, 2005b; Thompson et al., 2005b). It is likely that the continuum of behavior observed for psychophysical experiments also exists for the neural representation of search. As LIP and FEF appear to have the strongest representations for parallel and serial search, respectively, one could hypothesize that the continuum of behavioral results may arise from a continuum of relative importance of these two regions.

Our results are also in support of the premotor theory of attention (see General Introduction for an overview). This theory states that the direction of attention to a particular spatial location utilizes the same underlying neural framework association with making explicit saccades. Previous evidence for this theory comes from psychophysical studies in humans (Rayner et al., 1978; Deubel and Schneider, 1996a; Peterson et al., 2004) and through microstimulation of FEF in primates (Moore and Fallah, 2001, 2004). Our work provides a complementary result, demonstrating that post-saccadic motor responses can be used to follow the spotlight of attention as it moves before the saccade. Taken together, these results provide strong support for the role of the frontal eye fields in the direction of spatial attention.

Chapter 7: General Conclusions

The goal of this thesis was to investigate the role of parietal and frontal cortex in the control of attention. Visual attention is believed to fall under the control of two different sources: external stimuli can automatically grab one's attention, while internal motivations can also direct and maintain attention. Based on previous work it has been suggested that these two competing sources of control for attention originate in different manners – external capture is believed to propagate forward through the cortex while internal direction originates is thought to be a top-down process originate from frontal cortex and directing posterior cortex. However, although there is a large degree of circumstantial evidence there has been no direct evidence for the existence of such a dichotomy (Miller and D'Esposito, 2005).

The double dissociation we found in relative timing between frontal and parietal cortex supports this model. As information flows forward from parietal to frontal during pop-out, our data supports the theory that external capture of attention is produced through a bottom-up mechanism, one that likely relies on lateral inhibition in order to allow for the most salient stimulus to capture attention (Desimone and Duncan, 1995). During visual pop-out, when attention is captured by the saliency of the target stimulus, LIP carries a very strong representation of the target location. However, when a conjunctive search was used such that the target stimulus is no longer the most salient object in the visual field, LIP is no longer responsive to the target location. However, it is important to note that although LIP is not able to locate the target in the visual search condition, it is still involved in the task and is carrying task relevant information. Our results fit well with proposed models of parietal cortex maintaining a saliency map of the visual field (Bisley and Goldberg, 2003; Constantinidis and Steinmetz, 2005).

Prefrontal cortex is well situated anatomically to act as a top-down influence on posterior cortex and it has been shown to be carry information about cognitive parameters necessary for executing the task at hand (Miller and Cohen, 2001). Our results are the first direct evidence that regions within frontal cortex are the source of internal direction of attention. Although we only studied one behavior, it is likely that frontal cortex acts as the source of top-down influences in many cognitive behaviors. However, it is important to note that previous work comparing the relative roles of PFC and the basal

ganglia show that prefrontal cortex is not always the primary source of direction for all behaviors (Pasupathy and Miller, 2005). Understanding the distinction between behaviors reliant on prefrontal cortex as supposed to other integrative brain regions (such as the basal ganglia or hippocampus) will provide a more complete picture of the role each region plays in, and how they interact in order to support, behavior.

The synchrony between brain regions varied significantly between our two tasks in a frequency-dependent manner. Synchronization in the middle frequency band appears to be increased during visual search tasks, while synchronization is increased in the upper frequency band for visual pop-out. These differences may reflect differences in the integration of information between these two tasks. During pop-out the local information within parietal cortex has solved the task so that information should be both locally enhanced as well as passed forward into frontal cortex. In contrast, during visual search, a more integrative approach is needed in order to search the external visual array with the remember target identity. Under these conditions a lower frequency band is enhanced, which some have argued may aid in compensating for the larger timing disparities inherent in integrating across several regions (von Stein et al., 2000; von Stein and Sarnthein, 2000; Engel et al., 2001).

Synchrony for both frequency bands increases preferentially between electrodes with specific location preference relationships. The upper frequency band synchronizes most strongly with electrodes sharing a preferred direction, again supporting a possible role in enhancing the feed-forward flow of information. A different pattern was observed for the middle frequency band: although dIPFC and LIP LFPs continue to be preferentially synchronized between areas of shared preferred direction, FEF and LIP LFPs are synchronized the strongest when the preferred direction is offset by 90° . In light of our results suggesting a serial mechanisms underlying visual search, this may highlight the passage of information from the previously attended location from LIP into FEF. This may aid in the decision to shift attention towards the next location. Therefore, while the middle-frequency band may play a role in integrating information across more diverse brain regions, it is not exclusively used for top-down information. This is evident in the phase relationship between LIP and frontal cortex LFPs – the positive phase offset suggests LIP leads frontal cortex. The importance of the middle frequency band in visual search is also reflected in the positive correlation in the degree of coherence on single trials with the

reaction time of that trial. However, the expected result of a decrease in reaction time with increased coherence between LIP and frontal cortex is missing, suggesting more research needs to be done to further investigate the relationship of coherence to behavior.

The 4-8 Hz frequency band appears to act as a global reference signal to which all of the other frequency bands are able to synchronize. In addition to showing strong cross-band synchrony, it is also the frequency band for which the strongest spike-field synchrony was observed. A global clock may provide a very useful function to the brain by allowing it to occasionally resynchronize all of the frequency bands across all of the brain regions. This would reduce the degree of drift between frequency bands, something that may be important if neurons carrying different information are synchronized to different frequency bands. Not only must downstream neurons be able to integrate this information for processing, but exact temporal relationships may also be necessary for spike-timing dependent plasticity

We provide evidence for a serial mechanism underlying visual search and suggest activity in FEF directs the attentional spotlight from location to location. In contrast, visual pop-out appears to rely on a more parallel mechanism. Interestingly, dIPFC neurons are only responsive once the target has been attended in their receptive field. However, the source of this attention does not seem to matter – it can either be through top-down, serial mechanisms relying on FEF or bottom-up, parallel mechanisms through LIP. Furthermore, dIPFC was the only area to have significantly overlapping neural mechanisms between tasks, suggesting that it may, in part, be providing a common mechanism between the two tasks. One possible mechanism dIPFC may provide is to compare the currently attended target stimulus with the remembered sample stimulus in order to identify the target stimulus as correct and allow for a behavioral response.

The time to shift the attentional spotlight from location to location was estimated from both the behavior and from the neural activity to be approximately 40 ms. Therefore the animal is able to shift its attention at a rate of roughly 25 Hz, a value falling squarely in the middle frequency band found to be

Figure 64




Figure 64. Schematic of Control of Attention. During bottom-up, external direction of attention, selectivity flowed forward from parietal cortex into frontal cortex. In contrast, when attention was directed in a top-down, internal, manner selectivity flowed from the frontal cortex. These results support models of attention with both top-down and bottom-up influences, and suggests that top-down direction of behavior originates in the frontal cortex.

enhanced during visual search. The oscillations within the middle frequency band may allow for the synchronization of a diverse network of brain areas that work together to serially search the visual array.

Collecting our results, and relying heavily on previous work in the field, we can propose a working model of the control of visual attention as outlined in Figure 64. Salient objects in the visual field are able to capture attention as their representations propagate forward through the cortex. This propagation likely relies on lateral inhibition similar to a winner-take-all manner. During a visual pop-out task this process automatically selects the target location (reflected in LIP), eventually reaching dIPFC and then FEF for a behavioral response. This process happens quickly and in a fairly consistent manner (as indicated by the tighter reaction time distribution). It has been well documented that high frequency synchrony within visual cortex is associated with increase attention and that it can enhance downstream representations. In the current work we find similar effects within LIP (reaction time decreases as synchrony increases) and between LIP and frontal cortex (which is greater during pop-out). However, when internal direction of attention is necessary, either because there is no overtly salient objects in the visual field or because one desires to direct our attention in a specific manner, activity in frontal cortex is necessary. In the case of spatial direction of attention, the top-down signal appears to originate from the frontal eye fields although other prefrontal regions are likely to be involved in the top-down direction of other behaviors. Furthermore, the brain appears to oscillate at approximately 25 Hz, both allowing attention to shift from location to location every 40 ms and allowing for the necessary information to be integrated within a consistent temporal window.

