UNIVERSITY OF CALIFORNIA, SAN DIEGO

Global Effects of Attention in Human Visual Cortex

A dissertation submitted in partial satisfaction of the
requirements for the degree of Doctor of Philosophy
in
Neurosciences
by
Melissa Saenz

Committee in charge:

Professor Karen Dobkins, Chair
Professor Geoffrey Boynton
Professor Steven Hillyard
Professor John Reynolds
Professor Marty Sereno

2002

The dissertation of Melissa Saenz is approved, and it is acceptable in quality and form for publication on microfilm:

Chair

University of California, San Diego

2002

TABLE OF CONTENTS

LIST OF FIGURES AND TABLES

ACKNOWLEDGEMENTS

It has been my unique privilege to be Geoff Boynton's first graduate student. Being the only student in the lab was like being the only child in a family … I was spoiled! Now that his lab has grown, more students are benefiting from his natural talents as a scientist and mentor. I will always be grateful to Geoff for his mentorship and friendship.

Special mention must also go to my collaborator, Giedrius Buracas, whose enthusiastic support helped to keep this thesis project alive when early results lacked luster! In the true spirit of international scientific collaboration, Giedrius organized a fabulous vision conference in his home country of Lithuania that Geoff and I had the good fortune to attend.

I wish to thank Boynton lab members Rob Duncan, Ione Fine, Ed Hubbard, Eva Finney and August Tuan for their contribution to this project.

I wish to thank Professors Karen Dobkins, John Reynolds, Steve Hillyard, and Marty Sereno for all of their helpful insight as my thesis committee. I am also grateful to Sascha du Lac and Rich Krauzlis who welcomed me into their labs during first-year rotations.

Special thanks must also go to SBM, GDH, DDL, SC, EMT, AST, SJK, YJB, and DAN who are dispassionately referred to as "human subjects" in the following pages. This work depended on all of them!

Warm thoughts go to my family in San Antonio, TX who send me their love and support always.

Loving thoughts go to my fiancé, Christophe. Moving to San Diego meant moving 120 miles away from him. Despite the distance our relationship did not suffer!

This dissertation, in part, is a reprint of the material as it appears in the following:

1. Saenz M., Buracas, G.T., and Boynton, G.M. *Nature Neuroscience*, 5, 631-632, 2002.

2. Saenz M., Buracas, G.T., and Boynton, G.M. *Vision Research,* in press, submitted 2002.

VITA

1998	B.S. in Biology, California Institute of Technology
2002	Ph.D. in Neurosciences, University of California, San Diego

PUBLICATIONS

Saenz M., Buracas G.T., Boynton G.M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*. 5(7), 631-2.

Saenz M, Buracas GT, Boynton GM, (2002) “Global feature-based attention for motion and color”, *Vision Research*. in press.

Kornylo K., Dill N., Saenz M., Krauzlis R.J. (2002) “Canceling of pursuit and saccadic eye movements in humans and monkeys”. *Journal of Neurophysiology*. submitted.

Lewis T.B., Saenz M., O’Connell P., Leach R.J. (1994) “Localization of glucose-dependent insulinotropic polypeptide (GIP) to a gene cluster on chromosome 17q”. *Genomics* 19(3), 589-91.

ABSTRACTS

Saenz M., Boynton G.M. (2002). Feature-based attention modulates lateral interactions in contrast detection. *Society for Neuroscience,* 32nd annual meeting.

Saenz M., Buracas G.T., Boynton G.M. (2002). Global effects of feature-based visual attention to direction of motion and color. *Human Brain Mapping Conference,* Sendai, Japan.

Saenz M., Buracas G.T., Boynton G.M. (2002) Global effects of feature-based attention to direction of motion and color. *Vision Sciences Society*, 2nd annual meeting.

Saenz M., Buracas G.T., Boynton G.M. (2001) Feature-based attentional effects in early human visual cortex. *Society for Neuroscience*, 31st annual meeting.

Saenz M., Buracas G.T., Boynton G.M. (2001) Feature-based attentional effects in early human visual cortex. *Vision Sciences Society*, 1st annual meeting.

Buracas G.T, Saenz M., and Boynton G.M. (2000) Comparing fMRI responses to discrimination thresholds using population-based models. *Society for Neuroscience,* 30th annual meeting.

Saenz M., Krauzlis R.J. Countermanding Pursuit and Saccades. (1999) *Society for Neuroscience,* 29th annual meeting.

ABSTRACT OF THE DISSERTATION

Global Effects of Attention in Human Visual Cortex

by

Melissa Saenz

Doctor of Philosophy in Neurosciences

University of California, San Diego, 2002

Professor Karen Dobkins, Chair

At any moment in time more visual information reaches our eyes than our brains can fully process. What we see critically depends on how we are able to distribute attention within the visual scene. Early theories compared the distribution of visual attention to a moving “spotlight” enhancing the processing of stimuli within an attended location. This idea was based on substantial evidence that attention *locally* influences both neuronal and behavioral responses to visual objects presented within an attended location. Less well understood is how attention *globally* influences responses to visual objects located outside of the spatial focus of attention.

The experiments presented here test the hypothesis that visual attention globally enhances the processing of visual objects with behaviorally relevant features. This *feature-based* attention could be mediated by neurons in early stages of cortical visual processing that are tuned for elementary visual features including orientation, motion direction, and color. For example, this hypothesis predicts that when an observer searches a shelf for a red book, attention would sensitize neurons that are responsive to the color red with receptive field locations throughout the visual scene. A feature-based increase in neuronal sensitivity thus suggests a way in which attention could affect visual processing not just locally, but globally, throughout the visual field.

Chapter 1 presents the results of human functional magnetic resonance imaging (fMRI) experiments in which attention to a stimulus feature (motion direction or color) increased neuronal responses in multiple cortical visual areas to a spatially distant stimulus that was ignored, but shared the same feature. Chapter 2 presents complementary psychophysical studies in which attention to a stimulus feature facilitated the concurrent discrimination of a

spatially distant stimulus that had the attended feature. Chapter 3 presents a different psychophysical paradigm showing that feature-based attention (to orientation) influences well-known contextual interactions between stimuli with common features.

Together, these neuronal and behavioral findings support the hypothesis that feature-based attention globally enhances the processing of visual objects with behaviorally relevant features. This attentional mechanism could be highly useful to the visual system because the location of relevant visual objects is often not known in advance.

Introduction

What we see greatly depends on our ability to distribute attention within the visual field. Early theories compared the distribution of visual attention to a moving "spotlight" enhancing stimulus processing within its focus (Posner, Snyder et al. 1980; Eriksen and St James 1986). The spotlight metaphor was based on substantial evidence that attention *locally* modulates neuronal and behavioral responses to stimuli within an attended location (see below for review). Less well understood are the *global* effects of attention: how does attention affect responses to stimuli outside the focus of attention? The global distribution of attention may contribute significantly to our visual experience for a number of reasons including the following:

- Attention is not an all-or-none phenomenon. While attention is focused at one location an "ambient" level of visual attention may be available outside the focus (Braun and Julesz 1998; Nakayama and Joseph 1999).

- Attention may not be limited to a single focus. Psychophysical studies suggest that observers can simultaneously attend to objects at separate locations while ignoring irrelevant objects located between them (Kramer and Hahn 1995; Bichot, Cave et al. 1999).

- Attention is a dynamic process. While focusing attention at one location, the visual system must determine where to attend next. The global distribution of attention could direct subsequent shifts in the attentional focus and in eye position.

- Attention may select stimuli based on features rather than spatial location. Observes can selectively attend to a subset of overlapping stimuli with matching features (e.g. the color red) within the same location (Valdes-Sosa, Bobes et al. 1998; Alais and Blake 1999) or to one of two spatially overlapping objects (Duncan 1984; Blaser, Pylyshyn et al. 2000). Feature-based and object-based selection may allow for spatially complex distributions of attention within the visual field.

While certainly not an exhaustive list, these points are all motivations for studying the effects of attention outside of a single spatial focus. While the spotlight model is currently viewed as insufficient to describe the complexity of our visual experience, a majority of studies have still focused on how attention locally affects responses to *attended* stimuli. The work described in this dissertation is concerned with the global rather than local effects of visual attention including how attention effects responses to *ignored* stimuli outside the focus of attention.

The primary motivation for this work comes from the final point on the list: feature-based attention. Recent theories of visual attention propose that attention enhances the activity of cortical neurons that encode behaviorally relevant stimulus properties including, not only spatial location, but also features (Desimone and Duncan 1995; Duncan, Humphreys et al. 1997; Treue and Martinez Trujillo 1999). For example, these models predict that when an observer searches for a red book on a shelf, attention sensitizes neurons with receptive field locations throughout the visual scene that are tuned to the color red. A feature-based increase in signal strength thus suggests a way in which attention could affect visual processing not just locally but globally throughout the visual scene.

Critically, the feature-based model predicts that attention influences responses to unattended stimuli outside the focus of attention. In support of this prediction, Treue and Martinez-Trujillo (1999) found that attention to a visual object with a particular direction of motion increased responses of neurons tuned to that direction in macaque area MT, even among neurons with receptive fields outside the attended location. This effect of feature-based attention was far-reaching, affecting the responses of neurons with receptive field locations in the visual hemifield opposite of the attended location.

Here, we confirm and extend these findings in the human visual system. First we used functional magnetic resonance imaging (fMRI) to show that attention to a stimulus feature increased neuronal responses to a spatially distant stimulus that was ignored, but had the same feature compared to an opposing feature (Chapter 1). We found this global feature-based attention effect with two features for which cortical neurons show tuning: motion direction and color. FMRI responses were measured across the entire occipital lobe allowing us to explore multiple cortical visual areas where neurons are tuned to these features, including primary visual cortex (V1).

Second, we tested for a behavioral consequence of global feature-based attention (Chapter 2). If feature-based attention does indeed globally improve the processing of stimuli with attended features, this should profoundly impact what we see. In complementary psychophysical experiments, attention to a stimulus feature (motion direction or color) facilitated the concurrent discrimination of a spatially distant stimulus that had the same feature compared to the opposing feature. Both fMRI and psychophysical experiments were performed using similar stimuli in order to relate the neuronal and behavioral results as best as possible. A different psychophysical paradigm shows that feature-based attention to a third feature, orientation, influences well-known contextual interactions between stimuli with common orientations (Chapter 3).

Together, our results provide complementary neuronal and behavioral evidence that feature-based attention globally improves the processing of stimuli with behaviorally relevant features. Such a mechanism could be highly useful to the visual system because the locations of relevant stimuli are not always known in advance. Feature-specific signal enhancement could be useful in identifying relevant objects during a visual search (e.g. searching for a red book on a shelf) or in grouping stimuli with common features as part of a common object (e.g. tracking an animal moving behind trees in a forest).

This introduction continues with a brief review of the literature on spatially directed attention followed by a discussion of non-spatial, feature-based studies of attention. Finally, the scope of the thesis will be presented in more detail.

Spatial Attention

Experimental evidence confirms that attention determines what we see and how well we see it. Posner and colleagues were among the first to systematically study the behavioral costs and benefits of focusing attention on a limited region of visual space (Posner, Snyder et al. 1980). They developed a paradigm in which observers maintained eye position on a central fixation point and directed visual attention to a cued location in the periphery. Visual attention was thus *covertly* directed, without accompanying head or eye movements. Attending to a particular location reduced reaction times and lowered thresholds for detecting and discriminating stimuli subsequently presented there. Improved performance at attended locations was coupled with reduced performance at unattended locations. This basic result has been replicated many times (for excellent review see Pashler 1998).

The reduced responsiveness to unattended information can indeed be striking. Observers may fail to see highly salient but unexpected visual objects when engaged in an attentionally demanding visual task, a phenomenon aptly termed “inattentional blindness”. In one well-noted example, a majority of observers failed to notice a person in a gorilla suit walk through a scene (presented on video) while the observers performed a difficult ball-tracking task within the same scene ("Gorillas in our Midst", Simons and Chabris 1999). In other instances termed “change blindness”, observers fail to see abrupt changes in the visual scene that occur at unattended locations within their field of view (Chun and Marois 2002; Rensink 2002). These counter-intuitive but well-documented phenomena suggest that we only become aware of those visual objects that receive focused attention.

These behavioral effects are associated with clear neuronal effects of attention. Electrophysiological studies in non-human animals as well as ERP (event-related potential) and neuroimaging studies in the human show that visual cortical responses are strongly modulated by the attentive state of the observer. In a typical neurophysiological study of attention, the visual display is unchanged and observers are instructed to direct attention to one region of the display versus another while maintaining constant fixation (similar to the Posner paradigm). In order to maintain attention at a particular location, a demanding visual task is performed there. Because there is no change in the visual display or in eye position, any resulting changes in neuronal response across conditions can be attributed to non-retinal factors such as attention or task demands. Task demands and general levels of arousal must be controlled for across all conditions to isolate the effects of attention.

From the broad literature on spatial attention a number of conclusions can be drawn. First, attention locally modulates responses to attended stimuli. Under many conditions a neuron's response is increased when attention is directed to a single stimulus within its receptive field; however, greater attentional modulations are found when multiple stimuli compete for attentional selection (Moran and Desimone 1985; Haenny and Schiller 1988, Motter 1993; Luck, Chelazzi et al. 1997; Treue and Maunsell 1999). For example, Motter (1993) found that attentional modulation of responses in areas V1, V2, and V4 were more robust when attention was directed to a visual object in the presence of nearby distracting objects. The role of competition between multiple stimuli and the influence of attention on this competition will be discussed further in the next section.

Correspondingly, neuronal population responses in the human are also enhanced for attended stimuli and that enhancement is spatially localized to areas of cortex representing the attended location (Heinze, Mangun et al. 1994, for review see Kanwisher and Wojciulik 2000). Furthermore, both neuroimaging and neurophysiological studies report increases in the baseline activity of neurons with receptive field locations corresponding to the attended location, even when no stimulus is present (Luck, Chelazzi et al. 1997; Reynolds, Chelazzi et al. 1999; Kastner, Pinsk et al. 1999). Attended regions of the visual display may thus be primed for the processing of subsequently presented information.

Attention modulates responses at all cortical levels including the primary visual cortex (single-unit studies: Motter 1993; Roelfsema, Lamme et al. 1998; Ito and Gilbert 1999 and neuroimaging studies: Tootell, Hadjikhani et al. 1998; Brefczynski and DeYoe 1999; Gandhi, Heeger et al. 1999; Martinez, Anllo-Vento et al. 1999; Somers, Dale et al. 1999) and possibly the lateral geniculate nucleus (O'Connor, Fukui et al. 2002). However, the magnitude of attentional modulation is generally larger at later stages of cortical processing (Haenny and Schiller 1988; McAdams and Maunsell 1999). The delayed attentional modulation of stimulus-evoked responses in early visual areas (typically 50-100 ms after response onset in V1) may indicate the time course of feedback from higher cortical areas (Noesselt, Hillyard et al. 2002),

Several studies find that attention modulates neuronal responses without changing the tuning properties of a cell. In their elegant study, McAdams and Maunsell (1999) measured the effect of attention on the orientation tuning of individual neurons in macaque area V4. A grating of a range of different orientations was presented within the receptive field of a recorded neuron under two attentional conditions. The animal performed a visual task on either the grating within the receptive field or on another visual object presented elsewhere in the display. Attention caused an overall multiplicative gain increase in orientation tuning functions without changing cell selectivity for orientation (neither the tuning peak nor tuning width of the tuning functions were changed). This result was striking because a similar gain increase could be obtained by increasing the contrast of the stimulus (Tolhurst1973; Sclar and Freeman 1982) suggesting that attention effectively increased the contrast of the attended stimulus.

To test this prediction, Reynolds and colleagues measured contrast response functions of V4 neurons to gratings that were either attended or unattended (Reynolds, Pasternak et al. 2000). Typically, neuronal responses increase monotonically (towards saturation) with increasing stimulus contrast. An increase in contrast sensitivity would be reflected in a horizontal shift of the contrast response function (i.e. a reduction in the contrast needed to elicit a minimal response, a maximal increase in response to intermediate contrasts, and a minimal increase in response to contrasts above saturation). Their findings were consistent with the predicted shift in the contrast response function. Increased neuronal sensitivity to contrast also predicts that attention would reduce perceptual contrast detection thresholds, a result that has been reported for human observers (Zenger, Braun et al. 2000). Independent studies of direction-selective cells in macaque area MT also concluded that attention effectively increases stimulus contrast without changing the tuning properties of the cell (Martinez-Trujillo and Treue 2002).

In light of these data, attention cannot be understood as a process that simply increases responses to stimuli within an attended location. The size of the effect depends on how well the stimulus itself drives the cell. If the stimulus does not match the tuning properties of the cell (e.g. has a non-optimal orientation) or has already saturated the neuron’s contrast response function, the affect of attention may be minimal.

Guided by the advances made in the study spatial directed attention, we seek to show that a non-spatial, feature-based mechanism of attention also modulates both behavioral and neuronal responses to stimuli. We will test whether feature-based attention, like spatial attention influences multiple early stages of cortical visual processing (including area V1), whether the influence is greater at higher levels of cortical processing, and whether the influence is greater in the presence of distracting stimuli.

Feature-based attention

Just as knowing the location of a stimulus in advance improves its detection, advance knowledge about certain features of a stimulus improves its detection. For example, Sekuler and Ball (1977) instructed observers to detect a barely visible moving object that was briefly flashed on a visual display. Performance was greatly improved given advance knowledge of either the direction or speed of the moving object (2-AFC contrast detection). Similar facilitation has been shown with prior knowledge of the location, color, or spatial frequency of a visual object or the frequency of an auditory tone (Posner, Snyder et al. 1980; Davis, Kramer et al. 1983; Green 1961). In these examples, attentional selection is based on feature content rather than spatial location.

Here we use the term *feature-based attention* to describe these and other related phenomena that cannot be explained by spatial location alone. Feature-based attention can be studied independently of spatial attention with the used of spatially overlapping stimuli. Observers can selectively attend to a subset of stimuli within the same location that share a common feature (O'Craven, Rosen et al. 1997; Valdes-Sosa, Bobes et al. 1998; Alais and Blake 1999) or to one of two spatially overlapping objects (Duncan 1984; Blaser, Pylyshyn et al. 2000).

Feature-based selection may depend on the feature-tuning rather than the spatial-tuning (receptive field location) of neurons. In the previous section, evidence was presented that the effect of attention on a neuron’s response depends on how well the attended stimulus matches the feature selectivity of the neuron. The following section explores this feature dependence in more detail. We explore the possibility that feature-based selection allows for far-reaching and spatially complex distributions of attention throughout the visual field.

Feature-based attention – neurophysiology

In a groundbreaking study, Moran and Desimone demonstrated that attention modulated the responses of visual cortical neurons in a manner that depended on the feature selectivity of the cell (Moran and Desimone 1985). A pair of stimuli was placed within the receptive field of a recorded neuron in macaque areas V4 and IT. One matched the feature-selectivity of the cell and evoked a large response from the cell when presented alone, while the other stimulus did not. Neuronal firing rates increased attention was directed to the preferred stimulus of the pair and decreased when attention was directed to the non-preferred stimulus of the pair. These results clearly demonstrated that visual attention was more sophisticated than a simple spotlight that enhanced the responses of neurons with receptive fields falling within an attended location. Rather, the effect of attention on a neuron’s response (enhancement or suppression) depended on how the features of the attended stimulus matched the stimulus selectivity of the neuron. Similar results have been confirmed for a range of stimuli including color stimuli in V2, V4, and IT (Luck, Chelazzi et al. 1997; Reynolds, Chelazzi et al. 1999), motion stimuli in MT (Treue and Maunsell 1996; Treue and Maunsell 1999) and complex objects in V4 and IT (Chelazzi, Miller et al. 1993; Chelazzi, Duncan et al. 1998; Chelazzi, Miller et al. 2001).

A *biased competition model* has been proposed that accounts for this feature-specific attentional modulation (Desimone and Duncan 1995; Reynolds, Chelazzi et al. 1999) The model proposes that simultaneously presented visual objects activate competing populations of neurons. That competition may be biased in favor of neurons representing a particular object by either bottom-up sources (i.e. increasing the contrast of one of the objects) or by top-down attentional modulation (i.e. attending to one of the objects). The model proposes that a top-down attentional bias could be achieved by increased synaptic efficacy onto neurons tuned to the attended stimulus. Directing attention to one of two stimuli would result in a response enhancement for cells preferring that stimulus and in response suppression for cells preferring the unattended stimulus. In the absence of a competing stimulus, the top-down attentional bias may still enhance the activity of cortical neurons that encode behaviorally relevant stimulus properties.

The biasing effects of attention could apply to a range of stimulus properties for which cortical neurons are selective including spatial location (receptive field location), low-level features such as color and direction of motion, and even more complex object identities (e.g. neurons that respond well to individual faces have been identified in both macaques and humans Tovee, Rolls et al. 1994; Kanwisher, McDermott et al. 1997; Kreiman, Koch et al. 2000). In this way, the model offers a unified account of both spatial, feature-based, and perhaps object-based effects of attention. Importantly, feature-selective enhancement suggests a way in which attention could affect the processing of ignored stimuli throughout the visual field. Attention could sensitize neurons tuned to relevant stimulus features with receptive field locations *outside* the attended location.

Are the top-down biasing effects of feature-based attention far reaching enough to influence the processing of visual objects outside the attended location? In support of this, Chelazzi and colleagues showed that searching for a visual stimulus increased the firing rate of IT neurons tuned to that stimulus during a time period *prior* to stimulus presentation. This modulation of baseline firing rates was feature-driven and far-reaching because the exact location of the upcoming target was unknown (but the location was limited to a single visual hemifield). Motter (1994) trained monkeys to perform a discrimination task on bars of a cued color and found that responses of V4 neurons increased when the receptive field stimulus matched the cued color. In both cases, the effect depended on how the behaviorally relevant stimulus matched the feature-tuning of the recorded neuron and was not spatially driven.

Two recent studies have shown that the modulatory effects of feature-based attention are indeed far-reaching, influencing response to stimuli outside the attended location (Treue and Martinez Trujillo 1999; McAdams and Maunsell 2000). Treue and Martinez-Trujillo reported feature-specific attentional modulation of stimulus-evoked responses in macaque area MT. In their experiment an ignored random dots stimulus, moving coherently in the preferred direction, was presented inside the receptive field of a directionally tuned neuron. Attention was directed to a second stimulus, outside the receptive field, that either moved in the same or in the opposite direction. On average, neuronal responses to the ignored stimulus increased when the monkey attended the preferred direction and decreased when the monkey attended the opposing direction (compared to passive viewing trials). To account for these results, the authors proposed a *feature-similarity gain model* in which feature-based attention modulates the gain of cortical neurons that are selective for the behaviorally relevant stimulus property. The model emphasizes that the direction of the gain change (decrease or increase) depends on how well the attended properties (location or features) match the stimulus selectivity of the neuron and also emphasizes that the modulation will reach neurons with receptive field locations well outside the attended location.

McAdams and Maunsell also presented an attended and an ignored stimulus in opposite visual hemifields (McAdams and Maunsell 2000). The response to an ignored grating stimulus was increased when the monkey performed an orientation discrimination task on a second grating compared to when it performed a color discrimination task on a solid colored patch. This paradigm did not isolate the effects of feature-based attention because the compared conditions differed not only in the attended feature but also in the attended stimulus itself and the task involved. However, to our knowledge, this is the only other study to show that the response to an ignored stimulus outside the focus of attention depended on the features of the attended stimulus.

Feature-based attention - human neurophysiology

Event-related brain potentials (ERP) in the human are modulated differently by selective attention to location and to non-spatial features: spatial attention modulates early ERP components (P1 and N1) while attention to non-spatial features evokes a later response (selection negativity). Hillyard and colleagues tested whether the feature-based response could be evoked by stimuli outside the attended location (Hillyard and Munte 1984; Anllo-Vento and Hillyard 1996). Observers were instructed to respond to stimuli with a particular feature (a particular individual color or motion direction) when flashed at a cued location (left or right hemifield). Stimuli with the cued feature evoked the feature-specific ERP response when flashed at the attended location and *not* when flashed at the ignored location. These human results are quite different from the results of the two monkey physiology studies that did report effects of attention outside the attended location. Differences in stimuli and task demands may have contributed to the different outcomes. Here, we will test for global effects of feature-based attention in the human using stimuli that are more similar to those used in the monkey studies. Our experiments may help to resolve this difference in results across species.