Completing a search task, whether it relies on internal direction of attention or is automatically solved through external stimulus parameters, is a complex behavior. It requires a host of brain regions spanning from the back of the brain to the front and relies on the relationships between these regions to be flexible in order to ensure the most useful information is always available at the right time. This is why our work has focused on *relative* differences, whether they are timing differences between brain regions, synchrony differences between tasks, or even the commonality/disparity of networks underlying the task. We believe this is a fundamentally important distinction from the classic approach taken by neuroscience. Much of previous work in neuroscience has focused on a single brain region during a single task in order to determine what function is served by a particular brain region. In this

manner we have made great strides in understanding the brain – we have identified visual, auditory, somatosensory, motor, and cognitive regions and assigned each brain area a general role for it in behavior. Although this approach has been successful as a first-pass through the brain, there are reasons to believe that relationships *between* brain regions are important. The existence of strong feedback projections between cortical regions makes it difficult to study a brain region in isolation – outside of pure sensory areas the information flowing into a brain region is quickly a mix of feedforward and feedback projections. This not only emphasizes the need for capturing the leading edge of selectivity, but also demonstrates the need for relative studies across brain regions – when one records from a single area it is not clear what information is generated within that region as supposed to being reflections of computations done elsewhere. A common measure of the relative activity between brain regions is the study of synchrony across regions. Both previous work and the work presented here have demonstrated a direct relationship between behavior and the degree of synchrony between and within brain regions. These comparisons are not possible without the study of multiple brain regions simultaneously. We believe that the continued expansion of mutli-site recording technology will continue to push the field forward, providing greater insight into how individual brain regions interact in order to produce behavior.

The goal of this thesis was to better understand the relationships between parietal cortex (LIP) and prefrontal cortex (dlPFC and FEF) in the control of behavior. We believe our results have provided important pieces of information about the existence of multiple mechanisms underlying the control of attention, the origin of these mechanisms, and how these mechanisms might be implemented within a neural network structure by shifting the effective connectivity between regions. However, there are many questions about the control of attention (and more generally the control of behavior) left unanswered. One question of interest is how attention is directed to spatial locations and/or to objects and whether a common mechanism supports both forms of attention at some level. A separate open question is about the establishment of synchrony amongst brain regions – despite the apparent importance of synchronous relationships it is not clear how these relationships are established, maintained, and shifted in order to best suit the current task. There are several possible ways to address this question from both a correlative and causative direction – for example, utilizing the estimated coherence from a single trial it may be possible to demonstrate the existence of correlations between the neural activity in one brain region to the establishment of synchrony between two regions.

Additionally, the manner in the brain learns which synchronous relationships to enhance at which frequency bands while learning a new task is an open and interesting question. The use of multiple site recordings will play a major role in providing valuable insights into these, and other, questions of how attention is implemented and controlled within the brain. Answers to these questions are just the beginning and will eventually lead to a more complete model of how information is filtered, processed, and acted upon by the brain in order to support behavior.

Subjects

Two male rhesus monkeys, weighing 6 kg each, were used. All procedures followed the guidelines of the MIT Committee on Animal Care and NIH. Animals were implanted under general anesthesia with a titanium headpost to immobilize the head and titanium recording chambers for recording. Chambers were stereotaxically placed over frontal and parietal cortices (in the same hemisphere) using structural MRI scans. Novel software was developed in Matlab that produced three-dimensional models of each animal's skull and brain in stereotaxic coordinates. This allowed accurate placement of electrode penetrations.

Behavioral Task

The trial was initiated when the animal fixated a point at the center of the screen. Fixation was required within a 3.2 degree window (± 1.6 degrees centered on the fixation point). After a short fixation period (500 ms), the animal was presented with a sample, colored, oriented bar for 1000 ms, centered on fixation. The sample stimulus was removed and the monkey then maintained central fixation over a 500 ms memory delay, which ended with presentation of a visual search array. The array elements were identical in size and shape to the sample and appeared four degrees from fixation. One of the array items matched the sample in both color and orientation (the target). Monkeys needed to make a direct, linear, saccade from central fixation to the target and hold their gaze at the target for 150 ms to receive a juice reward. Any deviations from the correct saccade path, including saccades to non-target stimuli, were recorded as errors and not rewarded.

The colored, oriented bars were 0.16 degrees of visual angle in width and 1.6 degrees in height. Nine target stimuli were constructed each day from all possible combinations of three orientations and three different colors. The orientations and colors were chosen such that the differences between search items was consistent in both orientation and color space (stepping was fixed on a given day, but could range from 15 to 25 degrees in orientation and 2 to 5% change in hue). The number of search array items was varied from 2 to 4 for the psychophysical experiments that tested reaction time as a function of array size. Four search array items were used for all neurophysiological recording

experiments. Array items appeared at positions 45, 135, 225, and 315 degrees from the vertical meridian (see Figure 1). In pop-out, the non-targets (distractors) differed from the target by 90 degrees and were all the same color, which was opposite the target color on the color circle. In visual search, distractors differed from the target by either color or orientation alone. The difference in color and orientation between the target stimulus and distractors was the same as the difference between target stimuli. This allowed a target stimulus on one trial to be a distractor stimulus on the next.

The search and pop-out tasks were interleaved in blocks of approximately 35 trials each. The animals performed a minimum of 720 correct trials during recording sessions, ensuring at least 10 trials for each of the 9 possible targets (3 colors by 3 orientations) at each location and for each task. Data is presented from 25 recording sessions (10 in monkey S, 15 from monkey W). Behavioral control and the display of visual stimuli was done through CORTEX (<http://www.cortex.salk.edu>). An infrared based eye-tracking system monitored eye position at 240 Hz. As can be seen in Figure 65 there were no differences in the overall distribution of eye velocity between visual pop-out and visual search, nor were there differences in eye velocity between tasks for any given recording day. Unfortunately, a detailed analysis of micro-saccades is not possible due to the noise level inherent in the infrared based eye-tracking system. The eye-tracker does avoid eye coil surgeries and is more than adequate for monitoring fixation and normal saccadic eye movements, it is not qualified to study the subtleties of microsaccades. However, based on the overall distribution of eye velocities, we feel certain that while we are unable to study microsaccades directly, they do not differ greatly in prevalence between tasks.

Figure 65




Figure 65. Top figure shows the distribution of eye velocity during the search period for both the visual search (red) and pop-out (blue) tasks. There was no significance different, suggesting that minor eye movements were not different between tasks and therefore could not have contributed to any effects. The bottom figure shows that the difference in eye velocity was not significantly different from zero for any recording session.

Recording Locations and Isolation of Neural Activity

A total of up to 50 electrodes were implanted into parietal and frontal cortex simultaneously, with a maximum of 25 in each anatomical area. A total of 802 neurons were recorded across all three anatomical regions in two monkeys (274 from the lateral intraparietal area, LIP, 272 neurons from lateral prefrontal cortex, dlPFC, and 243 neurons from the frontal eye fields, FEF; 280 neurons were recorded from monkey S and 522 neurons from monkey W).

The lateral intraparietal region (LIP) recording well was placed at approximately -7 mm AP from the interaural plane and was placed using structural MRIs. To identify LIP neurophysiologically, we trained the animals on a delayed saccade task. During central fixation, a brief spot of light was flashed in the periphery. After a memory delay, the fixation point was extinguished and the animal made a saccade to the remembered location of the light spot. This has been used to isolate LIP from surrounding regions, as it is the only region in the parietal cortex that shows selectivity for target location during the memory delay (Barash et al., 1991). The animals performed the delayed saccade task at the beginning of every recording session. Electrodes were only considered to be within LIP for that session if a neuron isolated from that electrode showed memory delay activity selective for the remembered location.

The frontal recording well was placed at approximately 23 mm AP from the interaural plane. Microstimulation was used to demarcate the frontal eye fields from dorsolateral prefrontal cortex. Stimulation was delivered as a 200 ms train of bi-phasic pulses with a width 400 μ s and an inter-pulse frequency of 330 Hz using the same electrodes used for recording. Current level was started at 150 μ A and reduced to find the threshold at which an eye movement vector was elicited 50% of the time. Only sites that had thresholds of stimulation amplitudes less than 50 μ A were classified as belonging to the frontal eye fields (Bruce and Goldberg, 1985). Anterior sites were classified as belonging to the dlPFC. In general, stimulation at dlPFC sites did not elicit eye movements even at the highest current amplitude tested (150 μ A).

As noted in the Introduction, we were interested in utilizing simultaneous recordings from all three regions for several reasons. First, and foremost, absolute timing of neural activity can vary across time with further training/experience, between tasks, and with the statistical criterion used. Thus, our main interest was in the *relative* timing differences between areas. Simultaneous recording from multiple electrodes aids in detecting them because it reduces the influence of extraneous variables such as differences in performance across sessions. Furthermore, simultaneously recording from neurons in several brain regions allows for the study of relative timing effects between neurons and areas, such as synchrony.

In order to record from all three regions simultaneously we used a custom built microdrive assembly to lower electrodes in pairs from a single screw. The microdrive assembly was designed to allow for a high density of electrodes in order to maximize the number of simultaneously recorded neurons and field potentials. We used epoxy coated tungsten electrodes for recording as well as for microstimulation. The electrodes were acutely lowered through an intact dura at the beginning of every recording session and allowed to settle for a minimum of 2 hours before recording. This ensured stable isolation of single units over the session. After recording the electrodes were retracted and the microdrive assembly was removed from the well.

Spiking activity and local field potentials were recorded for each electrode simultaneously. Both spiking activity and local field potentials were referenced to ground (rather than to one of the electrodes). This eliminated the possibility that coherence was due to the reference itself having temporal structure. However, ground referencing induces a characteristic, condition independent, drop-off in coherence with frequency (see Figure 15). The signal from each electrode was divided into spiking activity and a local field potential by filtering between 154 Hz and 8.8 kHz for spikes and between 3.3 and 88 Hz for the local field potential. Waveforms from single neurons were sorted from the raw spiking activity signal off-line, using a combination of principal component analysis of waveform traces along with other properties of the recorded waveforms (amplitude, trough/peak latency, etc).

Due to the large number of simultaneously recorded neurons across all three areas there was no optimization of the stimulus parameters for recording. Likewise, neurons were not pre-selected for responsiveness. Rather, we randomly selected neurons for recording as best we could. Although this may yield neural responses that are below their potential maximum, this factor is consistent across all three regions and across behavioral tasks, and therefore cannot influence comparative results. Furthermore, doing so gives us a more realistic approximation of the way the brain encounters typical visual scenes and since no pre-selection occurred we have a more complete sampling of available cell types.

Analysis Techniques

The majority of analysis techniques and statistical tests used in the text are common amongst the field and are considered standard. All statistical tests were used when appropriate, with an attempt made to avoid normality assumptions when appropriate. In addition to the common tests, we used several analysis techniques in this study which are not necessarily common in neuroscience, and therefore are described in detail here. These statistics include the mutual information statistic, the coherence statistic, and circular statistics. All analysis code was written in Matlab and/or C.

For all of the analysis presented in this manuscript, we required each neuron to be recorded for a minimum of 30 trials for each target location. This yielded 249 LIP neurons, 248 dIPFC neurons, and 225 FEF neurons that had enough trials during the pop-out condition. During the search condition 247 LIP neurons, 251 dIPFC neurons, and 225 FEF neurons met this criterion. Similar results were obtained from each animal alone, so they are combined for presentation.

Mutual Information Statistic

Neural selectivity was assessed through the use of the mutual information statistic (Shannon and Weaver, 1949). The mutual information statistic reflects how well one can predict the variable of interest when given the neural activity (for example, the spiking activity from a single neuron, the spiking activity of multiple neurons, or the potential value from the local field potential). The amount of information carried by the neural activity can be calculated by estimating the reduction in uncertainty

about the variable of interest given the neural activity: $I(X|R) = H(X) - H(X|R)$, where I is the mutual information and H is the entropy, or uncertainty. X is our variable of interest, while R is the response of the neural measure. The entropy can be defined as $H(X) = \sum_{i=1}^N p(X_i) \log p(X_i)$, where $p(X_i)$ is the probability of observing occurrence i out of possible values in X . For example, if X is the possible locations of the target in our visual search display, then i can take the values of upper left, upper right, lower right, or lower left. In order to compute the information about variable X provided by the observed values in R , we need to compute both $H(X)$ and $H(X|R)$. While $H(X)$ is fairly easy to compute (and for our 4 item display, with each location occurring equally, is equal to 2 bits), $H(X|R)$ requires us to estimate the probability of observing $p(X_i|R_j)$. This is not directly observable, however through the use of Bayes' rule we can indirectly estimate it from $p(R_j|X_i)$, $p(R_j)$ and $p(X_i)$ which are directly observable. This is done through the equation $p(X_i|R_j) = \frac{p(R_j|X_i)p(X_i)}{p(R_j)}$.

The mutual information statistic is ideal for estimating the selectivity of a neural measure for a particular variable for several reasons. First, it requires no assumptions to be made about the underlying distributions of either the observed neural response or the variable of interest. This is an improvement over ANOVA tests as there is good evidence that many neural responses are not normal in their distributions. Therefore, by not making assumptions about the normality of the data, we can avoid spurious results and capture effects that more typical selectivity measures miss. A second advantage is that mutual information does not require the underlying distributions to be unimodal. This is an important advantage over other possible statistics such as ROC. Furthermore, mutual information handles multiple stimulus possibilities well. For example, in our 4 target location example, the mutual information statistic can calculate the selectivity of a single cell, while other statistical tests that require pairs of data such as ROC, would fail. Finally, the mutual information statistic is very amicable to randomization tests, making the determination of significance easily done. Randomization tests are done by shuffling the associations between the observed response, R_p , and the variable of interest, X_i . By repeating this process several thousand times, we can generate a null distribution to which we can compare the observed mutual information value. Similar to the mutual information statistic, randomization tests are more effective than more classic techniques as it does not require assumptions about the underlying distribution, does not require a unimodal distribution, and handles multiple comparisons easily.