A number of human neuroimaging studies (both PET and fMRI) have taken another approach to the study of feature-based attention. They have tested the hypothesis that performing a task on a particular feature type (e.g. color vs. motion) specifically activates brain areas specialized for processing that feature type (e.g. V4 or MT). Indeed, if attention to a stimulus feature modulates the responses of individual neurons tuned to that feature, it might be expected that the greatest net modulation of activity would be measured in brain areas with a high proportion of neurons sensitive to that feature. Corbetta and colleagues compared PET responses while subjects performed either a color, velocity, or shape discrimination task on different elements of an identical visual display (Corbetta, Miezin et al. 1990; Corbetta, Miezin et al. 1991). As predicted, each task condition resulted in a unique pattern of activation and activity during the motion task co-localized with area MT+. Additional support comes from fMRI studies finding greater activity in MT+ when subjects attended to motion compared to another feature on unchanged visual displays (O'Craven, Rosen et al. 1997; Huk and Heeger 2000).

From these studies it can be concluded that visual cortical responses depend not only on the stimulus itself but also on the task being performed. The modulated responses cannot be attributed to spatial attention due to an absence of spatial shifts in attention. However, the responses cannot be related to feature-based attention alone because the compared conditions differ not only in the attended feature but also in the task being performed. In our studies we will isolate the effects of feature-based attention by using spatially overlapping stimuli and by keeping task demands constant across compared conditions.

In addition, it is important to be clear about how we define a ‘feature’ in our experiments. In all of our experiments, ‘feature’ is defined as a property within a stimulus dimension. For example, upward and downward directions of motion are two opposing features within the stimulus dimension of motion direction, and red and green are opposing features within the stimulus dimension of color. In this way, observers can be instructed to perform the same task (e.g. speed discrimination) on either of two overlapping stimuli with different features (e.g. upward vs. downward motion) without changing the nature of the task being performed.

Summary and scope of the dissertation

In summary, the behavioral benefits associated with attention and the striking costs in the absence of attention make clear the importance of distributing attention as efficiently as possible. The accumulating evidence suggests that attention enhances responses to stimuli with behaviorally relevant properties including spatial location, features, and possibly object properties. Treue and Martinez-Trujillo presented strong evidence that the effects of feature-based attention can be far-reaching, modulating macaque MT responses to an ignored stimulus well outside the focus of attention. This dissertation addresses a number of important questions that remain concerning the global effects of feature-based attention:

1. Does feature-based attention globally modulate neuronal responses to stimuli outside the attended location in the human visual system?
2. If so, does feature-based attention affect the multiple areas of cortical visual processing, including primary visual cortex, where feature-tuned neurons are located?
3. Do the global effects of feature-based attention extend to other features, besides motion direction, for which cortical neurons show tuning such as color and orientation?
4. If feature-based attention globally improves the processing of stimuli with a behaviorally relevant feature, this may profoundly affect what we see. What are the behavioral consequences?
5. Do the effects of feature-based attention increase in the presence of competing stimuli (as is consistently seen with spatial attention)?

Chapter 1 presents human fMRI experiments that address the first three questions. In the first experiment, observers were presented with two stimuli, one to attend and one to ignore, placed to the left and right of a central fixation point. The attended stimulus was a circular aperture of two overlapping fields of upward and downward moving dots, and the ignored stimulus was a circular aperture of a single field of dots moving in either direction, up or down. On the attended side, subjects performed a speed discrimination task alternately on the upward and downward moving fields of dots. Because the fields of dots on the attended side were overlapping, either direction of motion could be attended without changing the stimulus or the spatial distribution of attention. We tested for attentional modulation in the response to the *ignored* stimulus in multiple visual areas representing the earliest stages of cortical visual processing (V1, V2, V3, V3A, and MT+). Additional experiments used color and orientation as the attended features.

Chapter 2 describes complementary psychophysical experiments that investigate a behavioral correlate to the physiological effect. If attention to a stimulus feature enhances the processing of other stimuli with that same feature, this should facilitate the distribution of attention across multiple stimuli with common features compared to opposing features. We employed a dual-task psychophysical experiment that required observers to make concurrent discrimination judgments on two spatially separate stimuli containing either the same feature (the same direction of motion or the same color) or opposing features (opposing directions of motion or opposing colors). We predicted that attending to stimuli with common features would facilitate their concurrent processing. These experiments were repeated without the overlapping fields of dots in order to assess the role of competing distractors. Similar stimuli were used in both the fMRI and psychophysical experiments in order to compare the results as best as possible.

Chapter 3 uses a different psychophysical paradigm to test whether attention to orientation facilitates the processing of another stimulus with the matching orientation. The second paradigm relates the effects of feature-based attention to contextual influences in visual processing.

Chapter 1

Global effects of feature-based attention in human visual cortex

Abstract

Functional MRI was used in human subjects to test whether a feature-based mechanism of attention globally influences cortical responses to visual stimuli outside the attended location. Focusing attention on a stimulus feature (direction of motion or color) increased the neuronal response in multiple cortical visual areas (V1, V2, V3, V3A, V4, and V5/MT+) to a spatially distant stimulus that was ignored, but shared the same feature. These results are consistent with a *feature-similarity gain model* of attention which proposes that attention to a stimulus feature increases the gain of cortical neurons tuned to that feature throughout the visual field.

1.1 Introduction

Functional magnetic resonance imaging (fMRI) allows non-invasive imaging of activity in the human brain. Here, we used fMRI to study the effects of feature-based attention on multiple early stages of cortical visual processing including primary visual cortex, extrastriate visual areas, and the motion processing area MT+/V5. This chapter begins with a brief description of how fMRI works followed by a description of our general fMRI methods. Following these introductory sections, the feature-based attention experiments are presented.

1.1.1 MRI basics

Magnetic resonance imaging (MRI) uses a nuclear magnetic resonance (NMR) signal to create detailed images of the brain and other body tissues. Hydrogen nuclei, abundant due to the high water content of the body, are the basis for the MR signal. Hydrogen nuclei, possessing an intrinsic quantum property called spin, behave like magnetic dipoles. Placed in a magnetic field *B*, the nuclei partially align with the field generating a net magnetization. This net magnetization is many orders of magnitude smaller than *B* and cannot be directly detected. To make it detectable, a radio frequency pulse of energy is transmitted to the brain at a resonant frequency proportional to the magnetic field strength (64 MHz for hydrogen nuclei in a 1.5 Tesla scanner). The nuclei absorb that energy causing the net magnetization to oscillate (or *precess*) around *B* at the resonant frequency. The oscillating magnetization induces a current in a nearby coil and this is the detected signal. The signal quickly decays and the equilibrium magnetization regrows at rates (*T1* and *T2*, respectively) that depend on the local magnetic environment at the source. As a result, the strength of the signal varies for different brain tissues (e.g. gray matter compared to white matter) creating the contrast in an MR image. Signal localization is achieved by applying magnetic field gradients that cause the frequency and phase of precession to vary with location. The image is constructed by applying the inverse Fourier transform to the raw signal. A typical MR image of the brain is shown in Figure 1.1.

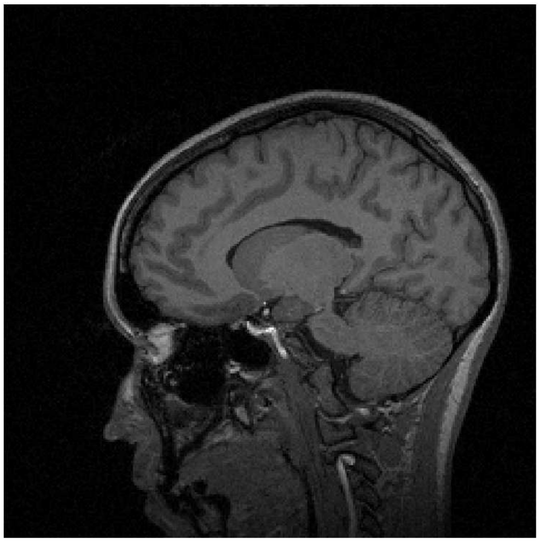

Figure 1.1 Typical MR image of the brain
Mid-sagittal slice view of the author's brain acquired with a T1-weighted MPRAGE pulse sequence (spatial resolution = 1 mm^3).

In *functional* MRI, small changes in the MR signal associated with brain activity are measured over time. The signal change depends on two serendipitous factors. First, neural activity is accompanied by a local increase in blood flow that supplies oxygen at a rate exceeding the rate of oxygen consumption. As a result of this oversupply, brain activity leads to a local increase in the concentration of oxygenated compared to deoxygenated hemoglobin in venous blood. Second, oxygenated hemoglobin attenuates the MR signal more slowly than deoxygenated hemoglobin. As a result, an increase in the ratio of oxygenated to deoxygenated hemoglobin leads to a local increase in the MR signal. This change is referred to as the blood oxygenation dependent signal (BOLD).

FMRI is a young technology. In 1990 Ogawa and colleagues first showed that blood oxygenation had a measurable effect on MR signals in the brain of the rat (Ogawa, Lee et al. 1990). In 1992, investigators demonstrated local signal increases in human primary visual and motor cortices after sensory stimulation (Kwong, Belliveau et al. 1992; Ogawa, Tank et al. 1992). With some modifications to allow for faster imaging, the MRI scanners possessed by many medical centers for conventional anatomical imaging could now be used for an altogether new application: functional imaging.

Many questions remain concerning how the BOLD signal quantitatively relates to neural activity. The ultimate goal is to understand the complex relationship between neuronal activity, oxygen metabolism, and blood flow.

For now, the accumulating evidence suggests that the BOLD response can be approximated as a linear function of neuronal activity. Tests of linearity of the BOLD response to visual stimulation demonstrate that it sums roughly linearly with stimulus duration (Boynton, Engel et al. 1996) and with repeated individual stimulation events (Dale and Buckner 1997). Another approach has been to compare to human BOLD responses to neural population response estimates made from monkey single-unit data. Rees and colleagues estimated that the average firing rate across the population of MT cells increased linearly with motion coherence and then showed that human MT+ responses also increased approximately linearly with motion coherence (Rees, Friston et al. 2000). Heeger and colleagues similarly estimated a V1 population contrast response function based on single-unit data and showed that human V1 contrast response functions scaled linearly with their estimate (Heeger, Huk et al. 2000). The linear model may thus be reasonable basis for the interpretation of BOLD data, particularly in visual cortex where it has been tested.

Another question to be answered concerns the type of brain activity that contributes to the BOLD response. In addition to the spiking rate of neurons, other factors that may contribute include subthreshold activity, neuronal inhibition, and changes in firing patterns that do not necessarily affect the spike rate (for review see Heeger and Ress 2002). A new direction of research comparing electrophysiological data from monkeys and high resolution fMRI data from the same animal will certainly shed further light on these issues (Logothetis, Pauls et al. 2001).

1.1.2 General fMRI methods

Data were acquired using a Siemens VISION 1.5 T scanner at the Thornton Hospital of the University of California, San Diego. High resolution (1x1x1 mm) anatomical images were collected using a volume head coil and a standard T1-weighted gradient echo pulse sequence (MPRAGE, magnetization prepared rapid gradient echo) to provide a three-dimensional reference anatomy for each subject. BOLD functional images were acquired using a small flex coil (localized over the occipital lobe) and a low-bandwidth T2*-weighted echo-planar imaging (EPI) pulse sequence. During all functional scans, 130 temporal frames were acquired over 260 seconds (TR=2 sec, flip angle=70 deg, 4x4 mm within-plane resolution, 16 coronal slices of 4mm thickness). The first ten temporal frames (20 sec) were discarded to avoid magnetic saturation effects. Typical scanning sessions consisted of ten functional scans. Each scanning session ended with a T1-weighted anatomical scan (1x1x1 mm resolution) acquired with the small flex coil and an MPRAGE pulse sequence. This anatomical data was manually aligned to the subject’s reference anatomy so that data from every scanning session was aligned into a common coordinate system.

During analysis of functional data, a linear trend in the time series of each voxel was removed. This standard procedure compensates for a drift in the signal that occurs slowly over time for unknown reasons. The time series of each voxel was then divided by its mean intensity to convert the data from arbitrary units of image intensity into units of percent signal change.

1.1.3 Stimulus presentation

Visual stimuli for all experiments were generated on a Macintosh PowerBook computer using Matlab v4.3 and the Psychophysics Toolbox (Brainard 1997; Pelli 1997). Stimuli were projected using a Proxima DP9300 LCD projector (max brightness = 1500 lumins; resolution = 1024 x 768; 60 Hz) onto a back-projection screen near the head inside the bore of the scanner. Subjects lay on their backs and looked up into an angled mirror to view the screen. A securely mounted dental-impression was used to minimize head movements. Subjects entered task responses onto a keypad positioned at their side.

1.1.4 Retinotopic mapping

Visual areas V1, V2, V3, V3A, and V4 were localized in each subject by their retinotopic organization. In retinotopically organized visual areas, neighboring neurons have nearby receptive fields creating 2-dimensional maps of visual space on the cortical surface. The retinotopic maps comprising each visual area were identified for each subject using standard phase encoding mapping methods (Sereno, McDonald et al. 1994; DeYoe, Carman et al. 1996; Engel, Glover et al. 1997). While maintaining fixation on a central point, subjects viewed a slowly rotating wedge (like the hand of a clock) and an expanding ring that map out the visual field in angle and eccentricity, respectively. These stimuli cause a traveling wave of activity across the cortical surface that peaks within a particular voxel when the stimulus passes through its associated region of visual space. The stimuli are cycled so that the phase of the fMRI response within each voxel reveals its retinotopic location. The fMRI response of each voxel is quantified as the phase of the (6.5 cycle/scan, 40 sec period) sinusoid that best fits the time-series of the voxel.

The wedges and rings were composed of flickering black and white checkerboard patterns (8 Hz counter phase flicker, mean luminance = 340 cd/m^2, contrast = 100%) and were presented on a mean gray background. During each scan, a wedge or ring was presented for six and a half 40 second cycles (260 seconds). In addition, alternating vertical and horizontal wedges (“hourglass” and “bowtie” shapes, width = 45 deg) of flickering checkerboard patterns were used to localize the horizontal and vertical meridians that define the borders of each visual area. The vertical and horizontal meridian stimuli were alternated in a block design every 20 seconds for six and a half 40 second cycles. For all scans, data from the first half cycle were discarded to avoid magnetic saturation effects.

The motion sensitive area MT+ was identified as an area on the lateral occipital surface that modulated in response to alternating moving (8 deg/sec radially inward and outward) vs. stationary random white dots on a black background. Human area MT+ (referred to as the MT complex) is likely to be the human homologue of macaque areas MT and MST (Watson, Myers et al. 1993). While MT is known to be retinotopically organized in the macaque visual system, large receptive fields make it difficult to measure retinotopic maps with the resolution of our functional mapping methods. A recent study shows that with repeated averaging it is possible to measure some degree of retinotopy in human MT that distinguishes it from neighboring MST (Huk, Dougherty et al. 2002).

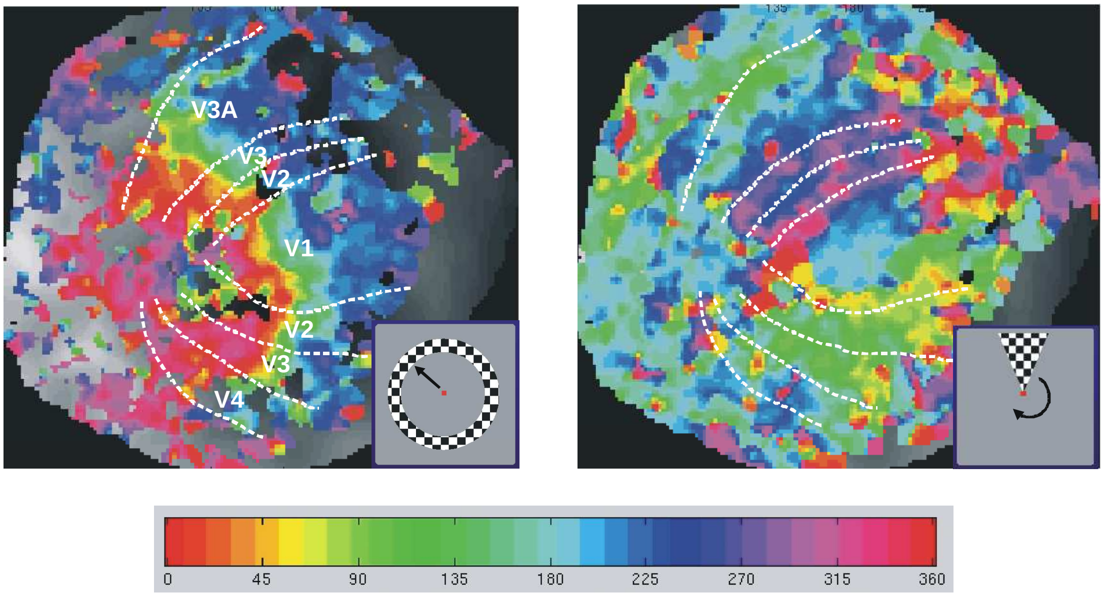

Figure 1.2 Localization of retinotopic visual areas
Functional MRI data in response to a cycled expanding ring and rotating wedge are projected onto a computationally flattened representation of the occipital cortical surface (right hemisphere). These stimuli cause a traveling wave of activity across retinotopically organized visual areas (V1, V2, V3, V4, and V3A). The temporal phase of the peak fMRI response within each voxel is color-encoded and visual area borders are labeled with white lines. Once localized, visual area coordinates are projected onto the 3-D anatomy (see Figure 1.3).

The 2-D retinotopic maps are best viewed when the functional data is projected onto a computationally flattened cortical surface. First, each subject's cortical surface (gray matter) was identified in their high-resolution reference anatomy using a Bayesian classification algorithm (Teo, Sapiro et al. 1997). Next, the occipital surface of each brain hemisphere was computationally flattened using a multidimensional scaling algorithm (Engel 1997). Software for both gray matter segmentation and cortical flattening is available online at http://white.stanford.edu. For illustration, fMRI responses to rotating wedge and expanding ring stimuli are presented on the experimenter's flattened occipital surface (Figure 1.2). Response phase (0 to 360) is color-encoded and the borders of the visual areas are labeled with white lines. The visual area borders are delineated by hand. These border estimations are typically made after comparing data from two repeats each of the wedge and ring scans and data from two meridian scans.

Once identified on the flattened maps, the visual area coordinates are saved and projected back into the 3-D reference coordinates. Figure 1.3 shows a 3-D rendering of the experimenter's brain with labeled visual areas. The visual cortical anatomy in different human subjects is qualitatively similar, but the size and shape of the visual areas varies significantly across individuals. For example, the size of human area V1 varies approximately 2-fold across individuals (Duncan and Boynton 2002). Considering this natural variability, our method of identifying visual areas based on functional activity is far superior to estimation based on anatomical landmarks.

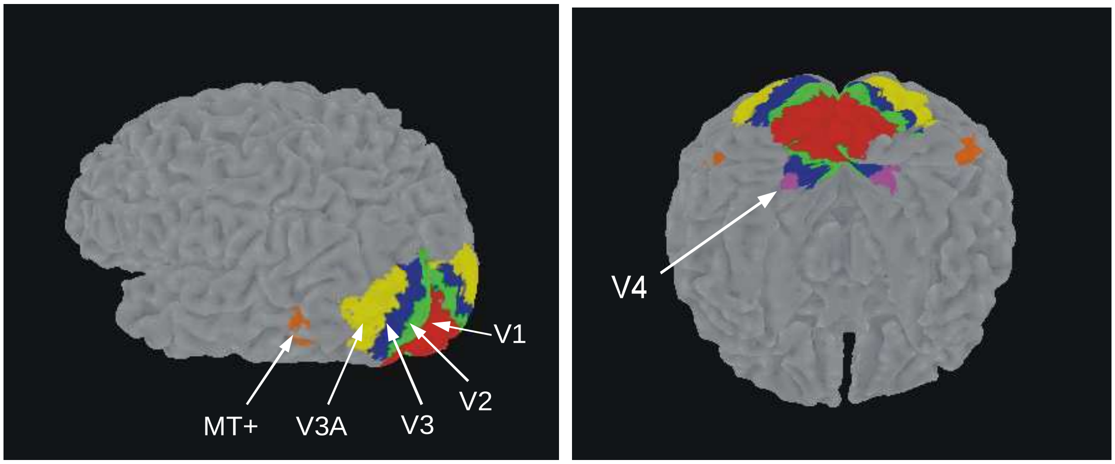

Figure 1.3 Three-dimensional view with labeled visual areas
Three-dimensional rendering of the cortical surface viewed from two angles. The bottom view (cerebellum removed) reveals the ventral surface of the brain and the symmetry across hemispheres.

1.2 Feature-based attention to direction of motion

1.2.1 Methods

Three human subjects (SBM, GDH, and DDL) with normal visual acuity participated. Subjects gave written informed consent and experiments were approved by the Salk Institutional Review Board. Human observers were presented with a stimulus to attend on one side of a central fixation point and a stimulus to ignore on the other side (Figure 1.4a). Stimuli were circular apertures (radius = 5 deg) of coherently moving random dots displayed in the lower visual field (centered 11 deg to the left and right of the fixation point, 2.5 deg below fixation, baseline speed 10 deg/s). Dots were white (560 cd/m^2) on a gray background (230 cd/m^2). The attended stimulus was composed of two overlapping fields of upward and downward moving dots. The ignored stimulus was a single field of dots moving in a single direction, either upward or downward. The two overlapping fields on the attended side perceptually segregated due to their opposing motion, allowing observers to selectively attend to one direction of motion at a time within the same location.

To direct their attention to one direction of motion or the other, subjects were instructed to perform a two-interval forced choice (2-IFC) speed discrimination task on either the upward or downward moving field of dots on the attended side in alternating blocks. The dots had limited lifetimes (200 ms) to prevent observers from attending to individual dots rather than the whole field. A cue (0.5 deg length line extending from the fixation point) instructed subjects to alternate the attended field of dots (the *target* field) every 20 seconds, so that each four minute fMRI scan consisted of six 40-second cycles in which attention alternated between both directions (Figure 1.4a). The field of dots on the ignored side did not change direction throughout the scan. By this design, conditions alternated every 20 seconds between the *same* condition in which the ignored stimulus matched the direction of the target field, and the *different* condition in which the ignored stimulus moved in the opposite direction of the target field.

Stimuli were presented in a series of trials initiating every 3.3 s. During each trial, all stimuli were presented simultaneously for two sequential 1000 ms intervals separated by a 100 ms interval in which only the fixation point was present. For each of the three fields of dots, the dots moved at a baseline speed during one interval and a just-detectable increment above this baseline during the other interval. Subjects indicated with a key press whether the target field moved faster during the first or second interval. Importantly, whether the dots moved faster during the first or second interval was independently randomized for each field of dots every trial. In addition, the baseline speed was randomly and independently jittered across trials in all three fields of dots, so that the non-target fields could not be used as a speed reference for comparison. These measures were taken to require subjects to direct attention only to the target field of dots. There was no information to be gained by attending to either non-target field (the overlapping distracting field and the ignored field in the opposite hemifield) and doing so could only hinder task performance. Furthermore, because a speed judgment requires observers to integrate information over time, the speed discrimination task encouraged subjects to continuously focus attention on the target field during the stimulus presentation intervals.

Subjects trained for several hours on this task until performance reached an asymptotic level prior to fMRI scanning. Speed increment thresholds were measured for each subject using a staircase procedure to determine the speed increment that predicted 79 percent correct performance. These increment thresholds were used during fMRI scanning. During scans, feedback was given during the inter-trial interval and all responses were recorded. Scans were counterbalanced for the attended side (left/right), the starting attended direction, and the direction of motion on the ignored side. Because the stimuli were presented in the periphery, responses to the left and right stimuli were separated into opposite brain hemispheres (MT+ is a possible exception, as discussed below).

Eye-monitoring was performed outside of the scanner using an identical stimulus display and a standard pupilometry eye tracking system (ISCAN Co, Burlington, MA). Mean horizontal and vertical eye-positions during stimulus presentation intervals were compared between the two conditions (same and different). Eye blink artifacts were removed from all eye position traces before statistical analysis. No statistical difference in mean horizontal or vertical eye position was found (t-test, $p>0.05$) excluding the possibility that eye position systematically varied with condition. This demonstrates that there was no difference in the location of the stimulus on the retina between conditions and no difference in the effective retinal velocity of the moving dot stimulus between conditions.