Figure 66

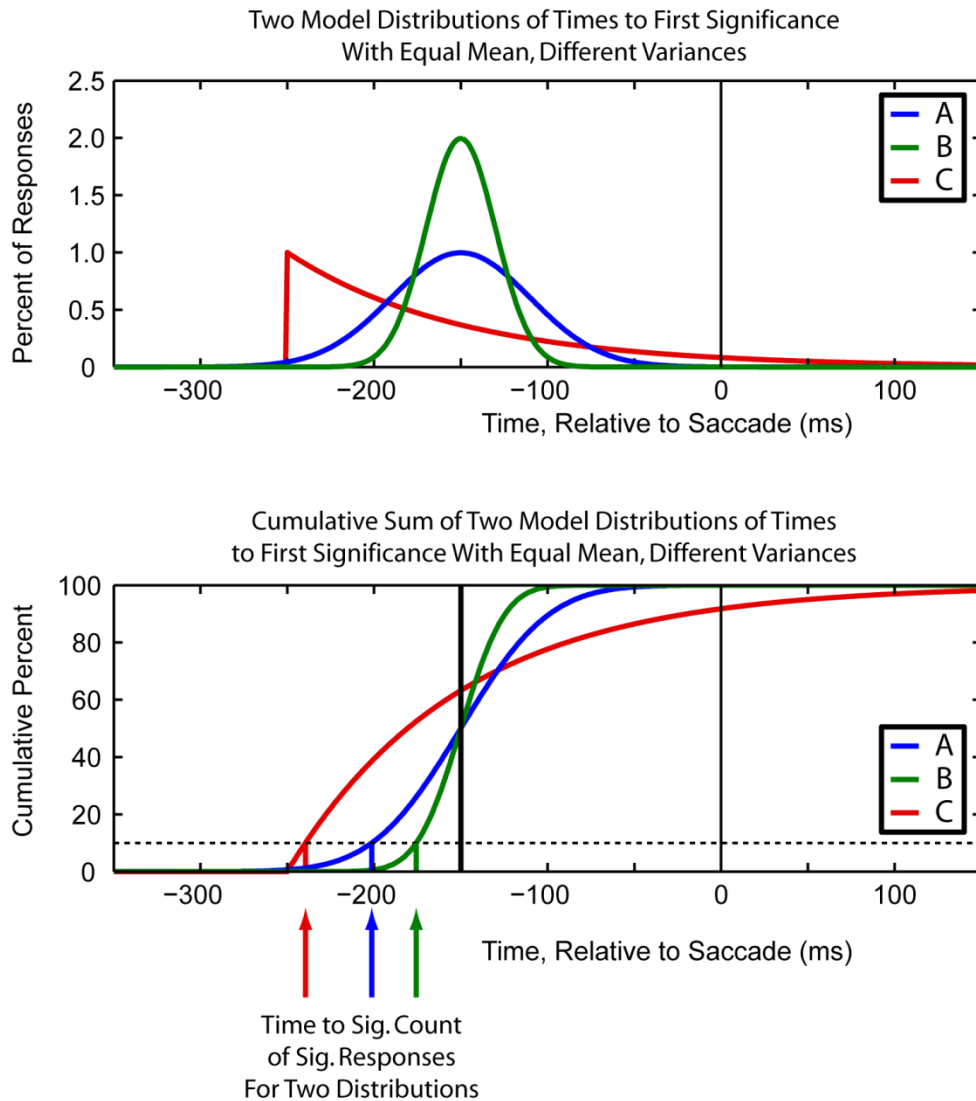


Figure 66. Three model distributions created to demonstrate the importance of capturing the leading edge of selectivity rather than the average. The top figure displays the probability density function and the bottom figure displays the cumulative density function. Both distribution A and B are normal distributions with a mean at -150 ms but differing variances. Distribution C is an exponential distribution also with a mean of -150 ms. These results highlight the fact that using the mean as a measure of the latency of a distribution will often disguise true differences in the distributions.

Comparing the Mean vs. Leading Edge of a Distribution

A crucial question underlying the results presented in this thesis is the manner in which to determine when a population of neurons (for example, like the ones recorded from an anatomical region) carries significant information about a variable of interest. Typically this question has been answered in one of two possible ways: through an estimation of the mean of the overall distribution or by estimating the leading edge of the distribution. We believe that the use of an arbitrary criterion to determine when significance is reached is fallible regardless of whether this threshold is the mean or a leading edge. In contrast, one should define the latency of response as the point at which a significant number of neurons carry (or have carried) a significant amount of information about our parameter of interest. This has the advantage of correcting for multiple comparisons and avoids the arbitrariness associated with selecting a threshold.

A related issue with selecting a threshold lies in the relationship between the shape of the distribution and when that distribution reaches the selected threshold. To demonstrate this point, three example distributions of when theoretical neurons carried significant information about our parameter of interest are shown in Figure 66. The top plot shows the probability distribution function, while the bottom plot shows the cumulative distribution function. All three distributions have equal means, but either differ in their variance (A and B) or differ in their underlying distribution (A/B and C). As we can see in the cumulative distribution plot, if we were to take the mean of the distribution as our arbitrary threshold all three distributions would show the same 'latency' even though this is clearly not the case. This problem becomes exacerbated when the variance of the distribution becomes large with respect to the selected period of interest – the mean will tend towards the middle regardless of the underlying distribution. An improved solution is determine the level at which a significant number of neurons are carrying significant information (here modeled as 10%, but this parameter can be set on a per-population basis to compensate for different N's and multiple comparisons). Now the latency differences between the populations better reflect the true differences. Furthermore, it is crucially important when investigating the flow of information to characterize a population by its leading edge as later-responding neurons may be responding to subsequent information received from other regions.

Regardless of the exact method used to determine the latency of a population’s response to a particular parameter of interest, this discussion highlights the difficulty in assigning an absolute neural latency to an area. Not only does the manner in determining the latency have a large impact on what that latency is (even if we improve this using a less arbitrary method) but many other variables will play a direct role – the exact paradigm used, the manner of training, the previous experience of an animal, etc. Many of these parameters are not only unreported, but to a large degree are unknowable. In order to overcome these difficulties, one must report the *relative* latencies of neuronal responses between simultaneously recorded areas, in the same animal, under the same paradigm, and with the same latency determination method. Although this will not completely remove all biases in the data, it will go a long way towards mitigating them, allowing one to observe and report latency differences of much finer detail.

Coherence Statistic

The coherence statistic was used to measure the synchrony between local field potentials. Coherence is a measure of the co-spectrum between two signals, and is normalized for the power in each signal alone. By normalizing it allows coherence values to be compared across conditions and ensures that increases in coherence cannot be the result of an increase in the amplitude of the underlying signal.

$$S_{XY} \approx FFT(D_X) \bullet FFT(D_Y)^*$$

$$C_{XY} = \frac{S_{XY}}{\sqrt{S_{XX} \bullet S_{YY}}}$$

$$Coh = |C_{XY}|$$

In the above equations, D_X is the recorded signal from source X, and S_{XY} is the co-spectrum between two signals from sources X and Y. Spectral estimates were made using a multi-tapered method and were averaged across trials. For our analyses, a smoothing level of 7.5 Hz was used to generate the discrete prolate spheroidal sequences used as tapers. A detailed description of this method of calculating coherence can be found in Jarvis and Mitra (Jarvis and Mitra, 2001). Since we typically used a

200 ms window to estimate the spectrum, frequencies below 10 Hz were difficult to estimate, and therefore were excluded from the current analysis. However, we were able to use the filtered LFP signals and circular statistics to investigate certain frequency bands of interest below 10 Hz (see below).

Coherence was observed for both the visual search and visual pop-out task, and the difference between the two tasks was tested for significance at each frequency independently. Significance was determined through a randomization test similar to the one used for mutual information: the trial assignments for each observed signal were randomly shuffled and the coherence level was recalculated in order to generate a null distribution. Significant differences in the coherence between signals were determined by calculating the difference in the randomly generated coherence values in order to generate a null distribution of the difference in coherence. This trial-shuffling procedure corrects for any trial-invariant differences between visual search and visual pop-out. Based on the randomly generated null distribution, a z-score was determined for the observed coherence by determining how far our observed coherence was relative to the null mean in units of the null distribution's standard deviation. To correct for multiple comparisons across frequencies we increased the significance requirement (22 different frequencies were used, therefore significance was measured at $p = 0.05/22$, or $\sim 2 \times 10^{-3}$). The z-score threshold in Figure 16B reflects this adjusted significance level. Non-parametric permutation tests were also used and found the same condition-dependent differences in the same frequency bands (similar to Nichols and Holmes, 2002).

Estimating Coherence of Single Trials

As coherence is a measure of the reliability of the phase relationship between two signals it inherently averages across multiple samples (in our case this tends to be trials). This makes it difficult to directly estimate on a single sample (trial) what the degree of coherence was between two signals. However there are many interesting questions that can be asked if the coherence on a single trial is known, including, but not limited, to those presented in this thesis, such as examining the relationship between coherence and reaction time or the relationship between the activity of single neurons and large-scale inter-/intra-areal coherence. Several techniques have been used to attempt to ask these questions, but we propose and utilize a new method in this thesis. One previously used option was to bin trials into two or more groups based on the variable of interest (for example, sorting trials into 'fast'

and 'slow' reaction time groupings, Womelsdorf et al., 2006b). However this method is often difficult when the number of trials is limited and when the temporal range is restricted, making estimation of the underlying spectra difficult to begin with. Another option is to attempt to leverage multi-site recordings to determine the coherence across all pairs of electrodes on a single trial. This method also suffers from difficulties in a large enough sample size (large pairs of electrodes in areas of interest can be difficult to acquire) and ignores the non-stationarity of coherence across particular sites (something that we have shown to be important with our analysis of coherence by preferred target location).

The method proposed and used in this thesis is to estimate the coherence on a single trial through a jackknife procedure. The jackknife is often used in order to estimate the bias and uncertainty of a particular statistic, but here we will use it to estimate the contribution of a single trial to the overall coherence. In general terms this is done by estimating the coherence across all trials and with a single trial removed. The difference between these two coherence statistics can be thought of as the 'contribution' of the removed trial to the overall coherence. This can be done for each individual trial in order to estimate the coherence between two signals across samples.