Echo-planar imaging (EPI) was performed using a Siemens Vision 1.5 T scanner (4X4X4 mm voxels, 16 slices, TR=2 s). We analyzed the blood oxygenation level-dependent (BOLD) response to the ignored stimulus in multiple brain areas representing the first stages of cortical visual processing: V1, V2, V3, V3A, and MT+. Analysis was restricted to pre-selected voxels within each visual area that responded to an on/off alternating reference stimulus (high-contrast counterphase modulating (8Hz) checkerboard pattern, 6X40 sec cycles alternated with a uniform gray background) presented within the same location in the visual field as the experimental stimulus. Voxels that responded to the reference stimulus with correlations above a liberal threshold ($r > 0.23$) were selected for further analysis. This step was taken to exclude the least responsive voxels within each visual area that either (1) did not correspond to the retinotopic location of the experimental stimulus or (2) contained a high proportion of white matter.

The fMRI response to the experimental stimulus was quantified as the phase and amplitude of the (6-cycle/scan, 40 sec period) sinusoid that best fit the time-series of voxel responses averaged across the pre-selected voxels within a given visual area. This calculated response was then projected onto a unit vector with an angle of 35 deg accounting for a 4.4-sec hemodynamic response latency (Heeger, Boynton et al. 1999).

1.2.2 Results

In all areas we found that the response to the ignored stimulus was stronger during the *same* condition, when its direction of motion matched the direction attended in the opposite visual hemifield. As illustrated for area MT+ (Figure 1.4b), the time series of the fMRI response to the ignored stimulus modulated as subjects alternately attended to the *same* and *different* directions of motion in the opposite hemifield (averaged across all 3 subjects, 24 repetitions per subject). In this and subsequent time series figures, the raw signal is smoothed by convolution with a Gaussian function. Response amplitudes were calculated for each visual area (Figure 1.4c) as the amplitude of the best fitting sine wave to the time series projected onto a predicted phase. The observed neuronal modulation occurred without changing the stimulus, eye position, task difficulty, or locus of spatial attention, and can therefore be attributed only to a feature-based allocation of attention.

Although multiple elements of the experimental design encouraged subjects to maintain attention on the target field of dots, we tested the possibility that the observed modulation was due to subjects inadvertently directing spatial attention to the ignored stimulus by examining both behavioral and fMRI responses to the attended stimulus. First, if spatial attention was diverted toward the ignored stimulus when it matched the direction of the target field, we would expect this withdrawal of attention from the target stimulus to result in poorer task performance on the attended stimulus during

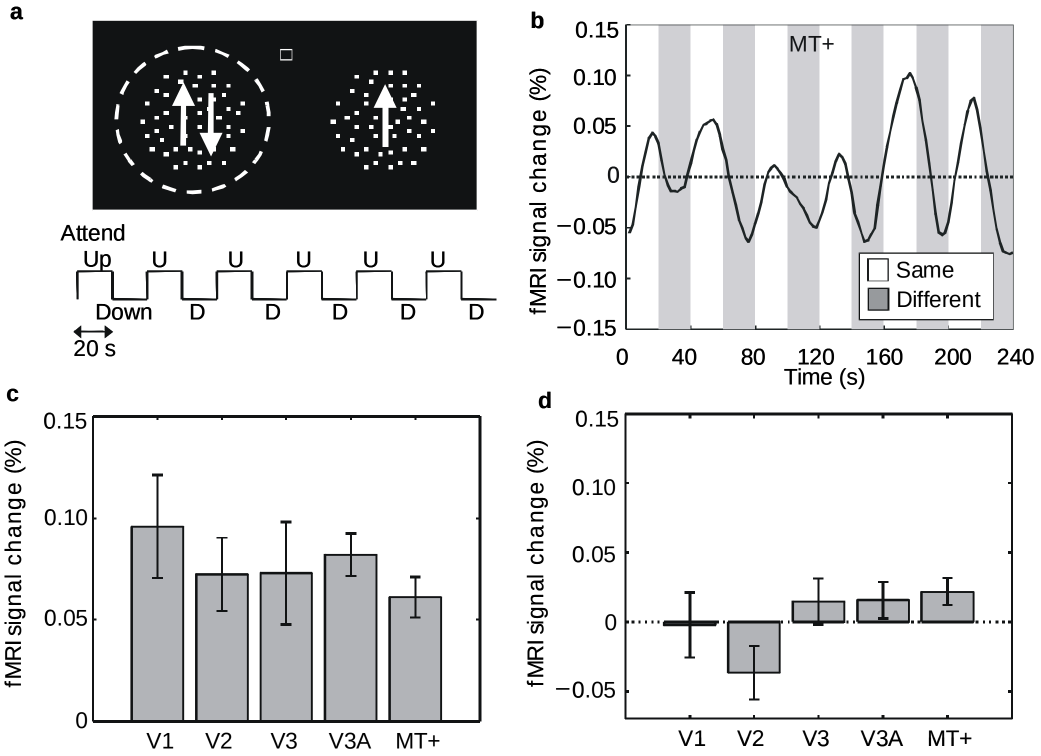

Figure 1.4 Feature-based attention to motion direction

(**a**) Stimulus diagram (not drawn to scale). Stimuli were circular apertures (radius 5 deg) of coherently moving random dots displayed in the lower visual field (centered 11 deg to the left and right of the fixation point, 2.5 deg below fixation, baseline speed 10 deg/s). Each arrow represents a field of 50 dots moving upward or downward. Dots were white (560 cd/m^2) on a gray background (230 cd/m^2). The dashed circle (not actually present in the stimulus) indicates the spatial focus of attention. Observers were cued to alternately attend to one of the two overlapping directions of motion as indicated by the example time course. (**b**) fMRI time series in response to the ignored stimulus for MT+, averaged across 3 subjects and 24 repetitions per subject. Response to the ignored stimulus modulated as subjects alternately attended to the same or different direction of motion in the opposite visual hemifield (as indicated by the background shading). (**c**) Projected response amplitudes to the ignored stimulus show that responses were greater during the *same* than during the *different* condition across multiple visual areas. (**d**) Responses to the attended stimulus, however, did not modulate largely or systematically between *same* and *different* conditions.

same than during *different* conditions (Posner, Snyder et al. 1980; Lee, Koch et al. 1999). However, there was no significant difference in performance between the *same* condition (mean 87.6% correct) and the *different* condition (mean 87.0% correct, p>0.05). In separate psychophysical trials, we verified that task performance was significantly impaired when subjects were instructed to divide attention across both left and right stimuli. Second, divided attention imaging studies (Vandenberghe, Duncan et al. 1997) predict that distribution of attention between two stimuli would result in a decrease in the fMRI response to the attended stimulus during the *same* condition, resulting in out-of phase modulation, or negative response amplitudes. However, the fMRI response to the attended stimulus did not modulate strongly or consistently across visual areas (Figure 1.4d).

In area MT+, we did observe a small but significant modulation that was in phase with the response to the ignored stimulus. This may be due to large receptive fields in areas MT/MST that include part of the ipsilateral visual field (Duffy and Wurtz 1991). Given such receptive fields, the responses to the attended and ignored stimuli may not be completely isolated to the contralateral hemisphere in area MT+. Therefore, a portion of the response to the ignored stimulus in area MT+ (Figure 1.4b) may be due to large receptive fields that include the attended stimulus.

1.2.3 Relative responses

The attentional response amplitudes measured in the first experiment (Figure 1.4c) reflect the difference in neuronal response to the ignored stimulus when subjects attended the same vs. different directions of motion in the opposite hemifield (*same – different*). These amplitudes are, however, difficult to interpret on their own. A small amplitude could be caused by a small attentional effect, or it could be that the stimulus evokes only a small response to begin with. We normalized the response amplitudes measured in the above experiment to the stimulus-evoked response elicited by cycling the ignored stimulus on and off. In a second experiment (diagrammed in Figure 1.5), the same three subjects alternately attended to the upward and downward motion of the overlapping fields of dots on the attended side, as in the first experiment. Meanwhile, however, the ignored stimulus was cycled ‘on’ when subjects attended its matching direction and cycled ‘off’ when subjects attended its opposing direction. Thus the fMRI response to the ignored stimulus reflected the difference in response when subjects attended the same direction vs. when the stimulus was turned off (*same – off*).

Figure 1.6 plots the attentional response as a percentage of the stimulus-evoked response calculated by dividing the modulation amplitudes from the first experiment by the modulation amplitudes from the second

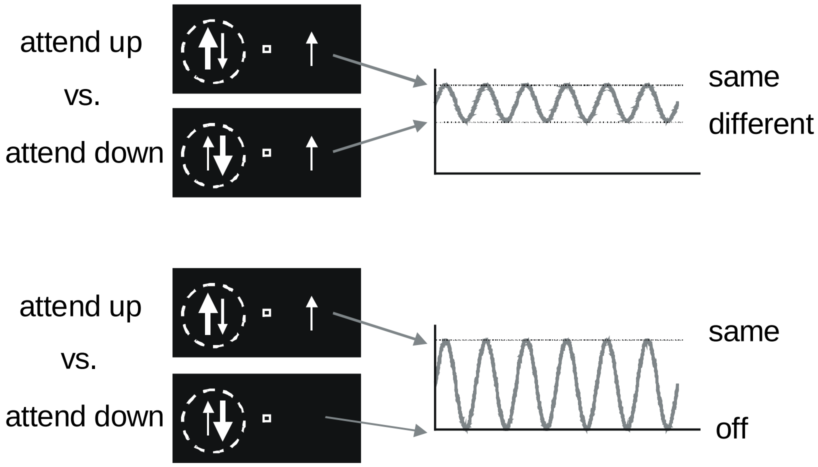

$$\text{Attentional Modulation (\%)} = \frac{\text{same - different}}{\text{same - off}} \times 100$$

Figure 1.5 Normalization of response amplitudes
(Top) Attentional response amplitudes from the first experiment reflect the difference in response to the ignored stimulus when subjects attended the same vs. different directions of motion (*same – different*). (Bottom) Response amplitudes from the second experiment reflect the difference in response to the ignored stimulus when subjects attended the same direction vs. when the stimulus was turned off (*same – off*). The ratio of these two response amplitudes gives a measure of attentional modulation as a percentage of stimulus-evoked response.

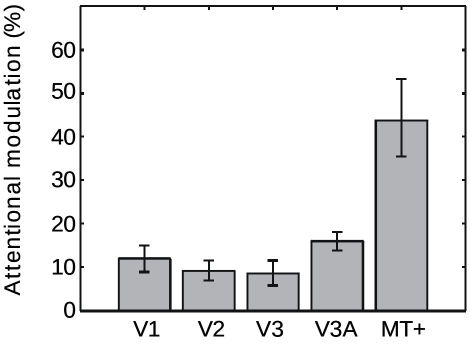

Figure 1.6 Normalized responses – motion direction experiment
Attentional modulation amplitudes as a percentage of stimulus evoked response across multiple visual areas, averaged across 3 subjects. Error bars indicate 1 SEM.

experiment: 100x(same-different)/(same-off). A 0% outcome would indicate no attentional modulation, while 100% would indicate that attention modulated the response to the ignored stimulus as much as cycling that stimulus on and off. The attentional modulation varies from about 10% in V1 to 50% in MT+, a substantial effect.

1.2.4 Changing the ignored direction of motion

In an alternate version of the experiment just described, a similar result was observed for a single subject (SBM) when the attended direction of motion was held constant and the ignored direction of motion was gradually rotated (the opposite manipulation of the first experiment). In this second experiment, subjects performed the same speed discrimination task as before on either the upward and downward moving dots of the attended stimulus, *without switching the attended direction* (Figure 1.7a). Meanwhile, the direction of motion of the ignored stimulus rotated counterclockwise by 30 deg every trial, so that it completed six full rotations during each four-minute scan.

Figure 1.7b plots the fMRI time series response to the ignored stimulus in area MT+ (1 subject, averaged across 36 repetitions). The response increased as the stimulus rotated into alignment with the attended direction of motion and decreased as it rotated away. The magnitude of the modulation was similar across visual areas (V1, V2, V3, V3A, & MT+) to the modulation observed for the same subject during the first experiment (Figure 1.7c). There was no systematic response modulation to the attended stimulus (Figure 1.7d). These results rule out the possibility that the feature-based attentional modulation reported above is related to the switching of attention between features because in this experiment there was no attentional switching.

1.3 Feature-Based Attention to Color

1.3.1 Methods

To address the question of whether this attentional effect generalizes to other features, the experiment was adapted to use color as the attended feature. Analogous to the first experiment, colored stimuli to be attended and ignored were presented to the left and right of a central fixation point (Figure 1.8a). Stimuli were circular apertures (radius 5 deg) of stationary red and green random dots displayed in the upper visual field (centered 11 deg to the left and right of the fixation point, 2.5 deg above fixation). The attended stimulus was composed of overlapping fields of stationary red and green dots and the ignored stimulus was composed of a single field of dots, either red or green. Dots had limited lifetimes to prevent subjects from attending to individual dots

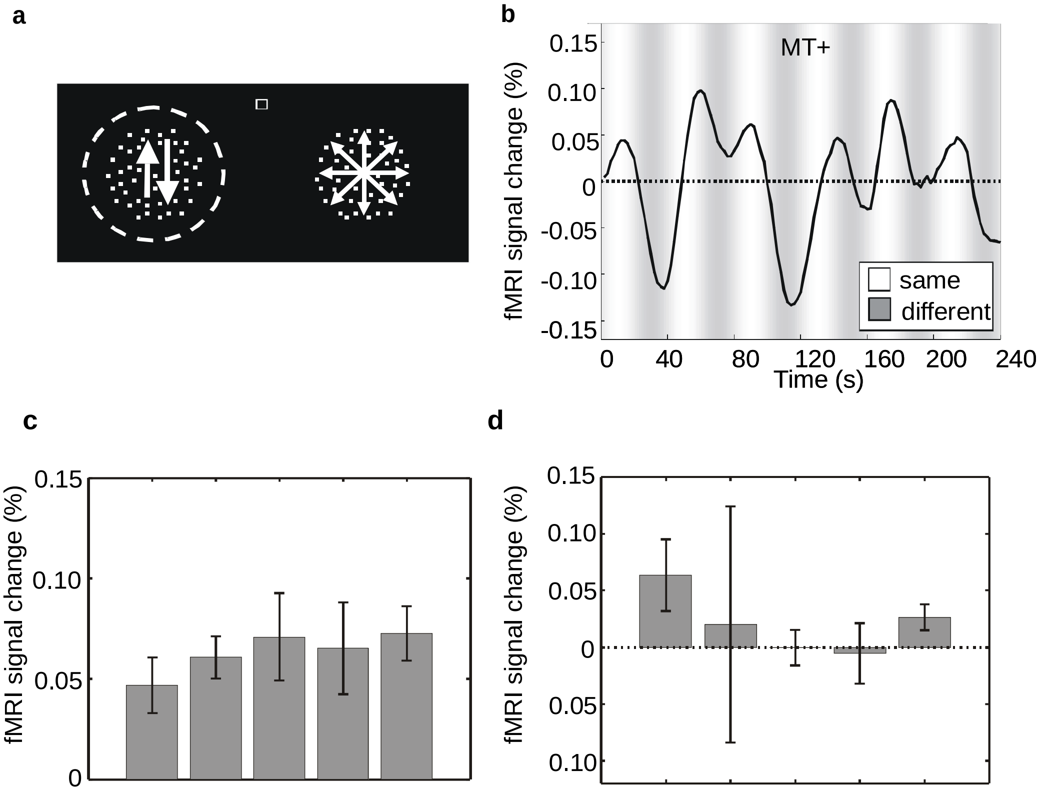

Figure 1.7 Changing the ignored motion direction
(**a**) Subjects attended one of the two overlapping fields of moving dots without alternating. The direction of motion of the ignored stimulus rotated counterclockwise by 30 deg every trial, so that it made a full rotation every 40s and completed 6 rotations during each four-minute fMRI scan. (**b**) fMRI time series in response to the ignored stimulus in MT+ for a single subject, averaged across 36 repetitions. The response modulated, increasing as the ignored stimulus rotated into alignment with the attended direction in the opposite visual hemifield and decreasing as it rotated away (as indicated by the background shading). (**c**) Projected amplitudes from the ignored stimulus demonstrate this modulation across multiple visual areas. (**d**) Responses to the attended stimulus, however, did not modulate largely or systematically between *same* and *different* conditions.

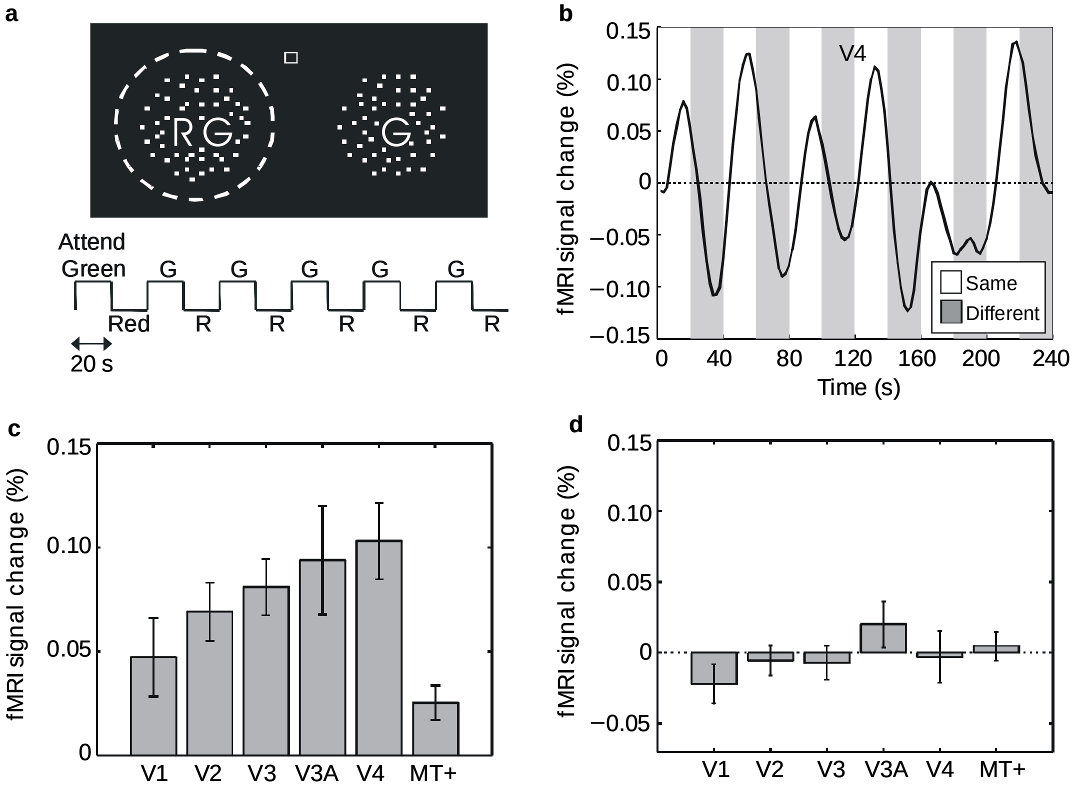

Figure 1.8 Feature-based attention to color
(**a**) Stimulus diagram (not drawn to scale). Stimuli were circular apertures (radius 5 deg) of stationary red and green random dots displayed in the upper visual field (centered 11 deg to the left and right of the fixation point, 2.5 deg above fixation). Each letter R and G represents a field of 50 red or green dots on a gray background. Dots had limited lifetimes (200 ms) to prevent the tracking of individual dots and the fields thus appeared to flicker. Observers were cued to alternately attend to the overlapping red and green fields of dots as indicated by the example time course. (**b**) fMRI time series in response to the ignored stimulus for V4, averaged across 3 subjects and 24 repetitions per subject. Response to the ignored stimulus modulated as subjects alternately attended to the same or different direction of motion in the opposite visual hemifield (as indicated by the background shading). (**c**) Projected amplitudes from the ignored stimulus show that responses were greater during the *same* condition than during the *different* condition across multiple visual areas. (**d**) Responses to the attended stimulus, however, did not modulate between *same* and *different* conditions.

and thus appeared to flicker. Stimuli were now placed in the upper visual field to include the suspected human homologue of area V4 in the analysis for which only a ventral region (upper visual field representation) has been identified (Tootell and Hadjikhani 2001).

Three subjects with normal visual acuity and color vision participated (SBM, DDL, and MTS) including two subjects from the motion experiment. Analogous to the motion experiment, subjects were instructed to perform a 2-IFC discrimination task at threshold on the red or green field of dots (the target field) on the attended side in alternating blocks. The color of the ignored field of dots did not change during each scan. Each four minute fMRI scan consisted of six 40-second cycles in which attention alternated between the *same* condition in which the color of the ignored stimulus matched the target field, and the *different* condition in which the color of the ignored stimulus was different from the target field (Figure 1.8a). During each trial, all stimuli were presented simultaneously for two sequential intervals (trial structure same as above). A threshold level luminance intensity change occurred between the two intervals for each field of dots on every trial. Subjects indicated with a key press whether the target dots were brighter during the first or second interval. Whether the dots were brighter during the first or second interval was independently randomized for each field of dots. Additionally, the baseline luminance was independently jittered across trials in all fields of dots, so that the non-target fields could not be used for comparison to the target field. Again, the ideal task strategy for subjects was to maintain attention on the target field of dots while ignoring the two non-target fields. All other stimulus presentation, training, eye monitoring, fMRI, and data analysis procedures were repeated as described above. No difference in mean eye position was found between same and different conditions (t-test, $p>0.05$).

1.3.2 Results

The fMRI response to the ignored stimulus was modulated by feature-based attention to color in multiple visual areas: V1, V2, V3, V3A, V4, and MT+. As illustrated for area V4 (Figure 1.8b), the time series of the fMRI response modulated as subjects alternately attended to the same and different colors in the opposite visual hemifield. Projected response amplitudes across visual areas indicate that the response to the ignored stimulus was stronger during the *same* condition, when its color matched the attended color (Figure 1.8c).

We again tested the possibility that the observed modulation was due to subjects inadvertently directing spatial attention to the ignored stimulus by examining both behavioral and fMRI responses to the attended stimulus (arguments explained above). There was no significant difference in task performance between the *same* condition (mean = 89.5%) and the *different* condition (mean = 87.2%, p>0.05). Additionally, there was no modulation in the fMRI response to the attended stimulus across visual areas (Figure 1.8d).

Analogous to the first experiment, we normalized the attentional response amplitudes to the stimulus-evoked response elicited by cycling the ignored colored stimulus on and off. The attentional modulation as a percentage of stimulus-evoked response varies from about 10% in V1 to 35% in MT+ (Figure 1.9). In both the color and direction of motion experiments, the attentional modulation in area MT+ was proportionately large compared to the stimulus-evoked response. These results are consistent with the visual attentional literature; studies generally report larger effects of attention at later stages of cortical processing (McAdams and Maunsell 1999).

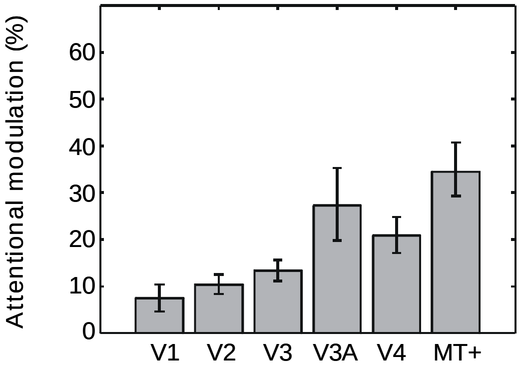

Figure 1.9 Normalized Responses – color experiment
Attentional response amplitudes as a percentage of stimulus-evoked response. Error bars indicate 1 SEM.

1.4 Unpublished Studies

The following two experiments yielded null results and are not published. It is difficult to draw new conclusions here because these null results may be due to an inability to measure the hypothesized effects.

1.4.1 Feature-Based Attention to Orientation

Our study was also adapted to use orientation as the attended feature. In the first experiment, sinusoidal gratings (one attended and one ignored) were presented to the left and right of a central fixation point (Figure 1.10a). The orientation of the attended grating was alternated between two orthogonal oblique angles (-45 deg and +45 deg) every 20 s during each 4-minute fMRI scan. To maintain attention, subjects performed a 2-IFC threshold level orientation discrimination task on the attended grating. Meanwhile, the orientation of the ignored grating did not change (either +45 deg or –45 deg) within a scan. We predicted that the response to the ignored grating would increase when subjects attended the matching orientation compared to the orthogonal orientation.