We can show this explicitly by following a similar logic to that used in deriving the general jackknife: suppose we have a random sample of n values, X_1, X_2, \dots, X_n , then we can calculate the sample mean in the typical fashion:

$$\bar{X} = \sum_i^n X_i/n$$

We can also determine the average when a single sample is left out:

$$\bar{X}_{-j} = (\sum_i^n X_i - X_j)/(n - 1)$$

These two equations can be combined to solve for X_j :

$$X_j = n\bar{X} - (n - 1)\bar{X}_{-j}$$

While this is an obviously trivial result if the sample values are already known, it can become a very powerful method of estimating the sample values when they are not known (as is the case with

coherence). This often happens when our statistic is some function of our sample values. The same derivation is applicable, but now our overall estimate of the statistic is $\hat{E} = f(X_1, X_2, \dots, X_n)$ and our estimate without the j -th trial is \hat{E}_{-j} . Therefore, by analogy with our derivation above, we can estimate the value of our function on a single trial:

$$\hat{E}_{-j}^* = n\hat{E} - (n - 1)\hat{E}_{-j}$$

Typically these values are used to estimate the bias and variability in our statistic in order to determine significance. For our analysis, we will use this method to estimate the coherence on a single trial in order to determine co-variations in single-trial coherence with reaction time.

Circular Statistics

As an extension of the results from the coherence analysis, we were interested in further investigating the role of the interesting frequency bands in the local field potential signal. For example, we were interested in the relationship between spiking activity and the frequency bands, as well as the relationship between multiple bands. In order to determine this relationship we wanted to determine if there was a significant phase relationship between a particular frequency band in the LFP and either spiking activity or another frequency band. Since the phase relationship by definition fits around a circle, we used circular statistics to determine these relationships. Although linear statistics are well developed, circular statistics are less well understood, and therefore only a few statistics are available for our use.

We are interested in determining whether the observed relationship between our frequency band-limited LFP and another variable deviates from the uniform distribution. The uniform distribution is our null distribution – if there is no specific relationship between the two variables then the phase relationship between these variables should fall randomly, therefore generating a uniform distribution around the circle. In order to determine these possible deviations, we used a combination of parametric and non-parametric tests. Similar to linear statistics, the parametric tests provide increased resolution, but require the data to fit the underlying assumptions. We used the Rayleigh test as a parametric test to determine whether the observed phase relationship differed from uniform. The Rayleigh test makes the assumption that the distribution is uniform. While this is a valid assumption for some of our data, it is

not always valid (for example, when examining the phase relationship between different frequency bands, a higher frequency band must peak several times within a lower frequency band, requiring the phase relationship to be multi-modal). When the unimodal distribution assumption cannot be reasonably made, we used a non-parametric test to determine deviation from uniformity.

In addition to determining the deviation from uniformity, it is also informative to investigate the average phase relationship between two signals. In order to determine these changes we used both the median and mean angle of deviation and generated confidence intervals.

Filtering Local Field Potentials

In order to determine LFP-LFP and Spike-LFP phase-locking using circular statistics, it is first necessary to filter the local field potentials into our frequency bands of interest. For the analysis presented here we used four main frequency bands: 4-8 Hz, 8-16 Hz, 18-34 Hz, and 35-65 Hz. Local field potentials were filtered using a digital IIR filter, consisting of 13 sections of Chebyshev, Type II filters (see Figure 67). The filter was of order 26 and was attenuated to at least -40 dB in the stop-band regions. The Type II Chebyshev filter avoided any rippling within the pass-band but did have a non-linear phase offset across frequencies. In order to compensate for this non-linear phase effect, we filtered the LFP signal both forward and backward in time. This ensures a zero phase shift. As this increases the effective order of the filter we filtered across the entire trial's signal, ensuring enough data to avoid edge effects.

Figure 68 shows a filtered LFP signal from an example trial (top) as well as the filtered LFP signal from a single trial averaged across all trials (bottom). In both figures one can see 'phase-jumps' – periods of time in which the phase of the filtered signal changed in a non-linear manner (denoted by black triangles, outlined in the filtered signal's color). Classically these are thought of as being the result of either a switch in the source of the signal, a restart of the source, or an interaction between two signal sources. Although it is not possible to determine which of these is the case from our current data, this is an interesting future consideration. This effect also highlights an important reason why the phase-locking approach can be more appropriate for synchrony measures than one relying on frequency

space (such as coherence). Since methods in frequency space rely on the signal being stationary, long periods of time are often hard to interpret – even from the two examples shown here, it is obvious that these signals are not stationary; rather they change over time, likely relating to the processes underlying behavior. These changes are important to capture and using a phase-locking approach allows us to do so. Phase relationships were determined by finding the local maxima and minima across time and fitting a linear phase distribution between local extrema. The phase relationship of any point process could then be easily determined, including, but not limited to, other filtered LFP signals, spiking activity, and behavioral responses.

Figure 67




Figure 67. The amplitude response curve for the four filters used throughout the manuscript. All filters were sectioned Chebyshev, Type II, IIR filter of order 26 (across 13 sections). They were designed to pass with no ripple in the pass-band and stop with at least -40 dB attention in the stop-bands. The phase response was continuous in the pass-band and although it has a non-linear phase distortion, this was corrected for using a zero-phase filtering approach.

Figure 68




Figure 68. Examples of filtered LFP signals from a single dIPFC electrode. Examples are taken from a single visual search trial (top) or averaged across all visual search trials (bottom). Both traces show non-linear changes in phase ('phase-jumps') as indicated by the black triangles.

References

- Aertsen AM, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: modulation of "effective connectivity". *J Neurophysiol* 61:900-917.
- Allman J, Miezin F, McGuinness E (1985) Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Annual Review of Neuroscience* 8:407-430.
- Amaral DG (1986) Amygdalohippocampal and amygdalocortical projections in the primate brain. *Adv Exp Med Biol* 203:3-17.
- Amaral DG, Price JL (1984) Amygdalo-cortical projections in the monkey (*Macaca fascicularis*). *Journal of Comparative Neurology* 230:465-496.
- Andersen RA, Asanuma C, Essick G, Siegel RM (1990) Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *J Comp Neurol* 296:65-113.
- Armstrong KM, Moore T (2007) Rapid enhancement of visual cortical response discriminability by microstimulation of the frontal eye field. *Proc Natl Acad Sci U S A* 104:9499-9504. Epub 2007 May 9421.
- Armstrong KM, Fitzgerald JK, Moore T (2006) Changes in visual receptive fields with microstimulation of frontal cortex. *Neuron* 50:791-798.
- Asaad WF, Rainer G, Miller EK (2000) Task-specific activity in the primate prefrontal cortex. *Journal of Neurophysiology* 84:451-459.
- Barash S, Bracewell RM, Fogassi L, Gnadt JW, Andersen RA (1991) Saccade-related activity in the lateral intraparietal area. I. Temporal properties; comparison with area 7a. *Journal of Neurophysiology* 66:1095-1108.
- Barbas H, Pandya DN (1989) Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *Journal of Comparative Neurology* 286:353-375.
- Barbas H, De Olmos J (1990) Projections from the amygdala to basoventral and mediodorsal prefrontal regions in the rhesus monkey. *J Comp Neurol* 300:549-571.
- Barbas H, Henion TH, Dermon CR (1991) Diverse thalamic projections to the prefrontal cortex in the rhesus monkey. *Journal of Comparative Neurology* 313:65-94.
- Bendiksbj MS, Platt ML (2006) Neural correlates of reward and attention in macaque area LIP. *Neuropsychologia* 44:2411-2420. Epub 2006 Jun 2416.
- Bichot NP, Rossi AF, Desimone R (2005) Parallel and Serial Neural Mechanisms for Visual Search in Macaque Area V4. *Science* 308:529-534.
- Bichot NP, Thompson KG, Chenthal Rao S, Schall JD (2001) Reliability of macaque frontal eye field neurons signaling saccade targets during visual search. *J Neurosci* 21:713-725.
- Bisley JW, Goldberg ME (2003) Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299:81-86.
- Bisley JW, Goldberg ME (2006) Neural Correlates of Attention and Distractibility in the Lateral Intraparietal Area. *Journal of Neurophysiology* 95:1696-1717.
- Blatt GJ, Andersen RA, Stoner GR (1990) Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *J Comp Neurol* 299:421-445.
- Bremmer F, Distler C, Hoffmann KP (1997) Eye Position Effects in Monkey Cortex. II. Pursuit- and Fixation-Related Activity in Posterior Parietal Areas LIP and 7A. *J Neurophysiol* 77:962-977.
- Bressler SL (1996) Interareal synchronization in the visual cortex. *Behavioural Brain Research* 76:37-49.