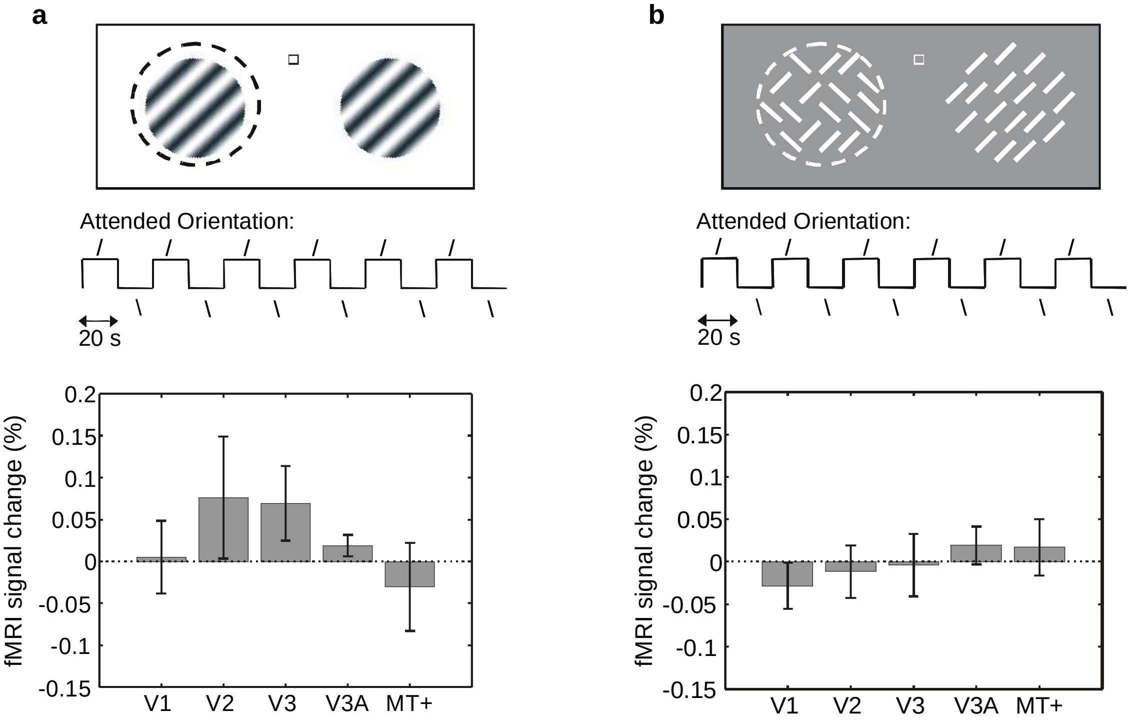

Figure 1.10: Feature-based attention to orientation

(**a**) Stimulus diagram (not drawn to scale): Stimuli were sinusoidal gratings presented to the left and right of a central fixation point. On the attended side, the grating alternated its orientation between -45 and +45 degrees as indicated by the example time course. Bar graph**:** Projected amplitudes show no significant modulation in the response to the ignored stimulus (averaged across 2 subjects, 8 repetitions per subject). Error bars indicate 1 SEM. (**b**) Stimulus diagram (not drawn to scale): Stimuli were composed of fields of oriented line segments (-45 and +45 deg). On the attended side, subjects alternately attended the two overlapping fields of oriented line segments as indicated by the example time course. Bar graph: Projected amplitudes show no significant modulation in the response to the ignored stimulus.

Projected response amplitudes to the ignored stimulus were determined as in the previous experiments and are plotted in Figure 1.10a across visual areas (averaged across 2 subjects, SBM and EMT, 8 scans per subject). No significant modulation was observed in any visual area. In a very similar experiment, McAdams and Maunsell also tested whether attention to an oriented grating in one visual hemifield affected the response to a task-irrelevant grating (same vs. orthogonal orientation) in the opposite visual hemifield in macaque V4. They also failed to find orientation-specific feature-based attentional modulation (personal communication, C. McAdams).

The paradigm here was different from that in our motion direction and color experiments in that subjects did not have to filter out an overlapping distracting stimulus in order to attend to the target stimulus. Rather, the orientation of the attended stimulus was physically changed with each alternated block. Multiple attention studies have found greater attentional effects when multiple stimuli compete for attentional selection (Motter 1993; Luck, Chelazzi et al. 1997; Treue and Martinez Trujillo 1999). To include a distractor, we next considered using an attended stimulus composed of overlapping gratings. However, the stationary overlapping gratings created a plaid pattern making it difficult for observers to selectively attend to either of the component gratings. Instead, we performed a second experiment using fields of oriented line segments more similar to the fields of random dots from our previous experiments.

In this second experiment, the attended side of the display was composed of two overlapping fields of obliquely oriented line segments (-45 deg and +45 deg) and the ignored side was composed of a single field of oriented line segments (either –45 or +45 deg). Line segments were randomly placed on a grid of non-overlapping positions. To direct attention, subjects performed a contrast discrimination task on each of the two overlapping fields in alternating blocks. Even though the task was not an orientation discrimination task, presumably attention to orientation was required to selectively attend to the target field in the presence of the overlapping distracting field. Projected response amplitudes to the ignored stimulus are plotted in Figure 1.10b across visual areas (averaged across 2 subjects, SBM and EMT, 8 scans per subject). Again, no significant modulation was observed in any visual area. This null result may be related to the inadequacy of this particular stimulus. The oriented line segments within each field did not perceptually group as strongly as did the moving and colored dots. It is likely that subjects performed the task by focusing on individual line segments rather than attending to the field as a whole, making the task less dependent on selective attention to orientation. Chapter 3 readdresses the issue of feature-based attention to orientation using an altogether new paradigm.

1.4.2 Combined effects of feature-based attention and spatial attention

We have shown neuronal response modulation due to shifts of feature-based attention in the absence of shifts of spatial attention. This modulation was found in multiple early stages of cortical visual processing (V1, V2, V3, V3A, V4, and MT+). Previous fMRI studies report neuronal modulation due to spatial attention within each of these cortical visual areas (Brefczynski and DeYoe 1999; Gandhi, Heeger et al. 1999; Martinez, Anllo-Vento et al. 1999). This suggests that a feature-based mechanism of attention may work in parallel with a spatial mechanism to influence visual processing in these early cortical areas. In support of this, Treue and Martinez-Trujillo (1999) found that when monkeys were instructed to shift both feature-based and spatial attention, their effects on MT responses summed remarkably well.

We also tested whether the effects of feature-based and spatial attention summed. The stimulus display was the same as in the *feature-based attention* to motion direction experiment: the stimulus on one side of fixation was composed of overlapping fields of upward and downward moving dots and the stimulus on the other side (the *test* stimulus) was composed of a single field of moving dots (Figure 1.11).

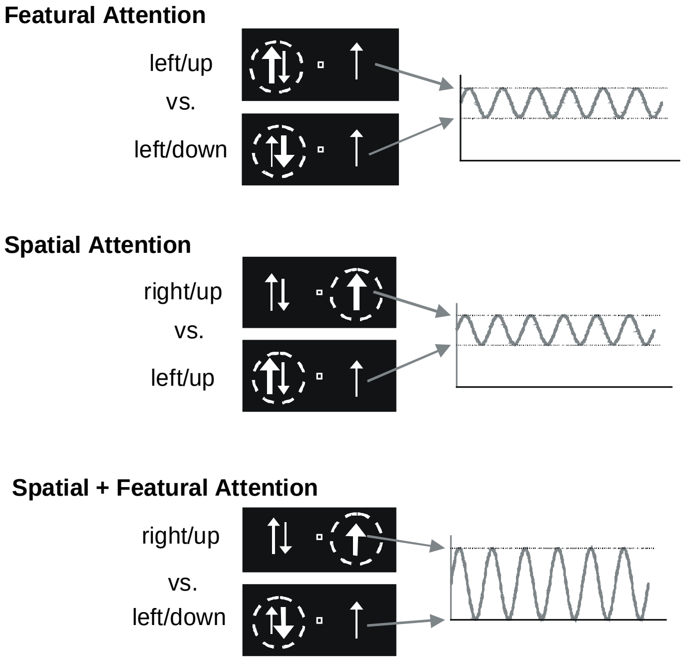

Figure 1.11 Combined spatial and feature-based attention
We expected a larger response modulation in the *spatial* + *feature-based* attention experiment than was measured in either the *feature-based* or *spatial* attention experiments. (The enlarged arrows indicate the attended fields of dots.)

During a new *spatial attention* experiment, observers alternately attended to the test stimulus on one side and to the field of dots moving in the same direction on the other side. By this design, modulation in response to the test stimulus reflects the difference in response between attending to the test stimulus vs. attending away to the same direction of motion. During a new *spatial* + *feature-based attention* experiment, observers alternately attended to the same test stimulus on one side and the field of dots moving in the opposite direction on the other side. Modulation measured in response to the test stimulus should now reflect the difference in response between attending to the test stimulus vs. attending away to a different direction of motion. Scans were counterbalanced for the side and direction of motion of the test stimulus. If the effects of spatial and feature-based attention combine additively, then the response measured in the *spatial* + *featural attention* experiment should be greater than the modulation caused by either manipulation alone.

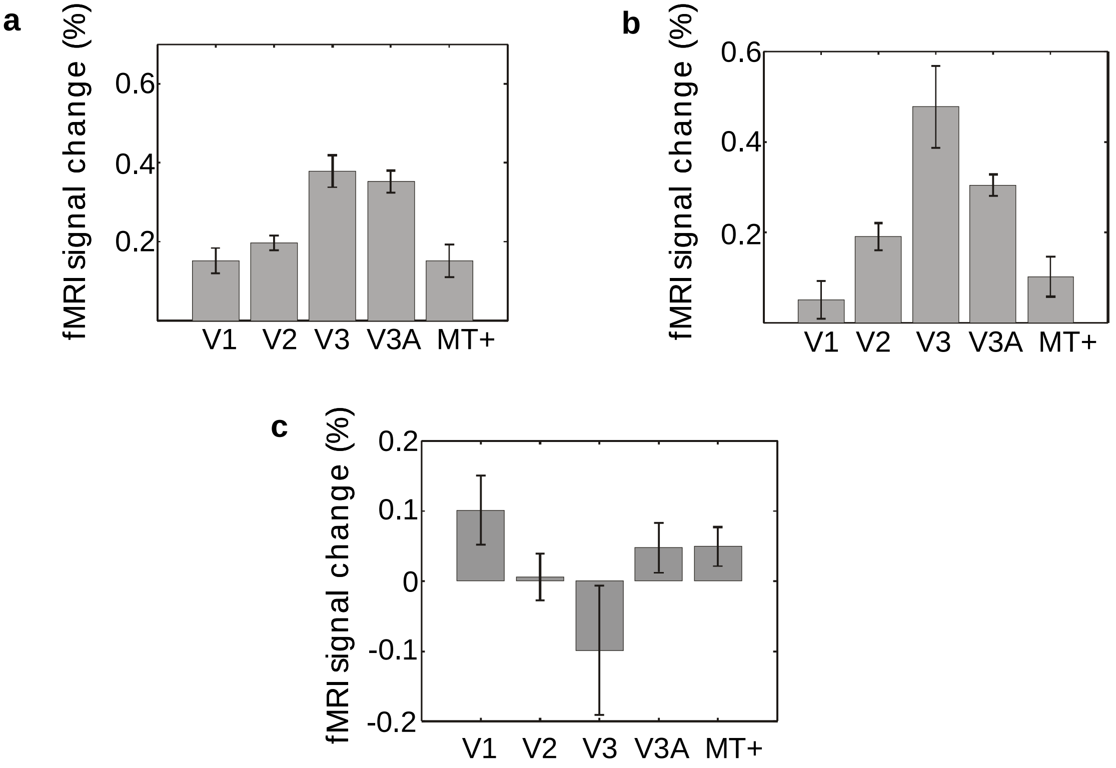

Figure 1.12 Combined spatial and feature-based attention results (**a**) Response amplitudes to the target stimulus in the *spatial* + *feature-based attention* experiment. (**b**) Response amplitudes to the target stimulus in the *spatial attention* experiment. (**c**) The differences in amplitudes (a-b) across visual areas were not consistently different from zero.

All three subjects (SBM, GDH, and DDL) from the original feature-based attention to motion study each participated in the new *spatial attention* and *spatial + feature attention* experiments. Mean projected amplitudes in response to the test stimulus are plotted across visual areas for both experiments in Figures 1.12a and b (each bar averaged across 3 subjects, 16 scans per subject). We did not find a larger modulation in the *spatial + feature attention* experiment than in the *spatial attention* experiment. Figure 1.12c plots the difference in amplitudes (b-a) across visuals areas. These differences were not consistently different from zero. The expected differences were the amplitudes of modulation measured in the feature-based attention experiment, ranging between 0.05 and 0.10 percent signal change. It is possible that we lacked sufficient signal to noise to measure such a small signal change as a difference between two independent measurements.

1.5 Discussion

Feature-based attention modulated neuronal responses to stimuli outside the focus of attention in the human brain. Attending to a stimulus feature (motion direction or color) increased the response to a spatially distant ignored stimulus with the matching feature compared to the opposing feature. This global modulation was found in multiple cortical visual areas (V1, V2, V3, V3A, V4, and MT+) representing the earliest stages of cortical visual processing. Consistent with the literature on spatial attention, the magnitude of the effect was overall larger at later stages of cortical visual processing.

These results are consistent with a previous report of global feature-based attentional modulation of directionally tuned neurons in macaque MT (Treue and Martinez-Trujillo 1999). Here, we confirm and extend these findings in the human visual system and across multiple cortical visual areas. We further show that this response generalizes to color, another feature for which cortical neurons show tuning. In addition, our experiment goes farther in ruling out the serious potential confound that observers shifted spatial attention towards the ignored stimulus. We show that neither the behavioral nor neuronal responses to the attended stimulus were modulated across compared conditions.

Global feature-based modulation was observed when the attended feature was motion direction or color, but not orientation. However, we cannot easily rule out global attention to orientation because there were differences in the experimental design that may have contributed to the null effect of orientation. Previous fMRI studies have shown neuronal modulation due to spatial attention in each of these cortical visual areas (DeYoe, Carman et al. 1996; Brefczynski and DeYoe 1999; Gandhi, Heeger et al. 1999; Martinez, Anllo-Vento et al. 1999). A feature-based mechanism of attention may thus work in parallel with a spatial mechanism to influence the earliest stages of cortical visual processing. Our attempts to sum the effects of spatial and feature-based attention were unsuccessful, but this could have been due to a lack of sufficient signal to noise.

1.5.1 Directional and color selectivity in visual cortex

Our results are consistent with a proposed mechanism, the *feature-similarity gain model*, whereby feature-based attention modulates the gain of cortical neurons tuned to the attended feature throughout the visual field (Treue and Martinez Trujillo 1999). Single-unit measurements commonly identify directionally-tuned neurons in all of the reported visual areas, with the highest proportion found in area MT (Zeki 1978). Neurons with a range of color-tuning properties have been identified in areas V1, V2, V3, and V4 (for review, Kiper, Levitt et al. 1999). However, color sensitivity in the motion processing area MT is a controversial topic.

A segregation of motion and color processing has been proposed (Livingstone and Hubel 1987), but accumulating evidence suggests that MT does have some access to color information. First, isoluminant chromatic motion evokes responses in macaque MT (Saito, Tanaka et al. 1989; Dobkins and Albright 1994) and those responses are direction selective (Seidemann, Poirson et al. 1999). Isoluminant motion evokes responses in human MT+ as well (Wandell, Poirson et al. 1999). Second, macaque MT receives input from all three cone types: L, M, and S (Chatterjee and Callaway 2002). Some degree of color selectivity could be achieved if cells varied in the strengths of their relative cone inputs and one study reports that MT cells do vary in the strength of their S cone inputs (Seidemann, Poirson et al. 1999). Color opponent cells in MT have yet to be identified.

Thus, the effect of feature-based attention to color in MT+ could be mediated by color-selective neurons located there. Alternately, color-tuned neurons in only certain visual areas could have identified the ignored stimulus and communicated more broadly with non-feature-specific neurons with receptive fields in the location of the ignored stimulus. This alternate mechanism (which could apply to the motion experiment as well) would be spatially crude in that it would enhance responses to any stimulus within the vicinity of the relevant ignored stimulus. However, we cannot rule out a contribution of this alternate mechanism with our fMRI methods, as we are unable to distinguish the subpopulation of neurons that contribute to the net response. It should also be noted that area MT+ is also likely to include human homologue of macaque area MST and possibly other nearby motion-sensitive areas where color selectivity could conceivably exist. Thus, areas nearby MT are likely to have contributed to the feature-based modulation in MT+.

1.5.2 Consideration of receptive field sizes across visual areas

The feature-based effect was spatially globalized because it modulated the responses of neurons with receptive fields that did not include the attended location. This assumes that receptive fields in the studied areas are not large enough to include an extended region of the ipsilateral visual field. In our experiments, attended and ignored stimuli were centered eleven degrees to the left and right of fixation (radii = 5 deg).

In the visual cortex of the macaque, receptive field sizes progressively increase at successively higher levels of visual processing (Zeki 1978). Receptive field sizes are smallest at the fovea, ranging from less than 1 degree in V1 to approximately 3 degrees in V3, V3A and V4 (for review see Smith, Singh et al. 2001). Receptive field sizes at lateral eccentricities are larger ranging up to approximately 8 degrees in areas V3/V4 (Gattass, Sousa et al. 1988). Assuming homology in the human, it is unlikely that receptive fields in these areas were large enough to include both the contralateral and ipsilateral stimuli in our experiments. However, large receptive fields in areas MT, and more so in MST, may extend 15 degrees or more into the ipsilateral visual field (Duffy and Wurtz 1991).

Consistent with the known monkey physiology, one study of human cortical retinotopy found that MT+, but not earlier visual areas, was activated by ipsilateral stimulation distant from the vertical meridian (Tootell, Mendola et al. 1998). Given this ipsilateral coverage in MT+, responses to the attended and ignored stimuli may not be completely separated into opposite brain hemispheres in that visual area. Therefore, a portion of the response to the ignored stimulus in MT+ may be due to large receptive fields that include the attended stimulus.

Neuronal mechanisms behind the global feature based attention effect are addressed further in the Discussion section of Chapter 2.

1.5.3 Behavioral consequences

If attention globally improves the processing of stimuli containing an attended feature, what are the behavioral consequences? We measured a neuronal response modulation due to feature-based attention across multiple early stages or cortical visual processing. The ubiquity and strength of this effect (up to 50% of stimulus evoked responses in MT+) suggests that it may profoundly influence what we see. A feature-based increase in signal strength could be expected to improve the detection, discrimination, and/or memory of visual objects with behaviorally relevant features, even if those objects are not the subject of focused attention. Global facilitation could be highly useful to the visual system because the location of relevant stimuli is not always known in advance. Feature-based facilitation may be relevant not only to vision but also to the processing of features in other modalities such as audition (e.g. auditory frequency).

The purpose of the experiments presented in Chapter 2 of this dissertation is to relate our fMRI results to a behavioral consequence using stimuli adapted from the fMRI experiments. A number of feature-specific effects already reported in the psychophysical literature may also be related and are discussed here.

1.5.4 Attentional Set (also referred to as mental or perceptual set)

Observers are better at detecting a visual object when they know something in advance about its features. In a classic example, Sekuler and Ball (1977) instructed observers to detect a barely visible moving object that was briefly flashed on a visual display (2-AFC contrast detection). Performance was greatly improved given advance knowledge of either the direction or speed of the moving object. Similar facilitation has been shown with prior knowledge of the location, color, or spatial frequency of a visual object or the frequency of an auditory tone (Posner, Snyder et al. 1980); Davis, Kramer et al. 1983; Green 1961). In these examples, the observer had to maintain an internal representation of the cued feature (an ‘attentional set’) that primed the visual system for the detection of the expected signal. At the neuronal level, this priming could be achieved by biasing the responses of neurons that encode the relevant feature *prior* to stimulus onset. Consistent with this possibility, Chelazzi and colleagues showed that searching for a visual stimulus increased the baseline firing rates of IT neurons tuned to that stimulus during a time period prior to stimulus presentation (Chelazzi, Miller et al. 1993; Chelazzi, Duncan et al. 1998).

In the behavioral studies just described, target location was certain. It would be most interesting to test whether facilitated detection persisted even if target location was uncertain. This result would suggest a *global* feature-based priming for detection. Similarly in our own experiments, the location of the ignored stimulus was unchanged. In future studies, we would like to test whether the feature-based response enhancement persisted for ignored stimuli at unpredictable locations. This result would also strengthen the evidence for a truly global attentional effect. An event-related fMRI paradigm would allow for the randomization of trials needed for this experiment.

1.5.5 Feature-based priming of attentional capture

Theories of attention commonly distinguish between voluntary goal-oriented (top-down) and automatic stimulus-driven (bottom-up) controls of attention. Our experiments deal primarily with the first case; observers voluntarily allocate attention to a particular feature as instructed. In stimulus-driven cases, a salient object or event (such as a sudden motion onset or a red object in a field of green objects) automatically captures attention. When exploring a visual environment, the interaction between both top-down and bottom-up influences determines what we pay attention to.

For example, stimuli which involuntarily attract attention do so more effectively when they share a feature in common with a stimulus that is being searched for (Folk, Remington et al. 1992). Folk and Remington reported that motion onset distractors at irrelevant locations captured attention when observers were searching for a motion onset target and not when searching for a color singleton (the reverse was also true). Thus, the stimulus-driven orienting of attention depends on how well the stimulus is related to feature-specific task demands.

In an unusual experiment (previously mentioned in the general introduction), a majority of observers failed to notice a man in a gorilla suit walk through a scene (presented on video) beating his chest while the observers performed an attentionally demanding task within the same scene (Simons and Chabris 2000). The task was to attend to one of two teams of people (in while or black shirts) passing a ball being between themselves and to keep track of the passes. The failure to detect the unexpected event indicated in a striking way that unattended visual objects may entirely escape awareness. However, quite interestingly, the investigators noted that observers were more likely to detect the black gorilla if they were attending to the team in black shirts. Further experiments on a more controlled visual display explored this relationship and concluded that observers were more likely to detect an unexpected visual object if its features were similar to those of attended objects in the display (Most, Simons et al. 2001).

Another type of feature-based bias in detection is observed in pop-out visual search tasks (e.g. searching for a red item within a background of green items). The hallmark of a pop-out task is that search times for a salient visual object are relatively short and are not affected by the number of background items. This pop-out visual search performance is improved by repetition of the same target feature across consecutive trials, as evidenced by improved reaction times and response accuracy (Maljkovic and Nakayama 1994; Bichot, Cave et al. 1999). (This is contrast to an inhibition of return with respect to repeated target location.)

Are these feature-specific effects related to a global priming of responses to stimuli with behaviorally relevant features? These findings could be extended and related to feature-based attention using our experimental paradigm. In our tasks with spatially overlapping stimuli, attention to a particular direction of motion or color was required in order to select the target field in the presence of the overlapping distracting field. In new experiments, we could test the ability of observers to detect test stimuli while they were concurrently engaged in the feature-based attention task. We could test whether reaction times to abrupt onset stimuli were reduced if the onset stimulus shared the attended feature. Similarly, we could test whether pop-out visual search performance was improved (also reduced reaction times) when the target stimulus shared an attended feature. Future experiments are discussed further in the *Conclusions* section of this thesis.

1.5.6 Inhibition of distractors

In some instances, the feature-based attention effect may work against the goals of the observer. Studies of selective attention generally find that distractors are more distracting when they share a feature in common with an attended stimulus. The classic ‘cocktail party’ example of selective attention refers to the ability to selectively attend to one voice in the presence of many background voices. When tested, subjects were less able to ignore distracting voices that were similar in pitch to the attended voice (Cherry 1953, Treisman 1964). In another example, Folk and colleagues instructed observers to monitor a centrally presented stream of letters for a target letter of a particular color (e.g. red) and to report the identity of that letter (Folk, Leber et al. 2002). Irrelevent distracting letters were simultaneously flashed in the periphery. Even though attention was focused on the center, distracting letters disrupted task performance, but only when the distractors matched the relevant color.

An everyday example may be experienced when a driver is stopped at a street intersection impatiently waiting for the traffic light to turn green. The onset of any green light, even in the wrong lane, captures attention and may cause the driver to make a false start.

1.5.7 Divided attention tasks

Our results are also consistent with a psychophysics study of feature-specific attention (Rossi and Paradiso 1995) in which observers performed a primary task of discriminating a feature of a foveal grating (spatial frequency or orientation) and a secondary task of detecting a near-threshold grating in the periphery. Observers were better at detecting the peripheral grating when its spatial frequency or orientation matched the attended feature in the primary task. This result suggests attention to a stimulus feature facilitates the processing of other non-overlapping stimuli with the same feature. This idea that feature-based attention facilitates the processing of multiple stimuli with common features is considered further in Chapter 2.

1.5.8 Object-based attention

Two fields of dots moving in the same direction could be perceived as part of a common object viewed through two apertures. Because of this, our findings could be attributed to an object-based rather than feature-based allocation of attention. Object-based models of attention propose that the visual system automatically segments the visual scene based on Gestalt –type principles (such as similarity of features) and that attention automatically spreads across identified visual objects or surfaces (Driver and Baylis 1998, Nakayama and Joseph 1998). The relationship between these two steps may be interactive: object identification may guide attention and attention may guide object identification.