- Bruce CJ, Goldberg ME (1985) Primate frontal eye fields: I Single neurones discharging before saccades. *Journal of Neurophysiology* 53:607-635.
- Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* 315:1860-1862.
- Cabeza R, Nyberg L (2000) Imaging cognition II: An empirical review of 275 PET and fMRI studies. *J Cogn Neurosci* 12:1-47.
- Carrasco M (2006) Covert attention increases contrast sensitivity: Psychophysical, neurophysiological and neuroimaging studies. *Prog Brain Res* 154:33-70.
- Cave KR, Wolfe JM (1990) Modeling the role of parallel processing in visual search. *Cognit Psychol* 22:225-271.
- Chao LL, Knight RT (1997) Prefrontal deficits in attention and inhibitory control with aging. *Cereb Cortex* 7:63-69.
- Chelazzi L, Miller EK, Duncan J, Desimone R (1993) A neural basis for visual search in inferior temporal cortex. *Nature* 363:345-347.
- Chelazzi L, Duncan J, Miller EK, Desimone R (1998) Responses of neurons in inferior temporal cortex during memory-guided visual search. *J Neurophysiol* 80:2918-2940.
- Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. *Annual Review of Neuroscience* 22:319-349.
- Committee on Quality Improvement SoA-DHD (2000) Clinical Practice Guideline: Diagnosis and Evaluation of the Child With Attention-Deficit/Hyperactivity Disorder. *Pediatrics* 105:1158-1170.
- Constantinidis C, Steinmetz MA (2005) Posterior Parietal Cortex Automatically Encodes the Location of Salient Stimuli. *J Neurosci* 25:233-238.
- Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3:201-215.
- Corbetta M, Miezin FM, Shulman GL, Petersen SE (1993) A PET Study of Visuospatial Attention. *J Neurosci* 13:1202-1226.
- Corbetta M, Miezin FM, Dobmeyer S, Shulman GL, Petersen SE (1990) Attentional modulation of neural processing of shape, color, and velocity in humans. *Science* 248:1556-1559.
- Corbetta M, Akbudak E, Conturo TE, Snyder AZ, Ollinger JM, Drury HA, Linenweber MR, Petersen SE, Raichle ME, Van Essen DC, Shulman GL (1998) A common network of functional areas for attention and eye movements. *Neuron* 21:761-773.
- Coull JT, Nobre AC (1998) Where and When to Pay Attention: The Neural Systems for Directing Attention to Spatial Locations and to Time Intervals as Revealed by Both PET and fMRI. *J Neurosci* 18:7426-7435.
- Coull JT, Frackowiak RS, Frith CD (1998) Monitoring for target objects: activation of right frontal and parietal cortices with increasing time on task. *Neuropsychologia* 36:1325-1334.
- Crick F, Koch C (1998) Consciousness and neuroscience. *Cereb Cortex* 8:97-107.
- Crick F, Koch C (2003) A framework for consciousness. *Nat Neurosci* 6:119-126.
- Croxson PL, Johansen-Berg H, Behrens TEJ, Robson MD, Pinski MA, Gross CG, Richter W, Richter MC, Kastner S, Rushworth MFS (2005) Quantitative Investigation of Connections of the Prefrontal Cortex in the Human and Macaque using Probabilistic Diffusion Tractography. *J Neurosci* 25:8854-8866.
- de Brecht M, Saiki J (2006) A neural network implementation of a saliency map model. *Neural Netw* 19:1467-1474.
- Deco G, Rolls ET (2005) Neurodynamics of Biased Competition and Cooperation for Attention: A Model With Spiking Neurons. *Journal of Neurophysiology* 94:295-313.
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annual Review of Neuroscience* 18:193-222.

- Desimone R, Schein SJ, Moran J, Ungerleider LG (1985) Contour, color and shape analysis beyond the striate cortex. *25:441-452*.
- Deubel H, Schneider WX (1996a) Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Res 36:1827-1837*.
- Deubel H, Schneider WX (1996b) Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research 36:1827-1837*.
- Duncan J, Humphreys GW (1989) Visual search and similarity. *96:433-458*.
- Duncan J, Humphreys G, Ward R (1997) Competitive brain activity in visual attention. *Current Opinion in Neurobiology 7:255-261*.
- Eblen F, Graybiel A (1995) Highly restricted origin of prefrontal cortical inputs to striosomes in the macaque monkey. *J Neurosci 15:5999-6013*.
- Eglin M, Robertson LC, Knight RT (1991a) Cortical Substrates Supporting Visual Search in Humans. *Cereb Cortex 1:262-272*.
- Eglin M, Robertson LC, Knight RT (1991b) Cortical substrates supporting visual search in humans. *1:262-272*.
- Engel AK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci 2:704-716*.
- Felleman DJ, Van Essen DC (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex 1:1-47*.
- Ferraina S, Pare M, Wurtz RH (2002) Comparison of Cortico-Cortical and Cortico-Collicular Signals for the Generation of Saccadic Eye Movements. *J Neurophysiol 87:845-858*.
- Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal cortex. *Nature 443:85-88*.
- Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science 291:312-316*.
- Friedman-Hill SR, Robertson LC, Treisman A (1995) Parietal contributions to visual feature binding: evidence from a patient with bilateral lesions. *Science 269:853-855*.
- Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci 9:474-480*.
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science 291:1560-1563*.
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology 61:331-349*.
- Fuster JM (1973) Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *Journal of Neurophysiology 36:61-78*.
- Fuster JM (1995) *Memory in the cerebral cortex*. Cambridge, MA: MIT Press.
- Fuster JM, Alexander GE (1971) Neuron activity related to short-term memory. *Science 173:652-654*.
- G R, SC R, Miller EK (1999) Prospective coding for objects in the primate prefrontal cortex. *Journal of Neuroscience 19:5493-5505*.
- Gattass R, Gross CG, Sandell JH (1981) Visual topography of V2 in the macaque. *Journal of Comparative Neurology 201:519-539*.
- Gattass R, Sousa AP, Gross CG (1988) Visuotopic organization and extent of V3 and V4 of the macaque. *8:1831-1845*.
- Gaymard B, Ploner CJ, Rivaud S, Vermersch AI, Pierrot-Deseilligny C (1998) Cortical control of saccades. *Exp Brain Res 123:159-163*.
- Gottlieb JP, Kusunoki M, Goldberg ME (1998) The representation of visual salience in monkey parietal cortex. *Nature 391:481-484*.

- Gross CG, Bender DB, Rocha-Miranda CE (1969) Visual receptive fields of neurons in inferotemporal cortex of the monkey. *Science* 166:1303-1306.
- Haier RJ, Jung RE, Yeo RA, Head K, Alkire MT (2004) Structural brain variation and general intelligence. *NeuroImage* 23:425-433.
- Hasegawa RP, Matsumoto M, Mikami A (2000) Search Target Selection in Monkey Prefrontal Cortex. *J Neurophysiol* 84:1692-1696.
- Hillis AE (2006) Neurobiology of unilateral spatial neglect. *Neuroscientist* 12:153-163.
- Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195:215-243.
- Iba M, Sawaguchi T (2003) Involvement of the Dorsolateral Prefrontal Cortex of Monkeys in Visuospatial Target Selection. *J Neurophysiol* 89:587-599.
- Ipata AE, Gee AL, Goldberg ME, Bisley JW (2006) Activity in the Lateral Intraparietal Area Predicts the Goal and Latency of Saccades in a Free-Viewing Visual Search Task. *J Neurosci* 26:3656-3661.
- Itti L, Koch C (2000) A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research* 40:1489-1506.
- Itti L, Koch C (2001) Computational Modeling of Visual Attention. *Nature Reviews Neuroscience* 2:194-203.
- Jarvis MR, Mitra PP (2001) Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput* 13:717-749.
- Jonides J, Yantis S (1988) Uniqueness of abrupt visual onset in capturing attention. 43:346-354.
- Kastner S, Pinsk MA, De Weerd P, Desimone R, Ungerleider LG (1999) Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron* 22:751-761.
- Knight RT (1997) Distributed cortical network for visual attention. *J Cognitive Neurosci* 9:75-91.
- Knight RT, Grabowecky MF, Scabini D (1995) Role of human prefrontal cortex in attention control. *Adv Neurol* 66:21-34; discussion 34-26.
- Koch K, McLean J, Berry M, Sterling P, Balasubramanian V, Freed MA (2004) Efficiency of information transmission by retinal ganglion cells. *Curr Biol* 14:1523-1530.
- Koch K, McLean J, Segev R, Freed MA, Berry MJ, 2nd, Balasubramanian V, Sterling P (2006) How much the eye tells the brain. *Curr Biol* 16:1428-1434.
- Kopell N, Ermentrout GB, Whittington MA, Traub RD (2000) Gamma rhythms and beta rhythms have different synchronization properties 10.1073/pnas.97.4.1867. *PNAS* 97:1867-1872.
- Kusunoki M, Gottlieb J, Goldberg ME (2000) The lateral intraparietal area as a salience map: the representation of abrupt onset, stimulus motion, and task relevance. *Vision Res* 40:1459-1468.
- Lange K, Roder B (2006) Orienting attention to points in time improves stimulus processing both within and across modalities. *J Cogn Neurosci* 18:715-729.
- Lee J, Williford T, Maunsell JH (2007) Spatial attention and the latency of neuronal responses in macaque area V4. *J Neurosci* 27:9632-9637.
- Legatt AD, Arezzo J, Vaughan HG, Jr. (1980) Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. *J Neurosci Methods* 2:203-217.
- Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* 38:317-327.
- Lewis JW, Van Essen DC (2000) Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *J Comp Neurol* 428:112-137.
- Li Z (2002) A saliency map in primary visual cortex. *Trends in Cognitive Sciences* 6:9-16.
- Liu T, Slotnick SD, Serences JT, Yantis S (2003) Cortical mechanisms of feature-based attentional control. *Cereb Cortex* 13:1334-1343.