In support of this idea, human performance data shows that observers are better at judging two attributes of the same object compared to two attributes of two different objects, even when the distance between the two attributes is controlled for (Duncan 1984; Baylis and Driver 1992; Egly, Driver et al. 1994; He and Nakayama 1995; Blaser, Pylyshyn et al. 2000). Duncan (1984) presented two roughly superimposed objects: the outline of box and a straight line. Observers were better at judging two box attributes (size and location of a gap) or two line attributes (texture and orientation) than judging one attribute from each. He and Nakayama (1995) demonstrated that it is difficult for observers to focus attention on locations at different stereoscopic depths, unless they are viewed as part of a common surface.

Visual objects are often defined by the sharing of common features. A global feature-based signal enhancement could thus contribute to the spread of attention across objects or surfaces. As an example, one can imagine tracking an animal running behind trees in a forest. The trees will break the view of the animal into disconnected elements sharing common features such as a particular color and/or direction of motion. A feature-based enhancement of signal strength may aid in identifying and attending the animal as a whole. This is not to say that feature grouping can explain all instances of object attention. In some cases it is not the sameness of features, but rather the continuity of features (e.g. a smooth transition of color) or even more complex relationships that define an object.

1.5.9 Sources of top-down attentional control

A number of functional imaging studies have searched the brain for sources of top-down attentional signals that modulate responses in the early visual areas reported on here. Higher level areas may be necessary for the processing and memory of task demands as well as for generating attentional control signals. A network of frontal and dorsal parietal areas is consistently activated in functional imaging studies of spatially directed attention. Directing attention to a region of visual space activates regions of the posterior parietal lobe (near the inferior parietal sulcus) and frontal cortex in what is possibly the human homologue of the macaque frontal eye fields, FEF (Corbetta, Akbudak et al. 1998; Hopfinger, Buonocore et al. 2000). Unlike responses in lower visual areas these responses are not strongly retinotopic and are sustained in the absence of visual stimulation. In a study of non-spatial attention, these same areas (as well as MT+) were activated when observers were told to expect the onset of a particular direction of coherent motion within a random dot display (Shulman, Ollinger et al. 1999). Whether these areas will be consistently activated by other non-spatial, feature-directed attention tasks is the topic of future study. Another region of posterior parietal cortex (near the temporoparietal junction, TPJ) is transiently activated by both spatial and non-spatial shifts in attention (Yantis, Schwarzbach et al. 2002). Based on these findings and others, Corbetta proposed two systems for attentional control, one primarily for the maintenance of attention and the other for the switching of attention between locations or objects (Corbetta and Shulman 2002).

1.5.10 Lateral connections

In addition to feedback from higher visual areas, intrinsic connections within visual areas may contribute to attentional modulation in early visual areas. Long-range horizontal projections between V1 pyramidal cells extend parallel to the cortical surface for distances of up to 8 mm allowing cells to integrate information from outside their classical receptive fields (Gilbert and Wiesel 1979). These horizontal connections (both excitatory and inhibitory) preferentially link cortical columns with similar orientation selectivity in area V1 (Gilbert and Wiesel 1989; Weliky, Kandler et al. 1995). This “like-to-like” connectivity could contribute to feature-specific contextual modulations in the cortex, allowing neurons to strongly influence other neurons with non-overlapping receptive fields tuned to the same feature.

The question of whether the “like-to-like” connectivity extends to features other than orientation and to visual areas outside of V1 demands further study, but some work has been done in that direction. Horizontal connections between V1 blobs have been reported to preferentially link cells with common color opponency in cat visual cortex (i.e., red/green opponency) (T’so and Gilbert 1989). Also, one study reports a tendency for lateral connections in MT to link cortical columns of similar directional preference in the owl monkey (Malach, Schirman et al. 1997). Also worth noting, a recent report compared V1 horizontal and V2-to-V1 feedback connections and found that, unlike the horizontal, the feedback connections did not link cells with similar orientation preference (Stettler, Das et al. 2002).

Chapter 3 of this dissertation explores the role of feature-based attention in modulating a well-studied psychophysical contextual interaction (collinear facilitation) that is likely to be mediated by horizontal connections.

1.5.11 Conclusions

Our fMRI results are important in demonstrating that feature-based attention affects neuronal responses to stimuli *outside* the focus of attention. A global feature-based increase in signal strength may contribute to a number of psychophysical reports indicating that unexpected or unattended stimuli are more likely to be processed if they share a feature in common with attended visual objects. Feature-based attention could also contribute to the spreading of attention across visual objects and surfaces. Behavioral consequences are explored further in Chapter 2.

This chapter, in part, is a reprint of the material as it appears in Saenz M., Buracas, G.T., and Boynton, G.M. *Nature Neuroscience*, 5, 631-632, 2002. The dissertation author was the primary investigator and author of this paper.

Chapter 2

Behavioral consequences of feature-based attention

Abstract

We used a divided attention psychophysical task to test the hypothesis that visual attention to a stimulus feature facilitates the processing of other stimuli sharing the same feature. Performance on a dual task was significantly better when human observers divided attention across two spatially separate stimuli sharing a common feature (same direction of motion or same color) compared to opposing features. This attentional effect was dependent upon the presence of competing stimuli. These results are consistent with a spatially global feature-based mechanism of attention that increases the response of cortical neurons tuned to an attended feature throughout the visual field.

2.1 Introduction

In Chapter 1, we demonstrated that the neuronal response to an unchanging, ignored visual stimulus was enhanced when observers attended a matching feature (motion direction or color) elsewhere in the visual scene. This global feature-based attention effect was observed in multiple human visual areas representing the earliest stages of cortical visual processing (V1, V2, V3, V3A, and V5/MT+). The ubiquity and strength of this effect suggests that it may profoundly influence what we see. Here we explore a behavioral consequence of this neuronal effect.

Our results suggest that global feature-based attention increases the processing of stimuli with behaviorally relevant features throughout the visual field. Feature-specific attention may thus profoundly impact our ability to process multiple stimuli in a complex visual scene. Specifically, if attention to a stimulus feature enhances the processing of other stimuli with that same feature, this should facilitate the distribution of attention across multiple stimuli with common features compared to opposing features. The aim of the experiments presented in this chapter was to test that prediction.

We employed a dual-task psychophysical experiment that required subjects to make concurrent discrimination judgments on two spatially separate stimuli containing either the same feature (the same direction of motion or the same color) or opposing features (opposing directions of motion or opposing colors). We predicted that attending to stimuli with common features would facilitate their concurrent processing.

The stimuli from our previous fMRI experiment were adapted for these experiments in order to relate the results as best as possible. Observers divided attention equally across two stimuli placed to the left and right of a central fixation point. In the first experiment, each stimulus was a circular patch consisting of two transparently overlapping fields of upward and downward moving dots (Figure 2.1a). Subjects simultaneously performed a discrimination task on one field of dots from each side, either moving in the same direction (up or down on both sides) or in different directions (up on one side and down on the other). Thus, without changing the visual display or the spatial distribution of attention, subjects divided attention across stimuli composed of either a common feature or opposing features. In a second experiment, we adapted the stimulus to use color as the attended feature. Stimuli were composed of transparently overlapping fields of red and green stationary dots (Figure 2.2a). Again, subjects simultaneously performed a discrimination task on one field of dots from each side. In both experiments, subjects performed significantly better on the dual-task when dividing attention between two fields of dots with the same feature (same direction of motion or same color) rather than opposing features (opposing direction of motion or opposing color). Furthermore, this attentional effect was diminished in additional experiments that eliminated the need to filter out distracting stimuli.

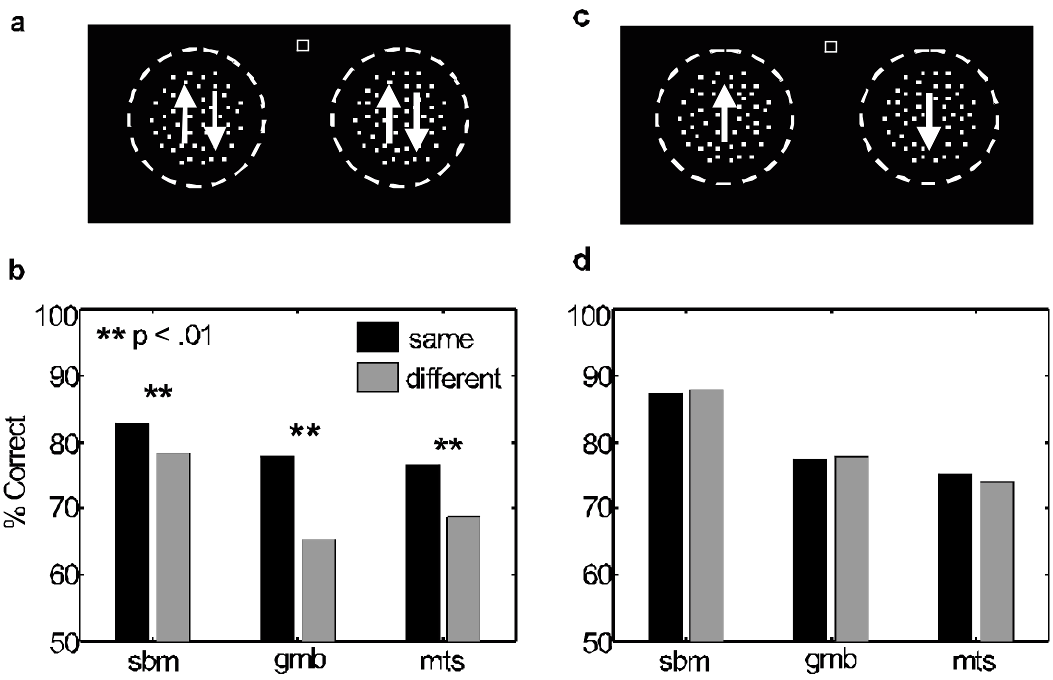

Figure 2.1 Motion dual-task experiment
(**a**) Stimulus diagram of the direction of motion experiment. Both left and right stimuli were composed of overlapping fields of upward and downward moving dots. While fixating, subjects concurrently performed a speed discrimination task on one field of dots from each side, either moving in the same direction (up or down on both sides) or in different directions (up on one side and down on the other). (**b**) Task performance was better when dividing attention across *same* vs. *different* directions for all subjects. (**c**) Stimulus diagram of the same experiment without distractors. Left and right stimuli were each composed of a single field of moving dots. Subjects concurrently performed a speed discrimination task on the single field of dots from each side, either moving in the same or different directions (only one example is diagrammed here). (**d**) The difference in task performance when dividing attention across *same* vs. *different* directions was reduced for all subjects.

2.2 Methods

2.2.1 Direction of motion experiment

The stimulus was composed of two spatially separate circular apertures (radius = 5 deg) of moving random dots centered 11 deg to the left and right of and 2.5 deg below a central fixation point (Figure 2.1a). Dots were white (560 cd/m^2) on a gray background (230 cd/m^2) and were each 0.6 deg of visual angle in width. The left and right sides of the display were identical and were each composed of two overlapping fields of upward and downward moving dots (50 dots per field). The dots within each field moved coherently and had limited lifetimes (200 ms) to prevent the tracking of individual dots. Because of their opposing motions, the overlapping fields of upward and downward moving dots on each side perceptually segregated allowing observers to selectively attend to a single direction of motion on each side. To direct their attention, subjects were instructed to perform a threshold level speed discrimination task on one field of dots from each side at the same time (dual-task). This was designed so that subjects could perform the dual-task on one field of dots from each side moving in either the same direction (both up or both down) or in different directions (one up and one down) without changing the visual stimulus, eye position, or the spatial distribution of attention.

The dual-task was performed in successive two-interval forced choice trials (2-IFC) initiating every 3.3 sec. During each trial, the stimulus was presented for two sequential 500 ms intervals separated by a 100 ms interval in which only the fixation point was present. Brief presentations of 500 ms were used to encourage subjects to perform the left and right discriminations simultaneously rather than sequentially. For each of the four fields of dots, a threshold level speed change occurred between the two intervals on 50% of trials. Specifically, on 50% of the trials, the dots moved at the baseline speed during one interval and at a slightly incremented speed during the other interval (in either order). On the other 50% of trials, there was no speed change across intervals; the dots moved at the baseline speed during both intervals. Whether a speed change occurred or not was independently randomized for each of the four fields of dots on every trial. At the end of each trial, the subjects' task was to report whether or not a speed change occurred within each of the two attended fields of dots (and speed changes that occurred in the distracting field of dots were to be ignored). Thus, there were four equally probable responses: change (on left)/change (on right), change/no change, no change/change, or no change/no change. Subjects indicated these responses by pressing 1 and 0 on a keypad in the following combinations: 11, 10, 01, and 00, respectively. Feedback was given during the inter-trial interval as a small ‘yes’ or ‘no’ appearing above the fixation point corresponding to each side.

It is important to note that the task was not to compare speeds across sides but rather to make an independent judgment on each side. On every trial, baseline speeds were different on each side so that observers could gain no benefit from comparing stimulus speeds across sides. If the baseline speed for the two fields of dots on the left was 10 deg/sec, then the baseline speed for the two fields of dots on the right was 20 deg/sec, and vice versa. Whether the higher baseline speed occurred on the left or right side was randomly determined for each trial. The difference in baseline speeds across sides was essential, as it would be a trivial result if subjects performed better when judging two fields of dots moving in the same direction because they benefited from comparing speeds across sides. There were four combinations of dot fields which could be attended per trial: up (on left)/up (on right), down/down, up/down, and down/up. Data were collected in blocks of 36 trials of each of the four trial types. At the start of a block of trials, a phrase presented on the screen instructed subjects which combination of dots to attend for that block (e.g. “Attend Up on the Left and Down on the Right”). Subjects each performed 9 interleaved blocks of each of the four trial types yielding a total of 1296 trials per subject.

Three subjects participated in this experiment. MTS and GMB were authors and SBM was a paid volunteer. Subjects (ages 25-36) had normal or corrected-to-normal visual acuity. All subjects gave written, informed consent. Before data collection, subjects trained equally on all four trial types until stable performance was achieved (minimum 1000 practice trials). Speed increments were chosen that resulted in a performance of approximately 80% correct on the dual task. The speed increments used for all subjects were 7.1 deg/sec for dots with a baseline speed of 10 deg/sec and 9.6 deg/sec for dots with a baseline speed of 20 deg/sec.

2.2.2. Direction of motion experiment without distractors

The same three subjects (MTS, GMB, and SBM) participated in a second version of this divided attention experiment that eliminated the need to filter out distracting motion. In this second experiment, only a single field of moving dots was presented on each side of the fixation point (Figure 2.1b). Subjects performed the same speed discrimination dual task as in the previous experiment on dots moving in either the same or in different directions of motion. Note that with only a single field of dots presented on each side, the stimulus was physically different during each of the four conditions: up (on left)/up (on right), down/down, up/down, and down/up. Subjects again performed 9 interleaved blocks of 36 trials of each of the four trial types, yielding a total of 1296 trials per subject. Without distractors, the task was less difficult and speed increments were reduced to maintain performance at approximately 80% correct. The speed increments used for all subjects were 5.7 deg/sec for dots with a baseline speed of 10 deg/sec and 7.9 deg/sec for dots with a baseline speed of 20 deg/sec.

Performing the dual task both with and without distractors allowed us to compare our results to previous studies reporting greater effects of attention when multiple stimuli compete for attentional selection within single neuronal receptive fields (Luck, Chelazzi, Hillyard, & Desimone, 1997; Moran & Desimone, 1985; Motter, 1993; Treue & Martinez Trujillo, 1999; Treue & Maunsell, 1999). If a performance difference between trial types found in the first experiment was due to attention, then we might expect the effect to be diminished by removing the overlapping distracting dots. Alternatively, if the performance difference was simply due to benefit gained from comparing speeds across sides, then we would expect that benefit to remain after removing the distractors.

2.2.3 Color experiment

We performed an analogous second experiment using color as the attended feature. The general methods were the same as in the first experiment and only the differences are emphasized here. The left and right sides of the display were each composed of two overlapping fields of stationary red and green random dots (50 dots per field) (Figure 2.2a). Whenever there were overlapping pixels between two dots in the display, those pixels were randomly assigned the color of one of the two overlapping dots so that neither field of dots appeared to be in front of the other. Stimuli were displayed in the upper visual hemifield (2.5 deg above fixation) to be consistent with our fMRI experiment involving feature-based attention to color. The dots had limited lifetimes (200 ms) and appeared to flicker. Subjects were instructed to perform a threshold level luminance discrimination task on one field of dots from each side at the same time. Under identical stimulus conditions, attention could thus be divided across two fields of dots with either the same color (both red or both green) or with different colors (one red and one green).

During each 2-IFC trial, the task was to report whether or not a threshold level luminance change occurred between the two intervals for each of the two attended fields of dots. As in the first experiment, whether a luminance change occurred or not was independently randomized for each of the four fields of dots every trial. There were four equally probable responses: change (on left)/change (on right), change/no change, no change/change, or no change/no change. Furthermore, baseline luminances on the two sides were randomized across trials so that subjects could not benefit from comparing luminances across sides.

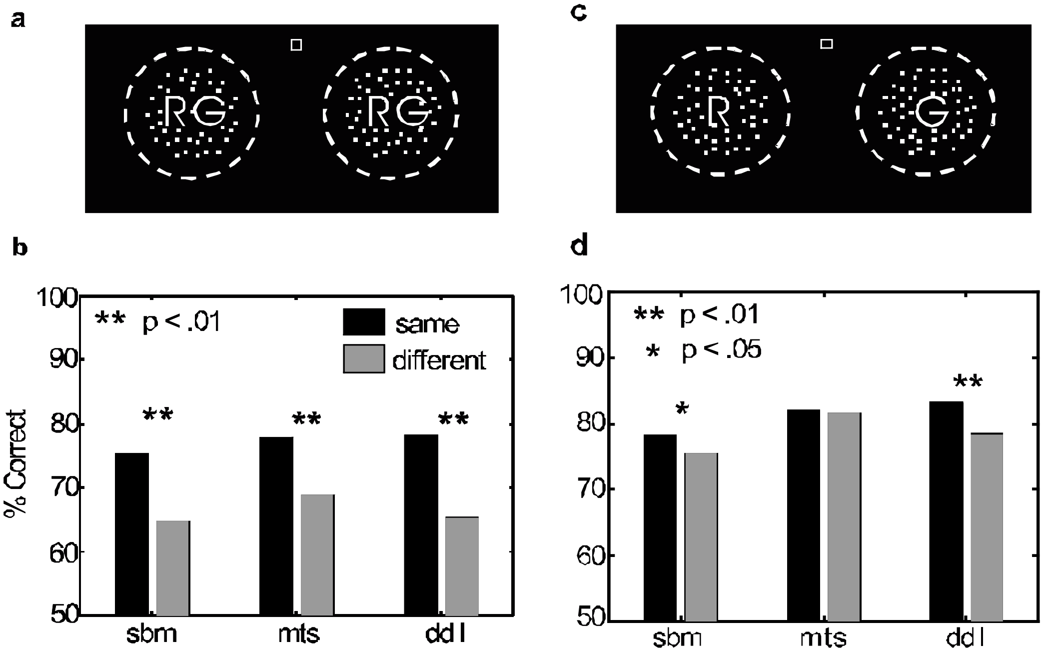

Figure 2.2 Color dual-task experiment
(**a**) Stimulus diagram of the color experiment. Both left and right stimuli were composed of overlapping fields of red and green stationary dots. While fixating, subjects concurrently performed a luminance discrimination task on one field of dots from each side, either of the same color (red or green on both sides) or of different colors (red on one side and green on the other). (**b**) Task performance was better when dividing attention across *same* vs. *different* colors for all subjects. (**c**) Stimulus diagram of the same experiment without distractors. Left and right stimuli were each composed of a single field of colored dots. Subjects concurrently performed a luminance discrimination task on the single field of dots from each side, either of the same color or of different colors (only one example is diagrammed here). (**d**) The difference in task performance when dividing attention across *same* vs. *different* colors was reduced for all subjects.

There were four combinations of dot fields which could be attended: red (on left)/red (on right), green/green, red/green, and green/red. Data were collected in blocks of 36 trials of each of the four trial types. Subjects each performed 9 interleaved blocks of each of the four trial types, yielding a total of 1296 trials per subject.

Subject MTS was an author and SBM and DDL were paid volunteers. Subjects (ages 25-27) had normal visual acuity and color vision. All subjects gave written, informed consent. Before data collection, subjects trained equally on all trial types until stable performance was achieved (minimum 1000 practice trials). Luminance increments were chosen that resulted in performance of approximately 80% correct on the dual task. The red and green dots were not equated for luminance so each had different baseline luminance values. Luminance increments used for all subjects were 15 cd/m^2 and 17 cd/m^2 for red dots with baseline luminances of 137 cd/m^2 and 153 cd/m^2, respectively. Luminance increments were 25 cd/m^2 and 28 cd/m^2 for green dots with baseline luminances of 225 cd/m^2 and 250 cd/m^2, respectively. (Weber fractions of 0.11).

2.2.4 Color experiment without distractors

The same three subjects (MTS, SBM, and DDL) performed a second version of the color experiment that eliminated the need to filter out distracting stimuli. Only a single field of dots was presented on each side of the fixation point (Figure 2.2b). The fields were either of the same color or of different colors: red (on left)/red (on right), green/green, red/green, or green/red. Data was collected in interleaved blocks of 36 trials of each of the four trial types, yielding a total of 1296 trials per subject. Surprisingly, the dual task was not noticeably easier without distractors and the same luminance increment thresholds were used as in the previous color experiment to maintain a task performance of approximately 80%.

2.2.5 Equipment and stimulus details

Stimuli for both experiments were generated on a Macintosh PowerBook computer using Matlab v4.3 and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli were displayed using an LCD projector (60 Hz frame rate) using a back-projection screen. Subjects sat in a darkened room in an upright position with their head in a chin-rest and viewed the screen at a distance of 18 cm. Stimuli were presented using an LCD projector instead of a CRT monitor in order to match the stimulus characteristics used in the previous fMRI experiment (Saenz et al., 2002).

2.3 Results

2.3.1 Direction of motion experiment

Subjects performed the dual task on two fields of dots moving in either the same direction (*same* trials) or different directions (*different* trials). The visual display was unchanged across *same* and *different* trials; only the attentional state of the observer differed. Figure 2.1a plots performance on *same* trials compared to *different* trials for each of three subjects. All subjects performed significantly better on *same* trials than on *different* trials (SBM, 82.9% vs. 78.3% correct; GMB, 78.8% vs. 65.4%; MTS, 76.6% vs. 68.8%; $p<0.01$ for each subject; $n = 1296$ trials/subject). Subjectively, subjects reported that attending to a particular direction of motion on one side seemed to make the dots moving in the same direction on the other side more salient, thereby facilitating task performance on same trials and interfering with task performance on different trials.

2.3.2 Direction of motion experiment without distractors

Subjects performed the dual task in the absence of distracting stimuli. As shown in Figure 2.1b, there was no difference in performance on *same* vs. *different* trials for all subjects (SBM, 87.4% vs. 87.9% correct; GMB, 77.5% vs. 77.8%; MTS, 75.1% vs. 74.0%, $p>0.05$ for each subject; $n = 1296$ trials/subject).

2.3.3 Short presentation trials

It could be argued that the 500 ms presentation intervals were not sufficiently short to guarantee simultaneous performance of the dual task. To address this concern, we reran the full experiment for one subject with 200 ms presentation intervals. Additional training was required for that subject to perform the dual-task with shorter presentation intervals. The speed increments used were 9.0 deg/sec for dots with a baseline speed of 10 deg/sec and 9.6 deg/sec for dots with a baseline speed of 20 deg/sec. With the shorter presentation times, performance on same trials remained near 80% correct while performance on different trials dropped to near chance (MTS, 80.1% vs. 59.6%; $p<0.001$). As with the 500 ms presentations, there was no statistical difference in performance across the same and different conditions when the distracting fields were removed (MTS, 79.6% vs. 77.5%, $p>0.05$).

2.3.4 Color experiment

Subjects performed the dual task on two fields of dots of either the same color (*same* trials) or different colors (*different* trials). Again, the visual display was unchanged across *same* and *different* trials; only the attentional state of the observer differed. Figure 2.2a plots performance on *same* trials compared to *different* trials for each of three subjects. Consistent with the direction of motion experiment, all subjects performed significantly better on *same* trials compared to *different* trials (subject SBM, 75.4% vs. 64.8% correct; DDL, 78.4% vs. 65.4%; MTS, 77.9% vs. 68.9%; $p<0.01$ for each subject; $n = 1296$ trials/subject).