- Luck SJ, Chelazzi L, Hillyard SA, Desimone R (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology* 77:24-42.
- Mansouri FA, Matsumoto K, Tanaka K (2006) Prefrontal Cell Activities Related to Monkeys' Success and Failure in Adapting to Rule Changes in a Wisconsin Card Sorting Test Analog. *J Neurosci* 26:2745-2756.
- Martinez-Trujillo JC, Treue S (2002) Attentional Modulation Strength in Cortical Area MT Depends on Stimulus Contrast. *Neuron* 35:365-370.
- Mayer JS, Bittner RA, Nikolic D, Bledowski C, Goebel R, Linden DE (2007) Common neural substrates for visual working memory and attention. *Neuroimage* 36:441-453.
- McAdams CJ, Maunsell JHR (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience* 19:431-441.
- McDonald S, Bennett KM, Chambers H, Castiello U (1999) Covert orienting and focusing of attention in children with attention deficit hyperactivity disorder. *Neuropsychologia* 37:345-356.
- McLeod P, Driver J, Crisp J (1988) Visual search for a conjunction of movement and form is parallel. *Nature* 332:154-155.
- Miller BT, D'Esposito M (2005) Searching for "the top" in top-down control. *Neuron* 48:535-538.
- Miller EK (2000) The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience* 1:59-65.
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal function. *Annual Review of Neuroscience* 24:167-202.
- Miller EK, Gochin PM, Gross CG (1993) Suppression of visual responses of neurons in inferior temporal cortex of the awake macaque by addition of a second stimulus. *616:25-29.*
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154-5167.
- Mishra J, Fellous JM, Sejnowski TJ (2006) Selective attention through phase relationship of excitatory and inhibitory input synchrony in a model cortical neuron. *Neural Networks* 19:1329-1346.
- Moldakarimov S, Rollenhagen JE, Olson CR, Chow CC (2005) Competitive Dynamics in Cortical Responses to Visual Stimuli. *J Neurophysiol* 94:3388-3396.
- Moore T, Fallah M (2001) Control of eye movements and spatial attention. *Proc Natl Acad Sci U S A* 98:1273-1276.
- Moore T, Armstrong KM (2003) Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421:370-373.
- Moore T, Fallah M (2004) Microstimulation of the frontal eye field and its effects on covert spatial attention. *J Neurophysiol* 91:152-162.
- Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. *Science* 229:782-784.
- Motter BC (1993) Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology* 70:909-919.
- Murthy A, Thompson KG, Schall JD (2001) Dynamic dissociation of visual selection from saccade programming in frontal eye field. *Journal of neurophysiology* 86:2634-2637.
- Nachev P, Husain M (2006) Disorders of visual attention and the posterior parietal cortex. *Cortex* 42:766-773.
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping* 15:1-25.
- Niebur E, Koch C (1994) A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons. *Journal of Computational Neuroscience* 1:141-158.
- Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223-226.

- Palva JM, Palva S, Kaila K (2005) Phase Synchrony among Neuronal Oscillations in the Human Cortex. *J Neurosci* 25:3962-3972.
- Pandya DN, Barnes CL (1987) Architecture and connections of the frontal lobe. In: *The Frontal Lobes Revisited* (Perecman E, ed), pp 41-72. New York: The IRBN Press.
- Pandya DN, Yeterian EH (1990) Prefrontal Cortex in Relation to Other Cortical Areas in Rhesus Monkey - Architecture and Connections. *85:63-94*.
- Pasupathy A, Miller EK (2005) Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 433:873-876.
- Peterson MS, Kramer AF, Irwin DE (2004) Covert shifts of attention precede involuntary eye movements. *Percept Psychophys* 66:398-405.
- Petrides M, Pandya DN (1999) Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. *European Journal of Neuroscience* 11:1011-1036.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233-238.
- Porrino LJ, Crane AM, Goldman-Rakic PS (1981) Direct and indirect pathways from the amygdala to the frontal lobe in rhesus monkeys. *J Comp Neurol* 198:121-136.
- Posner MI, Snyder CR, Davidson BJ (1980) Attention and the detection of signals. *J Exp Psychol* 109:160-174.
- Pouget A, Driver J (2000) Relating unilateral neglect to the neural coding of space. *Curr Opin Neurobiol* 10:242-249.
- Pribram KH, Mishkin M, Rosvold HE, Kaplan SJ (1952) Effects on delayed-response performance of lesions of dorsolateral and ventromedial frontal cortex of baboons. *J Comp Physiol Psychol* 45:565-575.
- Rayner K, McConkie GW, Ehrlich S (1978) Eye movements and integrating information across fixations. *J Exp Psychol Hum Percept Perform* 4:529-544.
- Reynolds JH, Desimone R (2003) Interacting Roles of Attention and Visual Saliency in V4. *Neuron* 37:853-863.
- Reynolds JH, Chelazzi L (2004) Attentional Modulation of Visual Processing. *Annual Review of Neuroscience* 27:611-647.
- Reynolds JH, Chelazzi L, Desimone R (1999) Competitive Mechanisms Subserve Attention in Macaque Areas V2 and V4. *J Neurosci* 19:1736-1753.
- Reynolds JH, Pasternak T, Desimone R (2000) Attention increases sensitivity of V4 neurons. *Neuron* 26:703-714.
- Rizzolatti G, Riggio L, Dascola I, Umiltà C (1987) Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia* 25:31-40.
- Rubia K, Overmeyer S, Taylor E, Brammer M, Williams SC, Simmons A, Bullmore ET (1999) Hypofrontality in attention deficit hyperactivity disorder during higher-order motor control: a study with functional MRI. *Am J Psychiatry* 156:891-896.
- Salinas E, Sejnowski TJ (2000) Impact of Correlated Synaptic Input on Output Firing Rate and Variability in Simple Neuronal Models. *J Neurosci* 20:6193-6209.
- Salinas E, Sejnowski TJ (2001) Correlated Neuronal Activity and the Flow of Neural Information. *Nature Reviews Neuroscience* 2:539-550.
- Sato T, Murthy A, Thompson KG, Schall JD (2001) Search efficiency but not response interference affects visual selection in frontal eye field. *Neuron* 30:583-591.
- Sato TR, Watanabe K, Thompson KG, Schall JD (2003) Effect of target-distractor similarity on FEF visual selection in the absence of the target. *Exp Brain Res* 151:356-363.

- Schall JD (2002a) The neural selection and control of saccades by the frontal eye field. *Philos Trans R Soc Lond B Biol Sci* 357:1073-1082.
- Schall JD (2002b) Decision Making: Neural Correlates of Response Time. *Current Biology* 12:R800-R801.
- Schall JD, Hanes DP (1993) Neural basis of saccade target selection in frontal eye field during visual search. *Nature* 366:467-469.
- Schall JD, Bichot NP (1998) Neural correlates of visual and motor decision processes. *Curr Opin Neurobiol* 8:211-217.
- Schall JD, Hanes DP, Thompson KG, King DJ (1995) Saccade target selection in frontal eye field of macaque. I. Visual and premovement activation. *J Neurosci* 15:6905-6918.
- Schanze T, Eckhorn R (1997) Phase correlation among rhythms present at different frequencies: spectral methods, application to microelectrode recordings from visual cortex and functional implications. *Int J Psychophysiol* 26:171-189.
- Schiller PH, Tehovnik EJ (2001) Look and see: how the brain moves your eyes about. *Prog Brain Res* 134:127-142.
- Schwartz ML, Goldman-Rakic PS (1984) Callosal and intrahemispheric connectivity of the prefrontal association cortex in rhesus monkey: Relation between intraparietal and principal sulcal cortex. *Journal of Comparative Neurology* 226:403-420.
- Schwarz G (1978) Estimating the Dimension of a Model. *The Annals of Statistics* 6:461-464.
- Selemon L, Goldman-Rakic P (1988) Common cortical and subcortical targets of the dorsolateral prefrontal and posterior parietal cortices in the rhesus monkey: evidence for a distributed neural network subserving spatially guided behavior. *J Neurosci* 8:4049-4068.
- Sereno AB, Maunsell JHR (1998) Shape selectivity in primate lateral intraparietal cortex. *Nature* 395:500-503.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916-1936.
- Shannon C, Weaver W (1949) *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois.
- Sharma J, Dragoi V, Tenenbaum JB, Miller EK, Sur M (2003) V1 neurons signal acquisition of an internal representation of stimulus location. *Science* 300:1758-1763.
- Smith EE, Jonides J (1997) Working memory: a view from neuroimaging. *Cognit Psychol* 33:5-42.
- Solomon JA, Lavie N, Morgan MJ (1997) Contrast discrimination function: spatial cuing effects. *J Opt Soc Am A Opt Image Sci Vis* 14:2443-2448.
- Stanton GB, Goldberg ME, Bruce CJ (1988) Frontal eye field efferents in the macaque monkey: I. Subcortical pathways and topography of striatal and thalamic terminal fields. *J Comp Neurol* 271:473-492.
- Stanton GB, Bruce CJ, Goldberg ME (1993) Topography of projections to the frontal lobe from the macaque frontal eye fields. *J Comp Neurol* 330:286-301.
- Stanton GB, Bruce CJ, Goldberg ME (1995) Topography of projections to posterior cortical areas from the macaque frontal eye fields. *J Comp Neurol* 353:291-305.
- Szabo M, Almeida R, Deco G, Stetter M (2004) Cooperation and biased competition model can explain attentional filtering in the prefrontal cortex. *European Journal of Neuroscience* 19:1969-1977.
- Tallon-Baudry C, Bertrand O, Peronnet F, Pernier J (1998) Induced gamma-band activity during the delay of a visual short-term memory task in humans. *Journal of Neuroscience* 18:4244-4254.
- Tehovnik EJ, Sommer MA, Chou IH, Slocum WM, Schiller PH (2000) Eye fields in the frontal lobes of primates. *Brain Res Brain Res Rev* 32:413-448.
- Thomas NWD, Pare M (2007) Temporal Processing of Saccade Targets in Parietal Cortex Area LIP During Visual Search. *J Neurophysiol* 97:942-947.

- Thompson KG, Bichot NP (2005a) A visual salience map in the primate frontal eye field. *Prog Brain Res* 147:251-262.
- Thompson KG, Bichot NP (2005b) A visual salience map in the primate frontal eye field. *Prog Brain Res* 147:251-262.
- Thompson KG, Bichot NP, Schall JD (1997) Dissociation of Visual Discrimination From Saccade Programming in Macaque Frontal Eye Field. In: *Am Physiological Soc*.
- Thompson KG, Bichot NP, Sato TR (2005a) Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. *J Neurophysiol* 93:337-351. Epub 2004 Aug 2018.
- Thompson KG, Bichot NP, Sato TR (2005b) Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. *J Neurophysiol* 93:337-351.
- Thompson KG, Biscoe KL, Sato TR (2005c) Neuronal Basis of Covert Spatial Attention in the Frontal Eye Field. *J Neurosci* 25:9479-9487.
- Tiesinga PH, Fellous JM, Jose JV, Sejnowski TJ (2002) Information transfer in entrained cortical neurons. *Network* 13:41-66.
- Tiesinga PHE, Sejnowski TJ (2004) Rapid Temporal Modulation of Synchrony by Competition in Cortical Interneuron Networks. *Neural Computation* 16:251-275.
- Tomita H, Ohbayashi M, Nakahara K, Hasegawa I, Miyashita Y (1999) Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature* 401:699-703.
- Toth LJ, Assad JA (2002) Dynamic coding of behaviourally relevant stimuli in parietal cortex. *Nature* 415:165-168.
- Treisman A (1996) The binding problem. *Curr Opin Neurobiol* 6:171-178.
- Treisman A (1998) Feature binding, attention and object perception. *Philos Trans R Soc Lond B Biol Sci* 353:1295-1306.
- Treisman A, Gormican S (1988) Feature analysis in early vision: evidence from search asymmetries. *95:15-48*.
- Treisman AM, Gelade G (1980) A feature-integration theory of attention. *Cognitive Psychology* 12:97-136.
- Treisman AM, Kanwisher NG (1998) Perceiving visually presented objects: recognition, awareness, and modularity. *Current Opinion in Neurobiology* 8:218-226.
- Treue S, Maunsell JHR (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382:539-541.
- Treue S, Trujillo JCM (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575-579.
- van Swinderen B (2007) Attention-like processes in *Drosophila* require short-term memory genes. *Science* 315:1590-1593.
- von Stein A, Sarnthein J (2000) Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *International Journal of Psychophysiology* 38:301-313.
- von Stein A, Chiang C, Konig P (2000) Top-down processing mediated by interareal synchronization. *PNAS* 97:14748-14753.
- Wallis JD, Miller EK (2003) Neuronal activity in the primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *European Journal of Neuroscience* 18:2069-2081.
- Wallis JD, Anderson KC, Miller EK (2000) Neuronal representation of abstract rules in the orbital and lateral prefrontal cortices (PFC). *Soc Neurosci Abs* in press.
- Watanabe M (1996) Reward expectancy in primate prefrontal neurons. *Nature* 382:629-632.

- White IM, Wise SP (1999) Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research* 126:315-335.
- Williams ZM, Elfar JC, Eskandar EN, Toth LJ, Assad JA (2003) Parietal activity and the perceived direction of ambiguous apparent motion. *Nat Neurosci* 6:616-623.
- Williford T, Maunsell JHR (2006) Effects of Spatial Attention on Contrast Response Functions in Macaque Area V4. *J Neurophysiol* 96:40-54.
- Wolfe JM, Horowitz TS (2004) What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience* 5:495-501.
- Wolfe JM, Cave KR, Franzel SL (1989) Guided search: an alternative to the feature integration model for visual search. *J Exp Psychol Hum Percept Perform* 15:419-433.
- Womelsdorf T, Fries P (2007) The role of neuronal synchronization in selective attention. *Curr Opin Neurobiol* 17:154-160.
- Womelsdorf T, Anton-Erxleben K, Pieper F, Treue S (2006a) Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nat Neurosci* 9:1156-1160.
- Womelsdorf T, Fries P, Mitra PP, Desimone R (2006b) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439:733-736. Epub 2005 Dec 2021.
- Womelsdorf T, Schoffelen J-M, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P (2007) Modulation of Neuronal Interactions Through Neuronal Synchronization. *Science* 316:1609-1612.
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* advanced online publication.
- Yantis S, Serences JT (2003) Cortical mechanisms of space-based and object-based attentional control. *Curr Opin Neurobiol* 13:187-193.
- Zang YF, Jin Z, Weng XC, Zhang L, Zeng YW, Yang L, Wang YF, Seidman LJ, Faraone SV (2005) Functional MRI in attention-deficit hyperactivity disorder: evidence for hypofrontality. *Brain Dev* 27:544-550.