2.3.5 Color experiment without distractors

Subjects performed the dual task in the absence of distracting stimuli. As plotted in Figure 2.2b, the difference in performance between *same* vs. *different* trials was reduced or eliminated for all subjects (SBM, 78.9 vs. 75.6% correct, $p<0.05$; DDL, 83.3 vs. 78.6%, $p<0.01$, MTS, 82.1 vs. 81.6%, $p>0.05$; $n = 1296$ trials/subject).

2.3.6 Effects of learning

Before data collection, subjects trained on all trial types until stable performance was achieved. We confirmed that the amount of training was adequate by separately analyzing the data from the first and second halves of data collection in all experiments. In the direction of motion and color experiments with distractors, all subjects performed better on *same* trials than on *different* trials in both the first and second halves of the data ($p<0.05$ for each subject in each half of each experiment). In the direction of motion experiment without distractors, all subjects showed no significant performance difference between *same* and *different* trials in both halves of the data ($p>0.05$ for each subject in each half). In the color experiment without distractors, MTS showed no performance difference in either half ($p>0.05$), SBM showed a performance difference that was not significant in the first half but was significant during the second half ($p<0.05$), and DDL showed a performance difference in both halves ($p<0.01$). This analysis suggests that additional training would not have changed the outcome of the comparisons in any of the experiments.

2.4 Combined effects of feature-based attention to direction of motion and color

The dual task experiment was adapted to include both features (direction of motion and color) at the same time. These new experiments were designed to address the following questions: Would performance be more strongly affected when observers divided attention across stimuli with *two* same vs. *two* different features compared to the previous experiments when only one feature was used? Would performance be at an intermediate level when one feature was the same and the other feature different? Would these feature-based attention effects depend on the task used? Unfortunately, the results did not provide a clear answer to these questions. There were individual differences in the overall pattern of performance that likely reflected differences in the task strategy. These data are unpublished.

2.4.1 Methods

The left and right side of the display were each composed of two overlapping fields of colored (red or green), moving (upward or downward) dots (50 dots per field). Stimuli were displayed in the upper visual hemifield. Observers attended one field of dots from each side either with two common features (same color /same motion direction), two different features (different color/different motion direction), or a combination of features (same color/different motion direction or different color/same motion direction). To direct attention, subjects performed either the speed discrimination task (from the motion experiment) or luminance discrimination task (from the color experiment).

There were 16 possible combinations of dot fields to be attended (red-up/red-up, red-up/red-down, red-up/green-up, red-up/green-down, etc.) that grouped into the four main conditions listed above. At the start of each 36 trial block, subjects were instructed on which combination of dot fields to attend. Data were collected in interleaved blocks of each of the 16 combinations with the speed discrimination task (648 trials of each of the four conditions per subject). The experiment was then repeated with the luminance discrimination task. Three subjects (SBM, DDL, and MTS) participated and their previously determined discrimination thresholds were used.

2.4.2 Results

Individual subject data is plotted in Figure 2.3. The left column shows performance on speed discrimination trials and the right column shows performance on luminance discrimination trials. Each bar graph also plots the

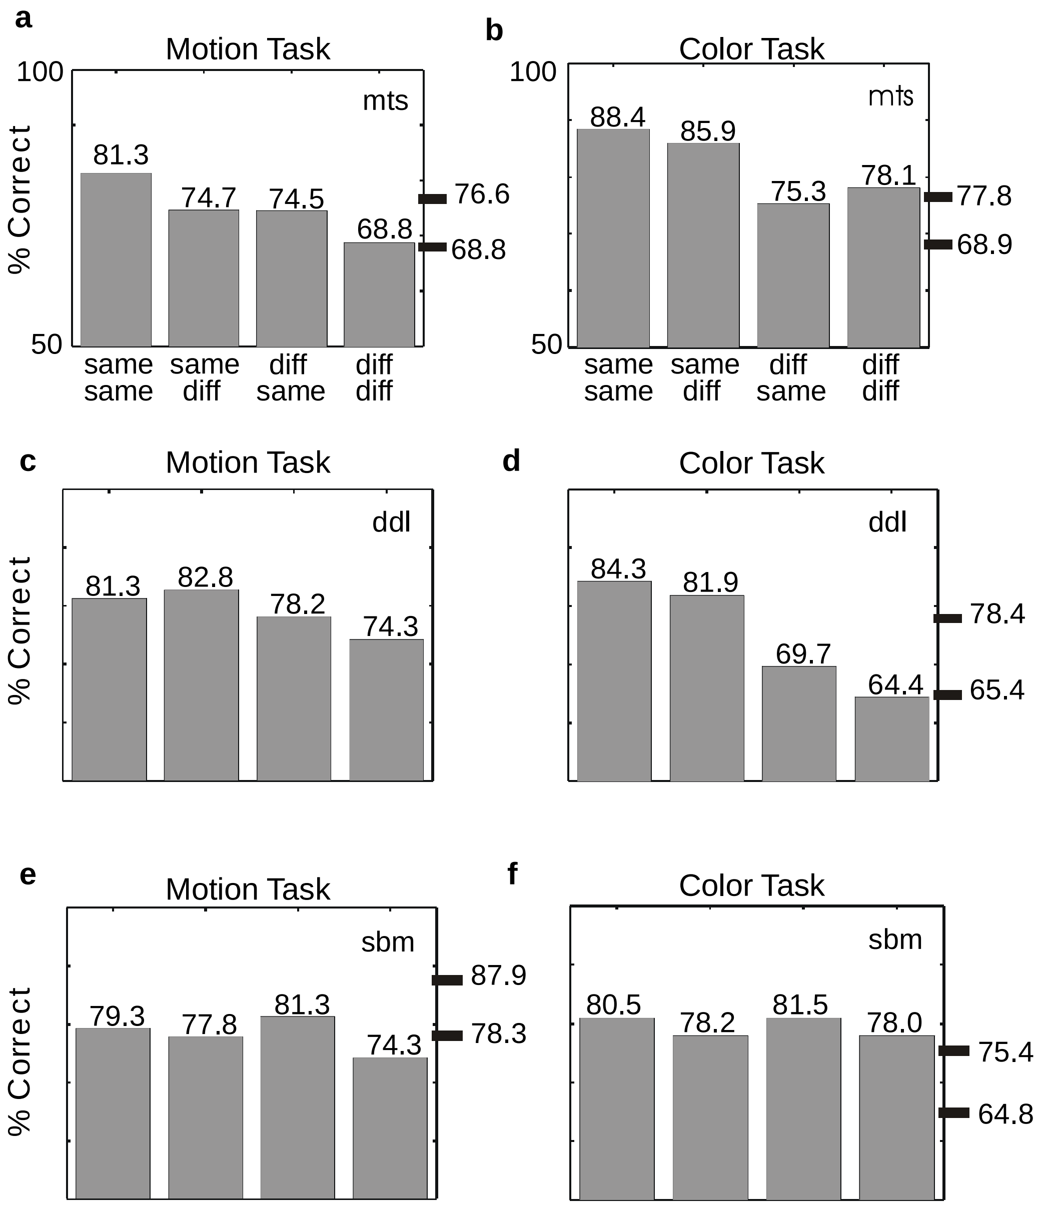

Figure 2.3 Combined features dual-task experiment
Each plot shows one subject's performance when attending fields with the same color/same direction, same color/different direction, different color/same direction, and different color/different direction. The left and right columns plot speed and luminance discrimination task performance, respectively. Each subject's performance on same vs. different feature trials from the single feature experiment are also plotted on the right vertical axes for comparison.

subject's previous performance on the single feature experiments for comparison (see right vertical axes). All subjects performed better when dividing attention across two same features compared to two different features (same/same vs. diff/diff). However, the data reveals individual differences in the exact patterns of performance. In some cases, the performance difference with two featureswas clearly larger than it had been in the single feature experiments (see Figure 2.3d) and in some cases it was not (see Figure 2.3e). In some cases, intermediate performance levels were seen in the combined feature conditions (see Figure 2.3a) and in some cases were not (see Figure 2.3f). The complexity of this two-feature experiment likely allowed subjects to employ different task strategies. For example, to select an upward moving red field, subjects could have selected it based on its color, its motion, or a combination of both. We had initially considered adapting our fMRI paradigm to include both features. However, considering the inter-subject variability seen here, we decided against pursuing the fMRI study.

2.4 Discussion

2.4.1 Summary

We found that observers were better able to concurrently discriminate spatially separate stimuli when those stimuli had common features compared to opposing features. This effect was demonstrated for the two features tested, direction of motion and color. We used overlapping stimuli that were identical in all conditions so that differences in task performance could not be confounded with changes in the stimulus itself or with changes in the spatial distribution of attention. Furthermore, the attentional effect was reduced when the need to filter out overlapping distractors was eliminated.

Additional experiments were directed at combining the effects of feature-based attention to direction of motion and color. However, due to variable performance on this more complex paradigm, no additional conclusions are drawn from these results.

Our findings are consistent with our previously reported fMRI findings in which attention to a particular feature of one stimulus was found to increase cortical responses to a spatially separate ignored stimulus sharing the attended feature (Saenz et al., 2002). This feature-specific response enhancement was observed in multiple early cortical visual areas and suggests that attention improves the processing of stimuli sharing the attended feature throughout the visual field.

If feature-based attention improves globally improves the processing of stimuli with the attended feature, we reasoned that this should facilitate the distribution of attention across multiple stimuli with common features compared to opposing features. Consistent with that interpretation, subjects in the present experiment reported subjectively that attending to a particular direction of motion or color on one side of the display seemed to make that feature more salient on the other side. Correspondingly, performance was facilitated when observers divided attention across matching features and performance was impaired when observers divided attention across opposing features. Together, the results from our fMRI and psychophysical studies provide complementary physiological and behavioral evidence that feature based attention does indeed improve the processing of stimuli throughout the visual field that share the attended feature.

2.4.2 Role of the distractors

Interestingly, the difference in task performance depended on the need to filter out competing stimuli. In the direction of motion experiment, the performance difference was eliminated for all subjects in trials without distractors. In the color experiment, the difference was reduced or eliminated for all subjects in trials without distractors. The difference in results obtained with and without distractors is not related to overall task difficulty because task performance was in the same range across both sets of trials. Rather, this result is consistent with neurophysiology studies reporting greater effects of attention on individual neurons when multiple stimuli compete for attentional selection within the receptive field (Luck et al., 1997; Moran & Desimone, 1985; Motter, 1993; Treue & Martinez Trujillo, 1999; Treue & Maunsell, 1999). These studies suggest that the role of attention in target selection is greatest in the presence of nearby distractors.

The weakening of the attentional effect in the absence of distractors rules out a potential confounding factor. It would be a trivial finding if subjects performed better when judging matching features simply because they benefited from comparing those features across sides. This was not the case. If such a benefit existed, it would have also been evident on trials without distractors. In the color experiment without distractors, the performance difference was greatly reduced but not eliminated for two out of three subjects. This remaining difference may be interpreted as an estimate of the size of the effect that could be attributed to other factors such as comparing luminances across sides. However, it may also be the case that in the color experiment, feature-specific attention facilitated the discrimination of stimuli with common features even in the absence of distractors.

2.4.3 Simultaneous vs. sequential task performance

In all experiments, subjects were instructed to divide attention equally to the left and right sides of the display and perform the two tasks concurrently.

However, it is difficult to rule out the possibility that subjects shifted spatial attention back and forth between the two sides and performed the tasks sequentially. The increased performance difference obtained with shorter presentation times (200 ms compared to 500 ms) is consistent with the hypothesis that observers were better able to divide their attention *concurrently* across stimuli with common features compared to opposing features. The shorter presentation time should have been more effective in preventing the observer from switching attention between the two stimuli to avoid this limitation.

2.4.4 Possible neuronal mechanisms

Our psychophysical results and previous fMRI results are consistent with a neuronal mechanism by which attention enhances the activity of cortical neurons that encode behaviorally relevant stimulus properties. A *biased competition model* predicts this type of feature-specific attentional modulation (Desimone & Duncan, 1995; Reynolds, Chelazzi, & Desimone, 1999). The model proposes that multiple stimuli activate competing populations of neurons and attention biases the competition in favor of neurons that encode the features of the attended stimulus. Multiple studies have shown that when a pair of stimuli with different features is presented within a visual cortical neuron’s receptive field, the response of the neuron is determined by which of the two stimuli is attended. Attending to the preferred stimulus of the pair increases the neuron’s firing rate and attending to the non-preferred stimulus decreases the firing rate. Thus, the effect of attention on a neuron’s response (enhancement or suppression) depends on how the features of the attended stimulus match the stimulus selectivity of the neuron. This result has been confirmed for a range of stimuli and visual areas including color stimuli in V2, V4, and IT (Luck et al., 1997; Moran & Desimone, 1985; Reynolds et al., 1999), motion stimuli in MT/V5 (Treue & Maunsell, 1996; Treue & Maunsell, 1999) and complex objects in V4 and IT (Chelazzi et al., 1998; Chelazzi et al., 1993; Chelazzi, Miller, Duncan, & Desimone, 2001).

Based on these findings, we can speculate about the neuronal mechanisms that mediated our behavioral results. In our divided attention experiments with distractors, overlapping fields of dots with opposing features were presented, presumably activating neurons tuned to both of those features (i.e. upward and downward direction selective neurons or red and green color selective neurons). Attending to one of the fields would have increased the responses of neurons encoding the features of the attended field and suppressed the responses of neurons encoding the features of the overlapping distracting field. Attending to the same feature on both sides of the display may have mutually enhanced the responses of neurons throughout the visual field tuned to the attended feature and suppressed neurons tuned to the opposing feature. This mutual enhancement and suppression may have aided the selection of target fields on our ‘same’ feature trials, facilitating task performance. Attending to opposing features may have initiated competing effects of enhancement and suppression in both populations of neurons. This interference may have made target selection more difficult on ‘different’ feature trials, hindering task performance. Thus, a combination of neuronal facilitation and suppression due to attention may have contributed to our psychophysical results.

When the need to filter out (or suppress) overlapping stimuli was removed, the competition between opposing neuronal populations would have been reduced. A subset of the studies listed above also measured the effects of attention when only a single stimulus was presented inside the receptive field (Luck et al., 1997; Moran & Desimone, 1985; Treue & Maunsell, 1999). Another study also compared the effects of attention on responses (V1, V2, and V4) to a target stimulus in the presence or absence of nearby distractors (Motter, 1993). In all cases attentional modulation was reduced in the absence of competing distractors and, in some cases, was eliminated (Luck et al., 1997; Moran & Desimone, 1985). Consistent with these results, in our experiments the performance difference was reduced in the absence of distractors.

Our interpretation requires that the top-down biasing effects of attention be far-reaching enough to affect the processing of a visual object located in the opposite visual hemifield. In support of this, Chelazzi and colleagues showed that searching for a visual stimulus increased the firing rate of IT neurons tuned to that stimulus during a time period *prior* to stimulus presentation (Chelazzi et al., 1998; Chelazzi et al., 1993). This modulation of baseline firing rates was feature-driven and far-reaching because the exact location of the upcoming target was unknown (but the location was limited to a single visual hemifield). Other studies have shown that the modulatory effects of feature-based attention do indeed extend into the opposite visual hemifield (McAdams & Maunsell, 2000; Treue & Martinez Trujillo, 1999).

In particular, Treue and Martinez-Trujillo reported feature-specific attentional modulation of stimulus-evoked responses in macaque area MT/V5. In their experiment an ignored random dots stimulus, moving coherently in the preferred direction, was presented inside the receptive field of a directionally tuned neuron. Attention was directed to a second stimulus, outside the receptive field, that either moved in the same or in the opposite direction. On average, neuronal responses to the ignored stimulus increased when the monkey attended the preferred direction and decreased when the monkey attended the opposing direction (compared to passive viewing trials). To account for these results, the authors proposed a *feature-similarity gain model* in which feature-based attention modulates the gain of cortical neurons that are selective for the behaviorally relevant stimulus property. The model emphasizes that the direction of the gain change (decrease or increase) depends on how the attended properties (location or features) match the stimulus selectivity of the neuron and also emphasizes that the modulation will reach neurons with receptive field locations well outside the attended location. Our previous fMRI results as well as the present psychophysical results are consistent with a spatially non-specific mechanism of feature-specific neuronal modulation.

Another explanation for the effectiveness of the distractors, besides competition within receptive fields, may also have to do with task strategy. In the experiments with distractors, observers were required to select one of two overlapping fields of dots with a particular feature and perform a discrimination task on the selected field. In the motion experiment, observers selected a field with a particular direction of motion in order to perform a speed discrimination task and in the color experiment observers selected a field of a particular color in order to perform a luminance discrimination task. Hence, it was primarily the selection of the target field in the presence of the distracting field, rather than the task itself, that required feature-based attention to either direction of motion or color. When the need to filter out the overlapping stimulus was removed, target selection may have been less dependent on feature-based attention. This may have contributed to the reduction of the attentional effect in both experiments without distractors. The amount of feature-based attention that remained after the removal of the distractors may have been different for the two tasks (speed discrimination and luminance discrimination) which could have contributed to the different degrees of effect reduction found in the motion and color experiments without distractors.

2.4.5 Related human psychophysical studies

The results of our divided attention study indicate that attention to a stimulus feature facilitates the concurrent processing of other stimuli sharing that same feature. This interpretation is consistent with previous psychophysical studies suggesting that observers have a limited ability to attend to more than one spatial frequency at a time (Shulman & Wilson, 1987; Sperling & Melchner, 1978). Our results are also consistent with a study of feature-specific attention (Rossi & Paradiso, 1995) in which observers performed a primary task of discriminating a feature of a foveal grating (spatial frequency or orientation) and a secondary task of detecting a near-threshold grating in the periphery. Although the tasks were not performed concurrently, observers were better at detecting the peripheral grating when its spatial frequency or orientation matched the attended feature in the primary task.

Lee, Koch, and Braun asked a related question of whether the ability to perform simultaneous tasks depends on the similarity of the two tasks involved. Observers performed a dual-task that involved discriminating dissimilar stimulus dimensions (e.g. form vs. motion) compared to similar stimulus dimensions (e.g. motion vs. motion) (Lee, Koch, & Braun, 1999). They concluded that while it was more difficult to perform two tasks compared to one, it did not matter whether those two tasks were similar or dissimilar. Another recent study (Morrone, Denti, & Spinelli, 2002) reported that performing concurrent tasks on the same stimulus dimension was more difficult than on different stimulus dimensions for tasks involving color vs. luminance contrast discrimination. Our findings are not inconsistent with these results. In our dual-task experiment, subjects always discriminated the same stimulus dimension at a time (either direction of motion or color). What varied was whether the simultaneous tasks involved the same vs. opposing features within a particular stimulus dimension. The competitive neuronal mechanisms described above could apply most specifically to neurons encoding opposing directions of a particular stimulus dimension.

2.4.6 Conclusions

Using a dual-task psychophysical paradigm, we found that subjects were better at detecting changes in a pair of spatially separated stimuli when they share a common feature, such as a direction of motion or color, than when they did not share a common feature. Our results are consistent with our fMRI findings of Chapter 1 and with a proposed mechanism, called the *feature-similarity gain model*, in which feature-based attention modulates the gain of cortical neurons tuned to the attended feature throughout the visual field (Treue & Martinez Trujillo, 1999). This global feature-based mechanism of attention could play an important role in the process of selecting the location of relevant stimuli for further processing. A feature-based increase in the saliency of regions of visual space would be useful in identifying relevant peripheral stimuli during visual search for guiding eye-movements, or in grouping stimuli with common features as part of the same object.

This chapter, in part, is a reprint of the material as it will appear in Saenz M., Buracas, G.T., and Boynton, G.M. *Vision Research,* in press, submitted 2002. The dissertation author was the primary investigator and author of this paper.

Chapter 3

Feature-based attention influences contextual interactions

Abstract

Collinear visual flankers can facilitate the detection of an oriented visual target. Freeman, Sagi and Driver (*Nature Neuroscience*, 4:1032-6 (2001)) recently reported that these lateral interactions are not only passive but can be influenced by attention. In their experiment, a foveal Gabor patch was simultaneously flanked by high contrast collinear Gabors on the collinear axis and by orthogonal Gabors on the orthogonal axis. Detection thresholds for the foveal Gabor were decreased when human observers attended to the collinear flankers compared to the orthogonal flankers, even though visual stimulation was identical in both conditions. Here we propose that this facilitation is not limited to collinear lateral interactions, but rather can be accounted for by a feature-based mechanism of visual attention that facilitates the processing of stimuli matching an attended feature (in this case, orientation). To test this hypothesis we replicated the Freeman experiment and added an additional condition in which the flanking stimuli were co-oriented but placed on the orthogonal axis (thus co-oriented without being collinear). We found that detection of the central target was facilitated by attending co-oriented flankers on either axis (collinear or orthogonal) compared to attending orthogonal flankers on either axis. Again, visual stimulation was unchanged across all compared conditions. These results are consistent with a mechanism of visual attention whereby attention to a stimulus feature facilitates the processing of other stimuli in the visual scene sharing that feature.

3.1 Introduction

In chapter 1, fMRI experiments revealed that attention to a visual object enhanced neuronal responses to a spatially separate object that was ignored, but had the same feature. This global feature-based neuronal modulation was detected when the attended feature was motion direction or color, but not orientation. However, we cannot rule out global effects of orientation because there were differences in the experimental paradigm with orientation that could have contributed to this null result. In chapter 2, complementary psychophysical results revealed that attention to a visual object improved the discrimination of a spatially separate object with the matching feature. This psychophysical paradigm was not tested with orientation. In the present chapter, we re-explore the role of orientation using an altogether new psychophysical paradigm. In this new experiment, attention to stimulus orientation facilitated the *detection* of another stimulus with the matching orientation.

These new results are also significant because they show that attention modulates a well-known contextual interaction in an unexpected way. Previous psychophysical studies have shown that that detection of an orientated visual target is facilitated by the presence of collinear flankers (see details below). Co-oriented flankers that are not collinear did not cause the same facilitation. In these previous studies, the flankers were task-irrelevant and assumed to be passively viewed. Here we show that when the flankers are made behaviorally relevant, co-oriented flankers facilitate detection as much as collinear flankers. This is an exciting result because it demonstrates that contextual interactions may be dynamically modified with changing task demands to meet the behavioral needs of the observer. This attentional modulation was consistent with a global feature-based mechanism of attention.

This introduction continues with background information on the relevant contextual influences on both behavioral and neuronal responses, new topics in this dissertation.

3.1.1 Lateral interactions

Contrast detection thresholds for a foveal Gabor patch are decreased by up to a factor of two by the presence of high contrast collinear flanking Gabors (Figure 3.1a) (Polat and Sagi 1993). This collinear facilitation quickly declines with increased orientation, spatial frequency or position offset of the flankers relative to the central target. This result has been confirmed by others using similar Gabor stimuli (Yu and Levi 2002) as well as oriented line segments (Figure 3.1b) (Kapadia, Ito et al. 1995). In contrast, when the same co- oriented flanking stimuli are presented along an axis orthogonal to the

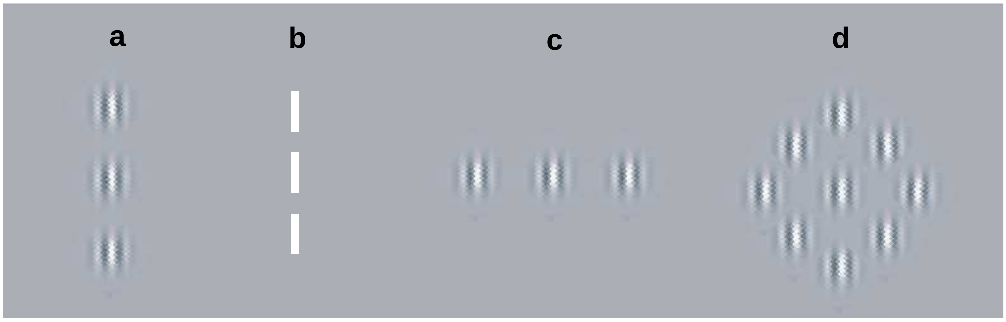

Figure 3.1 Collinear lateral interactions
(**a**) The presence of collinear flankers facilitates target detection. Facilitation is maximal when the flankers match the central target in orientation, phase, and spatial frequency (Polat and Sagi 1993). (**b**) Similar collinear facilitation occurs with oriented line segments (Kapadia, Ito et al. 1995). (**c**) In contrast, the co-oriented Gabors at other locations facilitate detection much less than collinear flankers or not at all (Polat and Sagi 1994; Solomon and Morgan 2000; Yu and Levi 2002). (**d**) Co-oriented surrounds do not facilitate detection either (Solomon and Morgan 2000). In all experiments, the flankers were task-irrelevant and assumed to be passively viewed.

orientation of the central target (Figure 3.1c), facilitation is much reduced (Polat and Sagi 1994; Solomon and Morgan 2000) or eliminated (Yu and Levi 2002). Full co-oriented surrounds also eliminate facilitation even though collinear flankers may still be present (Figure 3.1d) (Solomon and Morgan 2000); Yu and Levi 2002). In all of these experiments, the flankers or surrounds were assumed to be passively viewed. Under these passive viewing conditions, the contextual facilitation is most effective in the special case of collinearity and is not caused by co-orientation in general. This specificity has led to the proposal that collinear lateral interactions contribute to contour detection in low-level vision.

Contextual modulation of neuronal responses in primary visual cortex suggests a cellular basis for the psychophysical observations. Single-unit responses in cat V1 to low-contrast Gabor stimuli are increased by the presence of collinear flankers (Polat, Mizobe et al. 1998) and macaque V1 responses to oriented line segments are enhanced by the addition of a collinear flanking line segment (Kapadia, Ito et al. 1995). In both experiments, the flankers were located outside the classical receptive field of the recorded neuron and did not evoke responses themselves. Several properties of this neuronal facilitation covaried with those of the psychophysical facilitation; facilitation decreases with increased offset in the orientation, phase, spatial frequency or position of the flankers.

3.1.2 V1 horizontal connections

It has been proposed that the contextual modulation in V1 is mediated by long-range horizontal connections. Horizontal connections between V1 pyramidal cells extend parallel to the cortical surface for distances up to 8 mm allowing cells to integrate information from outside their classical receptive fields (Gilbert and Wiesel 1979). These projections are known to preferentially link cortical columns with similar orientation selectivity in multiple species (Gilbert and Wiesel 1989; Weliky, Kandler et al. 1995). Consistent with this preferential “like-to-like” connectivity, cross-correlation analysis reveals that cells in different orientation columns with similar tuning preferences show correlated spiking in cat V1 (T’so and Gilbert 1989, Gray et al 1989).

The link between collinear lateral interactions and V1 horizontal connections was strengthened by reports that the projections from a cortical column do not extend equally in all directions along the cortical surface. The projections in some species are anisotropic, extending preferentially along an axis in the map of visual space parallel to a column’s preferred orientation (Bosking, Zhang et al. 1997; Schmidt, Goebel et al. 1997; Sincich and Blasdel 2001). However, a recent study reports that V1 horizontal connections are not anisotropic in the macaque (Angelucci, Levitt et al. 2002).

3.1.3 Attention modulates lateral interactions

Recently, Freeman, Sagi and Driver (2001) showed that the psychophysically reported collinear interactions are not only passive but are influenced by attention. In a clever experiment, a foveal Gabor was flanked simultaneously by high contrast iso-oriented Gabors on the collinear axis and by cross-oriented Gabors on the orthogonal axis, as shown in Figure 3.2a. Contrast detection thresholds for the central Gabor were decreased when observers *attended* to the collinear flankers compared to when they attended to the cross-oriented flankers. Visual stimulation was identical in both conditions. This was a dual-task experiment in that observers performed a discrimination task on the flankers while simultaneously performing a detection

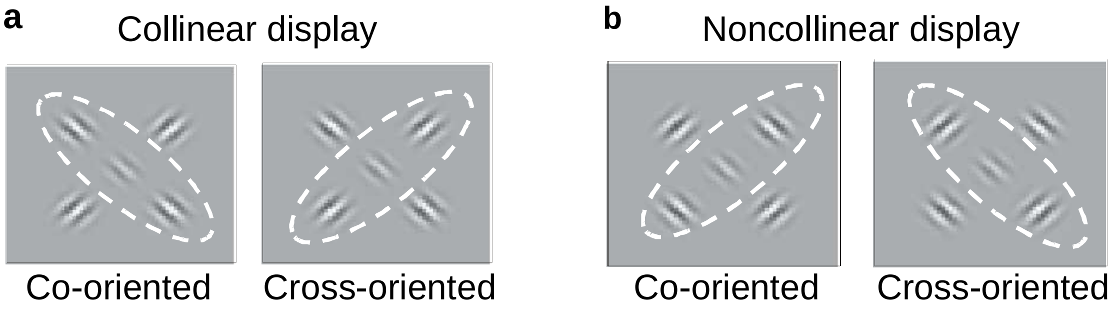

Figure 3.2 Four main stimulus conditions
(**a**) Freeman and colleagues showed that lateral interactions are modulated by attention. Contrast detection thresholds for the central Gabor were lower when observers attend to the collinear flankers compared to the cross-oriented flankers. We replicated the Freeman experiment and added an additional stimulus configuration (**b**) in which the co-oriented flankers were placed on the orthogonal axis. Hence, the flankers were still co-oriented but not collinear.

task on the center. Interestingly, the detection thresholds measured while attending to either pair of flankers were the same as detection thresholds measured on separate trials when only the attended flankers were present. In other words, ignoring a pair of flankers was equivalent to physically removing it from the display.

Here, we propose that this attentional facilitation is *not* limited to the special case of collinearity. Rather, it may be related to a feature-based mechanism of attention that facilitates the processing of stimuli matching an attended feature (in this case, orientation). To test our hypothesis we replicated the Freeman experiment and added an additional condition in which the flanking stimuli were co-oriented but placed along the orthogonal axis (thus co-oriented without being collinear) as show in Figure 3.2b. We predicted that target detection would be facilitated by attention to co-oriented flankers along either collinear or orthogonal axes.

3.2 Methods

3.2.1 Subjects

All subjects had normal or corrected-to-normal visual acuity. Three subjects (MST, GMB, and AST) were aware of the purposes of the experiment and three subjects were naïve (DAN, YGB, and SJK). All subjects gave written, informed consent.

3.2.2 Stimuli

Stimuli were presented on a 10-bit gamma corrected CRT display (1024X768 resolution, 85 Hz frame rate, background luminance 32 cd/m^2). Stimuli were generated using Matlab v4.3 and the Psychophysics Toolbox (Brainard 1997; Pelli 1997). Observers sat in a darkened room with their heads stabilized on a chin rest 60 cm from the display. The stimulus display was designed to replicate the Freeman experiment: a centrally located Gabor was surrounded by four flanking Gabors. The flankers were positioned on axes collinear and orthogonal to the orientation of the central Gabor. All Gabors had a wavelength and Gaussian distribution equal to 0.15 deg of visual angle. Spatial frequency (1/wavelength) was 6.7 cycles per degree. The center-to-center distance between the central target and flankers was 4 wavelengths (1 deg). A fixation diamond was continuously present and did not occlude the central target.

Figure 3.2 shows the two different stimulus displays: collinear and non-collinear. In the collinear display (replication of Freeman experiment), co-oriented flankers were placed on the collinear axis and cross-oriented flankers were placed on the orthogonal axis. In the non-collinear display (new to this experiment), co-oriented flankers were placed on the orthogonal axis and cross-oriented flankers were placed on the collinear axis. All Gabors had an orientation of either +45 or –45 deg with small deviations in flanker orientation as explained below. In half of the trials, the central Gabor orientation was –45 degrees, as depicted in the figure). In the other half of trials, the central Gabor orientation of +45. Thus, there were 8 different conditions: 4 main conditions (depicted in Figure 3.2) X 2 possible central orientations.

3.2.3 Procedure

Stimuli were presented in 2-interval forced choice trials (2-IFC). The trial sequence is shown in Figure 3.3. Each trial consisted of two consecutive 80 ms stimulus presentations, both of which contained the flankers and only one of which contained the central target. The stimulus presentation intervals were separated by a 500 ms interval in which only the fixation diamond was present. Observers performed two tasks: an orientation discrimination task on one flanker pair and a contrast detection task on the central target. The brief presentation intervals (80 ms) encouraged subjects to perform the two tasks simultaneously. At the end of each trial, subjects entered two unhurried responses and the next trial commenced 500 ms after the second response was received. Our flanker task differed from the Vernier offset task used in the Freeman experiment. We chose to use an orientation discrimination task in order to maximize attention to orientation rather than to spatial phase or position.

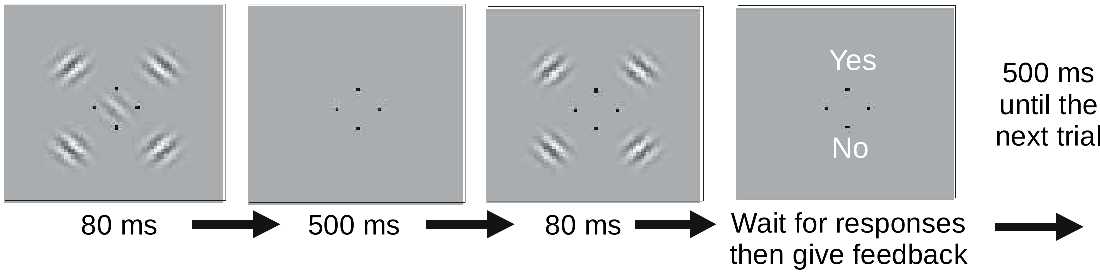

Figure 3.3 Trial sequence
Subjects performed two simultaneous tasks during each 2-interval forced choice trial. Task 1 was an orientation discrimination task on the flankers. Task 2 was a contrast discrimination task on the center. Here, the contrast of the central target is enhanced for visibility.

During each trial, the flanker pairs were rotated slightly clockwise around the diagonal axis during one interval and were rotated slightly counterclockwise during the other interval. Whether the flankers were rotated clockwise during the first or second interval was randomly and independently determined for each of the two flanker pairs, but the two members of a flanker pair always rotated in unison. Subjects indicated with a key press (“1” or “2”) in which interval the attended flankers were rotated clockwise. Threshold level orientation changes were pre-determined for each subject (range: 0.85 to 1.6 deg rotation) using a standard staircase procedure. During data collection, flanker orientation changes were fixed at the pre-determined threshold level (predicting 80% performance) and flanker contrast was fixed at 100%.

For the contrast detection task, subjects indicated with a second key press (“1” or “2”) which interval contained the central target. Contrast detection thresholds were measured under the 8 different conditions described above using a method of constant stimuli. During blocks of 75 trials of a given condition, the contrast of the central target was varied over 5 levels (1.7, 3.0, 5.2, 9.3, and 16.6% contrast) in randomized order. Psychometric curves were fit with a Weibull function using a maximum likelihood procedure to determine the contrast that predicts 80% performance. Instructions about which flanker pair to attend were displayed on the screen at the start of each block. During each experimental session (which lasted approximately 45 minutes), subjects performed eight blocks (one of each condition) in randomized order. Subjects each performed 6-10 sessions.

In order to emphasize attention to the flankers, subjects were told to consider the flanker task as the primary task and the central task as the secondary task. Subjects were informed that a minimum performance of 70% correct on the flanker task was required for each block (80% performance on average was expected). Flanker scores on all blocks were checked at the end of each session; if performance was below 70% the data was discarded and subjects repeated the block.

3.3 Results

3.3.1 Individual subject results

Figure 3.4 plots individual subject contrast detection thresholds for the four main stimulus conditions. Thresholds were typically lower when observers attended co-oriented flankers on the collinear axis compared to cross-oriented flankers (replication of Freeman experiment). Thresholds were also typically lower when observers attended co-oriented flankers on the orthogonal axis compared to cross-oriented flankers (new result). Threshold reduction was modest and ranged from 0.5 to 3% contrast, within the range reported in the Freeman experiment. The effect was not significant for all subjects and one subject (YGB) showed a reverse trend on the non-collinear display.

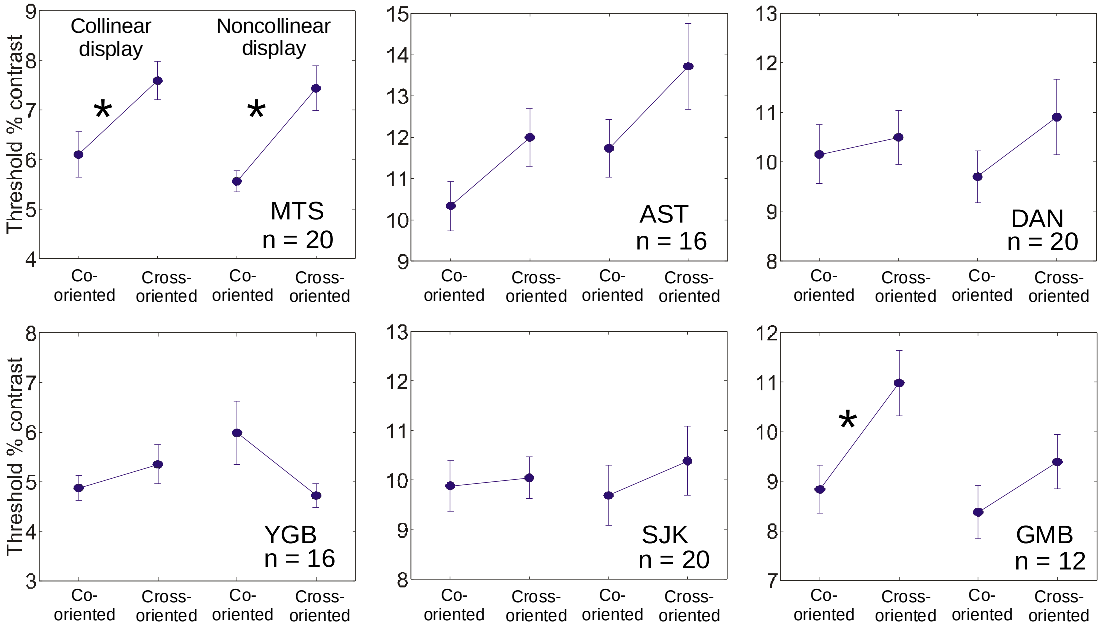

Figure 3.4 Individual results
Contrast detection thresholds measured in the four stimulus conditions for each of 6 subjects. Each data point was determined by an average of *n* independent threshold measurements. Error bars indicate SEM. * denotes a significance level of $p<0.05$.

3.3.2 Group results

The results are better appreciated after averaging the data across all subjects. Figure 3.5 plots the mean relative thresholds. To calculate relative thresholds the data was normalized across subjects and displays: each threshold measurement was divided by the observer's mean threshold on the given display (collinear or noncollinear). Thus, we would expect the normalized values to be either above or below 1. Each data point was determined by an average of 104 independent threshold measurements (104 blocks, 7800 trials). The difference in mean relative thresholds was significantly differently (t-test, $p < 0.05$) between the two main conditions on both displays.

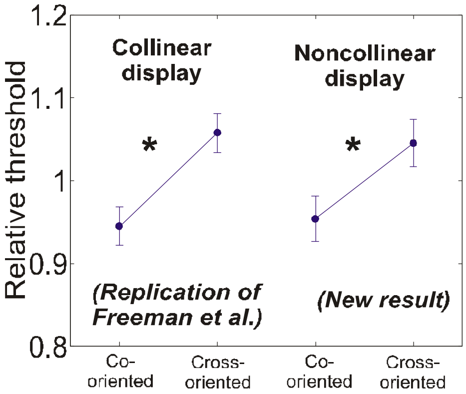

Figure 3.5 Group Results
Mean relative thresholds averaged across all six subjects. To normalize results, each threshold measurement was divided by the subject's mean performance on that display (see text for details). Each data point was determined by an average of 104 independent threshold measurements (7800 trials). Error bars indicate SEM. * denotes a significance level of $p<0.05$.

3.3.3 Flanker task results

We analyzed performance on the flanker task. Mean performance on the orientation discrimination task for all subjects was 80.1% (s.e.m. = 0.3%) indicating that rotation thresholds were well chosen. Co-orientation of the flankers with the central target had no effect on task performance (mean performance on co-oriented trials = 80.2% vs. cross-oriented trials = 80.0%; $p>0.05$) indicating that there was not a trade-off in performance between the

Table 3.1 Performance on flanker task

	Mean percent correct +/– standard error		
	All trials	Co-oriented trials	Cross-oriented trials
MTS	78.1 +/– 0.5%	78.2 +/– 0.8%	78.0 +/– 0.7%
AST	82.4 +/– 0.7%	82.2 +/– 0.9%	82.1 +/– 1.0%
DAN	79.5 +/– 0.6%	79.3 +/– 0.8%	79.6 +/– 0.9%
YGB	80.0 +/– 0.7%	81.2 +/– 1.0%	78.9 +/– 1.0%
SJK	81.2 +/– 0.5%	80.6 +/– 0.8%	81.7 +/– 0.7%
GMB	80.1 +/– 0.7%	80.4 +/– 1.0%	79.8 +/– 1.2%

two tasks. Individual subject performance is shown in Table 3.1. Performance did depend on target contrast. Performance was slightly but significantly impaired on high contrast trials compared to low contrast trials. This effect matches observers' reports that when the central target was clearly visible it captured attention and interfered with performance of the flanker task.

3.4 Discussion

Central target detection was facilitated by attention to co-oriented flankers on either axis (collinear or orthogonal) compared to attention to cross-oriented flankers on either axis. Hence, attentional facilitation was not limited to a special case of collinearity. These results are consistent with a mechanism of feature-based attention by which attention facilitates the processing of stimuli sharing a behaviorally relevant feature (in this case, orientation).

The magnitude of this facilitation was modest (approximately 10% threshold reduction on both displays) and was not significant in all subjects. The size of the effect reported by Freeman was similarly modest, so we feel confident that we are comparing the same effect. We are concerned that in both displays the attentional facilitation was greater for non-naïve compared to naïve subjects. Neither of the differences shown in Figure 3.5 is significant if only data from naïve observers is considered, although a trend in the same direction is still observed. Additional naïve observers should be added before submitting these results for publication.

3.4.1 Relation to lateral interactions

The new results presented here would *not* be predicted by the psychophysical literature on lateral interactions. In previous studies, co-oriented flankers along the orthogonal axis (like ours) facilitate target detection much less than collinear flankers (Polat and Sagi 1994, Solomon and Morgon 2000) or not at all (Yu and Levi 2002). In these previous studies only a single flanker pair was presented at a time and was assume to be passively viewed. Our results show that when the flankers are made behaviorally relevant, orthogonal co-oriented flankers facilitate detection as much as collinear flankers. These results suggest that passive contextual interactions are not the whole story. Rather, contextual interactions may be dynamically modified with changing task demands to meet the current needs of the observer.

In support of this possibility, one study clearly shows that contextual interactions in V1 are modulated by attention (Ito and Gilbert 1999). In that experiment, an oriented line segment was placed within the receptive of a recorded macaque V1 neuron and facilitating collinear line segments were placed outside the receptive field. Responses were compared when attention was either focused on the receptive field of the recorded neuron or more broadly distributed. This manipulation of attentional state did not strongly influence responses to the central line segment itself but did strongly influence the facilitatory effect of the flankers. In some cases, facilitation was greatest under focused attention and for others facilitation was greatest under distributed attention. Overall, the effect was variable across the two animals tested which may have reflected differences in training. These findings pose more questions than answers but do point to an exciting direction of research. Contextual effects, even at the level of primary visual cortex, could be influenced by cognitive factors such as attentional state, expectation, and learning.

What are the neuronal circuits underlying contextual effects and their modulation? A large degree of flexibility in contextual modulation could potentially be mediated by horizontal connections. Horizontal connections in V1, which preferentially link regions with similar orientation preferences, are both excitatory and inhibitory (the latter depending on inhibitory interneurons) (Weliky, Kandler et al. 1995). Top-down attentional signals could alter the balance between excitatory and inhibitory influences. Contextual interactions could also be mediated by feedback connections between visual areas. Two recent studies have compared V2-to-V1 feedback connections to V1 horizontal connections (Stettler, Das et al. 2002) Feedback connections are spatially more extensive than horizontal connections and would presumably allow neurons to incorporate information from larger regions of visual space (Angelucci, Levitt et al. 2002). However, unlike horizontal connections, the feedback connections do not appear to preferentially link regions with similar orientation preference (Stettler, Das et al. 2002).

3.4.2 Feature-based attention to orientation

Our results are consistent with a mechanism by which attention to a stimulus feature globally facilitates the processing of other stimuli sharing the same feature. In Chapter 2 of this dissertation, we showed that observers were better able to concurrently process spatially separate stimuli when those stimuli had common features compared to opposing features. This effect was shown for two features tested: motion direction and color. Here we extend this finding to another feature for which cortical neurons show tuning: orientation.

The behavioral effects reported here are relatively subtle which may reflect differences in experimental design. Among the many differences, there was a large change in spatial scale. Here, the visual display was relatively small with each Gabor being less than 1 degree of visual angle in diameter. This particular display was not designed to optimally study the global effects of feature-based attention. Rather, its utility was that it allowed a comparison of the attentional effects with an extensive literature on lateral interactions.

It may also be the case that attention to orientation has a relatively small global effect compared to motion direction and color. In Chapter 1, our fMRI experiment failed to measure a global modulation of responses due to attention to orientation, although we did show clear effects due attention to color and motion direction. In a possibly relevant study, Motter and Belky (1998) analyzed macaque eye movements during a visual conjunction search for color and orientation (i.e. visual search for a red bar tilted to the right). The monkeys used color rather than orientation to guide visual search; they successively fixated on objects with the target color until finding the target orientation. This finding suggests a greater access to color information than to orientation information outside the focus of attention.

3.4.3 Motion

In my proposal for the advancement to candidacy, I suggested using this paradigm to explore lateral interactions for motion. To adapt this paradigm for motion, the Gabor stimuli could simply be set into a direction of motion orthogonal to their orientations. After adding motion, the orthogonal axis becomes an axis of *motion-collinearity*. I suggested that, under passive viewing conditions, co-oriented flankers on the orthogonal axis would cause facilitation if they were moving in the same directions as the central target. This hypothesis was tested and confirmed (Kubodera and Sato 2002). A recent study found that a single pair of passively viewed co-oriented flankers facilitated target detection when placed on either the orthogonal axis (motion-collinearity) or on the collinear axis (orientation-collinearity). Facilitation occurred only if the flankers moved in the same direction as the target.

3.4.5 Conclusions

Attention to stimulus orientation facilitated the detection of another stimulus with the matching orientation. This result is consistent with a global feature-based mechanism of attention and demonstrates that contextual interactions are modulated by the behavioral goals of the observer. In natural images visual objects are rarely viewed in isolation and the critical role of context in visual processing is only beginning to be understood.

Conclusions and Future Studies

Our neuronal and behavioral findings support the hypothesis that feature-based attention globally enhances the processing of visual objects that have behaviorally relevant features. Our main conclusions are listed, in the following order, to address the motivating questions posed in the introduction:

(1) Feature-based attention modulated neuronal responses to stimuli outside the focus of attention in the human brain. Attending to a stimulus feature (motion direction or color) increased the response to a spatially distant ignored stimulus with the matching feature compared to the opposing feature. This response modulation is consistent with a *feature-similarity gain model* of attention which proposes that attention to a stimulus feature increases the gain of cortical neurons tuned to that feature throughout the visual field.

(2) Global feature-based attention modulated responses in multiple cortical visual areas where feature-tuned neurons are found (V1, V2, V3, V3A, V4, and MT+). These areas represent the earliest stages of cortical visual processing. Consistent with the spatial attention literature, the magnitude of the effect was overall larger at later stages of cortical visual processing.

(3) Global feature-based neuronal modulation was observed when the attended feature was motion direction or color, but not orientation. However, we cannot easily rule out global attention to orientation because there were differences in the experimental design that may have contributed to the null result with orientation.

(4) Psychophysical results suggested a behavioral consequence of global feature-based attention. Attention to stimulus feature at one location facilitated the concurrent processing of a spatially separate stimulus when it had the matching feature (same direction of motion or same color) compared to when it had the opposing feature (opposing direction of motion or opposing color).

(5) This behavioral consequence of feature-based attention (to both motion direction and color) was reduced in the absence of competing stimuli.

(6) A second psychophysical paradigm revealed that attending to stimulus orientation facilitated the detection of a stimulus with the matching orientation. This result also demonstrated that attention modulated a well-know contextual interaction.

Together, these findings show that the brain's response to a stimulus does not depend solely on the physical characteristics of the stimulus. Rather, the processing of a visual stimulus also depends on how the properties of that stimulus relate to other objects in the visual field and to the behavioral goals of the observer. In other words, we have studied the effects of both external physical context and internal mental context. This contextual dependence of neuronal responses was feature-specific, spatially global in its reach, and was correlated with clear behavioral consequences.

Here, we conclude with a description of suggested future studies that are designed to better explore the global extent of feature-based attention and to better quantify its neuronal effects. Future studies also aim to further explore the behavioral relevance of global feature based attention.

Ignored stimuli at unpredictable locations

In the present fMRI experiments, the ignored stimulus was always symmetrically positioned in the visual hemifield opposite to the attended stimulus. If the feature-based attention effects are truly global, we would expect to see similar results for ignored stimuli presented in (1) different and (2) unpredictable locations. This prediction could be tested with an event-related fMRI paradigm that allows us to vary the location of the ignored stimulus on a trial-by-trial basis. To start, the attended stimulus could be foveally located with task-irrelevant stimuli randomly presented at one of four equidistant locations, one in each visual quadrant. With this configuration, we would expect the four ignored stimuli to have similar cortical representations and to evoke similar baseline responses. We would expect these responses to be modulated by a feature-based attentional bias that was independent of spatial position.

Tuning of feature based attention

In the present motion fMRI experiment, the motion of the ignored stimulus was always in the direction same or opposite to the direction of the attended stimulus. We do not know how quickly the response enhancement would decline with an increasing difference in motion direction. In a new experiment, we could present the ignored stimulus with a range of discrete intermediate motion directions to measure the tuning of feature-based attention. This data could be compared to the known tuning properties of visual cortical neurons in the macaque. This experiment could be combined with the experiment just described – ignored stimuli could be presented at unpredictable locations with unpredictable directions of motion.

Contrast response functions

Reynolds and colleagues measured the effect of spatial attention on cortical contrast response functions (Reynolds, Pasternak et al. 2000). Responses of macaque V4 neurons were measured to oriented gratings (of varied contrasts) when attention was either focused on the grating or focused elsewhere in the display. The difference between the two conditions was a horizontal shift in the contrast response function indicating that attention effectively increased stimulus contrast. A separate study confirmed this result in macaque MT (Martinez-Trujillo and Treue 2002).

Would feature-based attention similarly cause an effective increase in the signal strength of a stimulus *outside* the focus of attention? In a new experiment, we could measure the contrast response function to an ignored motion stimulus under two conditions, when its direction of motion was either the same or opposite to the attended direction of motion. Measurement of human contrast response functions with fMRI is known to be feasible. Boynton and colleagues previously measured human contrast response functions in early cortical visual areas and showed that they increased monotonically (Boynton, Demb et al. 1999).

In this new experiment, the attended side of the display would consist of overlapping fields of upward and downward moving dots (as in the present fMRI experiments). The ignored side of the display would consist of a single moving grating, the contrast of which could be varied. An event related fMRI experiment would allow us to vary the contrast of the ignored grating on a trial-by-trial basis. We would first need to establish that attention to the moving dots stimulus modulated responses to this different type of ignored motion stimulus (an interesting question in and of itself).

Role of distractors

In the present fMRI experiments, observers selectively attended to a target field of dots in the presence of an overlapping distracting field of dots. We do not know if the effect depended on the presence of these competing stimuli. In our complementary behavioral study (Chapter 2), the effects of attention were reduced in the absence of competing stimuli. Furthermore, electrophysiological studies generally report stronger effects of spatial attention when multiple stimuli compete for selection within single neuronal receptive fields (Moran & Desimone, 1985; Motter, 1993; Luck et al., 1997; Treue & Martinez Trujillo, 1999; Treue & Maunsell, 1999). In light of these findings, we might expect that removing the distracting field of dots would reduce the response modulation in our fMRI experiments.

On the other hand, Treue and Martinez-Trujillo (1999) demonstrated a global effect of feature-based attention to motion without nearby competing stimuli. It would be most interesting to repeat our fMRI experiments without the overlapping field of dots to compare results. This experiment is currently underway with the moving dots stimulus.

Behavioral Relevance

The ubiquity and strength of the global feature-based attention effect (across multiple cortical visual areas) suggests that it may profoundly influence what we see. A feature-based increase in signal strength might improve the detection, discrimination, and/or memory of visual objects with behaviorally relevant features, even if those objects are not the subject of focused attention. A series of psychophysical experiments could be run to test these predictions.

Task-performance on peripheral stimuli could be assessed while observers were engaged in our feature based attention task (with overlapping stimuli) at the fovea. We could test whether performance was facilitated for stimuli matching the attended feature compared to the ignored feature in the following visual tasks:

(1) detection of low contrast stimuli at unpredictable locations (2-AFC detection task),
(2) detection of supra-threshold stimuli at unpredictable locations (reaction time analysis),
(3) detection of pop-out visual targets (reaction time analysis), and
(4) memory of task-irrelevant stimuli including words (semantic priming).

The strength of our general paradigm is that it allows us to compare performance on conditions in which the physical stimulus, eye position, task difficulty, and spatial distribution of attention are held constant. Only feature-based attention is changed across conditions.

References

1. Alais, D. and R. Blake (1999). Neural strength of visual attention gauged by motion adaptation. *Nat Neurosci, 2,* 1015-8.

2. Angelucci, A., J. B. Levitt, E. J. Walton, J. M. Hupe, J. Bullier and J. S. Lund (2002). Circuits for local and global signal integration in primary visual cortex. *J Neurosci, 22,* 8633-46.

3. Anllo-Vento, L. and S. A. Hillyard (1996). Selective attention to the color and direction of moving stimuli: electrophysiological correlates of hierarchical feature selection. *Percept Psychophys, 58,* 191-206.

4. Baylis, G. C. and J. Driver (1992). Visual parsing and response competition: the effect of grouping factors. *Percept Psychophys, 51,* 145-62.

5. Bichot, N. P., K. R. Cave and H. Pashler (1999). Visual selection mediated by location: feature-based selection of noncontiguous locations.PG - 403-23. *Percept Psychophys, 61*.

6. Blaser, E., Z. W. Pylyshyn and A. O. Holcombe (2000). Tracking an object through feature space. *Nature, 408,* 196-9.

7. Bosking, W. H., Y. Zhang, B. Schofield and D. Fitzpatrick (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J Neurosci, 17,* 2112-27.

8. Boynton, G. M., J. B. Demb, G. H. Glover and D. J. Heeger (1999). Neuronal basis of contrast discrimination. *Vision Res, 39,* 257-69.

9. Boynton, G. M., S. A. Engel, G. H. Glover and D. J. Heeger (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci, 16,* 4207-21.

10. Brainard, D. H. (1997). The Psychophysics Toolbox. *Spat Vis, 10,* 433-6.

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

12. Burock, M. A., R. L. Buckner, M. G. Woldorff, B. R. Rosen and A. M. Dale (1998). Randomized event-related experimental designs allow for

extremely rapid presentation rates using functional MRI. *Neuroreport, 9,* 3735-9.

13. Chatterjee, S. and E. M. Callaway (2002). S cone contributions to the magnocellular visual pathway in macaque monkey. *Neuron, 35,* 1135-46.

14. Chelazzi, L., J. Duncan, E. K. Miller and R. Desimone (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *J Neurophysiol, 80,* 2918-40.

15. Chelazzi, L., E. K. Miller, J. Duncan and R. Desimone (1993). A neural basis for visual search in inferior temporal cortex. *Nature, 363,* 345-7.

16. Chelazzi, L., E. K. Miller, J. Duncan and R. Desimone (2001). Responses of neurons in macaque area V4 during memory-guided visual search. *Cereb Cortex, 11,* 761-72.

17. Cherry, E.C. (1953). Some experiments on the recognition of speech, with one and with two ears. *J Acoustical Society of America*, 25, 975-79.

18. Chun, M. M. and R. Marois (2002). The dark side of visual attention. *Curr Opin Neurobiol, 12,* 184-9.

19. Corbetta, M., E. Akbudak, T. E. Conturo, A. Z. Snyder, J. M. Ollinger, H. A. Drury, M. R. Linenweber, S. E. Petersen, M. E. Raichle, D. C. Van Essen and G. L. Shulman (1998). A common network of functional areas for attention and eye movements. *Neuron, 21,* 761-73.

20. Corbetta, M., F. M. Miezin, S. Dobmeyer, G. L. Shulman and S. E. Petersen (1990). Attentional modulation of neural processing of shape, color, and velocity in humans. *Science, 248,* 1556-9.

21. Corbetta, M., F. M. Miezin, S. Dobmeyer, G. L. Shulman and S. E. Petersen (1991). Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. *J Neurosci, 11,* 2383-402.

22. Corbetta, M. and G. L. Shulman (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci, 3,* 201-15.

23. Dale, A.M., Bucker, R.L. (1997). Selective averaging of rapidly presented individual trials using fMRI. *Hum Brain Mapp* 5, 329-40.

24. Davis, E. T., P. Kramer and N. Graham (1983). Uncertainty about spatial frequency, spatial position, or contrast of visual patterns. *Percept Psychophys, 33,* 20-8.

25. Desimone, R. and J. Duncan (1995). Neural mechanisms of selective visual attention. *Annu Rev Neurosci, 18,* 193-222.

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

27. Dobkins, K. R. and T. D. Albright (1994). What happens if it changes color when it moves?: the nature of chromatic input to macaque visual area MT. *J Neurosci, 14,* 4854-70.

28. Driver, J. and Baylis, G. (1998). Attention and visual object segmentation. In *The Attentive Brain,* edited by R. Parasuraman, 299-325. Cambridge, MA: MIT Press.

29. Duffy, C. J. and R. H. Wurtz (1991). Sensitivity of MST neurons to optic flow stimuli. I. A continuum of response selectivity to large-field stimuli. *J Neurophysiol, 65,* 1329-45.

30. Duncan, J. (1984). Selective attention and the organization of visual information. *J Exp Psychol Gen, 113,* 501-17.

31. Duncan, J., G. Humphreys and R. Ward (1997). Competitive brain activity in visual attention. *Curr Opin Neurobiol, 7,* 255-61.

32. Duncan, R.O. and G.M. Boynton (2002). Cortical magnification factor in human primary visual cortex correlates with Vernier acuity thresholds, *Soc Neurosci*, 32nd annual meeting.

33. Egly, R., J. Driver and R. D. Rafal (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *J Exp Psychol Gen, 123,* 161-77.

34. Engel, S. A., G. H. Glover and B. A. Wandell (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex, 7,* 181-92.

35. Eriksen, C. W. and J. D. St James (1986). Visual attention within and around the field of focal attention: a zoom lens model. *Percept Psychophys, 40,* 225-40.

36. Folk, C. L., A. B. Leber and H. E. Egeth (2002). Made you blink! Contingent attentional capture produces a spatial blink. *Percept Psychophys, 64,* 741-53.

37. Folk, C. L., R. W. Remington and J. C. Johnston (1992). Involuntary covert orienting is contingent on attentional control settings. *J Exp Psychol Hum Percept Perform, 18,* 1030-44.

38. Freeman, E., D. Sagi and J. Driver (2001). Lateral interactions between targets and flankers in low-level vision depend on attention to the flankers. *Nat Neurosci, 4,* 1032-6.

39. Gandhi, S. P., D. J. Heeger and G. M. Boynton (1999). Spatial attention affects brain activity in human primary visual cortex. *Proc Natl Acad Sci U S A, 96,* 3314-9.

40. Gattass, R., A. P. Sousa and C. G. Gross (1988). Visuotopic organization and extent of V3 and V4 of the macaque. *J Neurosci, 8,* 1831-45.

41. Gilbert, C. D. and T. N. Wiesel (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J Neurosci, 9,* 2432-42.

42. Gray, C. M., P. Konig, A. K. Engel and W. Singer (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature, 338,* 334-7.

43. Green, D.M. (1961) Detection of auditory sinusoids of uncertain frequency. *Journal of the Acoustical Society of America*, 33, 897-303.

44. Haenny, P. E. and P. H. Schiller (1988). State dependent activity in monkey visual cortex. I. Single cell activity in V1 and V4 on visual tasks. *Exp Brain Res, 69,* 225-44.

45. He, Z. J. and K. Nakayama (1995). Visual attention to surfaces in three-dimensional space. *Proc Natl Acad Sci U S A, 92,* 11155-9.

46. Heeger, D. J., G. M. Boynton, J. B. Demb, E. Seidemann and W. T. Newsome (1999). Motion opponency in visual cortex. *J Neurosci, 19,* 7162-74.

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

48. Heeger, D. J. and D. Ress (2002). What does fMRI tell us about neuronal activity? *Nat Rev Neurosci, 3,* 142-51.

49. Heinze, H. J., G. R. Mangun, W. Burchert, H. Hinrichs, M. Scholz, T. F. Munte, A. Gos, M. Scherg, S. Johannes, H. Hundeshagen and et al. (1994). Combined spatial and temporal imaging of brain activity during visual selective attention in humans. *Nature, 372,* 543-6.

50. Hillyard, S. A. and T. F. Munte (1984). Selective attention to color and location: an analysis with event-related brain potentials. *Percept Psychophys, 36,* 185-98.

51. Hopfinger, J. B., M. H. Buonocore and G. R. Mangun (2000). The neural mechanisms of top-down attentional control. *Nat Neurosci, 3,* 284-91.

52. Huk, A. C., R. F. Dougherty and D. J. Heeger (2002). Retinotopy and functional subdivision of human areas MT and MST. *J Neurosci, 22,* 7195-205.

53. Huk, A. C. and D. J. Heeger (2000). Task-related modulation of visual cortex. *J Neurophysiol, 83,* 3525-36.

54. Ito, M. and C. D. Gilbert (1999). Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron, 22,* 593-604.

55. Kanwisher, N., J. McDermott and M. M. Chun (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci, 17,* 4302-11.

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

57. Kapadia, M. K., M. Ito, C. D. Gilbert and G. Westheimer (1995). Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neuron, 15,* 843-56.

58. Kastner, S., M. A. Pinsk, P. De Weerd, R. Desimone and L. G. Ungerleider (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron, 22,* 751-61.

59. Kreiman, G., C. Koch and I. Fried (2000). Category-specific visual responses of single neurons in the human medial temporal lobe. *Nat Neurosci, 3,* 946-53.

60. Kwong, K. K., J. W. Belliveau, D. A. Chesler, I. E. Goldberg, R. M. Weisskoff, B. P. Poncelet, D. N. Kennedy, B. E. Hoppel, M. S. Cohen, R. Turner and et al. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc Natl Acad Sci U S A, 89,* 5675-9.

61. Kubodera, T. and T. Sato, (2002). Non-classical receptive field structure for motion mechanisms revealed by lateral masking, *Vision Sciences Society*, 2nd annual meeting, 396.

62. Lee, D. K., C. Koch and J. Braun (1999). Attentional capacity is undifferentiated: concurrent discrimination of form, color, and motion. *Percept Psychophys, 61,* 1241-55.

63. Livingstone, M. S. and D. H. Hubel (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *J Neurosci, 7,* 3416-68.

64. Logothetis, N. K., J. Pauls, M. Augath, T. Trinath and A. Oeltermann (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature, 412,* 150-7.

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

66. Malach, R., T. D. Schirman, M. Harel, R. B. Tootell and D. Malonek (1997). Organization of intrinsic connections in owl monkey area MT. *Cereb Cortex, 7,* 386-93.

67. Maljkovic, V. and K. Nakayama (1994). Priming of pop-out: I. Role of features. *Mem Cognit, 22,* 657-72.

68. Martinez, A., L. Anllo-Vento, M. I. Sereno, L. R. Frank, R. B. Buxton, D. J. Dubowitz, E. C. Wong, H. Hinrichs, H. J. Heinze and S. A. Hillyard

(1999). Involvement of striate and extrastriate visual cortical areas in spatial attention. *Nat Neurosci, 2,* 364-9.

69. Martinez-Trujillo, J. and S. Treue (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron, 35,* 365-70.

70. McAdams, C. J. and J. H. Maunsell (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. *J Neurophysiol, 83,* 1751-5.

71. McAdams, C. J. and J. H. R. Maunsell (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci, 19,* 431-41.

72. Moran, J. and R. Desimone (1985). Selective attention gates visual processing in the extrastriate cortex. *Science, 229,* 782-4.

73. Morrone, M. C., V. Denti and D. Spinelli (2002). Color and luminance contrasts attract independent attention. *Curr Biol, 12,* 1134-7.

74. Most, S. B., D. J. Simons, B. J. Scholl, R. Jimenez, E. Clifford and C. F. Chabris (2001). How not to be seen: the contribution of similarity and selective ignoring to sustained inattentional blindness. *Psychol Sci, 12,* 9-17.

75. Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J Neurophysiol, 70,* 909-19.

76. Motter, B. C. and E. J. Belky (1998). The guidance of eye movements during active visual search. *Vision Res, 38,* 1805-15.

77. Nakayama K. and Joseph J.S. (1998) Attention, pattern recognition, and pop-out in visual search. In *The Attentive Brain,* edited by R. Parasuraman, pp. 279-298. Cambridge, MA: MIT Press.

78. Noesselt, T., S. A. Hillyard, M. G. Woldorff, A. Schoenfeld, T. Hagner, L. Jancke, C. Tempelmann, H. Hinrichs and H. J. Heinze (2002). Delayed striate cortical activation during spatial attention. *Neuron, 35,* 575-87.

79. O'Connor, D. H., M. M. Fukui, M. A. Pinsk and S. Kastner (2002). Attention modulates responses in the human lateral geniculate nucleus. *Nat Neurosci, 5,* 1203-9.

80. O'Craven, K. M., B. R. Rosen, K. K. Kwong, A. Treisman and R. L. Savoy (1997). Voluntary attention modulates fMRI activity in human MT-MST. *Neuron, 18,* 591-8.

81. Ogawa, S., T. M. Lee, A. R. Kay and D. W. Tank (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci U S A, 87,* 9868-72.

82. Ogawa, S., D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle and K. Ugurbil (1992). Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci U S A, 89,* 5951-5.

83. Pashler H. (1998) *The Psychology of Attention.* Cambridge, MA: MIT Press.

84. Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis, 10,* 437-42.

85. Polat, U., K. Mizobe, M. W. Pettet, T. Kasamatsu and A. M. Norcia (1998). Collinear stimuli regulate visual responses depending on cell's contrast threshold. *Nature, 391,* 580-4.

86. Polat, U. and D. Sagi (1993). Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. *Vision Res, 33,* 993-9.

87. Polat, U. and D. Sagi (1994). The architecture of perceptual spatial interactions. *Vision Res, 34,* 73-8.

88. Posner, M. I., C. R. Snyder and B. J. Davidson (1980). Attention and the detection of signals. *J Exp Psychol, 109,* 160-74.

89. Rees, G., K. Friston and C. Koch (2000). A direct quantitative relationship between the functional properties of human and macaque V5. *Nat Neurosci, 3,* 716-23.

90. Rensink, R. A. (2002). Change detection. *Annu Rev Psychol, 53,* 245-77.

91. Reynolds, J. H., L. Chelazzi and R. Desimone (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *J Neurosci, 19,* 1736-53.

92. Reynolds, J. H., T. Pasternak and R. Desimone (2000). Attention increases sensitivity of V4 neurons. *Neuron, 26,* 703-14.

93. Roelfsema, P. R., V. A. Lamme and H. Spekreijse (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature, 395,* 376-81.

94. Rossi, A. F. and M. A. Paradiso (1995). Feature-specific effects of selective visual attention. *Vision Res, 35,* 621-34.

95. Saenz, M., G. T. Buracas and G. M. Boynton (2002). Global effects of feature-based attention in human visual cortex. *Nat Neurosci, 17,* 17.

96. Saenz, M., G. T. Buracas and G. M. Boynton (2002). “Global feature-based attention for motion and color”, *Vision Research*. in press.

97. Saito, H., K. Tanaka, H. Isono, M. Yasuda and A. Mikami (1989). Directionally selective response of cells in the middle temporal area (MT) of the macaque monkey to the movement of equiluminous opponent color stimuli. *Exp Brain Res, 75,* 1-14.

98. Schmidt, K. E., R. Goebel, S. Lowel and W. Singer (1997). The perceptual grouping criterion of colinearity is reflected by anisotropies of connections in the primary visual cortex. *Eur J Neurosci, 9,* 1083-9.

99. Sclar, G. and R. D. Freeman (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp Brain Res, 46,* 457-61.

100. Seidemann, E., A. B. Poirson, B. A. Wandell and W. T. Newsome (1999). Color signals in area MT of the macaque monkey. *Neuron, 24,* 911-7.
101. Sekuler, R. and K. Ball (1977). Mental set alters visibility of moving targets. *Science, 198,* 60-2.

102. Sereno, M. I., C. T. McDonald and J. M. Allman (1994). Analysis of retinotopic maps in extrastriate cortex. *Cereb Cortex, 4,* 601-20.

103. Shulman, G. L., J. M. Ollinger, E. Akbudak, T. E. Conturo, A. Z. Snyder, S. E. Petersen and M. Corbetta (1999). Areas involved in encoding and applying directional expectations to moving objects. *J Neurosci, 19,* 9480-96.

104. Shulman, G. L. and J. Wilson (1987). Spatial frequency and selective attention to local and global information. *Perception, 16,* 89-101.

105. Simons, D. J. and C. F. Chabris (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception, 28,* 1059-74.

106. Sincich, L. C. and G. G. Blasdel (2001). Oriented axon projections in primary visual cortex of the monkey. *J Neurosci, 21,* 4416-26.

107. Smith, A. T., K. D. Singh, A. L. Williams and M. W. Greenlee (2001). Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cereb Cortex, 11,* 1182-90.

108. Solomon, J. A. and M. J. Morgan (2000). Facilitation from collinear flanks is cancelled by non-collinear flanks. *Vision Res, 40,* 279-86.

109. Somers, D. C., A. M. Dale, A. E. Seiffert and R. B. Tootell (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proc Natl Acad Sci U S A, 96,* 1663-8.

110. Sperling, G. and M. J. Melchner (1978). The attention operating characteristic: examples from visual search. *Science, 202,* 315-8.

111. Stettler, D. D., A. Das, J. Bennett and C. D. Gilbert (2002). Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron, 36,* 739-50.

112. Teo, P. C., G. Sapiro and B. A. Wandell (1997). Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Trans Med Imaging, 16,* 852-63.

113. Tolhurst, D. J. (1973). Separate channels for the analysis of the shape and the movement of moving visual stimulus. *J Physiol, 231,* 385-402.

114. Tootell, R. B., J. D. Mendola, N. K. Hadjikhani, A. K. Liu and A. M. Dale (1998). The representation of the ipsilateral visual field in human cerebral cortex. *Proc Natl Acad Sci U S A, 95,* 818-24.

115. Tovee, M. J., E. T. Rolls and P. Azzopardi (1994). Translation invariance in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque. *J Neurophysiol, 72,* 1049-60.

116. Treisman, A. (1964) Verbal cues, language and meaning in selective attention. *American Journal of Psychology*, 77, 206-19.

117. Treue, S. and J. C. Martinez Trujillo (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature, 399,* 575-9.

118. Treue, S. and J. H. Maunsell (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature, 382,* 539-41.

119. Treue, S. and J. H. Maunsell (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *J Neurosci, 19,* 7591-602.

120. Ts'o, D. Y. and C. D. Gilbert (1988). The organization of chromatic and spatial interactions in the primate striate cortex. *J Neurosci, 8,* 1712-27.

121. Valdes-Sosa, M., M. A. Bobes, V. Rodriguez and T. Pinilla (1998). Switching attention without shifting the spotlight object-based attentional modulation of brain potentials. *J Cogn Neurosci, 10,* 137-51.

122. Vandenberghe, R., J. Duncan, P. Dupont, R. Ward, J. B. Poline, G. Bormans, J. Michiels, L. Mortelmans and G. A. Orban (1997). Attention to one or two features in left or right visual field: a positron emission tomography study. *J Neurosci, 17,* 3739-50.

123. Watson, J. D., R. Myers, R. S. Frackowiak, J. V. Hajnal, R. P. Woods, J. C. Mazziotta, S. Shipp and S. Zeki (1993). Area V5 of the human brain: evidence from a combined study using positron emission tomography and magnetic resonance imaging. *Cereb Cortex, 3,* 79-94.

124. Weliky, M., K. Kandler, D. Fitzpatrick and L. C. Katz (1995). Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron, 15,* 541-52.

125. Yantis, S., J. Schwarzbach, J. T. Serences, R. L. Carlson, M. A. Steinmetz, J. J. Pekar and S. M. Courtney (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nat Neurosci, 5,* 995-1002.

126. Yu, C., Klein, S.A. and Levi, D.M. (2002). Facilitation of contrast detection by cross-oriented surround stimuli and its psychophysical mechanisms. *J Vision*, 2(3), 243-55.

127. Zeki, S. M. (1978). Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex. *J Physiol, 277,* 273-90.